

VMware SD-WAN Administration Guide

VMware SD-WAN 4.5

You can find the most up-to-date technical documentation on the VMware website at:

<https://docs.vmware.com/>

VMware, Inc.
3401 Hillview Ave.
Palo Alto, CA 94304
www.vmware.com

Copyright © 2021 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

Contents

1	About VMware SD-WAN Administration Guide	12
2	What's New	13
3	Overview	15
	VMware SD-WAN Routing Overview	16
	Solution Components	26
	SD-WAN Edge Performance and Scale Data	27
	Capabilities	34
	Tunnel Overhead and MTU	37
	Network Topologies	41
	Branch Site Topologies	44
	Roles and Privilege Levels	46
	User Role Matrix	48
	Key Concepts	51
	Supported Browsers	55
	Supported Modems	55
4	User Agreement	57
5	Log in to VMware SD-WAN Orchestrator Using SSO for Enterprise User	58
6	Monitor Enterprises	60
	Monitor Navigation Panel	60
	Network Overview	60
	Monitor Edges	64
	Overview Tab	66
	QoE Tab	67
	Transport Tab	70
	Applications Tab	72
	Sources Tab	74
	Destinations Tab	75
	Business Priority Tab	76
	System Tab	77
	VMware SD-WAN Orchestrator Data Retention	78
	Monitor Network Services	82
	Monitor Routing	85
	PIM Neighbors View	85

Monitor Alerts	86
Monitor Events	86
Auto Rollback to the Last Known Good Configuration	87
Monitor Reports	88

7 Monitor Enterprise using New Orchestrator UI 90

Monitor Network Overview	91
Monitor Edges	93
Monitor overview of an Edge	95
Monitor QoE	96
Monitor Links of an Edge	97
Monitor Path Visibility	99
Monitor Edge Applications	101
Monitor Edge Sources	102
Monitor Edge Destinations	104
Monitor Business Priorities of an Edge	106
Monitor System Information of an Edge	107
Monitor Network Services	108
Monitor Non SD-WAN Destinations through Gateway	109
Monitor Cloud Security Service Sites	110
Monitor Edge Clusters	110
Monitor Edge VNFs	111
Monitor Routing Details	112
Monitor Multicast Groups	112
Monitor PIM Neighbors	113
Monitor BGP Edge Neighbor State	114
Monitor BFD	115
Monitor BGP Gateway Neighbor State	115
Monitor Alerts	116
Monitor Events	118
Enterprise Reports	119
Create a New Enterprise Report	120
Create Customized Report	122
Monitor Enterprise Reports	128
View Analytics Data	131

8 Configure Segments 134

9 Configure Segments with new Orchestrator UI 137

10 Configure Network Services 140

About Edge Clustering	142
How Edge Clustering Works	143
Configure Edge Clustering	149
Troubleshooting Edge Clustering	151
Configure a Non SD-WAN Destination	153
VPN Workflow	154
Configure Non SD-WAN Destinations via Gateway	158
Configure a Non SD-WAN Destinations via Edge	206
Azure Virtual WAN IPsec Tunnel Automation	219
VMware SD-WAN in Azure Virtual WAN Hub Deployment	252
About VMware SD-WAN in Azure Virtual WAN Hub Deployment	252
Cloud Security Services	266
Configure a Cloud Security Service	266
Configure Cloud Security Services for Profiles	273
Configure Cloud Security Services for Edges	275
Configure Business Policies with Cloud Security Services	285
Monitor Cloud Security Services	287
Monitor Cloud Security Services Events	290
Configure DNS Services	292
Configure Netflow Settings	293
IPFIX Templates	295
Private Network Names	319
Configure Private Networks	319
Delete a Private Network Name	319
Configure Authentication Services	319
Configure Cloud Subscriptions	320
11	Configure Profiles 322
12	Manage Profiles with New Orchestrator UI 326
Create Profile with new Orchestrator UI	327
Configure Profiles with new Orchestrator UI	328
13	Configure a Profile Device 330
Configure a Device	330
Assign Segments in Profile	331
Configure Authentication Settings	333
Configure DNS Settings	333
Configure Netflow Settings for Profiles	335
Configure Syslog Settings for Profiles	337
Configure Cloud VPN for Profiles	343

Configure Multicast Settings	361
Configure VLAN for Profiles	363
Configure the Management IP Address	366
IPv6 Settings	366
Configure Device Settings	368
Configure Wi-Fi Radio Settings	392
Activate Multi-Source QOS	392
Configure Layer 2 Settings for Profiles	393
Configure SNMP Settings for Profiles	394
Configure NTP Settings for Profiles	396
Configure Visibility Mode	398
Assign Partner Gateways	399
Assign Controllers	402

14 Configure Business Policy 404

Configure Business Policy for Profiles	404
Configure Business Policy for Edges	405
Create Business Policy Rules	406
Configure Network Service for Business Policy Rule	413
Configure Link Steering Modes	415
Configure Policy-based NAT	421
Overlay QoS CoS Mapping	422
Tunnel Shaper for Service Providers with Partner Gateway	423

15 Configure Firewall 425

Configure Firewall for Profiles	427
Configure Firewall for Edges	429
Configure Firewall Rules	436
Configure Stateful Firewall Settings	440
Configure Network and Flood Protection Settings	441
Configure Edge Access	445
Troubleshooting Firewall	447

16 Provision an Edge 449

Provision a New Edge	449
Provision a New Edge with Analytics	452
Enable Analytics for an Existing Edge	454
Configure an Analytics Interface on an Edge	456
Configure Analytics Endpoint Settings	457
Activate SD-WAN Edges	457
Activate SD-WAN Edges Using Zero Touch Provisioning	458

Activate SD-WAN Edges Using Email	461
Request RMA Reactivation	468
Manage Edges	470
Assign Software Image	473
Reset Edges to Factory Settings	474
Manage Edges with New Orchestrator UI	474
Configure Edges with new Orchestrator UI	477
17 View or Modify Edge Information	480
18 Edge Device Configurations	487
Configure DSL Settings	489
Configure ADSL and VDSL Settings	489
Configure GPON Settings	493
Configure Netflow Settings for Edges	498
LAN-side NAT Rules at Edge Level	499
Configure Syslog Settings for Edges	507
Configure Static Route Settings	509
Configure ICMP Probes/Responders	511
Configure VRRP Settings	512
Monitor VRRP Events	515
Configure Cloud VPN and Tunnel Parameters at the Edge level	516
Configure VLAN for Edges	518
Loopback Interfaces Configuration	524
Loopback Interfaces—Benefits	525
Loopback Interfaces—Limitations	525
Configure a Loopback Interface for an Edge	526
Loopback Interfaces—Field References	528
Configure Orchestrator Management Traffic for Edges	529
Configure Device Settings	529
Configure DHCP Server on Routed Interfaces	529
Enable RADIUS on a Routed Interface	534
Configure Edge LAN Overrides	535
Configure Edge WAN Overrides	536
Configure Edge WAN Overlay Settings	536
SD-WAN Service Reachability via MPLS	549
Configure Class of Service	553
Configure Hot Standby Link	555
Configure Wi-Fi Radio Overrides	559
Security VNFs	560
Configure VNF Management Service	563

	Configure Security VNF without HA	568
	Configure Security VNF with High Availability	572
	Define Mapping Segments with Service VLANs	575
	Configure VLAN with VNF Insertion	576
	Monitor VNF for an Edge	578
	Monitor VNF Events	579
	Configure VNF Alerts	580
	Configure Layer 2 Settings for Edges	582
	Configure SNMP Settings for Edges	583
	Configure NTP Settings for Edges	585
	Configure Edge Activation	586
19	Edge Software Image Management	589
	Edge Software Image Management Overview	589
	Enable Edge Software Image Management	589
	Edge Image Assignment and Access	590
	Upgrade SD-WAN Edges	591
20	SD-WAN Gateway Migration	593
	SD-WAN Gateway Migration - Limitations	594
	Migrate Quiesced Gateways	595
	What to do When Switch Gateway Action Fails	596
21	Object Groups	598
	Configure Address Groups	598
	Configure Port Groups	600
	Configure Business Policies with Object Groups	600
	Configure Firewall Rules with Object Groups	602
22	Site Configurations	605
	Data Center Configurations	606
	Configure Branch and Hub	606
23	Configure Dynamic Routing with OSPF or BGP	615
	Enable OSPF	615
	Route Filters	618
	Configure BGP	619
	Configure BGP from Edge to Underlay Neighbors	619
	Configure BGP Over IPsec from Edge to Non SD-WAN Neighbors	631
	Configure BGP over IPsec from Gateways	642
	Monitor BGP Sessions	646

Monitor BGP Events	647
Troubleshooting BGP Settings	647
OSPF/BGP Redistribution	649
BFD Settings	650
Configure BFD	650
Configure BFD for BGP	652
Configure BFD for OSPF	653
Configure BFD for Gateways	654
Monitor BFD Sessions	656
Monitor BFD Events	658
Troubleshooting BFD	659
Overlay Flow Control	659
Configure Global Routing Preferences	661
Configure Subnets	662
Overlay Flow Control	664

24 Configure Alerts 668

25 Testing and Troubleshooting 673

Remote Diagnostics	673
Run Remote Diagnostics with new Orchestrator UI	675
Performing Remote Diagnostics Tests	676
Remote Actions	710
Perform Remote Actions with new Orchestrator UI	711
Diagnostic Bundles	713
Request Packet Capture Bundle	713
Request Diagnostic Bundle	714
Download Diagnostic Bundle	715
Delete Diagnostic Bundle	716
Diagnostic Bundles with new Orchestrator UI	716
Request Diagnostic Bundle	717
Request Packet Capture Bundle	718
Download Diagnostic Bundle	719
Delete Diagnostic Bundle	720

26 Enterprise Administration 721

System Settings	721
Configure Enterprise Information	721
Configure Enterprise Authentication	725
Manage Admin Users	750
Create New Admin User	751

Configure Admin Users	752
Roles	755
Functional Roles	755
Composite Roles	756
Role Customization	759
Edge Licensing	803
Example of Edge Licensing	805

27 Configure High Availability on SD-WAN Edge 808

How SD-WAN Edge High Availability (HA) Works	808
Failure Scenarios	809
High Availability Deployment Models	809
Standard HA	810
Enhanced HA	814
Mixed-Mode HA	821
Split-Brain Condition	822
Split-Brain Detection and Prevention	823
Support for BGP Over HA Link	824
Selection Criteria to Determine Active and Standby Status	825
VLAN-tagged Traffic Over HA Link	825
Configure High Availability (HA)	826
Deploying High Availability on VMware ESXi	826
Prerequisites	833
Activate High Availability	834
Wait for SD-WAN Edge to Assume Active	834
Connect the Standby SD-WAN Edge to the Active Edge	835
Connect LAN and WAN Interfaces on Standby SD-WAN Edge	835
Deactivate High Availability (HA)	836
HA Event Details	837

28 VMware Virtual Edge Deployment 838

Deployment Prerequisites for VMware Virtual Edge	838
Special Considerations for VMware Virtual Edge deployment	840
Cloud-init Creation	841
Install VMware Virtual Edge	843
Enable SR-IOV on KVM	843
Install Virtual Edge on KVM	845
Enable SR-IOV on VMware	849
Install Virtual Edge on VMware ESXi	851

29 Appendix 857

Enterprise-Level Orchestrator Alerts and Events	857
Supported VMware SD-WAN Edge Events for Syslogs	893

About VMware SD-WAN Administration Guide

1

The VMware SD-WAN™ (*formerly known as* VMware SD-WAN™ by VeloCloud®) Administration Guide provides information about VMware SD-WAN Orchestrator and the core VMware configuration settings, including how to configure and manage Network, Network Services, Edges, Profiles, and Customers who use the SD-WAN Orchestrator.

Intended Audience

This guide is intended for network administrators, network analysts, and IT administrators responsible for deploying, monitoring and managing Enterprise branch network.

Beginning with Release 4.4.0, VMware SD-WAN is offered as part of VMware SASE. To access SASE documentation for Cloud Web Security and Secure Access, along with Release Notes for version 4.4.0 and later, see [VMware SASE](#).

Here's a quick walkthrough 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

2

What's New in Version 4.5.0

Feature	Description
Enhanced support for IPv6 Address	<p>Allows to configure IPv6 addresses for Static Route Settings, BGP, BFD, and Loopback Interfaces. See:</p> <ul style="list-style-type: none">■ IPv6 Settings■ Configure Static Route Settings■ Configure BGP from Edge to Underlay Neighbors■ Configure BFD■ Loopback Interfaces Configuration
Configure Settings using New Orchestrator UI	<p>Allows the Enterprise user to configure the following settings using the new redesigned portal.</p> <ul style="list-style-type: none">■ Configure OSPF and BFD for Profiles. See Configure Profiles with new Orchestrator UI.■ Configure BFD, ICMP Probes, ICMP Responders, and High Availability for Edges. See Configure Edges with new Orchestrator UI.■ Generate and download Diagnostic bundles for Edges. See Diagnostic Bundles with new Orchestrator UI.
Support for Edges Without Wifi Modules	<p>VMware supports Edge models 510, 610, 620, 640, and 680 without WiFi modules for the following releases: 3.4.6, 4.2.2, 4.3.0, 4.3.1, and 4.5.0. For more information see:</p> <ul style="list-style-type: none">■ Configure Device Settings■ Chapter 17 View or Modify Edge Information
Zscaler Integration Enhancements	<p>At the Edge level, VMware SD-WAN and Zscaler integration supports:</p> <ul style="list-style-type: none">■ Zscaler GRE Tunnel Automation■ Zscaler Location/Sub-Location Configuration <p>See Configure Cloud Security Services for Edges and Configure a Cloud Security Service.</p>

Feature	Description
SD-WAN Gateway Migration	VMware SD-WAN 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. See Chapter 20 SD-WAN Gateway Migration .
Support to activate or deactivate USB port access on Edges	You can now choose to activate or deactivate USB port access on Edges. See Configure Edge Access .

Previous VMware SD-WAN Versions

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


Overview

3

VMware 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 VMware SD-WAN solution components. The components are described in more detail in the following sections.



To become familiar with the basic configuration and Edge activation, see [Activate SD-WAN Edges](#).

Read the following topics next:

- [VMware SD-WAN Routing Overview](#)
- [Solution Components](#)
- [SD-WAN Edge Performance and Scale Data](#)
- [Capabilities](#)
- [Tunnel Overhead and MTU](#)
- [Network Topologies](#)
- [Branch Site Topologies](#)
- [Roles and Privilege Levels](#)
- [User Role Matrix](#)
- [Key Concepts](#)
- [Supported Browsers](#)
- [Supported Modems](#)


VMware SD-WAN Routing Overview

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

Overview

VMware 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 in a manner 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 depicts 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)).



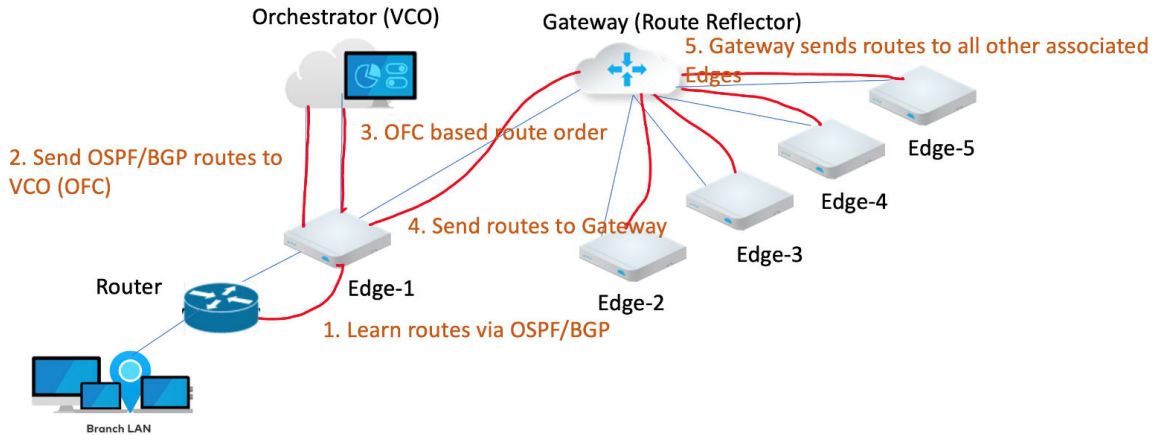
SD-WAN Components

VMware SD-WAN routing uses three components: Edge, Gateway, and Orchestrator.

- The **SD-WAN Edge** is an Enterprise-class device or virtualized cloud instance that provides secure and optimized connectivity to private, public and hybrid applications, and virtualized services. In SD-WAN routing the Edge is a **Border Gateway**. An Edge can function as a regular Edge (with no Hub configuration), as a Hub by itself or as part of a cluster, or as a Spoke (when Hubs are configured).
- The **SD-WAN 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 **SASE Orchestrator** is a multi-tenant cloud-based configuration and monitoring portal. In SD-WAN routing the Orchestrator manages routes for all enterprises and can override default routing behavior.



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 learnt 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.

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

Table 3-1. Order of Route Types

1. Longest Prefix Match
2. Connected Local
3. Static LAN/WAN Local
4. Connected Remote
5. Static LAN/WAN Remote
6. Static Non SD-WAN Destination

Table 3-1. Order of Route Types (continued)

7. Static Partner Gateway
8. Overlay Flow Control (OFC) Driven Route Order

Note Between local and remote routes of the same type, VMware SD-WAN will prefer the local over the remote. For example, a local connected route is preferred over a remote connected route. Similarly, for a local static route versus a remote static route, the local static route is preferred.

Connected and Static Routes

This section includes essential information regarding connected and static routes. A connected route is a route to a network that is directly attached to the interface. Information about static routes can be found in [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 which the connected route is configured.

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 **Advertise** flag must be set on the Edge interface for which the static route is configured.
 - The static route configuration must have **Preferred** and **Advertised** checked.
- Static Routes can forward traffic to the WAN Underlay for Global Segments, and to the LAN or WAN Underlay for Non-Global Segments.
- 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 blackholing traffic.

- 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.

The Preferred option is not available for an IPv6 address type.

When the **Advertise** check box is selected, the static route is advertised over VPN routes and other SD-WAN Edges in the network will have access to the resource.

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 Advertise option is not available for an IPv6 address type.

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 and require 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". When looking in an Edge's logs, a user would see these Self Routes with the route flag "s".

A Self Route is distinguished 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 a VMware 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 VMware 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 expected to be activated for all customers. This feature will automatically be activated for new customers in an upcoming release. For more information about DCC including best practices, see [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

Table 1-2 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 3-2. Dynamic Route Types

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

Table 3-2. Dynamic Route Types (continued)

Edge	Partner Gateway / Hosted Gateway
OSPF OE2	
Uplink BGP	

Table 3-3. Route Type Meanings

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

Note Non SD-WAN Destination (NSD) support with OFC is available from Release 4.3.0 and forward. For more information on NSDs, see [#unique_20](#).

Each route type has a preference value, 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 1-4 lists the default preference value for each route type.

Table 3-4. Preference Values

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

Table 3-4. Preference Values (continued)

Device	Route Type	Default Preference
Hub	OSPF OE2	1001007
Hub	BGP Uplink	1001008

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 covers **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.


The screenshot shows the VMware SD-WAN configuration interface. The top navigation bar includes 'Monitor', 'Configure' (selected), 'Diagnostics', and 'Settings'. The left sidebar lists 'Edge Configuration' items: Edges, Profiles, Object Groups, Segments, Overlay Flow Control (selected), Network Services, and Alerts & Notifications. The main panel is titled 'Overlay Flow Control' and has tabs for 'IPv4' (selected) and 'IPv6'. Under 'VRF Global Routing Preferences', there is a section for 'Preferred VPN Exits' with an 'EDIT' link. Below this is a table with two columns: 'Order' and 'Header'. The table lists five categories: 1. NSD, 2. Edge, 3. Partner Gateway, 4. Router, and 5. Hub. At the bottom, there is a link to 'Global Advertise Flags'.

Order	Header
1.	NSD
2.	Edge
3.	Partner Gateway
4.	Router
5.	Hub

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 / 2 or BGP Uplink type route.
- **Hub:** Any external site that is learned on an Edge 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:** Router represents any route prefix learned by an Edge with BGP or OSPF and determines the preference that is assigned to a dynamic route. Typically, all exit points above **Router** in the VPN Exit are assigned a low preference value and are thus more preferred, while all exit points below **Router** are assigned a higher preference value and are thus 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.



The screenshot shows the VMware SD-WAN configuration interface. The left sidebar is under the 'Configure' tab, with 'Overlay Flow Control' selected under 'Edge Configuration'. The main panel is titled 'Overlay Flow Control' and shows the 'IPv4' tab. Under 'VRF Global Routing Preferences', there is a section for 'Preferred VPN Exits'. Below this, there is a 'Default Priority' section with an 'EDIT' link. A table lists the preferred VPN exits in order:

Order	Header
1.	NSD
2.	Edge
3.	Partner Gateway
4.	Router
5.	Hub

The 'Router' entry in the table is highlighted with a green box. Below the table, there is a link to 'Global Advertise Flags'.

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. 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 click the **Pin Learned Route Preference** option.
 - b Modifying the order of the **Preferred VPN Exits** by clicking **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 more information on pinning a route, consult [Configure Subnets](#).

Tie-Breaking Scenarios for Dynamic 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 VMware 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.

Solution Components

This section describes VMware solution components.

VMware SD-WAN Edge

A thin “Edge” that is zero IT touch provisioned from the cloud for secured, optimized connectivity to your apps and virtualized services. The SD-WAN 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. SD-WAN 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.

VMware SD-WAN Orchestrator

The VMware SD-WAN 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.

VMware SD-WAN Gateways

VMware 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 VMware 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.

For more information about the VMware SD-WAN Gateways functionality and resiliency, see <https://kb.vmware.com/s/article/71374>.

SD-WAN Edge Performance and Scale Data

This section covers the performance and scale architecture of the VMware SD-WAN 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.

VMware SD-WAN Edge

VMware SD-WAN 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. VMware SD-WAN 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 3-1. FIGURE 1: Throughput performance test topology for devices 1 Gbps or lower

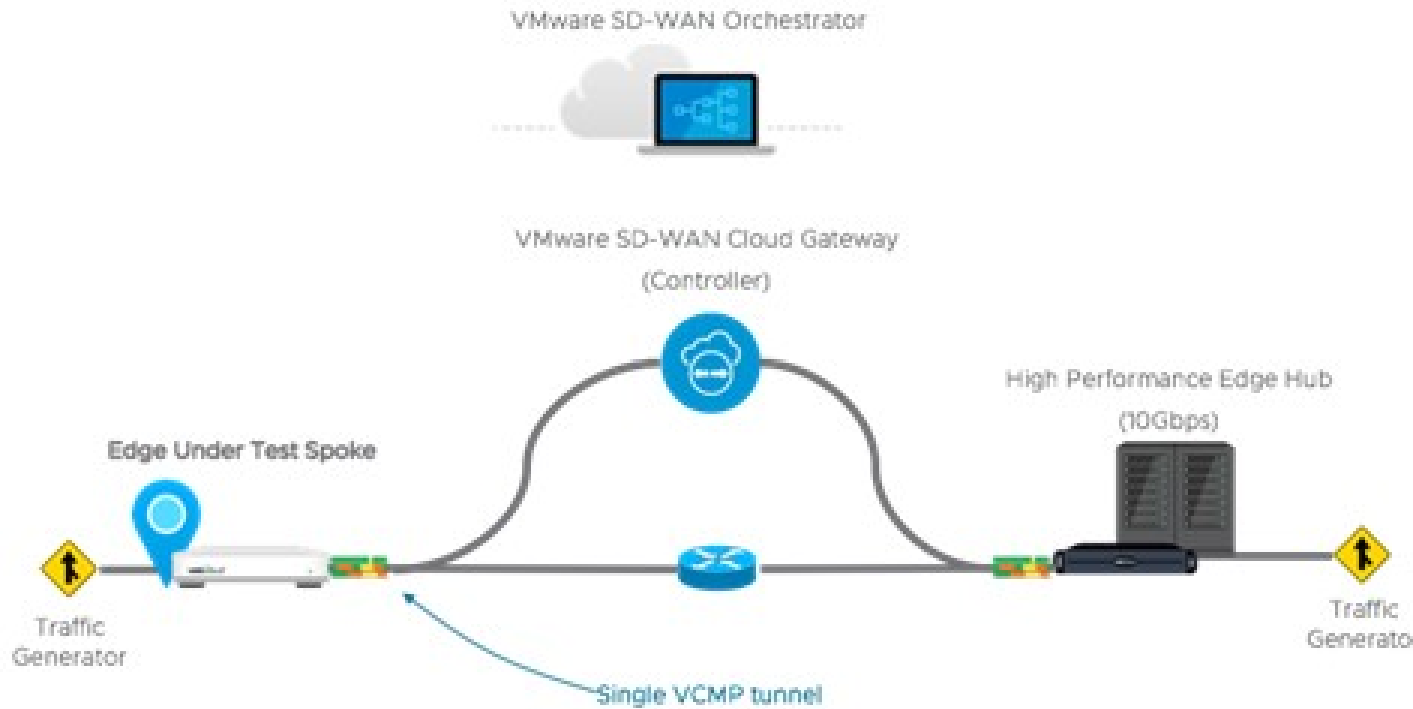
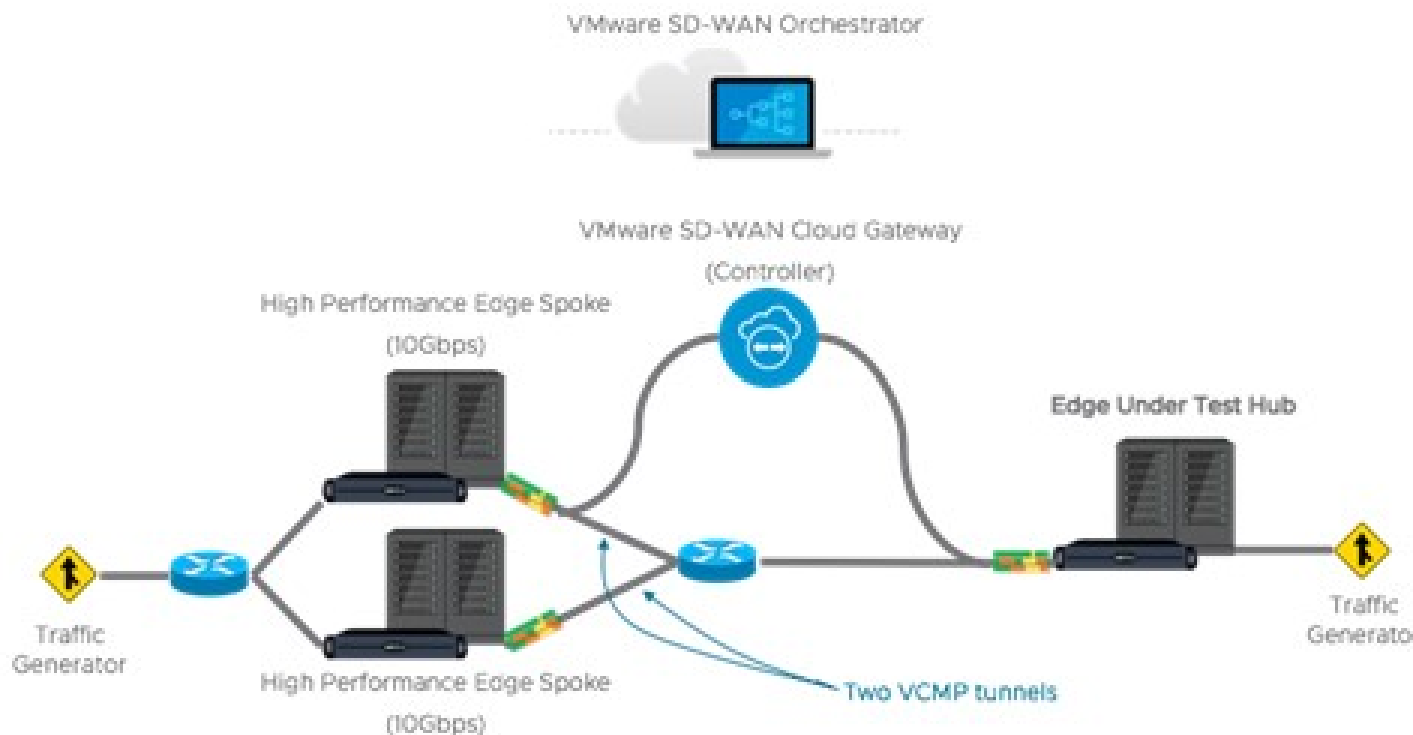


Figure 3-2. FIGURE 2: Throughput performance test topology for devices above 1 Gbps



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

VMware SD-WAN Edge Performance and Scale Results

Performance metrics are based on the Test Methodology detailed above.

Switched Port Performance: VMware 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, VMware 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 3-5. Physical Edge Appliances

VMware SD-WAN Edge	510, 510N	510-LTE	520	520V	540
Maximum Throughput Large Packet (1300-byte)					
Routed Mode All Ports	350 Mbps	350 Mbps	350 Mbps	350 Mbps	1 Gbps
Switched Mode All Ports	200 Mbps	200 Mbps	200 Mbps	200 Mbps	650 Mbps
Maximum Throughput Internet Traffic (IMIX)					
Routed Mode All Ports	200 Mbps	200 Mbps	200 Mbps	200 Mbps	500 Mbps
Switched Mode All Ports	80 Mbps	80 Mbps	80 Mbps	80 Mbps	200 Mbps
Other Scale Vectors					
Maximum Tunnel Scale	50	50	50	50	100
Flows Per Second	2,400	2,400	2,400	2,400	4,800
Maximum Concurrent Flows	240K	240K	240K	240K	480K
Maximum Number of Routes	100K	100K	100K	100K	100K
Maximum Number of Segments	128	128	128	128	128
Maximum Number of NAT Entries	80K	80K	80K	80K	150K

Table 3-6.

VMware SD-WAN Edge	640, 640C, 640N	680, 680C, 680N	840	2000	3400, 3400C
Maximum Throughput Large Packet (1300-byte)					
Routed Mode All Ports	3 Gbps	6 Gbps	4 Gbps	10 Gbps	7 Gbps
Switched Mode All Ports	1 Gbps	1 Gbps	1 Gbps	1.2 Gbps	1.2 Gbps
Maximum Throughput Internet Traffic (IMIX)					
Routed Mode All Ports	1 Gbps	2 Gbps	1.5 Gbps	5 Gbps	2.5 Gbps

Table 3-6. (continued)

VMware SD-WAN Edge	640, 640C, 640N	680, 680C, 680N	840	2000	3400, 3400C
Switched Mode All Ports	350 Mbps	350 Mbps	350 Mbps	350 Mbps	900 Mbps
Other Scale Vectors					
Maximum Tunnel Scale	400	800	400	6,000	4,000
Flows Per Second	19,200	19,200	19,200	38,400	38,400
Maximum Concurrent Flows	1.9M	1.9M	1.9M	1.9M	1.9M
Maximum Number of Routes	100K	100K	100K	100K	100K
Maximum Number of Segments	128	128	128	128	128
Maximum Number of NAT Entries	650K	650K	650K	960K	960K

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.

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 between any two Edges or an Edge and a Gateway is 16, regardless of Edge model or type.

Each WAN link (public or private) an Edge uses establishes two tunnels with another Edge or Gateway: one inbound, and one outbound. For example, if an Edge had four WAN links and connected to another Edge with four WAN links, that would be a total of 16 tunnels. In this example, no additional tunnels could be established between the two Edges if an additional WAN link is added to either Edge.

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

Table 3-7. Edge Maximum Throughput When a Firewall VNF is Actively Service Chained:

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

Table 3-8. Enhanced High-Availability (HA) Link Performance

Edge Model	510, 510N	510-LTE	520, 520v	610, 610C, 610N	610-LTE
Maximum Throughput (IMIX) Across Enhanced HA Link	90 Mbps	90 Mbps	100 Mbps	200 Mbps	200 Mbps

Edge Model	640, 640C, 640N	680, 680C, 680N	840	2000	3400, 3400C
Maximum Throughput (IMIX) Across Enhanced HA Link	800 Mbps	800 Mbps	800 Mbps	800 Mbps	800 Mbps

Important Performance with Edge Network Intelligence enabled:

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

Virtual Edge

Table 3-9. Private Cloud (Hypervisors)

Edge Device	Maximum Throughput	Maximum Number of Tunnels	Flows Per Second	Maximum Concurrent Flows	Maximum Number of Routes
ESXi Virtual Edge (2-core, VMXNET3)	2 Gbps (1300-byte) 800 Mbps (IMIX)	50	2400	240K	35K
KVM Virtual Edge (2-core, Linux Bridge)	500 Mbps (1300-byte) 200 Mbps (IMIX)	50	2400	240K	35K
KVM Virtual Edge (2-core, SR-IOV)	1.25 Gbps (1300-byte) 600 Mbps (IMIX)	50	2400	240K	35K
ESXi Virtual Edge (4-core, VMXNET3)	2 Gbps (1300-byte) 1.5 Gbps (IMIX)	400	19200	1.9M	35K
ESXi Virtual Edge (4-core, SR-IOV)	2 Gbps (1300-byte) 1.5 Gbps (IMIX)	400	19200	1.9M	35K
KVM Virtual Edge (4-core, Linux Bridge)	1 Gbps (1300-byte) 350 Mbps (IMIX)	400	4800	480K	35K
KVM Virtual Edge (4-core, SR-IOV)	2 Gbps (1300-byte) 1 Gbps (IMIX)	400	19200	1.9M	35K
ESXi Virtual Edge (8-core, VMXNET3)	5 Gbps (1300-byte) 2.5 Gbps (IMIX)	800	38400	1.9M	35K

Table 3-9. Private Cloud (Hypervisors) (continued)

Edge Device	Maximum Throughput		Maximum Number of Tunnels	Flows Per Second	Maximum Concurrent Flows	Maximum Number of Routes
ESXi Virtual Edge (8-core, SR-IOV)	Version 3.4 or older: 5 Gbps (1300-byte) 2.5 Gbps (IMIX)	Version 4.0 or newer: 9 Gbps (1300-byte) 4 Gbps (IMIX)	800	38400	1.9M	35
KVM Virtual Edge (8-core, SR-IOV)	Version 3.4 or older: 3.5 Gbps (1300-byte) 1 Gbps (IMIX)	Version 4.0 or newer: 9 Gbps (1300-byte) 3 Gbps (IMIX)	800	38400	1.9M	35

	2 vCPU	4vCPU	8vCPU	10vCPU
Minimum Memory (DRAM)	8 GB	16GB	32 GB	32 GB
Minimum Storage	8 GB	8 GB	16 GB	16 GB
Supported Hypervisors	Software Version 3.4 or older: <ul style="list-style-type: none"> ■ ESXi 6.0, 6.5U1, 6.7U1 ■ KVM Ubuntu 14.04 LTS or 16.04 Software version 4.0 and above: <ul style="list-style-type: none"> ■ 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 Instruction set: <ul style="list-style-type: none"> ■ AES-NI ■ AVX2 or AVX512 ■ SSE3, SSE4, and RDTSC instruction sets Hyper-threading disabled			

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

Public Cloud

Table 3-10. Amazon Web Services (AWS)

AWS Instance Type	c5.large	c5.xlarge	c5.2xlarge
Maximum Throughput	100 Mbps (1300-byte) 50 Mbps (IMIX)	200 Mbps (1300-byte) 100 Mbps (IMIX)	7 Gbps (1300-byte) 2.4 Gbps (IMIX)
Maximum Tunnels	50	400	800

Table 3-10. Amazon Web Services (AWS) (continued)

AWS Instance Type	c5.large	c5.xlarge	c5.2xlarge
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

Note c5.2xlarge performance and scale numbers are based on AWS Enhanced Networking (ENA SR-IOV drivers) being 'enabled'.

Table 3-11. Microsoft Azure

Azure VM Series	D2d v4	D4d v4	D8d v4
Maximum Throughput	100 Mbps (1300-byte) 50 Mbps (IMIX)	200 Mbps (1300-byte) 100 Mbps (IMIX)	1 Gbps (1300-byte) 450 Mbps (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

Note Azure Accelerated Networking is supported with a limited availability. Please contact your sales representative for details.

Use of DPDK on VMware SD-WAN Edges

To improve packet throughput performance, VMware 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 more 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.

Capabilities

This section describes VMware SD-WAN capabilities.

Dynamic Multi-path Optimization

VMware 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 VMware 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.

Multi-source Inbound QoS

VMware classifies 3000+ applications enabling smart control. Out-of-the-box defaults set the multi-source inbound Quality of Service (QoS) parameters for different application types with IT required only to establish application priority. Knowledge of network requirements for different application types, automatic link capacity measurements and dynamic flow monitoring enables automation of QoS configurations and bandwidth allocations.

Firewall

VMware 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 (e.g., disable 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 VMware Solution supports a platform to host multiple virtualized network functions to eliminate single-function appliances and reduce branch IT complexity. VMware 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

SD-WAN 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 SD-WAN 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 SD-WAN 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 SD-WAN Edge. The Overlay Flow Control (OFC) enables enterprise-wide route visibility and control for ease of programming and for full and partial overlay.

OSPF

VMware 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

VMware supports inbound/outbound filters and the filter can be set to Deny, or optionally adding/changing 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.

VMware enables 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, SD-WAN Edge learns routes from adjacent routers through OSPF or BGP. The SD-WAN Orchestrator maintains all the dynamically learned routes in a global routing table called the Overlay Flow Control. 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 more information, see [Chapter 23 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.

Tunnel Overhead and MTU

VMware, 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 VMware, 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.

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 VMware regardless of whether NAT is present or not (NAT-T is always activated).

VMware Tunnel Overhead

To support Dynamic Multipath Optimization™ (DMPO), VMware 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.

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

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

VMware 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.

Network Topologies

This section describes network topologies for branches and data centers.

Branches to Private Third Party (VPN)

Customers with a private data center or cloud data center often want a way to include it in their network without 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 SD-WAN Gateway to the customer's existing router or firewall. All the SD-WAN Edge that need to talk to the site will connect to the same SD-WAN Gateway to forward packets across the tunnel, simplifying the overall network configuration and new site bring up.

SD-WAN Edges

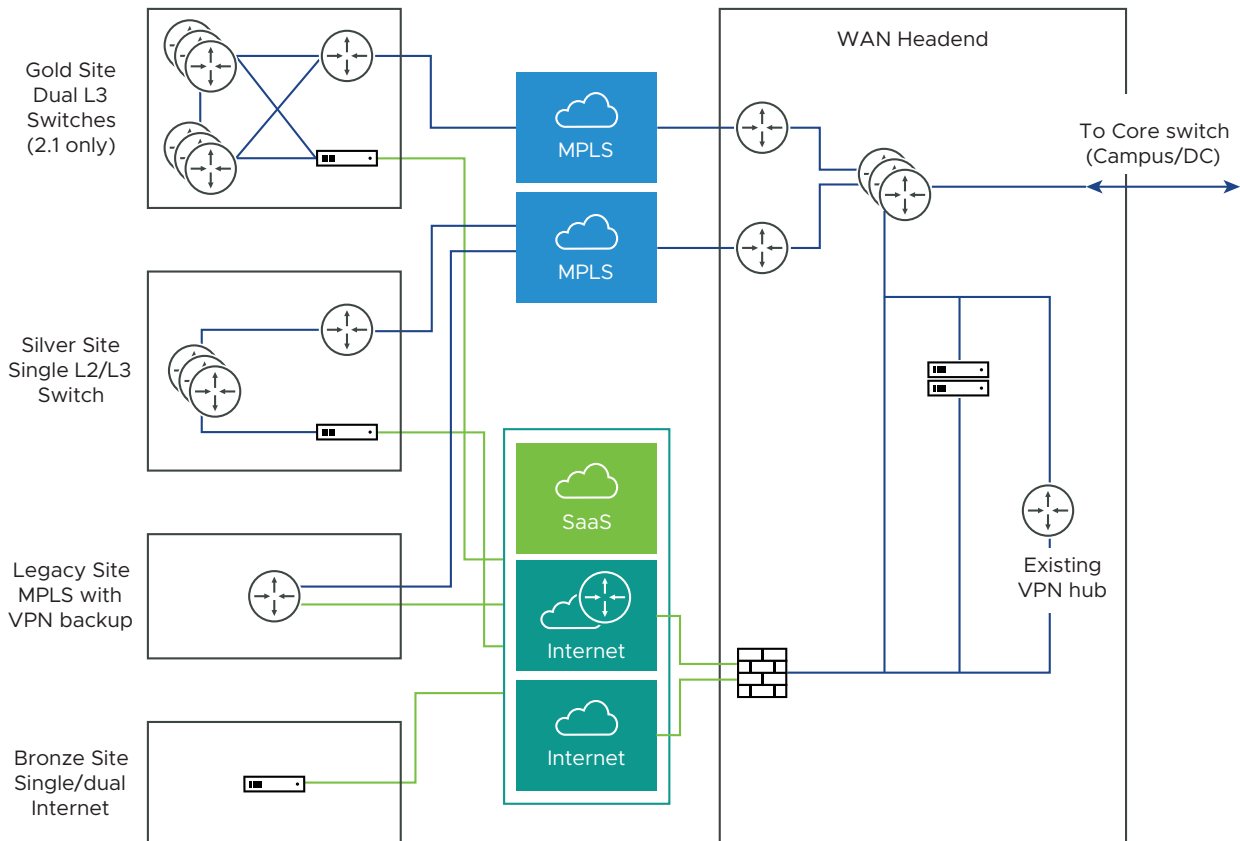


VMware simplifies the branch deployment and delivers enterprise great application performance or public/private link for cloud and/or on-premise applications.


Data Center Network Topology

The Data Center Network topology consists of two hubs and multiple branches, with or without SD-WAN 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).



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.



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 VMware 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. VMware customer-premise equipment, SD-WAN Edge, aggregates multiple broadband links (e.g., Cable, DSL, 4G-LTE) at the branch office, and sends the traffic to SD-WAN Gateways. Using cloud-based orchestration, the service can connect the branch office to any of type of data center: enterprise, cloud, or Software-as-a-Service.


SD-WAN 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.

Branch Site Topologies

The VMware service defines two or more different branch topologies designated as Bronze, Silver, and Gold. In addition, pairs of SD-WAN Edge 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 SD-WAN Edge. The following figure shows an overview of the Bronze 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 a MPLS link, which is terminated on a CE and is accessible through the L3 switch. In this case, the SD-WAN Edge goes between the L3 switch and Internet (replacing existing firewall/router).

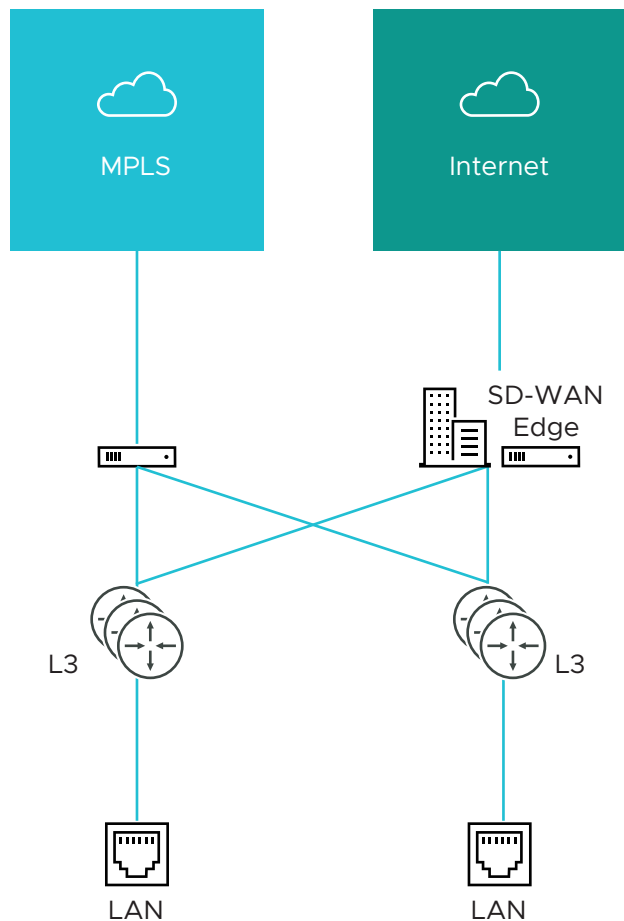


The second variant includes MPLS and Internet routers deployed using HSRP with an L2 switch on the LAN side. In this case, the SD-WAN Edge replaces the L2 switch.




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.




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.



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 VMware High Availability configuration using two SD-WAN Edge, one active and one standby.



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

Roles and Privilege Levels

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

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

Administrator

The Administrator configures, monitors, and administers the VMware service operation. There are three 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 SD-WAN Edge physical installation and activation with the VMware 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:

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:

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.

User Role Matrix

This section describes feature access according to VMware user roles.

Operator-level SD-WAN Orchestrator Features User Role Matrix

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

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

SD-WAN 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

SD-WAN Orchestrator Feature	Operator: Superuser Operator	Operator: Standard Operator	Partner: Business Specialist	Partner: Customer Support Operator	Super User	Standard Admin	Business Specialist	Customer Support
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 SD-WAN Orchestrator Features User Role Matrix

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

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

SD-WAN 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

SD-WAN Orchestrator Feature	Partner: Superuser	Partner: Standard Admin	Business Specialist	Customer Support
Settings	RW	R	R	R
Gateway Pools	RW	RWD	NA	R
Gateways	RW	RW	NA	R

Enterprise-level SD-WAN Orchestrator Features User Role Matrix

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

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

SD-WAN 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
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

SD-WAN Orchestrator Feature	Enterprise: Super User	Enterprise: Standard Admin	Customer Support	Read Only
Administration > System Settings	RW	RW	RW	NA
Administration > Administrators	RW	R	R	NA

Note Operator users have complete access to the SD-WAN Orchestrator features.

Key Concepts

This section describes the key concepts and the core configurations of SD-WAN Orchestrator.


Configurations

The VMware service has four core configurations that have a hierarchical relationship. Create these configurations in the SD-WAN Orchestrator.

The following table provides an overview of the configurations.

Configuration	Description
Network	Defines basic network configurations, such as IP addressing and VLANs. Networks can be designated as Corporate or Guest and there can be multiple definitions for each network.
Network Services	Define several common services used by the VMware 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.



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 VMware 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 SD-WAN Edge.



Note When using non-overlapping addresses, the SD-WAN 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 SD-WAN 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

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

Scenario	Description
SaaS	Used for Edges that do not require VPN connections between Edges, to a Non SD-WAN Destination, or to a VMware SD-WAN 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.
VMware SD-WAN Site VPN	Used for Edges that require VPN connections to a VMware SD-WAN 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 VMware SD-WAN Sites are defined in the profile.

For each scenario, perform the configurations in the SD-WAN 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 VMware SD-WAN Site or Non SD-WAN Destination. The final step is to create a new Edge and activate it.

Quick Start Configuration Steps	SaaS	Non SD-WAN Destination VPN	VMware SD-WAN 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 VMware SD-WAN 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 more information, see [Activate SD-WAN Edges](#).

Supported Browsers

The SD-WAN Orchestrator supports the following browsers:

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

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

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

Supported Modems

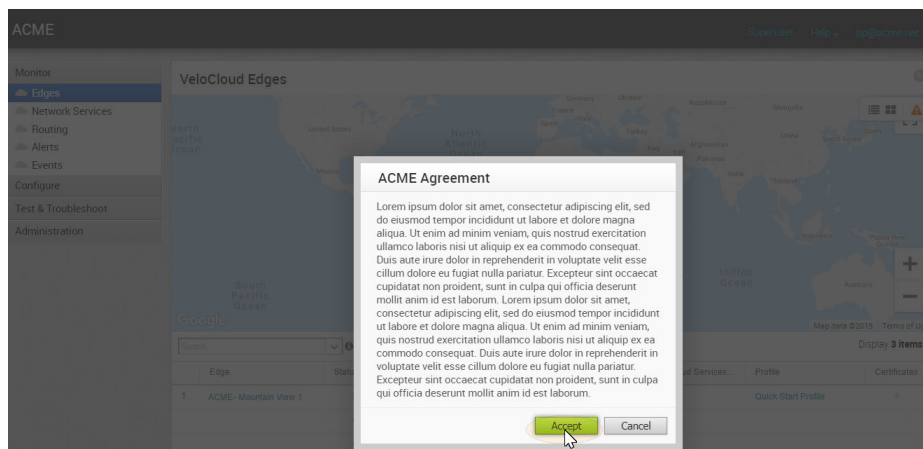
This section describes how to get a list of supported modems.

For a detailed list of supported modems, see <https://sdwan.vmware.com/get-started/supported-modems>.

User Agreement

4

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



Log in to VMware SD-WAN Orchestrator Using SSO for Enterprise User

5

Describes how to log in to VMware SD-WAN Orchestrator using Single Sign On (SSO) as an Enterprise user.

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

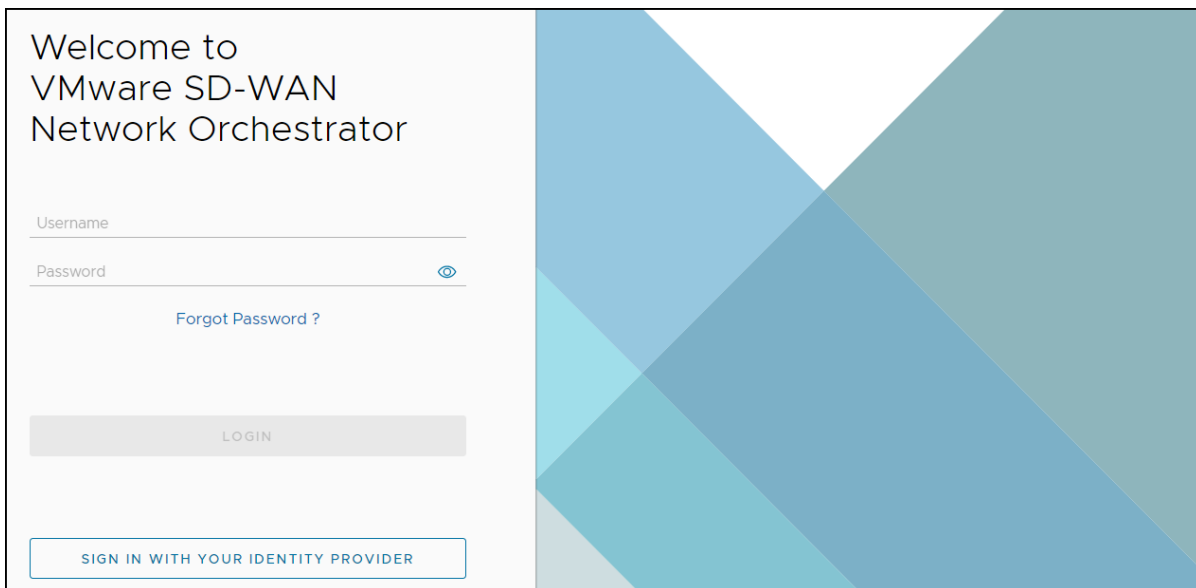
Prerequisites

- Ensure you have configured the SSO authentication in SD-WAN Orchestrator. For more information, see [Configure Single Sign On for Enterprise User](#).
- Ensure you have set up roles, users, and OIDC application for the SSO in your preferred IDPs. For more information, see [Configure an IDP for Single Sign On](#).

Procedure


- 1 In a web browser, launch the SD-WAN Orchestrator application as an Enterprise user.

The VMware SD-WAN Operations Console screen appears.



- 2 Click **Sign In With Your Identity Provider**.

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



Sign In using Identity Provider

VMware SD-WAN
Network Orchestrator

ACME

SIGN IN

BACK TO LOGIN PAGE

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

Note

- Once the users log in to the SD-WAN Orchestrator using SSO, they are not allowed to login again as native users.
 - The user can navigate to the Classic UI by clicking the **Open Classic Orchestrator** option located at the top right of the UI screen.
-

Monitor Enterprises

6

The SD-WAN Orchestrator provides monitoring functionality that allows you to observe various performance and operational characteristics of VMware SD-WAN Edges. Monitoring functionality is accessible in **Monitor** area of the navigation panel.

Read the following topics next:

- [Monitor Navigation Panel](#)
- [Network Overview](#)
- [Monitor Edges](#)
- [Monitor Network Services](#)
- [Monitor Routing](#)
- [Monitor Alerts](#)
- [Monitor Events](#)
- [Monitor Reports](#)

Monitor Navigation Panel

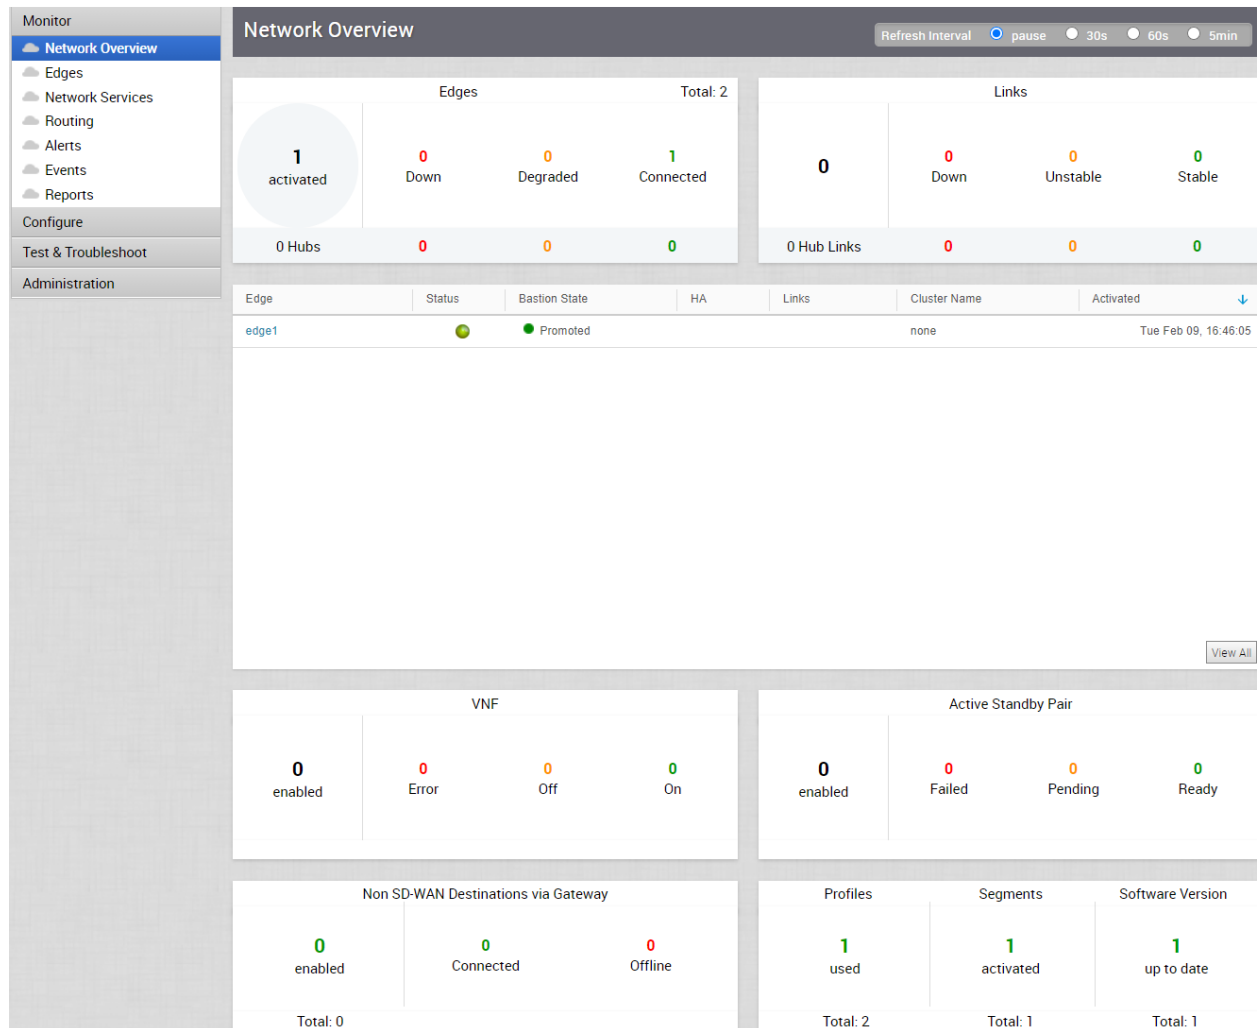
The following monitoring capabilities are displayed under **Monitor** in the navigation panel.

- [Network Overview](#)
- [Monitor Edges](#)
- [Monitor Network Services](#)
- [Monitor Routing](#)
- [Monitor Alerts](#)
- [Monitor Events](#)

Network Overview

The Network Overview feature helps to monitor networks by checking the Edge and Link (activated Edge) status summary. Clicking **Monitor > Network Overview** in the navigation panel opens the **Network Overview** screen, which provides a visual summary about the

enterprises running SD-WAN Edge devices, Non SD-WAN Destinations, profiles, segments, software versions, and their system configuration time and run time statuses.



The **Network Overview** screen presents the overall summary information about a network in three dashboard sections:

- SD-WAN Edge statistics - Includes the following information about the Edges and Links:
 - Total number of Edges
 - Total number of Edge Hubs
 - Total number of Links
 - Total number of Hub Links
 - Count of Edges/Edge Hubs (Connected, Degraded, and Down)
 - Count of Link/Hub Links (Stable, Unstable, and Down)

- Summary dashboard table - Includes a table that displays top ten Edges, or Edge Hubs, or Links, or Hub Links sorted by last contact time, based on the selected filter criteria in the SD-WAN Edge statistics section.
- Non-Edge statistics - Includes the following non-Edge related information:
 - Total number of Virtual Network Functions (VNFs)-activated Edges
 - Count of VNFs-activated Edges (Error, On, and Off)
 - Total number of VMware Active Standby Pair-activated Edges
 - Count of VMware Active Standby Pair-activated Edges (Failed, Pending, and Ready)
 - Total number of activated Non SD-WAN Destinations
 - Count of Non SD-WAN Destinations (Connected and Offline)
 - Count of used Profiles out of the total number of Profiles configured for the Enterprise.
 - Count of activated Segments out of the total number of Segments configured for the Enterprise.
 - Count of Edges with up-to-date Software version out of the total number of Edges configured for the Enterprise.

Note The minimum supported edge version is 2.4.0. You can change the target edge version against which the edges will be compared by using the system property `product.edge.version.minimumSupported`.

You can also get detailed information on a specific item in the **Network Overview** screen by clicking the link on the respective item or metric. For example, clicking the **Edge** link in the summary dashboard table takes you to the Edge detail dashboard for the selected Edge.

You can configure the refresh time interval for the information displayed in the Network Overview dashboard screen to one of the following options:

- pause
- 30s
- 60s
- 5min

SD-WAN Edge States and Transitions

Transitions are driven by Edge heartbeats (which occur under normal circumstances every 30 seconds), irrespective of the Links over which the heartbeats are received.

The following table describes the connection state types and transitions for a SD-WAN Edge.

Color	Edge State	Description
Green	Connected	<ul style="list-style-type: none"> ■ An Edge is in Connected state if a heartbeat has been received from the Edge in the last 60 seconds. ■ The Edge transitions from Connected to Degraded state when the Orchestrator determines that a heartbeat has not been received from the Edge for more than 60 seconds. ■ The Edge transitions from Connected to Offline state when the Orchestrator has not received two consecutive heartbeats from the Edge within a span of two minutes (120 seconds).
Amber	Degraded	<ul style="list-style-type: none"> ■ An Edge is in Degraded state if the Edge to Orchestrator connectivity appears to be impacted, possibly due to transient network conditions. ■ The Edge transitions from Degraded to Offline state when the Orchestrator determines that a heartbeat has not been received from the Edge for more than two minutes (120 seconds).
Red	Offline	An Edge is in Offline state if the Edge is unable to reach the Orchestrator due to some persistent network condition.

SD-WAN Orchestrator Link States and Transitions

SD-WAN Orchestrator drives state changes between the various Link states based on the most recent state change, taking into consideration the time when the Link was last active and the time when the event last occurred. Transitions are driven by a combination of:

- Edge-reported Link Stats values as received when the Edge pushes the Link Stats to the Orchestrator (occurs every 5 minutes).
- Edge-reported Events as received by Edge heartbeats (occurs every 30 seconds).

The following table describes the connection state types and transitions for a SD-WAN Orchestrator Link.

Color	Edge State	Description
Green	Stable	A Link is in Stable state if the Link conditions appear to be stable and the Orchestrator receives the Link Stats consistently.
Amber	Unstable	A Link is in Unstable state if an expected Link Stats push is not received, or Link is down, but has not yet been inactive for 10 minutes.
Red	Disconnected	A Link is in Disconnected state if the Link has been inactive for more than 10 minutes.

SD-WAN Edge Bastion States and Transitions

The following table describes the Bastion state types and transitions for a SD-WAN Edge.

Bastion State	Description
UNCONFIGURED	The initial state of a SD-WAN Edge before it is staged. The SD-WAN Edge is available only in Private Orchestrator.
STAGE_REQUESTED	An intermediate state before the SD-WAN Edge is staged to Public Orchestrator.
STAGED	The SD-WAN Edge is staged to Public Orchestrator.

Bastion State	Description
UNSTAGE_REQUESTED	An intermediate state before the SD-WAN Edge is removed from Public Orchestrator.
UNSTAGED	The SD-WAN Edge is removed from Public Orchestrator and available only in Private Orchestrator.
PROMOTION_REQUESTED	An intermediate state when a user has requested promotion of the SD-WAN Edge from Public Orchestrator to Private Orchestrator.
PROMOTION_PENDING	Configuration for the SD-WAN Edge to be promoted has been pushed to the Public Orchestrator and is waiting for the Edge to send heartbeat back to the Private Orchestrator.
PROMOTED	The SD-WAN Edge has been successfully promoted and currently heartbeats to the Private Orchestrator.

Note These states are available only if the Bastion Orchestrator feature is activated on the SD-WAN Orchestrator.

For more information, see *Bastion Orchestrator Configuration Guide* available at <https://docs.vmware.com/en/VMware-SD-WAN/index.html>.

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:

- 1 In the Enterprise portal, click **Monitor > Edges**.
- 2 The **Edges** page displays the Edges associated with the Enterprise.

3-site-public-cloud Open New Orchestrator UI Recently Viewed Operator Superuser Help super@velocloud.net

Monitor

- Network Overview
- Edges**
- Network Services
- Routing
- Alerts
- Events
- Reports

Configure

Test & Troubleshoot

Administration

Edges

Search... Cols Reset View Refresh CSV Display 3 items

	Edge	Status	HA	Links	VM Status	VNF	Edge Tunnels	Gateways	Profile
1	b1-edge1	●		↔ 3			↔ 3 ↔ 3	View	Quick Start Profile
2	b2-edge1	●		↔ 2			↔ 2 ↔ 2	View	Quick Start Profile
3	b3-edge1	●		↔ 2			↔ 2 ↔ 2	View	Quick Start Profile

The page displays the following options:

- **Table of edges** – Lists all edges provisioned in the network.
- **Search** – Enter a term to search for a specific detail. Click the drop-down arrow to filter the view by specific criteria.
- **Cols** – Click and select the columns to be shown or hidden in the view. By default, Edge and Status information are displayed.
- **Reset View** – Click to reset the view to default settings.
- **Refresh** – Click to refresh the details displayed with the most current data.
- **CSV** – Click to export all data to a file in CSV format.

Click the link to an Edge to view the details pertaining to the selected Edge. Click 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.

For each Edge, you can view the following details:

- [Overview Tab](#)
- [QoE Tab](#)
- [Transport Tab](#)
- [Applications Tab](#)
- [Sources Tab](#)
- [Destinations Tab](#)

- [Business Priority Tab](#)
- [System Tab](#)

Overview Tab

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:

Procedure

- 1 In the Enterprise portal, click **Monitor > Edges**.
- 2 Click the link to an Edge and the **Overview** tab is displayed by default.

Results

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



You can choose to view the Edge information live by selecting the **Stay in live mode** checkbox. When this mode is activated, 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 following details:

Option	Description
Links	The Interface and WAN links of the selected Edge.
Link Status	Connectivity status of the Link to the Gateway.
Interface (WAN Type)	The Interface connected to the Link.
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 the Orchestrator is 5 minutes.
Bandwidth	The maximum rate of data transfer across a given path. Displays both the upstream and downstream bandwidth details.
Pre-Notifications	Allows to activate or deactivate the alerts sent to the Operator. Click Edit to modify the notification settings.
Alerts	Allows to activate or deactivate the alerts sent to the Enterprise Customer. Click Edit to modify the notification settings.
Signal	Information on signal strength.
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.

The **Bandwidth Usage** section displays graphical representation of bandwidth and network usage of the following: Applications, Categories, Operating Systems, Sources, and Destinations of the Edges. Click **View Details** in each panel to navigate to the corresponding tab and view more details.

Hover the mouse on the graphs to view more 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.

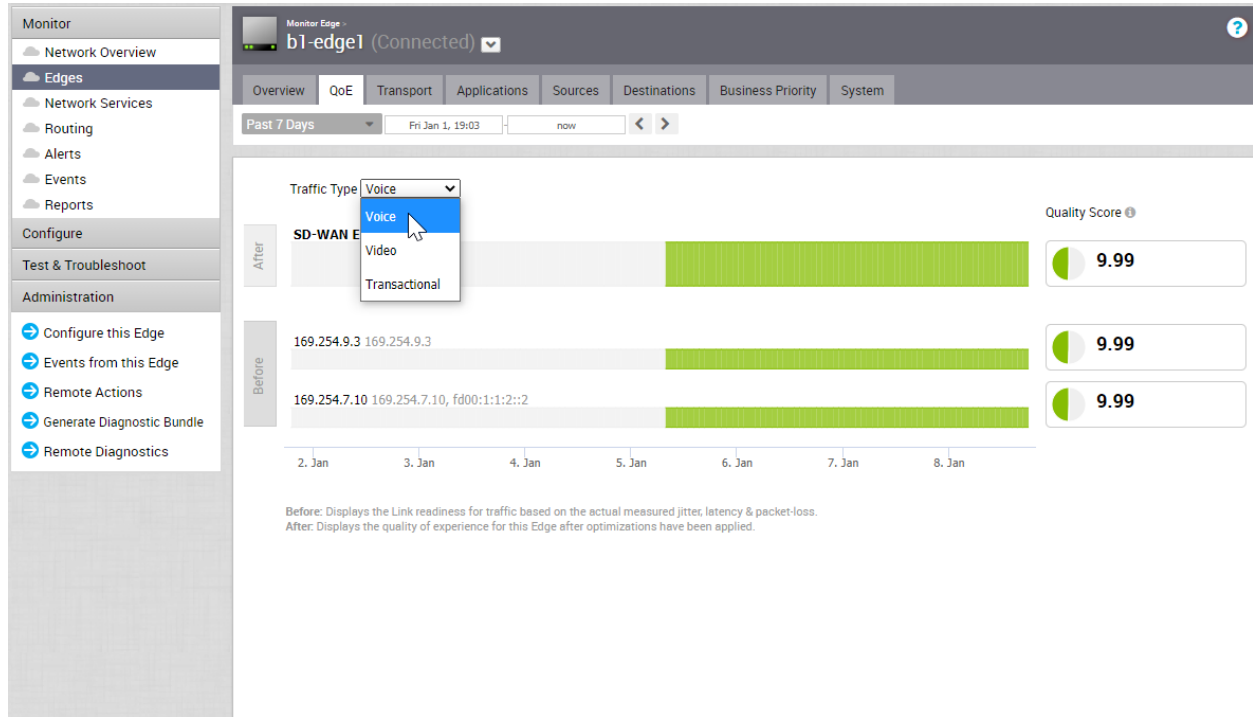
QoE Tab

The VMware **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.

Click the **Monitor > Edges > QoE** tab to view the following details.

Traffic Type

There are three different traffic types that you can monitor (Voice, Video, and Transactional) in the **QoE** tab. You can hover over a WAN network link, or the aggregate link to display a summary of Latency, Jitter, and Packet Loss.



Quality Score

The Quality Score rates an application's quality of experience that a network can deliver for a given time frame. Some examples of applications are video, voice, and transactional. QoE rating options are shown in the table below.

Rating Color	Rating Option	Definition
Green	Good	All metrics are better than the objective thresholds. Application SLA 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.


QoE Example

The following images show examples of QoE with before and after voice traffic scenario problems and how VMware solved them. The red numbers in the following images represent the scenario numbers in the table.

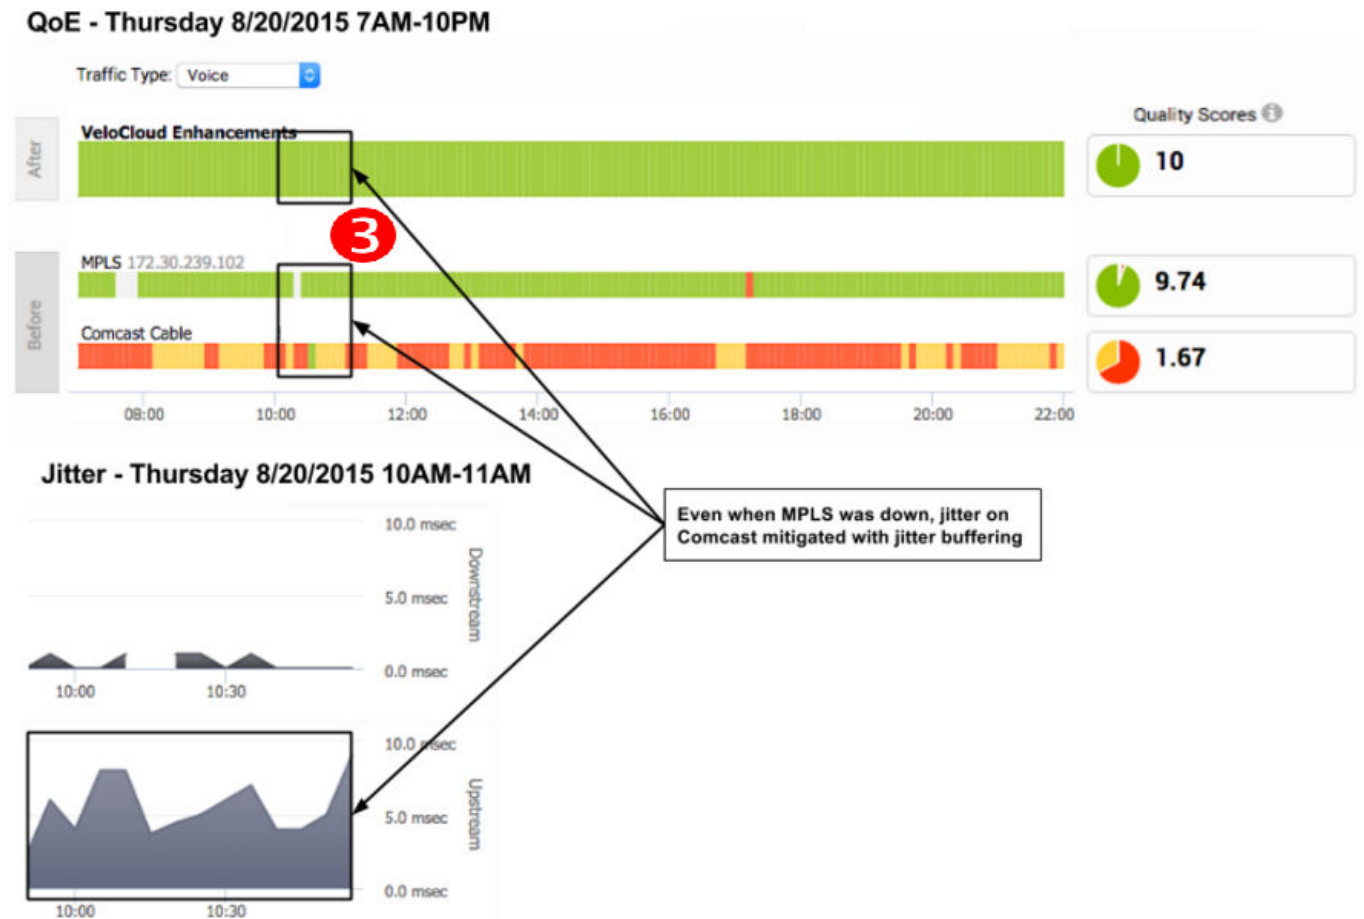
QoE Example Table

Scenario	Issue	VMware Solution
1	MPLS is down	Link steering
2	Packet loss	Forward error correction
3	MPLS is down; Jitter on Comcast	Link steering and jitter buffering

Scenario 1 and 2: Link Steering and Forward Error Correction Solution Example



Scenario 3: Link Steering and Jitter Buffering Solution Example



Transport Tab

You can monitor the WAN links connected to a specific Edge along with the status, interface details, and other metrics.

At any point of time, you can view which Link or Transport Group is used for the traffic and how much data is sent in the **Monitor > Edges > Transport** tab.

When you click the **Transport** tab, **Links** is screen is displayed by default. The screen displays Sent and Received data for your links. The links associated with an Edge are displayed in at the bottom of the screen under the Link column, along with the status for Cloud and VPN, WAN Interface, Application details, and details of Bytes.

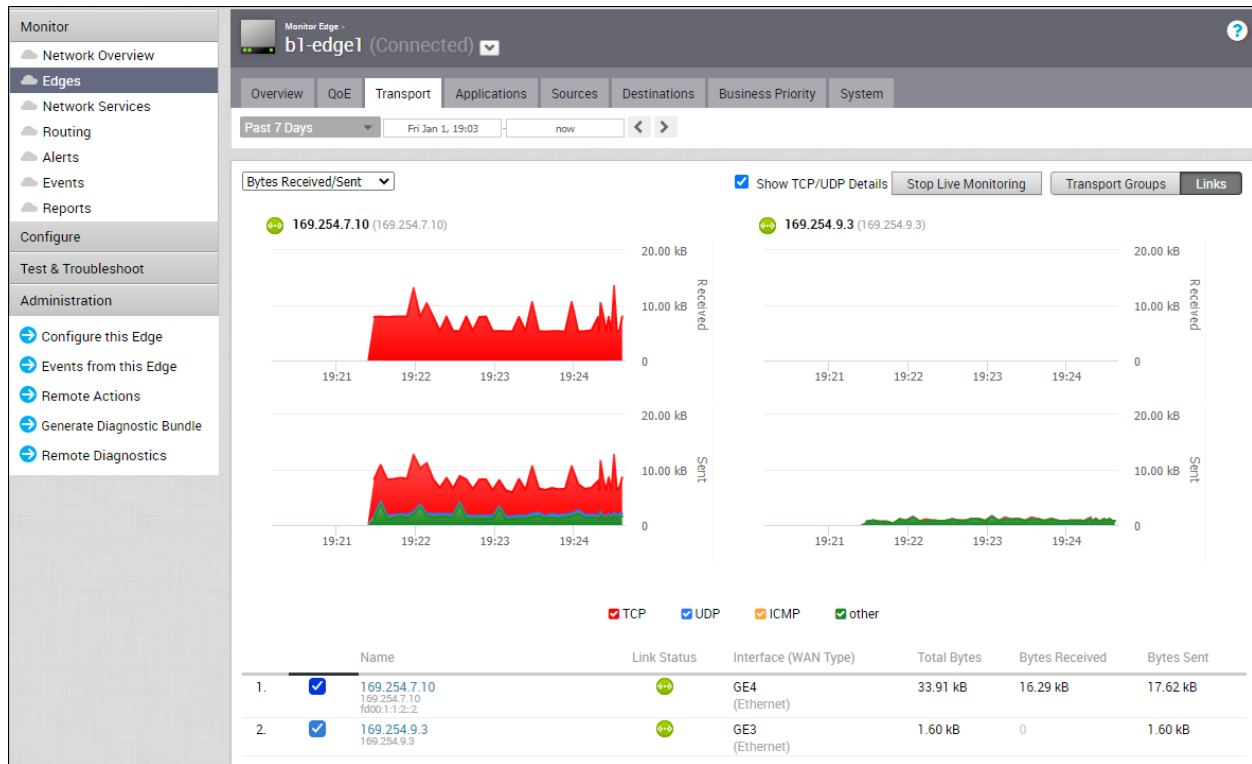
Hover the mouse on the graphs to view more details.

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

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

You can choose to view the information live by clicking the **Start Live Monitoring** option. When this mode is enabled, you can view live monitoring of the links and the transport groups. Live monitoring is useful for conducting active testing and calculating Average Throughput. It is also beneficial for troubleshooting security compliance and for monitoring how traffic policies are being leveraged in real time.

In the **Live Monitoring** screen, select the **Show TCP/UDP Details** checkbox to view protocol level link usage details.




By default, the **Scale Y-axis evenly** checkbox is enabled. This option synchronizes the Y-axis between the charts. If required, you can deactivate this option.

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.

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

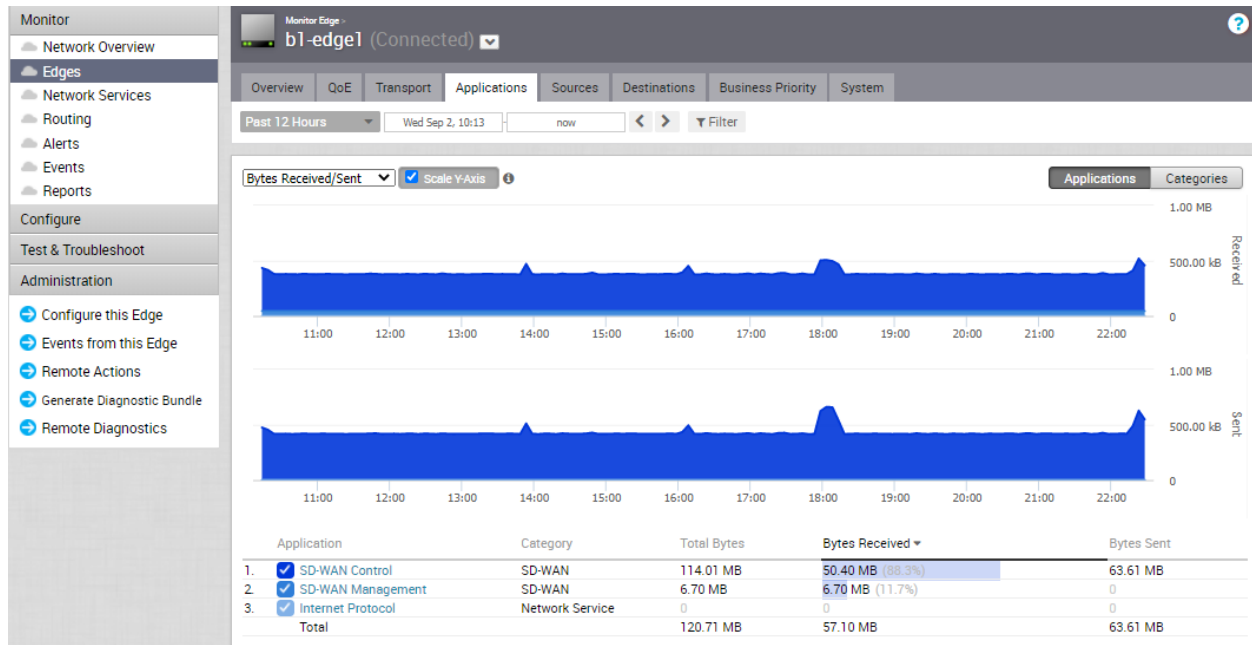
The following image shows a detailed report of transport groups with top applications.



Applications Tab

You can monitor the network usage of applications or application categories used by a specific Edge.

Click the **Monitor > Edges > Applications** tab to view the following:



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. For more information on the metrics parameters, see [Monitor Edges](#).

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.

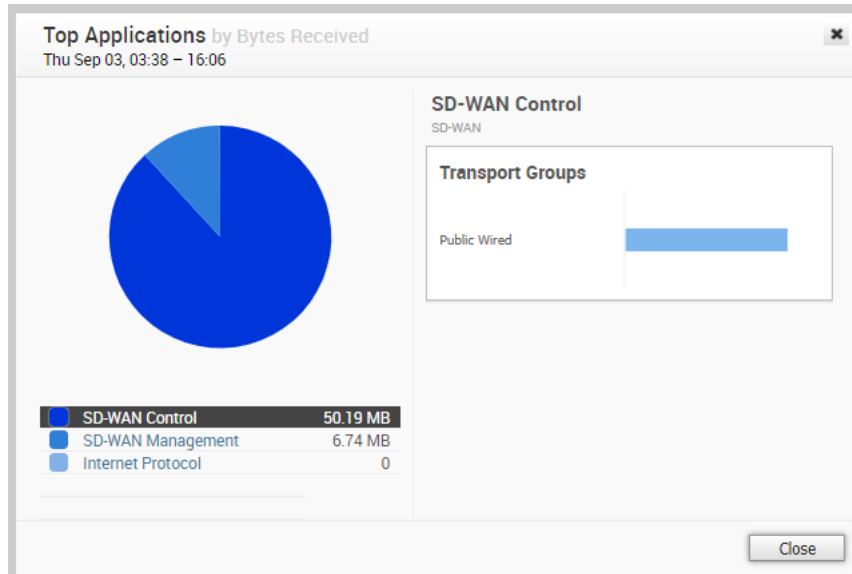
Hover the mouse on the graphs to view more details.

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

The bottom panel displays the details of the selected metrics for the applications or categories.

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

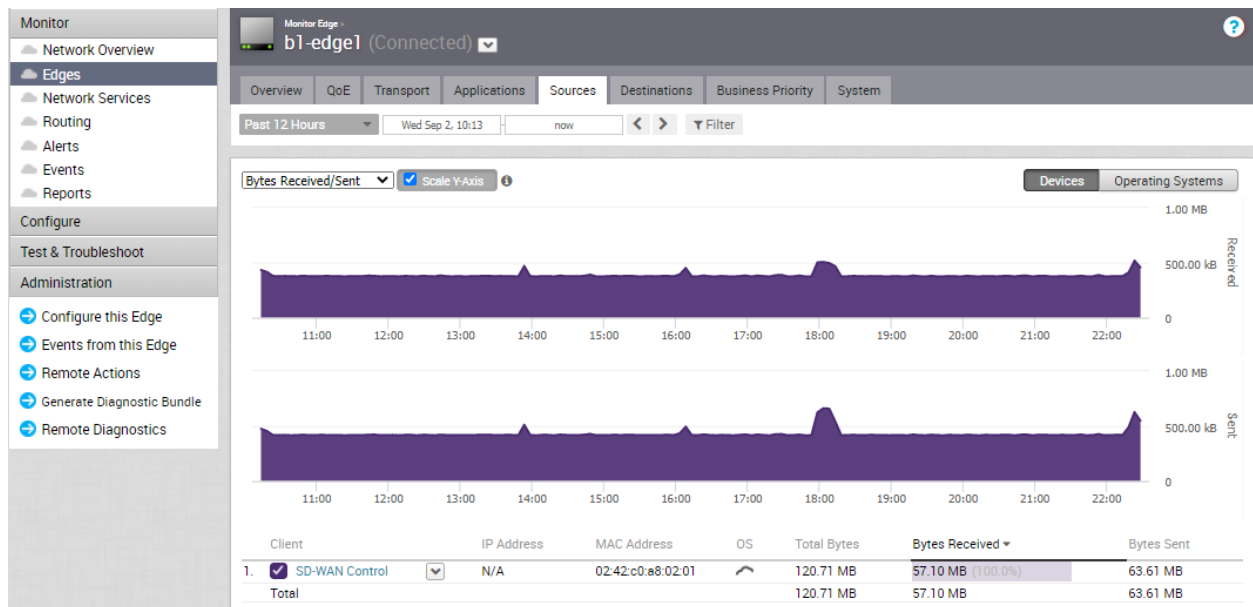
The following image shows a detailed report of top applications.



Sources Tab

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

Click **Monitor > Edges > Sources** to view the following:



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. For more information on the metrics parameters, see [Monitor Edges](#).

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 more details.

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

The bottom panel displays the details of the selected metrics for the devices or operating systems.

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

The following image shows a detailed report of top clients.

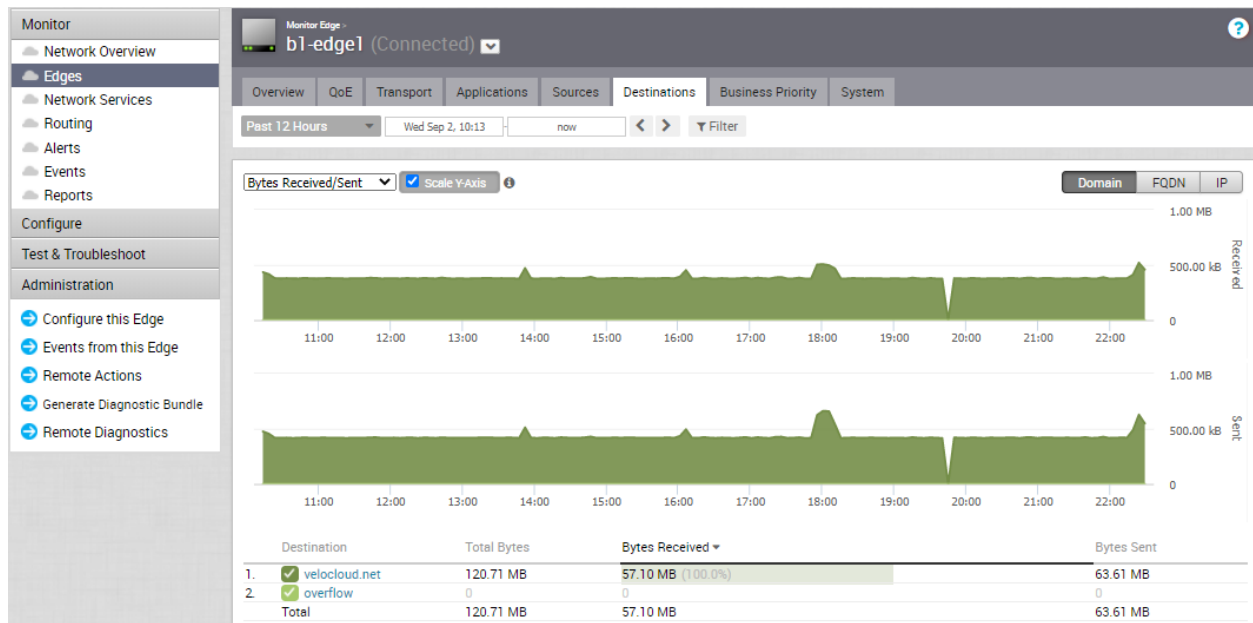


Click the arrows displayed next to **Top Applications** to navigate to the **Applications** tab.

Destinations Tab

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

Click the **Monitor > Edges > Destinations** tab to view the following:



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. For more information on the metrics parameters, see [Monitor Edges](#).

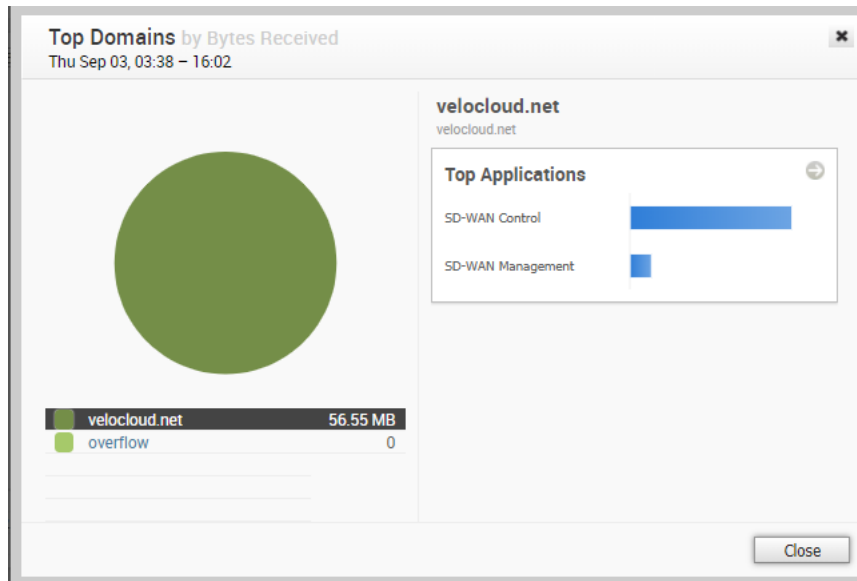
By default, the **Scale Y-axis evenly** check box is selected. This option synchronizes the Y-axis between the charts. If required, you can turn off this option.

Hover the mouse on the graphs to view more details.

The bottom panel displays the details of the selected metrics for the destinations by the selected type.

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

The following image shows a detailed report of top domains.

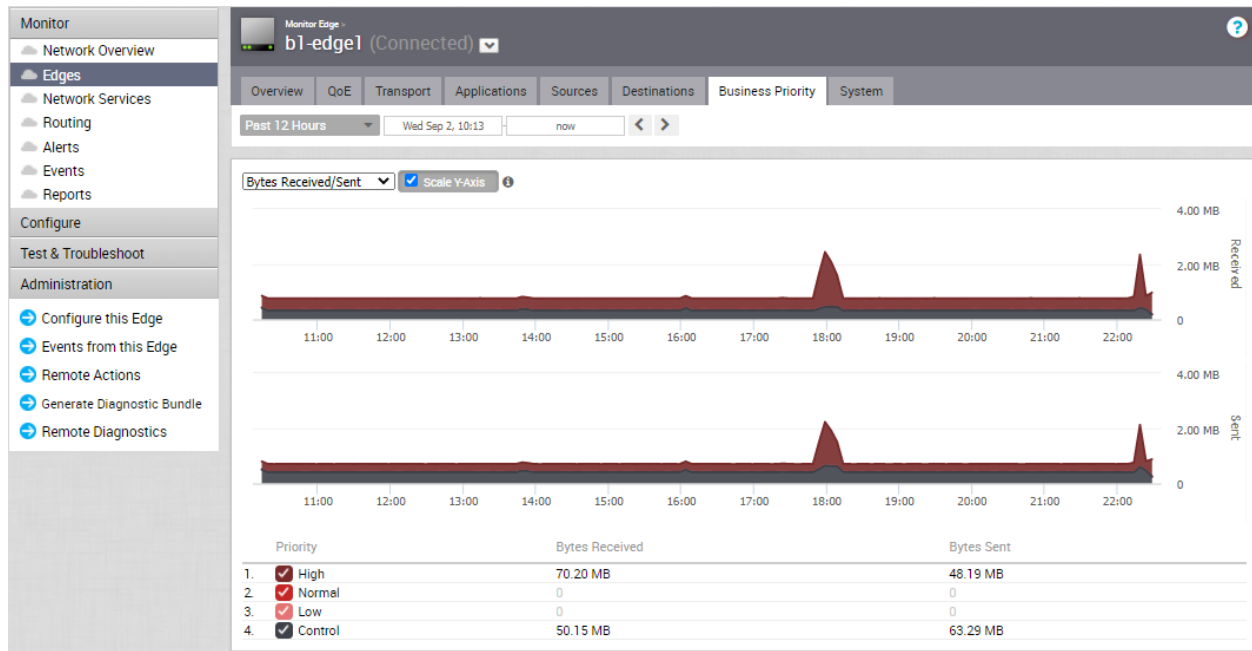


Click the arrows displayed next to **Top Applications** to navigate to the **Applications** tab.

Business Priority Tab

You can monitor the Business policy characteristics according to the priority and the associated network usage data for a specific Edge.

Click **Monitor > Edges > Business Priority** tab, to view the following:



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. For more information on the metrics parameters, see [Monitor Edges](#).

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 more details.

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

System Tab

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

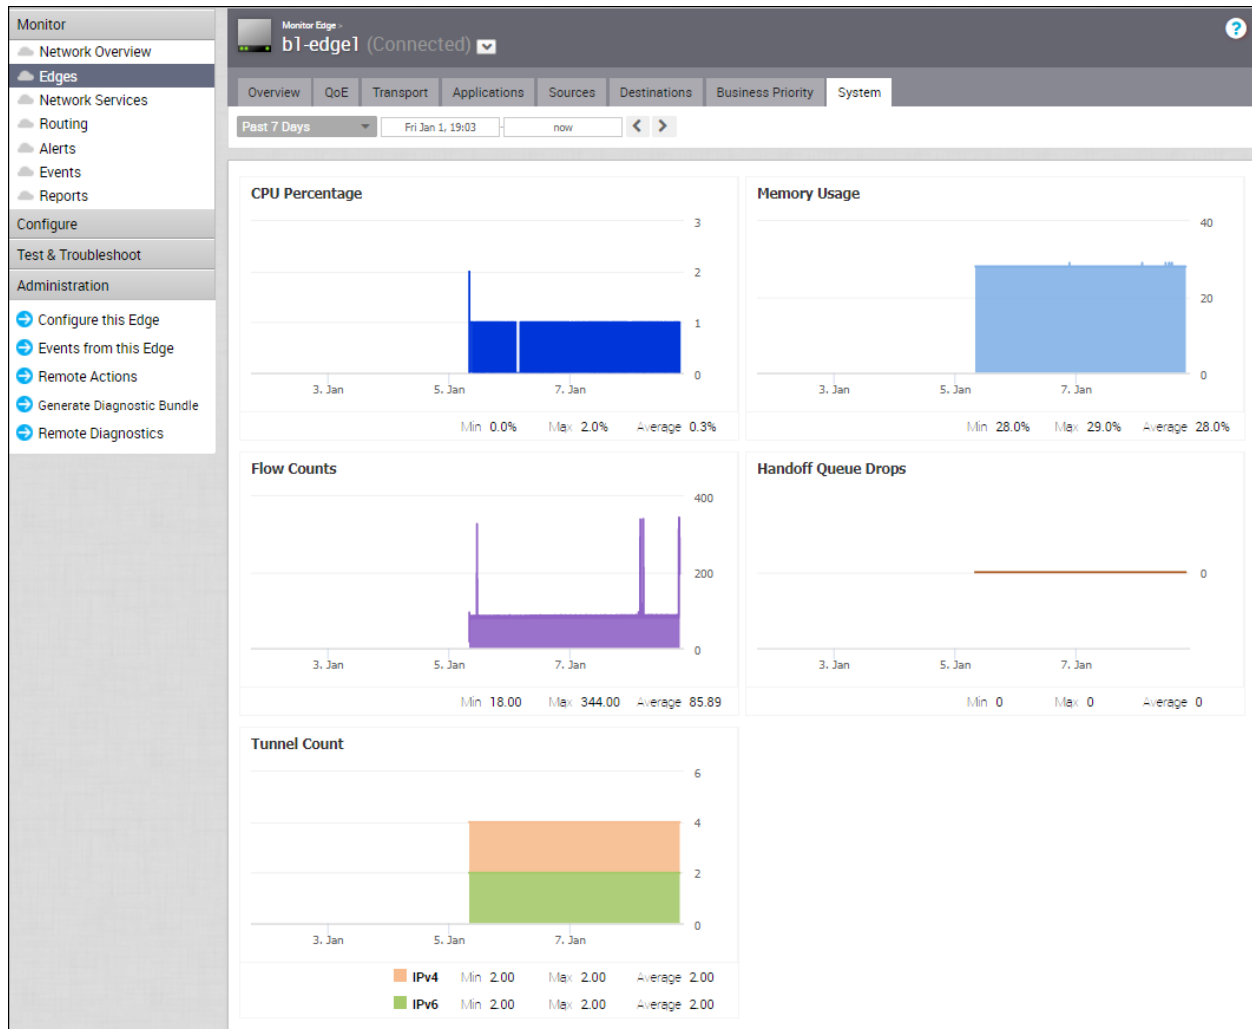
To view the details of system information:

Procedure

- 1 In the Enterprise portal, click **Monitor > Edges**.
- 2 Click the link to an Edge and click the **System** tab.

Results

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



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 Percentage** – Percentage of usage of CPU.
- **Memory Usage**– Percentage of usage of memory.
- **Flow Counts** – Count of traffic flow.
- **Handoff Queue Drops** – Count of packets dropped due to oversubscription of the Edge resources.
- **Tunnel Count** – Count of tunnel sessions.

Hover the mouse on the graphs to view more details.

VMware SD-WAN Orchestrator Data Retention

Describes the data retention policy for the VMware SD-WAN Orchestrator.

SD-WAN Data Retention

SD-WAN Data	System Property	Default	Maximum	Prior to 4.0 Release
Enterprise Events	retentionWeeks.enterpriseEvents	40 weeks	1 year	40 weeks
Enterprise Alerts	N/A	40 weeks	1 year	No policy
Operator Events	retentionWeeks.operatorEvents	40 weeks	1 year	40 weeks
Enterprise Proxy Events	retentionWeeks.proxyEvents	40 weeks	1 year	40 weeks
Firewall Logs	retentionWeeks.firewallLogs	Not supported	Not supported	40 weeks
Link Stats	retentionWeeks.linkStats	40 weeks	1 year	40 weeks
Link QoE	retentionWeeks.linkQuality	40 weeks	1 year	40 weeks
Path Stats	N/A	2 weeks	2 weeks	N/A
Flow Stats	retention.lowResFlows.months	1 year – 1 hour rollup	1 year – 1 hour rollup	1 year with rollup
	retention.highResFlows.days	2 weeks – 5 min	3 month – 5 min	
Edge Stats	retentionWeeks.edgeStats	2 weeks	2 weeks	No policy

Important Notes

- The 4.0 release flow statistics defaults will apply for the flow statistics retention post upgrade. To configure retention values for flow statistics, please follow the instructions below. The system properties pertaining to flow statistics in releases 3.3 and 3.4 will be deprecated upon upgrade to the 4.0 release.
- Currently, as per design, the Edge Stats data is present in monthly partitions. So if we were to truncate data older than 2 weeks, it essentially truncates data older than a month and the SD-WAN Orchestrator Monitor page displays the current and the immediate previous month's health stats data.

Changing the Flow Stats Retention Period

Operators can change the retention period by creating new System Properties. Follow the steps below to create new System Properties for high resolution and low resolution retention periods in days and months.

High Resolution Retention Period

High resolution flow stats retention can be configured anywhere between 1 and 90 days. Follow the steps below to create a new System Property for the high resolution retention period.

- 1 From the SD-WAN Orchestrator navigation panel, click **System Properties**.
- 2 In the **System Properties** screen, click the **New System Properties** button.

- 3 In the **New System Property** dialog box:
 - a Type `retention.highResFlows.days` in the **Name** text field.
 - b In the **Data Type** drop-down menu, choose **Number**.
 - c In the **Value** text field, enter the retention period in number of days.

Note High resolution retention period has a maximum of 90 days, and the resolution is 5 minutes.

New System Property...

Name:

Data Type:

Value:

Value is Password: ☐ Yes — ☒ No

Value is Read-only: ☐ Yes — ☒ No

Description:

Save Close

- 4 Click **Save**.

Low Resolution Retention Period

The low resolution flow stats can be configured to persist anywhere between 1 and 365 days. Follow the steps below to create a new System Property for the low resolution retention period.

- 1 From the SD-WAN Orchestrator navigation panel, click **System Properties**.
- 2 In the **System Properties** screen, click the **New System Properties** button.
- 3 In the **New System Property** dialog box:
 - a In the **Name** text field, type `retention.lowResFlows.months`
 - b In the **Data Type** drop-down menu, choose **Number**.

- c In the **Value** text field, enter the retention period in number of months.

Note The low resolution retention period has a maximum of 1 year, and the resolution is 1 hour.

New System Property...

Name:

Data Type:

Value:

Value is Password: ☐ Yes — ☒ No

Value is Read-only: ☐ Yes — ☒ No

Description:

Save **Close**

- 4 Click **Save**.

Changing the Flow Stats Query Interval

If Operators would like to view flow statistics for more than two weeks, the following system property must be enabled, `session.options.maxFlowstatsRetentionDays`. See the section below for instructions on how to enable this system property.

- 1 From the SD-WAN Orchestrator navigation panel, click **System Properties**.
- 2 In the **System Properties** screen, click the **New System Properties** button.
- 3 In the **New System Property** dialog box:
 - a In the **Name** text field, type `session.options.maxFlowstatsRetentionDays`
 - b In the **Data Type** drop-down menu, choose **Number**.
 - c In the **Value** text field, enter the retention period in number of days.

Modify System Property...

Name:

Data Type:

Value:

Value is Password: ☐ Yes — ☒ No

Value is Read-only: ☐ Yes — ☒ No

Description:

4 Click **Save**.

Monitor Network Services

You can view the details of configured network services for an enterprise from the **Monitor > Network Services** page in the Enterprise portal.

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

- **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, L7 health status, Number of profiles and Edges that use the Non SD-WAN Destination, Last contacted date and time, and Number of related state change Events.
- **Non SD-WAN Destinations via Edge** - Displays the configured Non SD-WAN Destinations along with 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.

- **BGP Gateway Neighbor State** – Displays the BGP neighbors connected to Gateways. The page displays the following details: Gateway name, IPv4 and IPv6 addresses of the BGP 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.
- **BGP Edge Neighbor State** – Displays the BGP neighbors connected to Edges. The page displays the following details: Edge name, IPv4 and IPv6 addresses 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.
- **Cloud Security Service Sites** – Displays the Cloud Security Services configured for the Enterprise along with the other configuration details such as Name, Type, IP address, Status of the Cloud Security Service, L7 health status, Status of the Edge using the Cloud Security Service, Date and Time of the status change, Number of related state change Events, and Deployment status.
- **Edge VNFs** – Displays the configured Edge VNFs along with other configuration details such as Name of the VNF Service, Number of Edges that use the VNF, and VM status.
- **Edge Clusters** – Displays the configured Edge clusters and the usage data along with other configuration details such as Name of the Edge cluster, Edges available in the cluster, Percentage of CPU and Memory utilization, Number of tunnels, Flow count, and Number of handoff queue drops.

Monitor

Network Overview

Edges

Network Services

Routing

Alerts

Events

Firewall Logs


Reports

Configure

Test & Troubleshoot

Administration

Network Services



Non SD-WAN Destinations via Gateway

	Name	Public IP	Status	Tunnel Status	Service Status	Used By	Last Contact	Events
1	NVS-CSR Cisco ISR	12.1.1.100 12.1.1.101	●	→ 1 → 1	N/A	4 Profiles 20 Edges	Sat Jul 10, 00:06:50	
2	NVS-CSR-Seg1 Cisco ISR	12.1.1.102 12.1.1.103	●	→ 1 → 1	N/A	4 Profiles 20 Edges	Sat Jul 10, 00:06:50	
3	NVS-CSR-Seg2 Cisco ISR	12.1.1.104 12.1.1.105	●	→ 1 → 1	N/A	4 Profiles 20 Edges	Sat Jul 10, 00:06:50	

Non SD-WAN Destinations via Edge

	Name	Public IP	Tunnel Status	Used By	Last Contact	Deployment Status
1	NVS-Edge Generic IKEv1 Router(Router Based VPN)	12.1.1.50	→ 1	1 Profile 0		N/A
2	NVS-EDGE4 Generic IKEv1 Router(Router Based VPN)	12.1.1.204 12.1.1.205	→ 2	2 Profiles 0		N/A

BGP Gateway Neighbor State Deletes Auto refresh Paused

	Gateway	Neighbor IP	State	State Changed Time	Msg Rec...	Msg Sent	Events	Up/Down	Prefix Rec...
<input type="checkbox"/>	gateway-5	169.254.0.89	● ESTABLISH...	Wed Jul 07, 19:09:18 2 days ago	20024	19072	36 View	2d04h58m	
<input type="checkbox"/>	gateway-5	169.254.0.85	● ESTABLISH...	Fri Jul 02, 13:43:57 7 days ago	23069	22234	34 View	2d12h45m	
<input type="checkbox"/>	Gateway3	12.1.1.105	REMOVED	Fri Jul 02, 13:42:33 7 days ago	0	0	5 View		
<input type="checkbox"/>	Gateway3	12.1.1.104	REMOVED	Fri Jul 02, 13:42:33 7 days ago	0	0	4 View		
<input type="checkbox"/>	Gateway3	12.1.1.103	REMOVED	Fri Jul 02, 13:42:33 7 days ago	0	0	4 View		
<input type="checkbox"/>	Gatewayv3	12.1.1.102							

BGP Edge Neighbor State Auto refresh Paused

	Edge Name	Neighbor IP	State	State Changed Time	Msg Rec...	Msg Sent	Events	Up/Down	Prefix Rec...
<input type="checkbox"/>	b1-hub2	fd00:172:21...	● ESTABLISH...	Fri Jul 09, 23:27:11 40 minutes ...	523	498	2 View	00:39:28	
<input type="checkbox"/>	b1-hub2	fd00:172:21...	REMOVED	Fri Jul 09, 23:18:54 an hour ago	0	1	4 View		
<input type="checkbox"/>	b1-hub2	172.21.41.1	● ESTABLISH...	Fri Jul 09, 23:16:56 an hour ago	2998	3098	120 View	00:50:30	
<input type="checkbox"/>	b1-hub2	10.1.1.25	● CONNECT	Fri Jul 09, 23:16:56 an hour ago	0	0	120 View	never	
<input type="checkbox"/>	b1-hub3	172.21.1.10	● IDLE	Fri Jul 09, 23:16:56 an hour ago	1885	5287	120 View	never	
<input type="checkbox"/>	b1-hub1	172.21.41.1							

Cloud Security Service Sites

	Name	Public IP	Status	Tunnel Status	Service Status	State Changed Time	Events	Deployment Status
1	Region1-Spoke	12.1.1.201 12.1.1.200	●	→ 1 → 1 → 2	4	Sat Jul 10, 00:04:53 2 minutes ...	480 View	N/A

Edge VNFs

	Service	Used By	Edge VM Status
1	CPM Check Point Firewall	1 Edge View	● Powered On (Insertion Enabled) 1 Edge

Monitor Routing

The Routing feature (**Monitor > Routing > Multicast** tab) displays Multicast Group and Multicast Edge information.

Routing

Multicast

Search... Refresh

Display 3 items.

Segment	Multicast Group	Source Address	RP	Multicast Edges	Created	Last Update
Global Segment	225.1.1.1	*	3.3.3.3	2 Edges	23 days ago	4 days ago
Global Segment	225.1.1.1	10.7.0.25	3.3.3.3	1 Edge	23 days ago	23 days ago
Global Segment	224.1.1.1	10.2.0.25	NA	1 Edge	4 days ago	4 days ago

Multicast Group Details - Global Segment / 225.1.1.1 / *

Search... Refresh

Display 2 items.

Multicast Edges	Upstream	Downstream
1 EDGE-2 View PIM Neighbors	GE6	GE6 br-network1
2 EDGE7 View PIM Neighbors	EDGE-3	br-network1 EDGE-3

PIM Neighbors View

The following figure shows the PIM neighbors of the selected Edge (per segment), the interface where the PIM neighbor was discovered, the neighbor IP address, and time stamps.

Multicast PIM Neighbors: EDGE7

Search... Cols

Display 1 items.

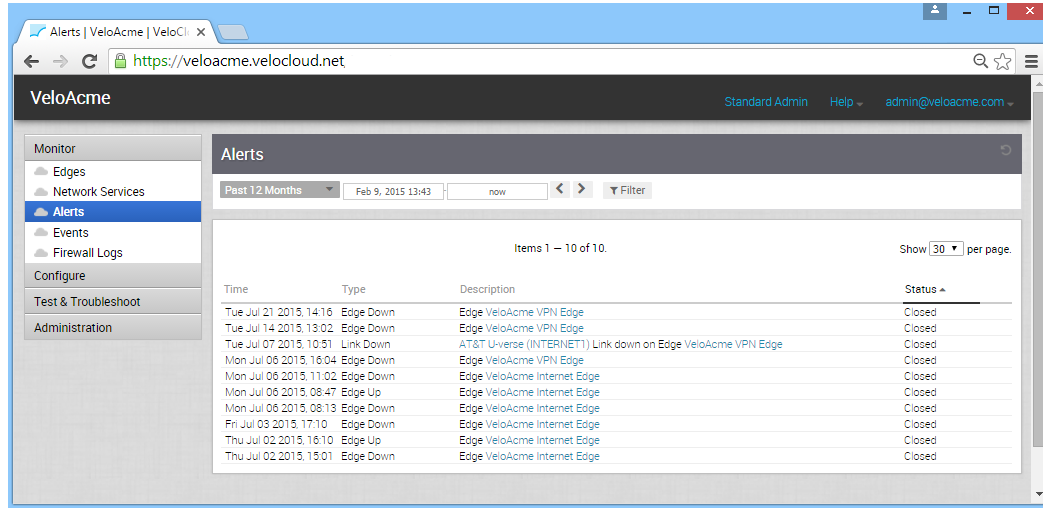
Segment	Edge Name	Interface	Address	Created	Last Update
1 Global Segment	EDGE-3	10.3.0.1	Sat Apr 07, 00:53:08	23 days ago	

Close

Monitor Alerts

SD-WAN Orchestrator provides an alert function to notify one or more Enterprise Administrators (or other support users) when a problem occurs. You can access this functionality by clicking **Alerts** under **Monitor** in the navigation panel.

You can send Alerts when a SD-WAN Edge goes offline or comes back online, a WAN link goes down, a VPN tunnel goes down, or when an Edge HA failover occurs. A delay for sending the alert after it is detected can be entered for each of the alert types. You can configure alerts in **Configure > Alerts and Notifications**.



Note If you are logged in using a user ID that has Customer Support privileges, you will only be able to view SD-WAN Orchestrator objects. You will not be able to create new objects or configure/update existing ones.

Monitor Events

The **Events** page in the navigation panel displays the events generated by the SD-WAN Orchestrator. These events can help you determine the operational status of the VMware system.

You can click the link to an Event link displayed in the **Events** page to view more details.

Time	Event	Segment	Edge	User	Severity	Message
Wed Aug 12, 17:16:23	MGD_DEVICE_C...		Az_Edge2		Warning	Inconsistent device settings d
Wed Aug 12, 17:16:23	Configuration appl...		Az_Edge2		Info	Applied new configuration for
Wed Aug 12, 17:15:58	Edge direct tunnel...	Global Segment		super@velocloud...	Info	Edge Direct Tunnels enabled
Wed Aug 12, 17:15:58	Profile updated			super@velocloud...	Info	profile [Quick Start Profile] ed
Wed Aug 12, 17:15:50	VPN Tunnel state ...				Error	Tunnel to [Velo NVS] - Failed
Wed Aug 12, 17:15:50	VPN Tunnel state ...				Notice	Tunnel to [Velo NVS] - Failed
Wed Aug 12, 17:15:24	MGD_DEVICE_C...		Az_Edge2		Warning	Inconsistent device settings d
Wed Aug 12, 17:15:24	Configuration appl...		Az_Edge2		Info	Applied new configuration for
Wed Aug 12, 17:15:24	Configuration appl...		Az_Edge2		Info	Applied new configuration for
Wed Aug 12, 17:15:05	Profile updated			super@velocloud...	Info	profile [Quick Start Profile] ed

The Events feature is useful for obtaining the following information:

- Audit trail of user activity [filter by user]
- Historical record of activity at a given site [filter by site]
- Record of outages and significant network events [filter by event]
- Analysis of degraded ISP performance [filter by time period]

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 SD-WAN 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.

Time	Event	Segment	Edge	User	Severity	Message
Wed Feb 13, 23:55:02	Bad configuration...		11-edge-102		Critical	Rolling back from device acti
Wed Feb 13, 23:55:11	Edge state received		11-edge-102		Warning	Error setting flow state: 10000

Monitor Reports

The Monitoring dashboard in the Enterprise portal allows to generate reports with overall network summary along with information on SD-WAN traffic and transport distribution. The reports enable the analysis of your network.

Note The reports focus on descriptive analytics and cannot be used for troubleshooting purposes. In addition, these reports are not dashboards that reflect the real-time data from the network.

In the Enterprise portal, click **Monitor > Reports**.

To create a new report:

- 1 In the **Reports** window, click **New Report**.
- 2 In the **New Report** window, enter a descriptive name for the report and choose the start and end dates.
- 3 Click **Create**.

New Report

Please note that reporting is limited in this release:

1. Reports have a 14-day maximum time frame
2. Reports will time out after 3 hours
3. Reports are not allowed with more than 600 Edges
4. Reporting will retain up to 10 reports in total

* Name:

* Start Date:

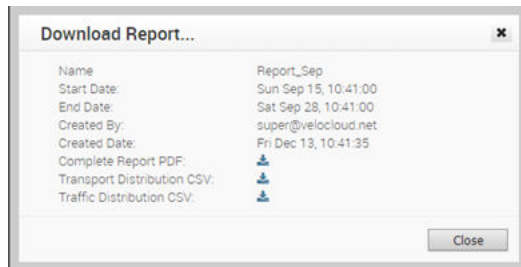
* End Date:

Note You can generate a report only for a duration of 14 days and for a maximum of 600 Edges. The report generation times out after 3 hours. The **Reports** table retains only the latest 10 reports at a time.

The **Status** of the report generation is displayed in the window. Once completed, you can download the report by clicking the **Completed** link.

Status	Name	Created By	Created Date	Report Period
In progress	Report_Oct	super@velocloud.net	Fri Dec 13, 10:43:42	Tue Oct 01, 10:42:00 - Mon Oct 14, 10:42:00
Completed	Report_Sep	super@velocloud.net	Fri Dec 13, 10:41:35	Sun Sep 15, 10:41:00 - Sat Sep 28, 10:41:00
Completed	Report_Aug	super@velocloud.net	Fri Dec 13, 10:41:09	Thu Aug 01, 10:40:00 - Wed Aug 14, 10:40:00

The **Download Report** window provides the following options:



You can download the report as a PDF that provides an overall summary of the traffic and transport distribution, represented as a pie chart. This report also provides the list of top 10 applications by the traffic and transport type.

You can choose to download the reports by transport or traffic distribution, as a CSV file.

- The transport distribution report displays the details of time, transport type, applications, name and description of the edges, and the bytes sent and received.
- The traffic distribution report displays the details of time, flow path, applications, name and description of the edges, and the bytes sent and received.

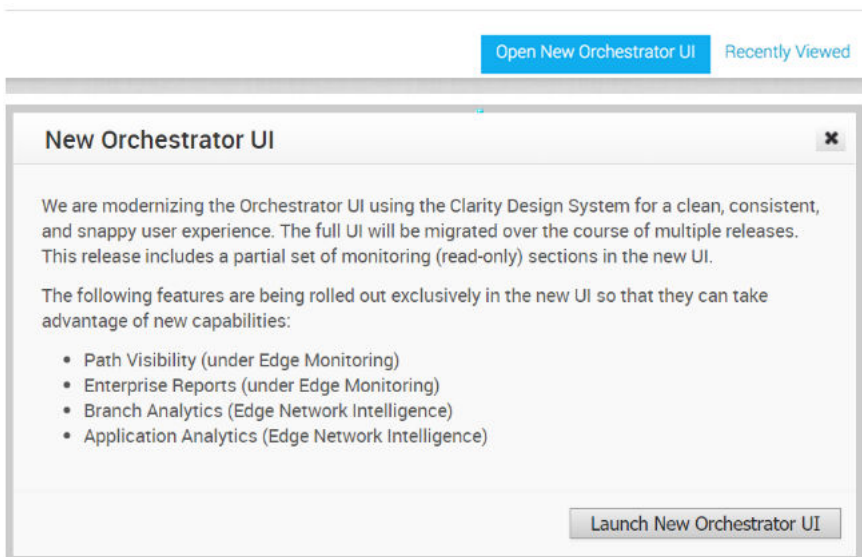
Monitor Enterprise using New Orchestrator UI

7

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

To access the new portal:

- 1 In the Enterprise portal, click **Open New Orchestrator UI**.



- 2 Click **Launch New Orchestrator UI** in the pop-up window.
- 3 The UI opens in a new tab displaying the monitoring and configuring options.
- 4 You can explore each monitoring option and click 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. Click the Filter Icon to filter the view by a specific criterion.
- **Column** – Click and select the columns to be shown or hidden in the view.
- **Refresh** – Click to refresh the details displayed with the most current data.

Read the following topics next:

- [Monitor Network Overview](#)

- [Monitor Edges](#)
- [Monitor Network Services](#)
- [Monitor Routing Details](#)
- [Monitor Alerts](#)
- [Monitor Events](#)
- [Enterprise Reports](#)
- [View Analytics Data](#)

Monitor Network Overview

The Network Overview displays the overall summary of the network like activated Edges, links, top applications, and other configured data.

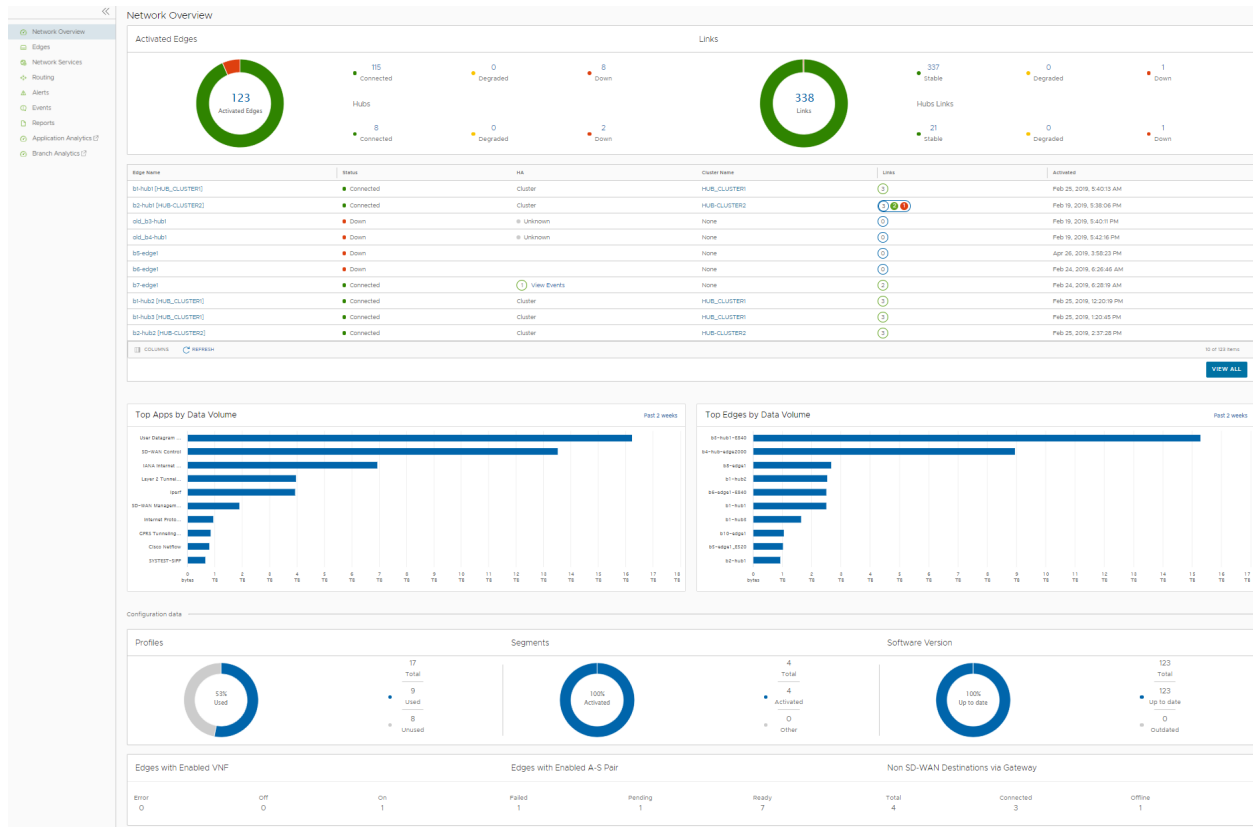
To view the Network Overview summary:

Procedure

- 1 In the Enterprise portal, click the **Open New Orchestrator UI** option available at the top of the Window.
- 2 Click **Launch New Orchestrator UI** in the pop-up window. The UI opens in a new tab.
- 3 Click **Monitor > Network Overview**.

Results

The **Network Overview** page displays the summary of the network.



The window displays the following details:

Option	Description
Activated Edges	<p>Displays the number of Edges and Hubs that are connected, degraded, and down, along with a graphical representation. Click the link to a number and details of the corresponding Edges or Hubs are displayed in the bottom panel.</p> <p>In the bottom panel, click the link to the Edge or the cluster name to navigate to the corresponding tabs.</p>
Links	<p>Displays the number of links and hub links that are stable, degraded, and down, along with a graphical representation. Click the link to a number and details of the corresponding links or Hub links are displayed in the bottom panel.</p> <p>In the bottom panel, click the link to the Hub name to navigate to the corresponding tab.</p>
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	Displays the details of used and unused profiles.
Segments	Displays the details of activated and other segments.

Option	Description
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 enabled with VNF, that are with status Error, Off, and On.
Edges with Enabled A-S Pair	Displays the number of Edges enabled 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 more details.

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:

- 1 In the Enterprise portal, click the **Open New Orchestrator UI** option available at the top of the Window.
- 2 Click **Launch New Orchestrator UI** in the pop-up window. The UI opens in a new tab.
- 3 Click **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.

Name	Status	HA	Links	VNF VM Status	VNF Type	Gateways	Last Contact
b1-hub1 [HUB_CLUSTER1]	Connected	Cluster	3			View	Dec 8, 2020, 10:28:46 PM
b2-hub1 [HUB_CLUSTER2]	Connected	Cluster	3			View	Dec 8, 2020, 10:28:41 PM
b7-edge1	Connected	1 View Events	2			View	Dec 8, 2020, 10:28:41 PM
b1-hub2 [HUB_CLUSTER1]	Connected	Cluster	3			View	Dec 8, 2020, 10:28:46 PM
b1-hub3 [HUB_CLUSTER1]	Connected	Cluster	3			View	Dec 8, 2020, 10:28:39 PM
b2-hub2 [HUB_CLUSTER2]	Connected	Cluster	2			View	Dec 8, 2020, 10:28:39 PM
b2-hub3 [CLUSTER3]	Connected	Cluster	2			View	Dec 8, 2020, 10:28:42 PM
b8-edge1	Connected	Unknown	4			View	Dec 8, 2020, 10:28:38 PM
b10-edge1	Connected	Standby ready	1			View	Dec 8, 2020, 10:28:40 PM
spoke-1-1-7	Connected		3			View	Dec 8, 2020, 10:28:47 PM
spoke-1-1-2	Connected		3			View	Dec 8, 2020, 10:28:27 PM

You can use the **Search** option to view specific Edges. Click the Filter Icon in the Search option to define a criterion and view the Edge details filtered by Edge Name, Status, Created Date, Serial Number, Custom Info, and so on.

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

You can click the link to **View** option in the **Gateways** column to view the details of Gateways connected to the corresponding Edge.

At the bottom of the page, you can click the following options to modify the view:

- **Columns:** Allows to select the columns to be shown or hidden in the view.
- **Refresh:** Displays the details in the page with the most current data.

Click the link to an Edge to view the details pertaining to the selected Edge. Click 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 list of metrics parameters. You can choose the metrics from the list to view the corresponding data. The following table lists the 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 SD-WAN 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.

For each Edge, you can view the following details:

- [Monitor overview of an Edge](#)
- [Monitor QoE](#)
- [Monitor Links of an Edge](#)
- [Monitor Path 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 click the **Shortcuts** option at the top to perform the following activities:

- **Configure** – Navigates to the Configuration tab of the selected Edge. See [Configure Edges with new Orchestrator UI](#).
- **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 with new Orchestrator UI](#).
- **Generate Diagnostic Bundle** – Allows to generate Diagnostic Bundle for the selected Edge. See [Diagnostic Bundles with new Orchestrator UI](#).
- **Remote Actions** – Allows to perform the Remote actions for the selected Edge. See [Perform Remote Actions with new Orchestrator UI](#).
- **View Profile** – Navigates to the Profile page, that is associated with the selected Edge.
- **View Gateways** – Displays the Gateways connected to the selected Edge.

Monitor overview of an Edge

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

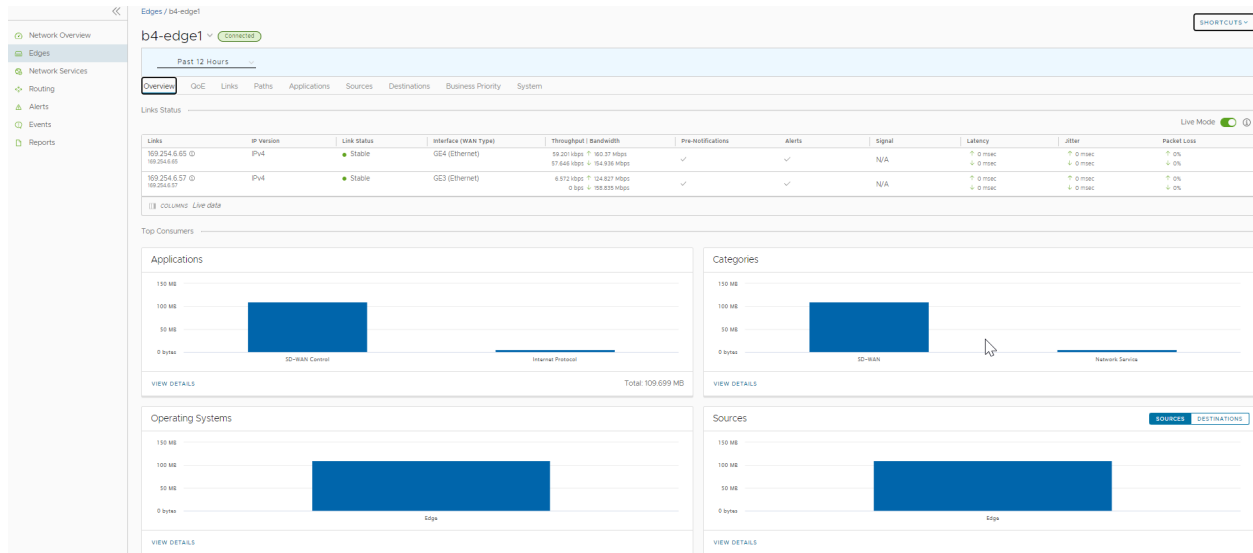
To view the information of an Edge:

Procedure

- 1 In the Enterprise portal, click the **Open New Orchestrator UI** option available at the top of the Window.
- 2 Click **Launch New Orchestrator UI** in the pop-up window. The UI opens in a new tab.
- 3 Click **Monitor > Edges** to view the Edges associated with the Enterprise.
- 4 Click the link to an Edge and the **Overview** tab is displayed by default.

Results

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



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, WAN Interface, Throughput, Bandwidth, Signal, Latency, Jitter, and Packet Loss. For more 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. Click **View Details** in each panel to navigate to the corresponding tab and view more details.

Hover the mouse on the graphs to view more 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.

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.

To view the QoE report of an Edge:

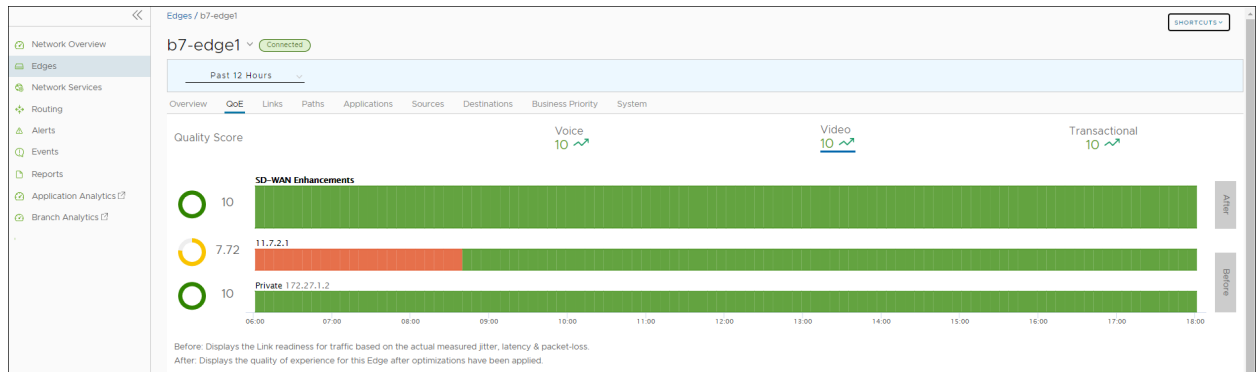
Procedure

- 1 In the Enterprise portal, click the **Open New Orchestrator UI** option available at the top of the Window.
- 2 Click **Launch New Orchestrator UI** in the pop-up window. The UI opens in a new tab.
- 3 Click **Monitor > Edges** to view the Edges associated with the Enterprise.

4 Click the link to an Edge and click the **QoE** tab.

Results

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



The following traffic types are supported: Voice, Video, and Transactional. Click the link to a traffic type displayed at the top, to view the corresponding data.

The QoE graphs display the quality scores of the selected Edge before and after the SD-WAN optimization.

You can hover the mouse on a WAN network link or an aggregate link to display a summary of Latency, Jitter, and Packet Loss.

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.

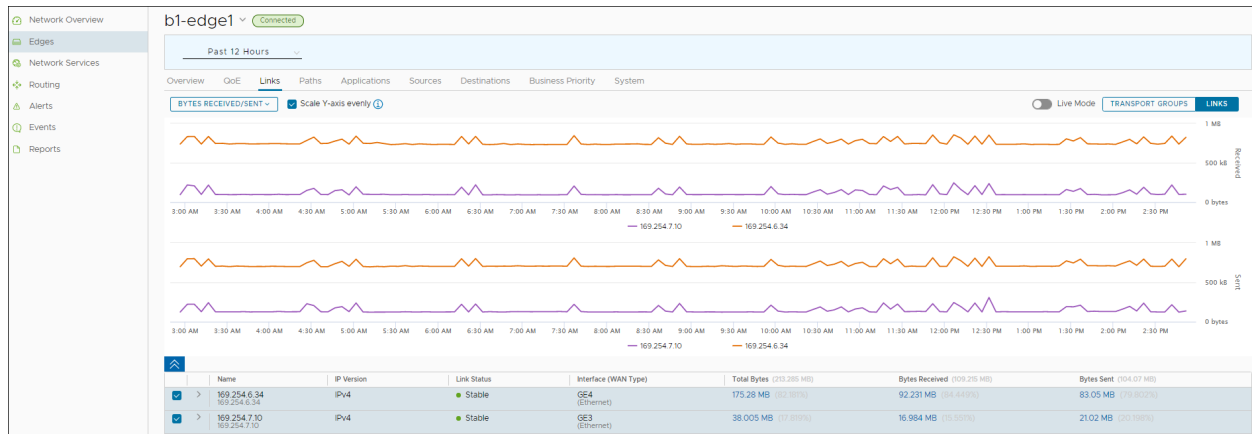
To view the details of Links and Transport groups used by the traffic:

Procedure

- 1 In the Enterprise portal, click the **Open New Orchestrator UI** option available at the top of the Window.
- 2 Click **Launch New Orchestrator UI** in the pop-up window. The UI opens in a new tab.
- 3 Click **Monitor > Edges** to view the Edges associated with the Enterprise.
- 4 Click the link to an Edge and click the **Links** tab.

Results

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

By default, the **Scale Y-axis evenly** checkbox is enabled. 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 more details.

Click **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. For more information on the metrics parameters, see [Monitor Edges](#).

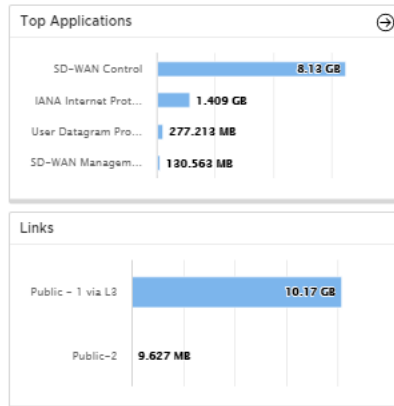

The bottom panel displays the details of the selected metrics for the links or the transport groups. You can view the details of a maximum of 4 links at a time.

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

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

Transport Groups by Total Bytes

Jul 1, 2020, 6:15:42 AM - Jul 2, 2020, 6:15:42 PM



CLOSE

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

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 SD-WAN peers. SD-WAN Orchestrator enables an Enterprise user to monitor the Path visibility using the monitoring dashboard.

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

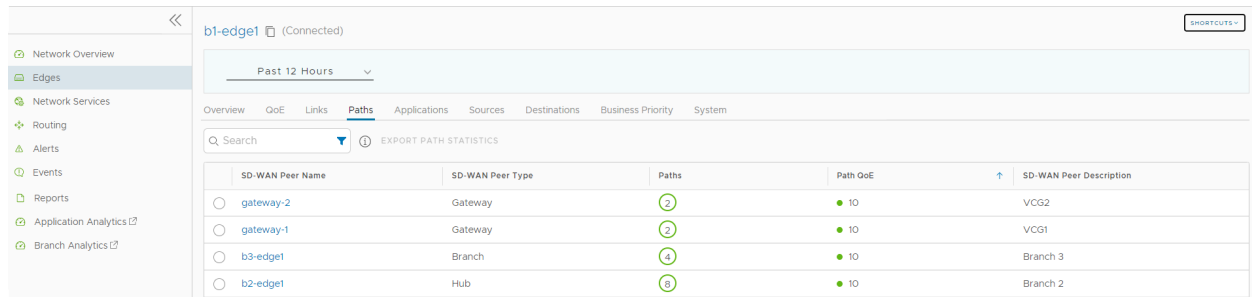
Procedure

- 1 In the Enterprise portal, click the **Open New Orchestrator UI** option available at the top of the Window.
- 2 Click **Launch New Orchestrator UI** in the pop-up window. The UI opens in a new tab.
- 3 Click **Monitor > Edges** to view the Edges associated with the Enterprise.
- 4 Click the link to an Edge and click the **Paths** tab.

Results

For the selected Edge, the **Paths** tab displays the details of 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.

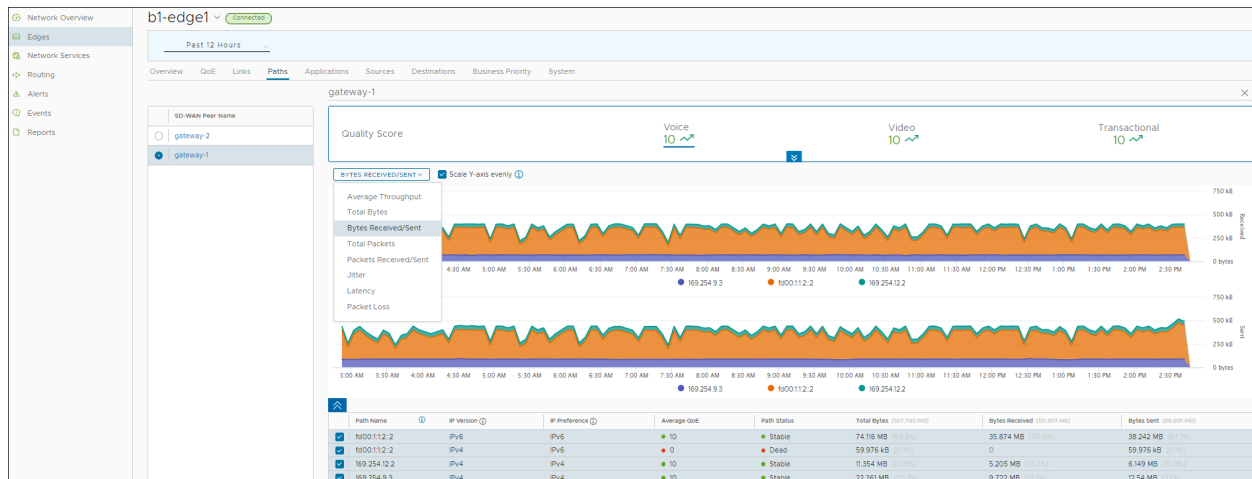


At the top of the page, you can choose a specific time period to view the details of the priorities for the selected duration.

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

Click the link to an SD-WAN peer to view the corresponding Path details as follows:

- All the 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 video, voice, transactional traffic
- Time series data for each path by metrics like: Throughput, Latency, Packet loss, Jitter, and so on. For more information on the parameters, see [Monitor Edges](#).

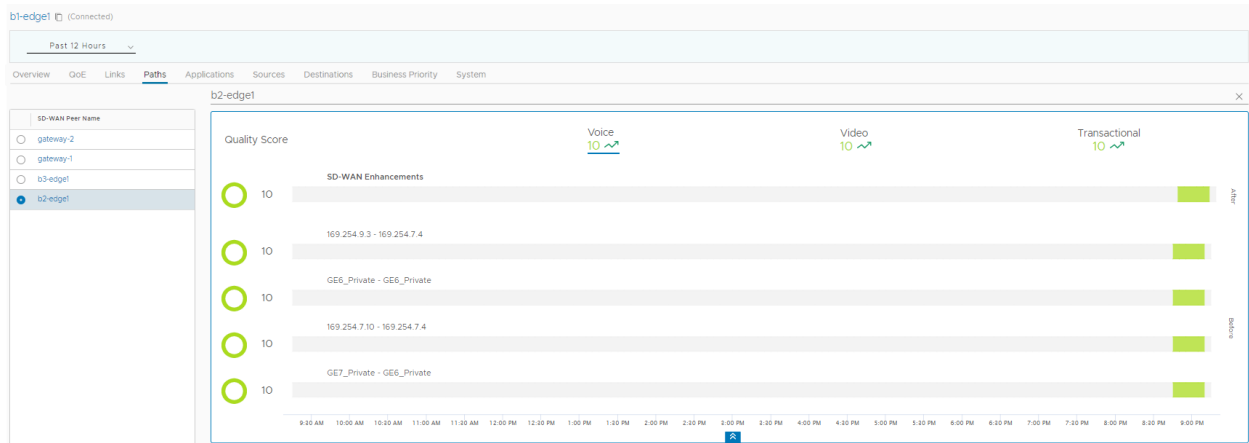


The metrics time-series data is displayed in graphical format. You can select and view the details of a maximum of 4 paths at a time.

Hover the mouse on the graphs to view more details.

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

Click the DOWN arrow in the **Quality Score** pane at the top, to view the Path score by the traffic types.



You can click an SD-WAN peer displayed at the left pane to view the corresponding Path details.

Monitor Edge Applications

You can monitor the network usage of applications or application categories used by a specific Edge.

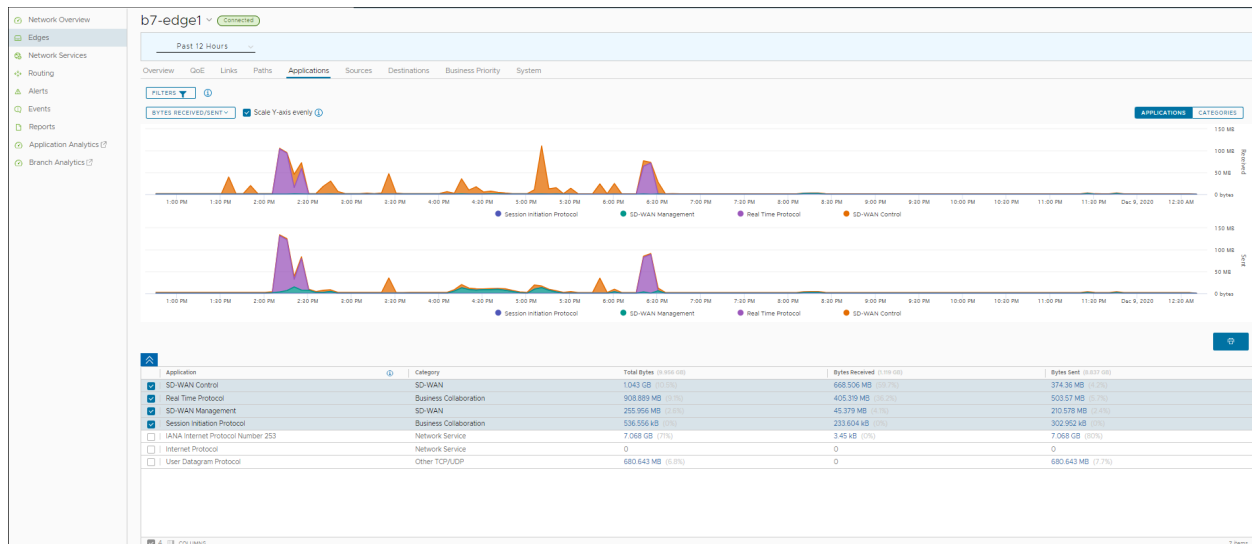
To view the details of applications or application categories:

Procedure

- 1 In the Enterprise portal, click the **Open New Orchestrator UI** option available at the top of the Window.
- 2 Click **Launch New Orchestrator UI** in the pop-up window. The UI opens in a new tab.
- 3 Click **Monitor > Edges** to view the Edges associated with the Enterprise.
- 4 Click the link to an Edge and click the **Applications** tab.

Results

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

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

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

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

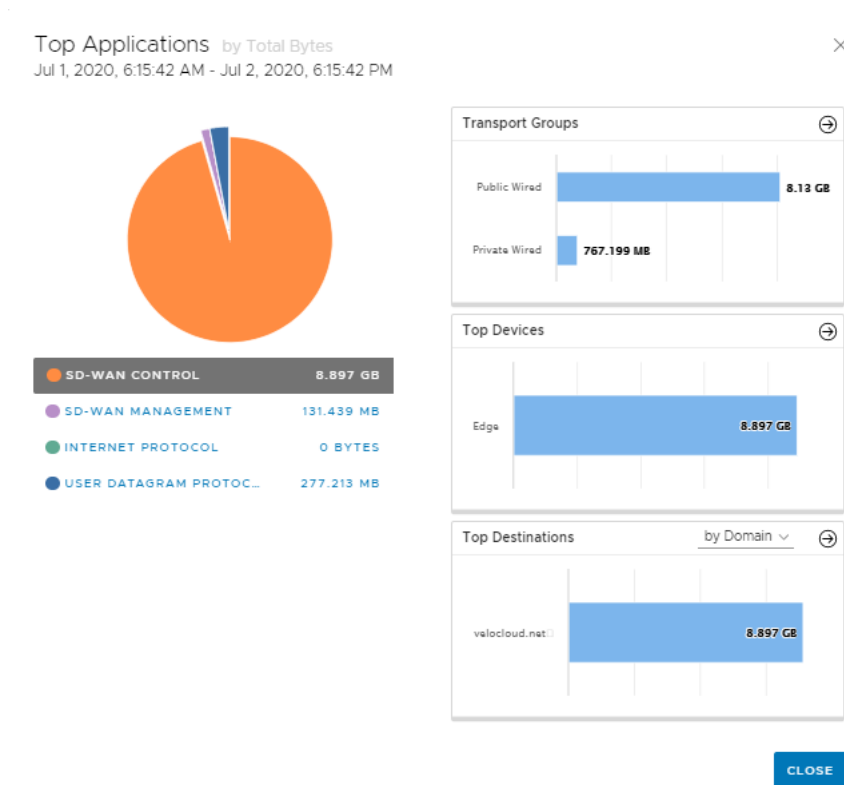
Hover the mouse on the graphs to view more details.

Choose the metrics from the drop-down to view the details related to the selected parameter. For more information on the metrics parameters, see [Monitor Edges](#).

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. Click **Columns** to select the columns to be shown or hidden in the view.

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

The following image shows a detailed report of top applications.



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

Monitor Edge Sources

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

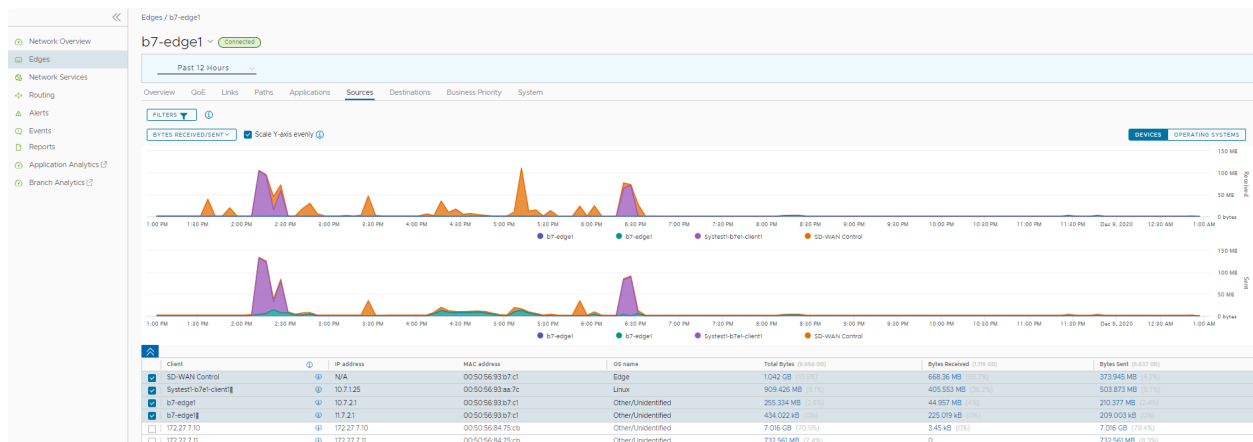
To view the details of devices and operating systems:

Procedure

- 1 In the Enterprise portal, click the **Open New Orchestrator UI** option available at the top of the Window.
- 2 Click **Launch New Orchestrator UI** in the pop-up window. The UI opens in a new tab.
- 3 Click **Monitor > Edges** to view the Edges associated with the Enterprise.
- 4 Click the link to an Edge and click the **Sources** tab.

Results

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

By default, the **Scale Y-axis evenly** checkbox is enabled. 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 more details.

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

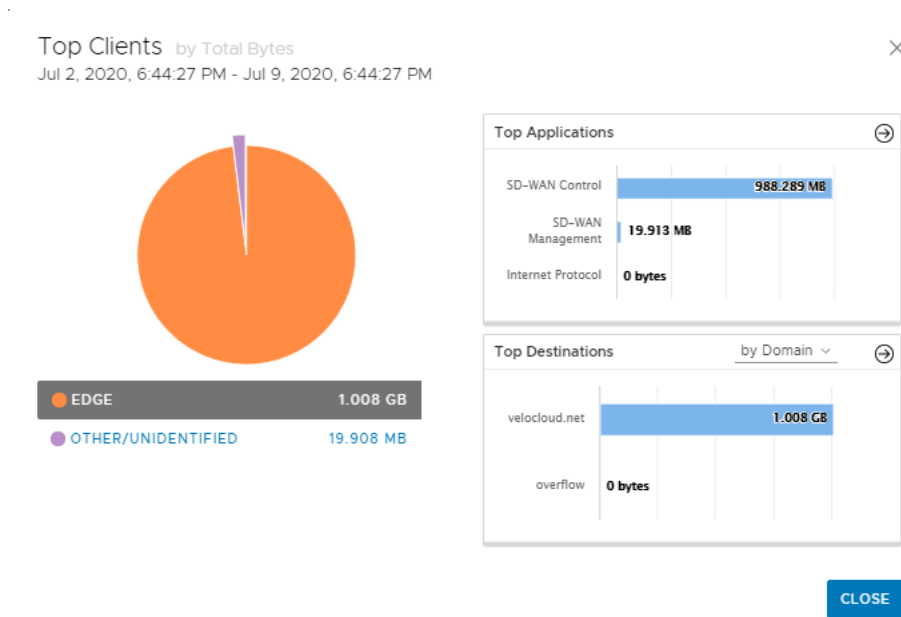
Click **Operating Systems** to view the report based on the Operating Systems used in the devices.

Choose the metrics from the drop-down to view the details related to the selected parameter. For more information on the metrics parameters, see [Monitor Edges](#).

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. Click **Columns** to select the columns to be shown or hidden in the view.

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

The following image shows a detailed report of top clients.



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

Monitor Edge Destinations

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

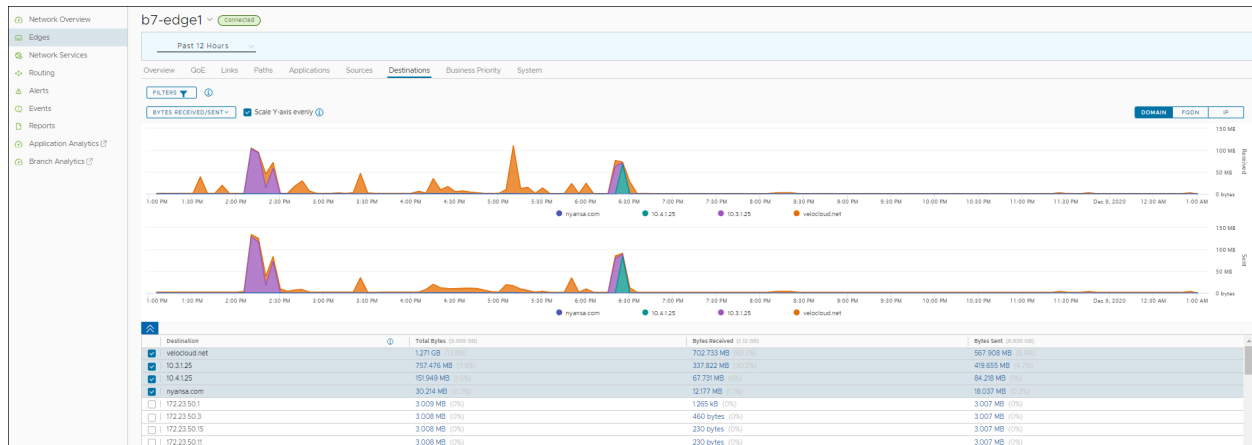
To view the details of destinations:

Procedure

- 1 In the Enterprise portal, click the **Open New Orchestrator UI** option available at the top of the Window.
- 2 Click **Launch New Orchestrator UI** in the pop-up window. The UI opens in a new tab.
- 3 Click **Monitor > Edges** to view the Edges associated with the Enterprise.
- 4 Click the link to an Edge and click the **Destinations** tab.

Results

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

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

By default, the **Scale Y-axis evenly** checkbox is enabled. 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. Click the relevant type to view the corresponding information.

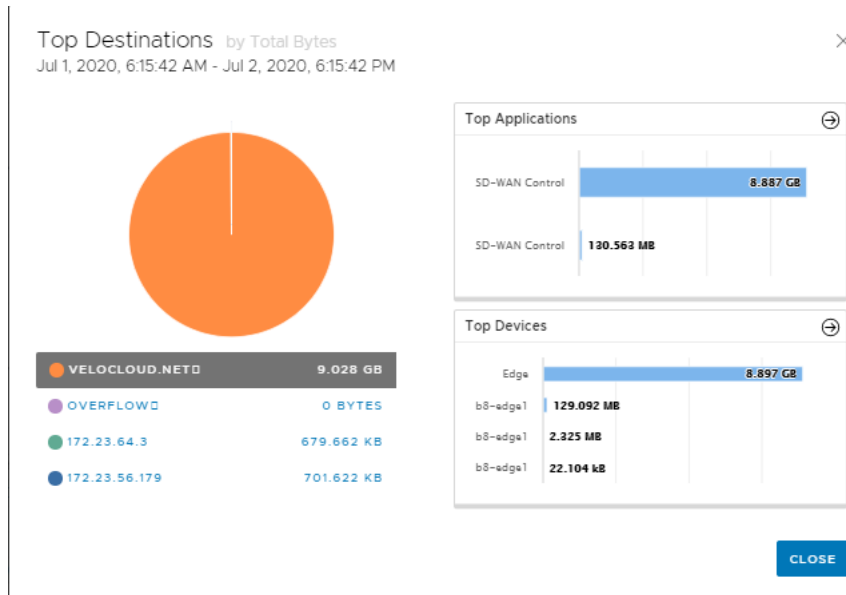
Hover the mouse on the graphs to view more details.

Choose the metrics from the drop-down to view the details related to the selected parameter. For more information on the metrics parameters, see [Monitor Edges](#).

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. Click **Columns** to select the columns to be shown or hidden in the view.

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

The following image shows a detailed report of top destinations.



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

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.

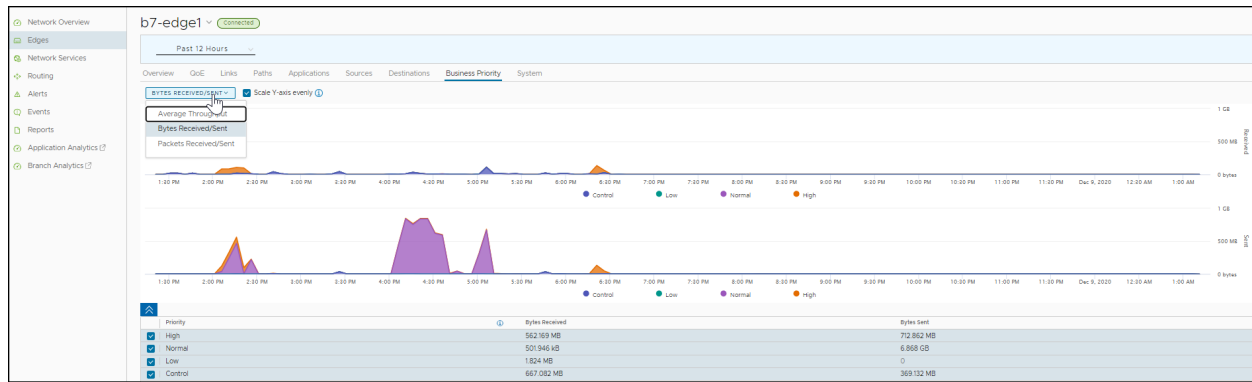
To view the details of business priorities of the network traffic:

Procedure

- 1 In the Enterprise portal, click the **Open New Orchestrator UI** option available at the top of the Window.
- 2 Click **Launch New Orchestrator UI** in the pop-up window. The UI opens in a new tab.
- 3 Click **Monitor > Edges** to view the Edges associated with the Enterprise.
- 4 Click the link to an Edge and click the **Business Priority** tab.

Results

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

Choose the metrics from the drop-down to view the details related to the selected parameter. For more information on the metrics parameters, see [Monitor Edges](#).

By default, the **Scale Y-axis evenly** checkbox is enabled. 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 more details.

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

Monitor System Information of an Edge

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

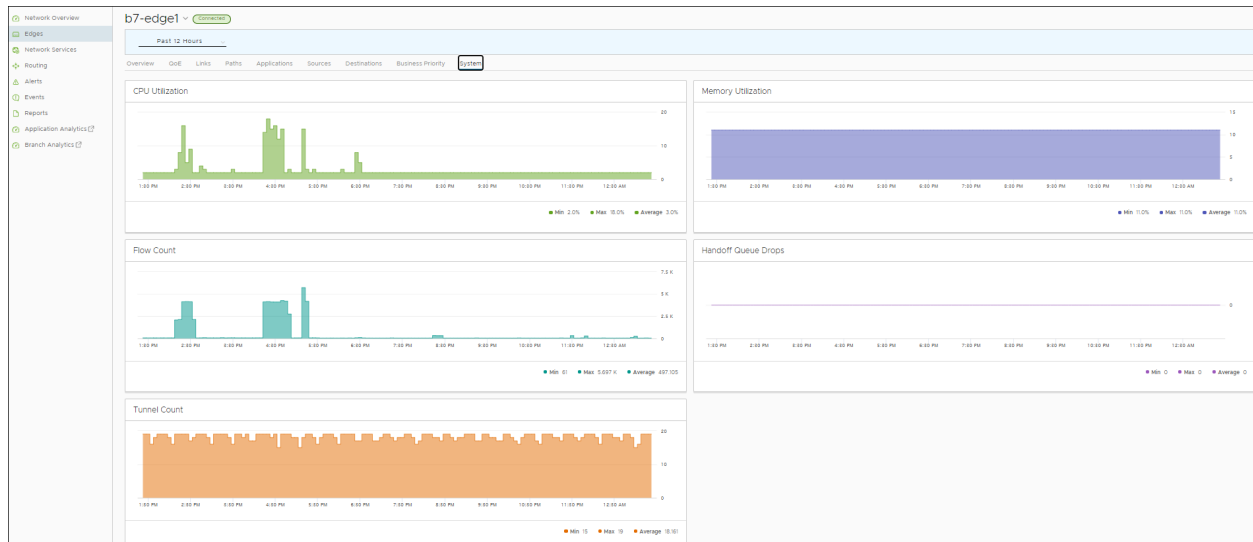
To view the details of system information:

Procedure

- 1 In the Enterprise portal, click the **Open New Orchestrator UI** option available at the top of the Window.
- 2 Click **Launch New Orchestrator UI** in the pop-up window. The UI opens in a new tab.
- 3 Click **Monitor > Edges** to view the Edges associated with the Enterprise.
- 4 Click the link to an Edge and click the **System** tab.

Results

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



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.
- **Memory Utilization**– Percentage of usage of memory.
- **Flow Count** – Count of traffic flow.
- **Handoff Queue Drops** – Total number of packets dropped from a queue since the Edge was last rebooted. Occasional drops are expected, usually caused by a large burst of traffic. However, a consistent increase in handoff queue drops usually indicates an Edge capacity issue.
- **Tunnel Count** – Count of tunnel sessions.

Hover the mouse on the graphs to view more details.

Monitor Network Services

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

To view the details of network services:

- 1 In the Enterprise portal, click the **Open New Orchestrator UI** option available at the top of the Window.
- 2 Click **Launch New Orchestrator UI** in the pop-up window. The UI opens in a new tab.
- 3 Click **Monitor > Network Services**.

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

- [Monitor Non SD-WAN Destinations through Gateway](#)
- [Monitor Cloud Security Service Sites](#)
- [Monitor Edge Clusters](#)

■ Monitor Edge VNFs

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:

Procedure

- 1 In the Enterprise portal, click the **Open New Orchestrator UI** option available at the top of the Window.
- 2 Click **Launch New Orchestrator UI** in the pop-up window. The UI opens in a new tab.
- 3 Click **Monitor > Network Services**. The **Non SD-WAN Destinations via Gateway** tab is displayed.

Results

The **Non SD-WAN Destinations via Gateway** tab displays the details of already configured Non SD-WAN Destinations. To configure the Non SD-WAN Destinations via Gateway, see [Configure a Non SD-WAN Destination](#).

Network Overview

Edges

Network Services

Routing

Alerts

Events

Reports

Application Analytics

Branch Analytics

Network Services

Non SD-WAN Destinations via Gateway

Cloud Security Service Sites

Edge Clusters

Edge VNFs

	Name	Public IP	Status	Tunnel Status	Used By	Last Contact
	NVS-CSR <small>Cloud IP</small>	12.11.100 12.11.101	Connected	Connected Connected	15 Profiles 26 Edges	Dec 9, 2020, 1:22:58 AM
	NVS-CSR-Seg1 <small>Cloud IP</small>	12.11.102 12.11.103	Connected	Connected Disconnected	14 Profiles 122 Edges	Dec 9, 2020, 1:22:58 AM
	NVS-CSR-seg2 <small>Cloud IP</small>	12.11.104 12.11.105	Connected	Connected Disconnected	14 Profiles 122 Edges	Dec 9, 2020, 1:22:58 AM

Columns

4 Items

General

IKE/IPSec Configuration

Events

Name

NVS-CSR

Type

Cisco ISR

Enable Tunnel(s):

☐

Location

Lat, Lng: 37.402889, -122.116859

Site Subnets

Subnet	Description	Advertise
12.12.0/25		<input type="checkbox"/>
192.168.93.0/24		<input type="checkbox"/>
192.168.93.1/32		<input type="checkbox"/>
192.168.93.2/32		<input type="checkbox"/>
192.168.93.3/32		<input type="checkbox"/>
192.168.93.4/32		<input type="checkbox"/>
192.168.93.5/32		<input type="checkbox"/>

Primary VPN Gateway

Public IP:

12.11.100

Secondary VPN Gateway

Public IP:

12.11.101

Redundant Cloud VPN

☒

- **IKE/IPSec Configuration** – Click 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** – Click the tab to view the events related to the selected Non SD-WAN Destination. Click the arrow displayed in the first column to view more details of an event.

Monitor Cloud Security Service Sites

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

To monitor the Cloud Security Services:

Procedure

- 1 In the Enterprise portal, click the **Open New Orchestrator UI** option available at the top of the Window.
- 2 Click **Launch New Orchestrator UI** in the pop-up window. The UI opens in a new tab.
- 3 Click **Monitor > Network Services > Cloud Security Service Sites**.

Results

The **Cloud Security Service Sites** tab displays the already configured Cloud Security Services. To configure a Cloud Security Service, see [Cloud Security Services](#).

The page displays the following details: Name, Type, IP address, Status of the Cloud Security Service, Status of the Edge using the Cloud Security Service, Date and Time of the status change, and the number of Events.

You can also sort the report by clicking 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.

Click a Cloud Security Service to view the related Events along with the IP address and State, in the bottom panel.

Monitor Edge Clusters

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

To view the details of Edge clusters:

Procedure

- 1 In the Enterprise portal, click the **Open New Orchestrator UI** option available at the top of the Window.
- 2 Click **Launch New Orchestrator UI** in the pop-up window. The UI opens in a new tab.
- 3 Click **Monitor > Network Services > Edge Clusters**.

Results

The **Edge Clusters** tab displays the details of already configured Edge clusters. To configure the clusters, see [Configure Edge Clustering](#).

🏠 Network Overview	Non SD-WAN Destinations via Gateway		Cloud Security Service Sites		Edge Clusters		Edge VNFs	
🏠 Edges								
🏠 Network Services								
⚙️ Routing	Cluster Name	Edges	CPU Utilization	Memory Utilization	# Tunnels	Flow Count	# Handoff Queue Drops	
🚨 Alerts	HUB_CLUSTER1	b1-hub1	4.00%	13.00%	251	220	1975	
		b1-hub2	4.00%	13.00%	182	2837	7921156	
		b1-hub3	5.00%	13.00%	214	5062	20901588	
📊 Events	HUB-CLUSTER2	b2-hub1	2.00%	14.00%	627	111	-	
📄 Reports		b2-hub2	-	13.00%	12	99	-	
		b2-hub3	-	13.00%	12	114	-	
📊 Application Analytics ⓘ								
📊 Branch Analytics ⓘ								

The page displays the following details: Name of the Edge cluster, Edges available in the cluster, percentage of CPU and Memory utilization, Number of tunnels, Flow count, and Number of handoff queue drops.

Monitor Edge VNFs

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

To view the Edge VNFs:

Procedure

- 1 In the Enterprise portal, click the **Open New Orchestrator UI** option available at the top of the Window.
- 2 Click **Launch New Orchestrator UI** in the pop-up window. The UI opens in a new tab.
- 3 Click **Monitor > Network Services > Edge VNFs**.

Results

The **Edge VNFs** tab displays the details of already configured VNFs. To configure VNF on an Edge, see [Security VNFs](#).

Network Overview

Edges

Network Services

Routing

Alerts

Events

Reports

Application Analytics

Branch Analytics

Non SD-WAN Destinations via Gateway

Cloud Security Service Sites

Edge Clusters

Edge VNFs

Service	Used By	Edge VM Status
<div><div></div><div>CPM</div><div>Check Point Security Firewall</div></div>	1 Edge	<div><div></div>Powered On (Insertion Enabled) 1 Edge</div>

COLUMNS

1 Items

VNF Edge Deployments

Edge Name	Edge VM Status
b6-edge1-EB40	<div><div></div>Powered On (Insertion Enabled)</div>

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

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

Monitor Routing Details

You can view the routing services configured in the Enterprise.

To view the details of configured routing services:

- 1 In the Enterprise portal, click the **Open New Orchestrator UI** option available at the top of the Window.
- 2 Click **Launch New Orchestrator UI** in the pop-up window. The UI opens in a new tab.
- 3 Click **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](#)

Monitor Multicast Groups

You can view the multicast groups configured for the Enterprise.

To view the multicast groups:

Procedure

- 1 In the Enterprise portal, click the **Open New Orchestrator UI** option available at the top of the Window.
- 2 Click **Launch New Orchestrator UI** in the pop-up window. The UI opens in a new tab.
- 3 Click **Monitor > Routing**. The **Multicast Groups** tab is displayed.

Results

The **Multicast Groups** displays the details of already configured multicast group settings. To configure multicast groups, see [Configure Multicast Settings](#).

The screenshot shows the 'Routing' section with the 'Multicast Groups' tab selected. The table lists several multicast groups, each with a segment, address, RP, and edges. Below the table, the 'Multicast Group Members' section shows the upstream and downstream edges for a selected group.

Segment	Multicast Group	Source Address	RP	Multicast Edges	Created	Last Update
Global Segment	224.0.1.40	*	1.1.1.1	2 Edges	7 days ago	10 hours ago
Global Segment	224.1.1.2	*	1.1.1.1	3 Edges	10 months ago	9 months ago
Global Segment	224.1.1.1	*	1.1.1.1	9 Edges	10 months ago	10 hours ago
Global Segment	227.1.1.1	*	1.4.1.1	2 Edges	10 months ago	9 months ago
Global Segment	227.1.1.9	*	1.4.1.1	2 Edges	10 months ago	9 months ago
Global Segment	227.1.1.10	*	1.4.1.1	2 Edges	10 months ago	9 months ago

Multicast Edges	Upstream	Downstream
b1-hub2 View PIM Neighbors	GE6	b8-edge1
b8-edge1 View PIM Neighbors	b1-hub2 - 1	

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.

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

Monitor PIM Neighbors

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

To view the PIM neighbors:

Procedure

- 1 In the Enterprise portal, click the **Open New Orchestrator UI** option available at the top of the Window.
- 2 Click **Launch New Orchestrator UI** in the pop-up window. The UI opens in a new tab.
- 3 Click **Monitor > Routing > PIM Neighbors**.

Results

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

Network Overview

Edges

Network Services

Routing

Alerts

Events

Reports

Application Analytics

Branch Analytics

Multicast Groups

PIM Neighbors

BGP Edge Neighbor State

BFD

BGP Gateway Neighbor State

Q Search

Edge Name

TEST

spoke-1-9-9

spoke-1-9-8

spoke-1-9-7

spoke-1-9-6

spoke-1-9-5

REFRESH

1 - 20 of 123 items

PIM Neighbors

Segment	Edge Name	Interface	Address	Created	Last Update
Global Segment	b1-hub1		10.1.1.1	Dec 8, 2020, 10:04:17 AM	Dec 8, 2020, 4:18:31 PM
Global Segment	b4-hub-edge2000		10.4.1.1	Dec 8, 2020, 5:50:20 PM	
Global Segment	b1-hub2		10.1.2.1	Dec 8, 2020, 2:46:29 PM	Dec 8, 2020, 4:17:23 PM

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.

Monitor BGP Edge Neighbor State

You can view the details BGP neighbors connected to Edges.

To view the BGP neighbors connected to Edges:

Procedure

- 1 In the Enterprise portal, click the **Open New Orchestrator UI** option available at the top of the Window.
- 2 Click **Launch New Orchestrator UI** in the pop-up window. The UI opens in a new tab.
- 3 Click **Monitor > Routing > BGP Edge Neighbor State**.

Results

The **BGP Edge Neighbor State** tab displays the Edges connected as BGP neighbors, when you have configured BGP settings on the Edges.

Network Overview

Edges

Network Services

Routing

Alerts

Events

Firewall Logs

Reports

Routing

Multicast Groups

PIM Neighbors

BGP Edge Neighbor State

BFD

BGP Gateway Neighbor State

Q Search

Edge Name

b10-edge1

b10-edge1

b10-edge1

b10-edge1

b1-hub2

b1-hub1

b1-hub2

b1-hub1

b1-hub2

b1-hub1

b1-hub2

b1-hub2

REFRESH

218 items

Edge Name	IPv4 Address	IPv6 Address	State	State Change Time	# Msg Rcvd	# Msg Sent	Up/Down	# Prefixes Rcvd
b10-edge1	172.30.11.1		Established	Jul 2, 2021, 1:42:23 PM 7 days ago	0	0	120	0
b10-edge1	172.30.11.10		Established	Jul 2, 2021, 1:42:23 PM 7 days ago	0	0	2	0
b10-edge1	172.30.12.1		Established	Jul 8, 2021, 2:04:00 AM 2 days ago	3,074	2,851	120	102
b10-edge1	192.168.5.25		Connect	Nov 11, 2020, 11:40:06 PM 8 months ago	0	0	Never	0
b1-hub2	fe00::172:21:1:1		Established	Jul 9, 2021, 8:52:47 PM 4 hours ago	949	958	2	01:19:02
b1-hub1	fe00::172:21:1:1		Established	Jul 9, 2021, 5:26:28 PM 7 hours ago	1,561	1,598	8	02:12:11
b1-hub2	fe00::172:21:4:1		Active	Jul 9, 2021, 8:17:06 PM 4 hours ago	0	0	Never	0
b1-hub1	fe00::172:21:4:1		Established	Jul 9, 2021, 5:26:28 PM 7 hours ago	7,753	7,942	4	02:12:11
b1-hub2	fe00::172:21:7:2		Established	Jul 9, 2021, 11:58:54 PM 1 hour ago	0	1	4	0
b1-hub1	fe00::172:21:7:2		Established	Jul 9, 2021, 10:22:36 PM 2 hours ago	1,609	1,597	10	02:12:09
b1-hub2	fe00::172:21:7:20		Established	Jul 9, 2021, 11:27:11 PM 1 hour ago	864	828	2	01:07:55

Related State Change Events

State Change Time	State	# Msg Received	# Msg Sent	Up/Down	# Prefixes Received
Jul 9, 2021, 8:52:47 PM 4 hours ago	Established	949	958	01:19:02	6
Jul 9, 2021, 8:17:06 PM 4 hours ago	Active	0	0	Never	0

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.

Click 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 click the Filter Icon next to the **Search** option to filter the details by Edge Name, Neighbor IP, Neighbor IP Type, and Status.

Monitor BFD

You can view the BFD sessions on Edges and Gateways.

To view the BFD sessions:

Procedure

- 1 In the Enterprise portal, click the **Open New Orchestrator UI** option available at the top of the Window.
- 2 Click **Launch New Orchestrator UI** in the pop-up window. The UI opens in a new tab.
- 3 Click **Monitor > Routing > BFD**.

Results

The **BFD** tab displays the details of already configured BFD sessions. To configure BFD, see [Configure BFD](#).

<ul style="list-style-type: none"> Network Overview Edges Network Services Routing Alerts Events Reports Application Analytics Branch Analytics 	Multicast Groups PIM Neighbors BGP Edge Neighbor State BFD BGP Gateway Neighbor State																									
	Edge BFD Sessions <input type="text" value="Search"/>																									
	Edge	Segment	Peer Address	Local Address	State	Remote Timers	Local Timers	Events																		
	b1-hub3	Global Segment	11.99.1	172.21.120	Down	rx: 300ms / tx: 300ms	rx: 300ms / tx: 300ms	110 View																		
	b1-hub2	Global Segment	11.99.1	172.21.110	Down	rx: 300ms / tx: 300ms	rx: 300ms / tx: 300ms	104 View																		
	b1-hub1	Global Segment	11.99.1	172.21.12	Down	rx: 300ms / tx: 300ms	rx: 300ms / tx: 300ms	120 View 12 minute(s), 29 second(s)																		
	b4-hub-edge2000	Global Segment	1.4.1.1	1.4.1.100	Down	rx: 1000ms / tx: 1000ms	rx: 300ms / tx: 300ms	87 View 22 second(s)																		
	b4-hub-edge2000	segment1	1.4.1.1	1.4.1.100	Down	rx: 1000ms / tx: 1000ms	rx: 300ms / tx: 300ms	33 View 15 minute(s), 10 second(s)																		
	b4-hub-edge2000	segment2	1.4.1.2.1	1.4.1.102	Down	rx: 1000ms / tx: 1000ms	rx: 300ms / tx: 300ms	120 View 21 hour(s), 56 minute(s), 55 second(s)																		
	b9-edge1_E540	Global Segment	1.9.1.1	1.9.1.100	Down	rx: 1000ms / tx: 1000ms	rx: 300ms / tx: 300ms	120 View 1 day(s), 14 hour(s), 44 minute(s), 33 second(s)																		
	b1-hub2	Global Segment	172.21.11	172.21.110	Down	rx: 1000ms / tx: 1000ms	rx: 300ms / tx: 300ms	120 View																		
Columns REFRESH 31 Items																										
Gateway BFD Sessions <input type="text" value="Search"/>																										
<table> <tr> <th>Gateway</th><th>Segment</th><th>Peer Address</th><th>Local Address</th><th>State</th><th>Remote Timers</th><th>Local Timers</th><th>Events</th><th>Session Time</th></tr> <tr> <td colspan="9">No BFD events available for selected enterprise</td></tr> </table>									Gateway	Segment	Peer Address	Local Address	State	Remote Timers	Local Timers	Events	Session Time	No BFD events available for selected enterprise								
Gateway	Segment	Peer Address	Local Address	State	Remote Timers	Local Timers	Events	Session Time																		
No BFD events available for selected enterprise																										

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.

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

Monitor BGP Gateway Neighbor State

You can view the details BGP neighbors connected to Gateways.

To view the BGP neighbors connected to Gateways:

Procedure

- 1 In the Enterprise portal, click the **Open New Orchestrator UI** option available at the top of the Window.
- 2 Click **Launch New Orchestrator UI** in the pop-up window. The UI opens in a new tab.
- 3 Click **Monitor > Routing > BGP Gateway Neighbor State**.

Results

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

Routing								
<div> <div>Network Overview</div> <div>Edges</div> <div>Network Services</div> <div>Routing</div> <div>Alerts</div> <div>Events</div> <div>Firewall Logs</div> <div>Reports</div> </div>								
<div> <div>Multicast Groups</div> <div>PIM Neighbors</div> <div>BGP Edge Neighbor State</div> <div>BFD</div> <div>BGP Gateway Neighbor State</div> </div>								
<div> <div>Q Search</div> <div>0</div> </div>								
Gateway	Neighbor IP	State	State changed time	# Msg Received	# Msg Sent	# Events	Up/Down	# Prefix Received
Gateway3	12.11.100	●	Jul 2, 2021, 1:42:33 PM 7 days ago	0	0	6		0
Gateway3	12.11.101	●	Jul 2, 2021, 1:42:33 PM 7 days ago	0	0	6		0
Gateway3	12.11.102	●	Jul 2, 2021, 1:42:33 PM 7 days ago	0	0	5		0
Gateway3	12.11.103	●	Jul 2, 2021, 1:42:33 PM 7 days ago	0	0	7		0
Gateway3	12.11.104	●	Jul 2, 2021, 1:42:33 PM 7 days ago	0	0	7		0
Gateway3	12.11.105	●	Jul 2, 2021, 1:42:33 PM 7 days ago	0	0	9		0
Gateway3	169.254.0.26	● Established	Apr 12, 2021, 2:37:38 AM 3 months ago	23,178	22,100	20	2d13h19m	0
gateway-5	169.254.0.30	● Established	Jun 9, 2021, 10:36:30 PM 1 month ago	23,217	22,100	32	2d13h19m	0
Gateway3	169.254.0.34	● Established	Apr 12, 2021, 2:37:38 AM 3 months ago	23,244	22,138	16	2d13h19m	0
gateway-5	169.254.0.38	● Established	Jun 9, 2021, 10:36:30 PM 1 month ago	23,244	22,100	32	2d13h19m	0
<div> <div>Columns</div> <div>Refresh</div> </div>								
<div> <div>Related State Change Events</div> <div> <div>State Changed Time</div> <div>State</div> <div># Msg Received</div> <div># Msg Sent</div> <div>Up/Down</div> <div># Prefix Received</div> </div> </div>								
<div> <div>Mar 31, 2021, 5:36:04 AM 3 months ago</div> <div>● Established</div> <div>4,456</div> <div>4,259</div> <div>11:45:04</div> <div>0</div> </div>								
<div> <div>Mar 31, 2021, 5:35:04 AM 3 months ago</div> <div>● Connect</div> <div>14,044</div> <div>13,335</div> <div>00:00:05</div> <div>0</div> </div>								
<div> <div>Mar 31, 2021, 11:01:53 PM 3 months ago</div> <div>● Established</div> <div>150</div> <div>142</div> <div>00:23:29</div> <div>0</div> </div>								
<div> <div>Mar 31, 2021, 10:51:25 PM 3 months ago</div> <div>● Active</div> <div>4,458</div> <div>4,265</div> <div>00:10:49</div> <div>0</div> </div>								
<div> <div>Mar 31, 2021, 10:50:26 PM 3 months ago</div> <div>● Connect</div> <div>4,458</div> <div>4,265</div> <div>00:00:19</div> <div>0</div> </div>								
<div> <div>Mar 30, 2021, 9:54:52 PM 3 months ago</div> <div>● Established</div> <div>14,035</div> <div>13,329</div> <div>03:13:12</div> <div>0</div> </div>								
<div> <div>Mar 30, 2021, 9:53:52 PM 3 months ago</div> <div>● Idle</div> <div>11,116</div> <div>10,562</div> <div>00:00:01</div> <div>0</div> </div>								

The page displays the following details: Gateway name, IP address of the BGP 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.

Click 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.

Monitor Alerts

SD-WAN Orchestrator allows to configure alerts that notify the Enterprise Administrators or other support users, whenever an event occurs.

Prerequisites

Ensure that you have configured the relevant alerts, along with the notification delay, in **Configure > Alerts & Notifications**. See [Chapter 24 Configure Alerts](#).

Procedure

- 1 In the Enterprise portal, click the **Open New Orchestrator UI** option available at the top of the Window.

2 Click **Launch New Orchestrator UI** in the pop-up window. The UI opens in a new tab.

3 Click **Monitor > Alerts**.

Results

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

Trigger Time	Notification Time	Category	Type	Description	Status
> Dec 8, 2020, 5:51:26 PM	Dec 8, 2020, 5:53:33 PM	Customer	Edge CSS Tunnel Down	11.4.2.1 Edge CSS Tunnel Down on Edge b4-hub-edge2000	Closed
> Dec 8, 2020, 5:50:15 PM	Dec 8, 2020, 5:51:23 PM	Customer	Edge HA Failover	Edge b4-hub-edge2000	Closed
> Dec 8, 2020, 5:11:59 PM	Dec 8, 2020, 5:14:03 PM	Customer	Edge CSS Tunnel Down	11.4.2.1 Edge CSS Tunnel Down on Edge b4-hub-edge2000	Closed
> Dec 8, 2020, 5:11:00 PM	Dec 8, 2020, 5:12:03 PM	Customer	Link Up	Private (SFP) Link Up on Edge b3-hub-E840	Closed
> Dec 8, 2020, 5:10:45 PM	Pending	Customer	Link Down	11.4.2.1 (GE4) Link Down on Edge b4-hub-edge2000	Closed
> Dec 8, 2020, 5:10:15 PM	Dec 8, 2020, 5:11:23 PM	Customer	Edge HA Failover	Edge b4-hub-edge2000	Closed
> Dec 8, 2020, 5:09:39 PM	Dec 8, 2020, 5:11:42 PM	Customer	Edge CSS Tunnel Down	11.2.1.1 Edge CSS Tunnel Down on Edge b3-hub-E840	Closed
> Dec 8, 2020, 5:08:40 PM	Dec 8, 2020, 5:10:43 PM	Customer	Edge CSS Tunnel Down	11.4.2.1 Edge CSS Tunnel Down on Edge b4-hub-edge2000	Closed
> Dec 8, 2020, 5:07:30 PM	Dec 8, 2020, 5:09:33 PM	Customer	Link Down	Private (SFP) Link Down on Edge b3-hub-E840	Closed
> Dec 8, 2020, 5:07:15 PM	Pending	Customer	Link Down	11.4.2.1 (GE4) Link Down on Edge b4-hub-edge2000	Closed
> Dec 8, 2020, 5:06:41 PM	Dec 8, 2020, 5:08:43 PM	Customer	Edge CSS Tunnel Down	11.3.1.1 Edge CSS Tunnel Down on Edge b3-hub-E840	Closed
> Dec 8, 2020, 5:05:33 PM	Dec 8, 2020, 5:07:33 PM	Customer	Edge CSS Tunnel Down	11.4.2.1 Edge CSS Tunnel Down on Edge b4-hub-edge2000	Closed
> Dec 8, 2020, 5:05:05 PM	Pending	Customer	Edge CSS Tunnel Down	11.3.1.1 Edge CSS Tunnel Down on Edge b3-hub-E840	Closed
> Dec 8, 2020, 4:42:55 PM	Dec 8, 2020, 4:45:03 PM	Customer	Edge CSS Tunnel Up	11.4.2.1 Edge CSS Tunnel Up on Edge b4-hub-edge2000	Closed
> Dec 8, 2020, 4:37:55 PM	Pending	Customer	Edge CSS Tunnel Down	11.4.2.1 Edge CSS Tunnel Down on	Closed

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

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

Click 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:

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. Click 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.

Monitor Events

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

To view the Events page:

Procedure

- 1 In the Enterprise portal, click the **Open New Orchestrator UI** option available at the top of the Window.
- 2 Click **Launch New Orchestrator UI** in the pop-up window. The UI opens in a new tab.
- 3 Click **Monitor > Events**.

Results

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

Events						
Past 12 Hours						
Q Search						
Event	User	Segment	Edge	Severity	Time	Message
Edge SSH login			spoke-1-1-6	Info	Dec 9, 2020, 1:54:48 AM	sshd[20663]: Accepted keyboard-interactive/pam for root from 172.18.0.254 port 56900 ssh2
Edge SSH login			spoke-1-1-5	Info	Dec 9, 2020, 1:54:48 AM	sshd[20654]: Accepted keyboard-interactive/pam for root from 172.18.0.254 port 44956 ssh2
Edge SSH login			spoke-1-1-5	Info	Dec 9, 2020, 1:54:48 AM	sshd[20661]: Accepted keyboard-interactive/pam for root from 172.18.0.254 port 44958 ssh2
BGP session established to edge neighbor			bf-hub1	Info	Dec 9, 2020, 1:54:48 AM	BGP session established for edge [bf-hub1] to neighbor: [1.199.1]
BGP session established to edge neighbor			bf-hub1	Info	Dec 9, 2020, 1:54:48 AM	BGP session established for edge [bf-hub1] to neighbor: [172.21.11.1]
BGP session established to edge neighbor			bf-hub1	Info	Dec 9, 2020, 1:54:48 AM	BGP session established for edge [bf-hub1] to neighbor: [172.21.12.1]
EDGE_BFD_NEIGHBOR_UP			bf-hub1	Info	Dec 9, 2020, 1:54:48 AM	BFD session up for edge [bf-hub1] to peer: [1.199.1]
Edge SSH login			spoke-1-1-4	Info	Dec 9, 2020, 1:54:47 AM	sshd[20707]: Accepted keyboard-interactive/pam for root from 172.18.0.254 port 53348 ssh2
Edge SSH login			spoke-1-1-4	Info	Dec 9, 2020, 1:54:47 AM	sshd[20713]: Accepted keyboard-interactive/pam for root from 172.18.0.254 port 53350 ssh2
Edge SSH login			bf-hub2	Info	Dec 9, 2020, 1:54:46 AM	sshd[14896]: Accepted keyboard-interactive/pam for root from 10.1.2.25 port 39991 ssh2
Edge SSH login			spoke-1-1-2	Info	Dec 9, 2020, 1:54:46 AM	sshd[21167]: Accepted keyboard-interactive/pam for root from 172.18.0.254 port 34702 ssh2
Edge SSH login			spoke-1-1-2	Info	Dec 9, 2020, 1:54:46 AM	sshd[21173]: Accepted keyboard-interactive/pam for root from 172.18.0.254 port 34704 ssh2
Edge SSH login			spoke-1-1-3	Info	Dec 9, 2020, 1:54:46 AM	sshd[20663]: Accepted keyboard-interactive/pam for root from 172.18.0.254 port 43172 ssh2
Edge SSH login			spoke-1-1-3	Info	Dec 9, 2020, 1:54:46 AM	sshd[20670]: Accepted keyboard-interactive/pam for root from 172.18.0.254 port 43174 ssh2
Edge SSH login			spoke-1-1-1	Info	Dec 9, 2020, 1:54:45 AM	sshd[22275]: Accepted keyboard-interactive/pam for root from 172.18.0.254 port 40642 ssh2
Edge SSH login			spoke-1-1-1	Info	Dec 9, 2020, 1:54:45 AM	sshd[22282]: Accepted keyboard-interactive/pam for root from 172.18.0.254 port 40644 ssh2
EDGE_BFD_NEIGHBOR_DOWN			bf-hub-edge2000	Info	Dec 9, 2020, 1:54:43 AM	BFD session down for edge [bf-hub-edge2000] to peer: [1.4.1.1]

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

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

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

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

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.

Option	Description
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.

Enterprise Reports

VMware SD-WAN allows you to generate Enterprise reports that enable 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 SD-WAN 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**.

To access the Enterprise reports:

- 1 In the Enterprise portal, click the **Open New Orchestrator UI** option available at the top of the Window.
- 2 Click **Launch New Orchestrator UI** in the pop-up window. The UI opens in a new tab.
- 3 Click **Monitor > Reports**.

Note You can also create and view the Reports in the **Monitor > Reports** page in the Enterprise portal. However, it is recommended to use the New Orchestrator UI to create reports with customizable options.

In the **Reports** page, you can create a new report, customize the report, and schedule report generation for a recurring period.



For more information, see [Create a New Enterprise Report](#).

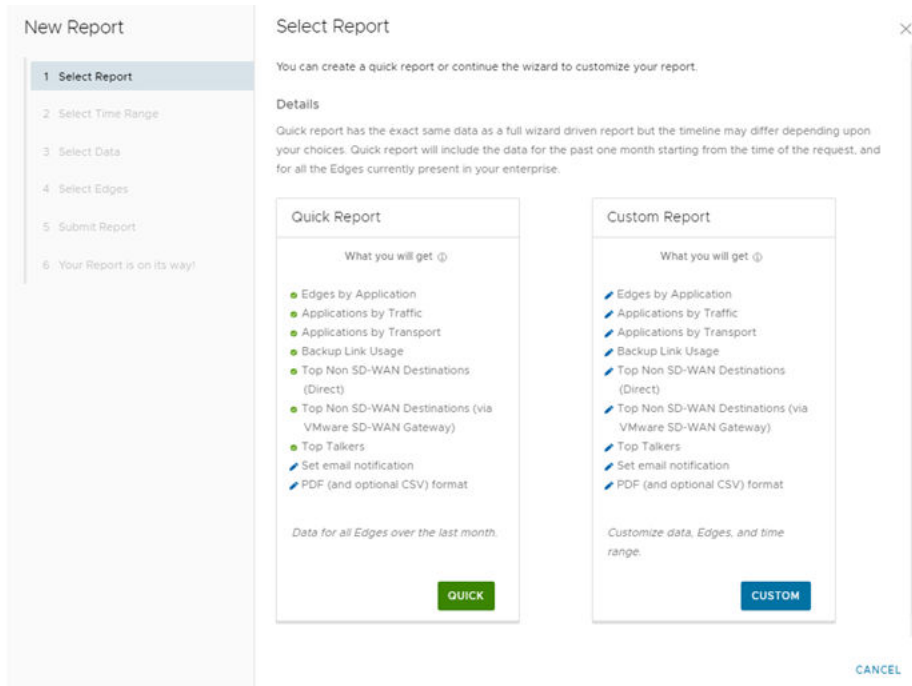
Create a New Enterprise Report

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

Procedure

- 1 In the New Orchestrator UI, click **Monitor > Reports**.
- 2 In the **Reports** page, click **New Report**.

- 3 In the **New Report** page, you can configure to generate a consolidated report or a customized report.



- 4 Click **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 VMware SD-WAN Edges with top 5 Edges for each destination.
 - Top Non SD-WAN destinations using VMware SD-WAN Gateways with top 5 Edges for each destination.
 - Top clients across Edges with top 5 applications for each client.
- 5 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](#).
- 6 In the window **Your Report is on its way** that appears, click **Done**.

Results

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

Name	Created By	Created Date	Report Period	Status
ACME Report	super@vnetcloud.net	Dec 1, 2020, 10:46:33 AM	Oct 31, 2020, 10:42:18 AM - Dec 1, 2020, 10:42:18 AM	In Progress

What to do next

Your report is generated and is displayed in the **Reports** page. See [Monitor Enterprise Reports](#).

To generate a customized report with specific values, see [Create Customized Report](#).

Create Customized Report

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

Procedure

- 1 In the New Orchestrator UI, click **Monitor > Reports**.
- 2 Click **New Report**.
- 3 In the **New Report** page, click **Custom**.

What to do next

Follow the instructions on the screen to select the configuration settings for the custom report. See [Select Time Range](#).

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.

Procedure

- 1 When you choose to customize the Enterprise report and click **Custom** in [Create Customized Report](#), the **Select Time Range** window appears.

New Report

- 1 Select Report
- 2 Select Time Range**
- 3 Select Data
- 4 Select Edges
- 5 Submit Report
- 6 Your Report is on its way!

Select Time Range [X]

Create a one-time report or schedule a recurring report. Select the time range to analyze.

Details
The report will include all data within your selected time period.

☒ Create a one-time report.
☐ Schedule a recurring report.

Time range is 05/24/2020 05:33 to 06/24/2020 05:33

OR choose from a pre-determined time range

- Past 31 Days (selected)
- Past 7 Days
- Past 2 Weeks
- Past 31 Days
- Past 6 Months
- Past 12 Months
- Not Selected

[CANCEL](#) [BACK](#) [NEXT](#)

- 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.

The screenshot shows a 'New Report' wizard with a sidebar on the left containing six steps: 1. Select Report, 2. Select Time Range (highlighted), 3. Select Data, 4. Select Edges, 5. Submit Report, and 6. Your Report is on its way!. The main panel is titled 'Select Time Range' and includes a close button (X) in the top right. Below the title, it says 'Create a one-time report or schedule a recurring report. Select the time range to analyze.' A 'Details' section follows, stating 'The report will include all data within your selected time period.' There are two radio button options: 'Create a one-time report.' (unselected) and 'Schedule a recurring report.' (selected). Under the selected option, there are three fields: 'Generate a report for the' with a dropdown menu showing 'Last Week', 'Repeat every week' (text label), and 'on' with a dropdown menu showing 'Monday'. Below these is an 'at' field with a time input '07:00' and a clock icon. At the bottom right of the panel are three buttons: 'CANCEL', 'BACK', and 'NEXT'.

- 4 Click **Next**.

What to do next

See [Select Data](#).

Select Data

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

Procedure

- 1 When you click **Next** after selecting the time range in [Select Time Range](#), the **Select Data** window appears.

New Report

- 1 Select Report
- 2 Select Time Range
- 3 **Select Data**
- 4 Select Edges
- 5 Submit Report
- 6 Your Report is on its way!

Select Data

Select the items to include in the report from the list below.

Details

Each report encapsulates unique insights into your network, click on each description title to know more.

Items	Brief Description
<input checked="" type="checkbox"/> EDGES BY APPLICATION	> Edges by Application
<input checked="" type="checkbox"/> APPLICATIONS BY TRAFFIC	> Applications by Traffic
<input checked="" type="checkbox"/> APPLICATIONS BY TRANSPORT	> Applications by Transport
<input checked="" type="checkbox"/> BACKUP LINK USAGE	> Backup Link Usage
<input checked="" type="checkbox"/> TOP NON SD-WAN DESTINATIONS (Direct)	> Top Non SD-WAN Destinations (Direct)
<input checked="" type="checkbox"/> TOP NON SD-WAN DESTINATIONS (via VMware SD-WAN Gateway)	> Top Non SD-WAN Destinations (via VMware SD-WAN Gateway)
<input checked="" type="checkbox"/> TOP TALKERS	> Top Talkers

CANCEL **BACK** **NEXT**

- 2 Select the checkboxes of the data that you want to include in the report from the following available options:
 - **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 VMware SD-WAN Edges with top 5 Edges for each destination.
 - **Top Non SD-WAN Destinations (via SD-WAN Gateway)** – List of top Non SD-WAN destinations via VMware SD-WAN Gateways with top 5 Edges for each destination.
 - **Top Talkers** – List of top clients across Edges with top 5 applications for each client.
- 3 Click **Next**.

What to do next

See [Select Edges](#).

Select Edges

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

Procedure

- 1 When you click **Next** after selecting the data to be included in the report in [Select Data](#), the **Select Edges** window appears.

- 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 click the Plus (+) Icon to include more conditions. After specifying the conditions, click **Apply** and the details of Edges selected according to the conditions are displayed at the right side.
- 4 Click **Next**.

What to do next

See [Submit Report](#).

Submit Report

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

Procedure

- 1 When you click **Quick** to create a Quick Report in [Create a New Enterprise Report](#), or click **Next** after selecting the Edges in [Select Edges](#), the **Submit Report** window appears.

- 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 checkbox 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 click **Submit**.
- 4 In the window **Your Report is on its way** that appears, click **Done**.

Results

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


What to do next

Your report is generated and is displayed in the **Reports** page. See [Monitor Enterprise Reports](#).

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 new Orchestrator UI, click **Monitor > Reports**. The page displays all the generated reports.



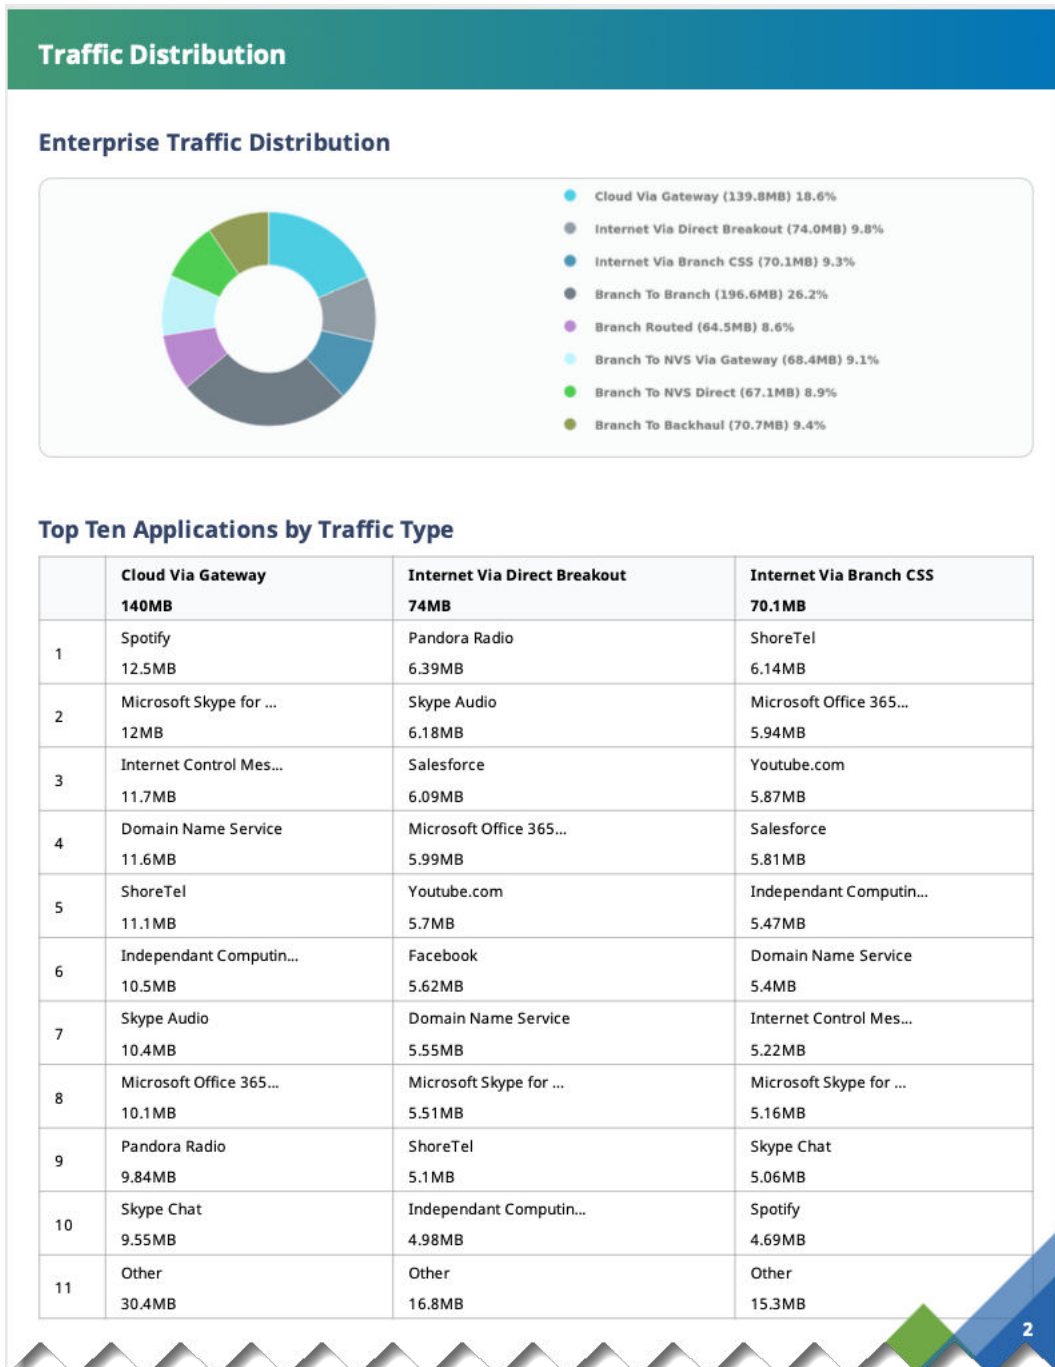
	Name	Created By	Created Date	Report Period	Status
<input checked="" type="checkbox"/>	Daily Report	super@velocloud.net	Dec 1, 2020, 3:00:00 PM	Nov 30, 2020, 2:30:00 PM - Dec 1, 2020, 2:30:00 PM	Completed
<input type="checkbox"/>	Weekly Report	super@velocloud.net	Dec 1, 2020, 2:00:00 PM	Nov 24, 2020, 1:30:00 PM - Dec 1, 2020, 1:30:00 PM	Completed
<input type="checkbox"/>	Monthly Report	super@velocloud.net	Dec 1, 2020, 12:30:00 PM	Nov 30, 2020, 12:30:00 PM - Dec 1, 2020, 12:30:00 PM	Completed
<input type="checkbox"/>	ACME Report	super@velocloud.net	Dec 1, 2020, 10:46:33 AM	Oct 31, 2020, 10:42:19 AM - Dec 1, 2020, 10:42:19 AM	Completed

To download a report, click 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:



The Enterprise Traffic distribution lists the following data:

- **Cloud Via Gateway:** Internet bound traffic that goes through the SD-WAN Gateway.
- **Internet Via Direct Breakout:** Internet bound traffic that breaks out directly from branch and does not go through VMware Tunnels.
- **Internet Via Branch CSS:** Traffic bound to Cloud Security Services directly from VMware branch.

- **Branch To Branch:** Traffic going through SD-WAN Gateway / SD-WAN Hub / dynamic SD-WAN Tunnels, directly between two VMware 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 SD-WAN 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 VMware SD-WAN 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 SD-WAN 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**:

	A	B	C	D	E	F	G	H	I	J	K	L
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 click **DELETE**.

To view the scheduled reports, click **RECURRING REPORTS**.

<input type="checkbox"/>	Name	Created By	Created Date	Recurrence	Recipients
<input type="checkbox"/>	Daily Report	super@velocloud.net	Dec 1, 2020, 12:13:03 PM	Every day at 3:00 PM	
<input type="checkbox"/>	Monthly Report	super@velocloud.net	Dec 1, 2020, 12:12:26 PM	Every month on day 1 at 12:30 PM	
<input type="checkbox"/>	Weekly Report	super@velocloud.net	Dec 1, 2020, 12:06:20 PM	Every week on Tuesday at 2:00 PM	admin@acme.com

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 click **DELETE**.

View Analytics Data

Once a SD-WAN Edge is provisioned with Analytics, the Analytics functionality collects data (application-specific Analytics or application and branch Analytics). The collected Analytics data are then sent directly from the SD-WAN Edge to the Cloud Analytics Engine. Operator Super User, Operator Standard Admin, Enterprise Super User, Enterprise Standard admin, Partner

Super User, and Partner Standard Admin can view the Analytics data for a specific customer in the Analytics portal (<https://app.nyansa.com>).

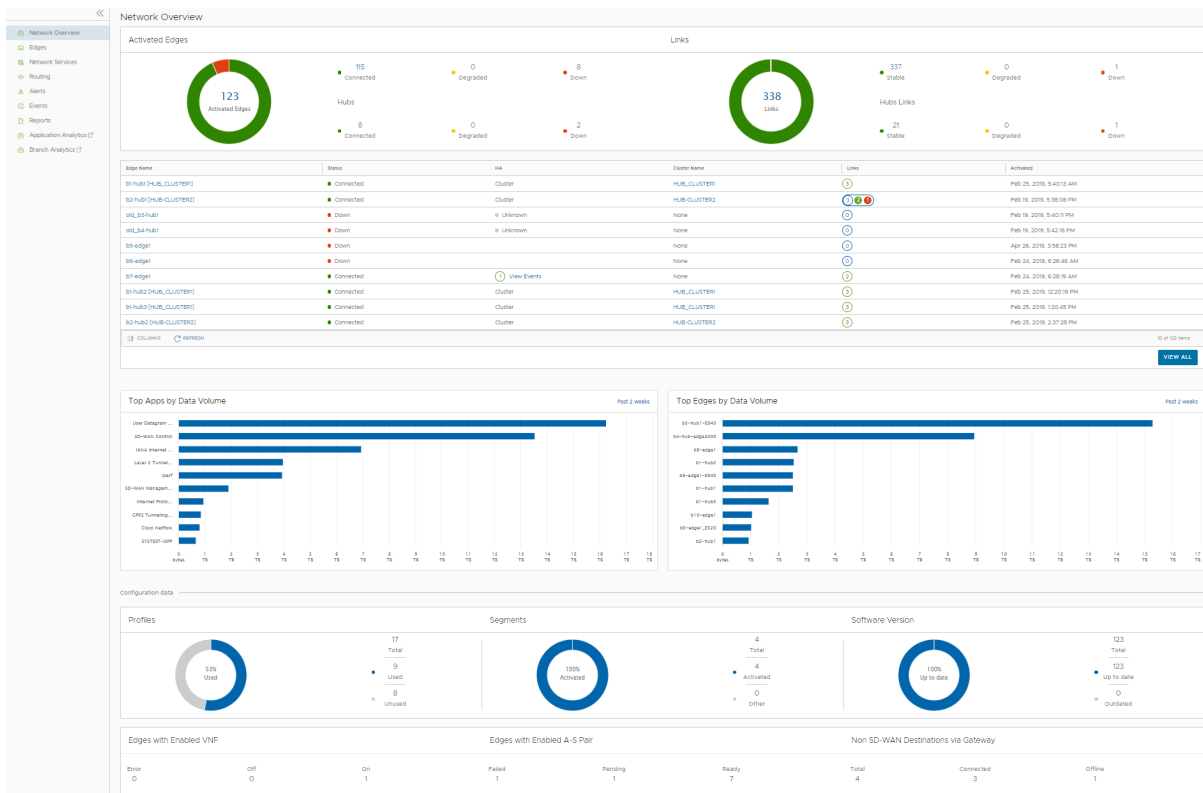
To view the Analytics data, perform the following steps.

Prerequisites

- Ensure that all the necessary system properties to enable Analytics are properly set in the SD-WAN Orchestrator. For more information, contact your Operator Super User.
- Ensure that you have access to the Analytics portal to view the Analytics data.

Procedure

- 1 In the Enterprise portal, click the **Open New Orchestrator UI** option available at the top of the Window.
- 2 Click **Launch New Orchestrator UI** in the pop-up window. The UI opens in a new tab.



- 3 To view the Application Analytics data for the selected enterprise, click **Application Analytics**.
- 4 To view Branch Analytics data, click **Branch Analytics**.

When the Analytics menu is clicked, the Analytics portal will be opened in a new browser tab, where you can view the Analytics data (Application and Branch) of all the Edges configured for a selected customer. Note that the Browser settings may prevent this action as popups. You need to allow it when browser shows notification.

What to do next

In the Analytics portal, you can configure additional data sources such as Wi-Fi and Wired metrics. For more information, see *VMware Edge Network Intelligence User Guide* available at <https://docs.vmware.com/en/VMware-Edge-Network-Intelligence/index.html>.

Configure Segments

8

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 important when traffic from different organizations and/or 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 more information about how to configure system properties, refer to the "System Properties" section in the VMware SD-WAN 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 more 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 VMware SD-WAN 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 SD-WAN 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 a new segment for an enterprise, perform the following steps:

- 1 From the SD-WAN Orchestrator navigation panel, go to **Configure > Segments**. The **Segments** page for the selected enterprise appears.

Segments						
Segment Name	Description	Type	Service VLAN	Delegate To Partner	Delegate To Customer	
Global Segment ⓘ	Default segment for traffic	Regular ⌵ ⓘ	<input type="text"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	ⓘ +
Guest	user flows hidden	Private ⌵	<input type="text"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	− +

- 2 Click the **+** button and enter the following details to configure a new segment.

Field	Description
Segment Name	The name of the segment (up to 256 characters).
Description	The description of the segment (up to 256 characters).
Type	<p>The segment type can be one of the following:</p> <ul style="list-style-type: none"> ■ Regular - The standard segment type. ■ Private - Used for traffic flows that require limited visibility in order to address end user privacy requirements. ■ CDE - VMware 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 VMware PCI certification. <p>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.</p>
Service VLAN	The service VLAN identifier. For information, see <i>Define Mapping between Segments and Service VLANs (Optional)</i> section in Security VNFs .
Delegate To Partner	By default, this checkbox is selected. If you unselect it, the Partner cannot change configs within the segment, including the interface assignment.
Delegate To Customer	By default, this checkbox is selected. If you unselect it, the Customer cannot change configs within the segment, including the interface assignment.

- 3 Click **Save Changes**.

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

- Does not upload user flow stats to Orchestrator except for VMware Control, VMware 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 in 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 VMware 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.

Configure Segments with new Orchestrator UI

9

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 enabled 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 configure the Segments using the new Orchestrator UI:

- 1 In the Enterprise portal, click the **Open New Orchestrator UI** option available at the top of the Window.
- 2 Click **Launch New Orchestrator UI** in the pop-up window.
- 3 The UI opens in a new tab displaying the monitoring and configuring options.

In the new Orchestrator UI, click the **Configure** tab.

Note The **Configure** tab is available only when the Operator has enabled the option. If the tab is not available for you, contact your Operator.

- 1 Click **Configure > Segments**.
- 2 The **Segments** page displays the existing Segments.



Segment Name	Description	Type	Service VLAN	Delegate To Partner	Delegate To Customer	Number of Profiles in Use
Global Segment	Default segment for traffic that is not ex...	Regular	Enter VLAN	Enabled	Enabled	2
segment1	Enter Description	Regular	Enter VLAN	Enabled	Enabled	1
segment2	Enter Description	Regular	Enter VLAN	Enabled	Enabled	1

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

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	<p>Choose the Segment type as one of the following:</p> <ul style="list-style-type: none"> ■ Regular - The standard segment type. ■ Private - Used for traffic flows that require limited visibility in order to address end user privacy requirements. ■ CDE - VMware 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 VMware PCI certification. <p>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.</p>
Service VLAN	Enter the service VLAN identifier. For more information, see Define Mapping Segments with Service VLANs .
Delegate To Partner	By default, this checkbox is selected. If this checkbox is not selected, the Partner cannot change the configurations within the segment, including the Interface assignment.
Delegate To Customer	By default, this checkbox is selected. If this checkbox is not selected, the Customer cannot change the configurations within the segment, including the Interface assignment.

4 Click **Save Changes**.

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

- Does not upload user flow stats to Orchestrator except for VMware Control, VMware 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 in 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 VMware 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 click **Delete**. You cannot delete a Segment used by a Profile.

For more information, see [Chapter 8 Configure Segments](#).

Configure Network Services

10

As an enterprise user, SD-WAN Orchestrator allows you to configure a number of network services such as Edge Cluster, Non SD-WAN Destinations, Cloud Security Service (CSS), VNFs and so on from **Configure > Network Services**.

Note If you are logged in using a user ID that has Customer Support privileges, you will only be able to view SD-WAN Orchestrator objects. You will not be able to create new objects or configure/update existing ones.

You can configure the following Network Services:

- Edge Cluster
- Cloud VPN Hubs
- Non SD-WAN Destinations via Gateway
- Non SD-WAN Destinations via Edge
- Cloud Security Service
- VNFs
- VNF Licenses
- DNS Services
- Netflow Settings
- Private Network Names
- Authentication Services
- Cloud Subscriptions

Note Configuring Network Services are optional and can be configured in any order.

NSD Dev

Open New Orchestrator UI Recently Viewed Operator Superuser Help super@velocloud.net

Monitor

Configure

Edges

Profiles

Object Groups

Segments

Overlay Flow Control

Network Services

Alerts & Notifications

Customer

Test & Troubleshoot

Administration

Services

Edge Cluster

New Cluster

Delete Cluster

Name	Location	Used in Profiles

Cloud VPN Hubs

Hub	Type	Used in Profiles	Segment	VPN Hub ⓘ	Backhaul Hub ⓘ

Non SD-WAN Destinations via Gateway

New...

Delete...

Actions

Name	Servers	Tunnels	Pre-Notifications ⓘ	Alerts ⓘ	Used By
<input type="checkbox"/> VM NVS	Type: Generic IKEv1 Router	Not enabled	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	1 Edge 0
<input type="checkbox"/> Generic IKEv2	Generic IKEv2 Router(Router Based VPN)				0

Cloud Security Service

New...

Delete...

Name	Type	Used By
<input type="checkbox"/> Zscaler CSS auto	Zscaler Cloud Security Service	1 Edge

VNFs

New...

Delete...

Name	Type	Used By

VNF Licenses

New...

Delete...

Name	Type	Used By

DNS Services

New...

Delete...

Name	Type	Servers	Used By
<input type="checkbox"/> OpenDNS	Public	Primary: 208.67.222.222 Backup: 208.67.220.220	0
<input type="checkbox"/> Google	Public	Primary: 8.8.8.8 Backup: 8.8.4.4	2 Profiles

Netflow Settings

Collector Name	Collector IP	Collector Port	Used By

Filter Name	Used By

Private Network Names

New...

Delete...

Name	Used By

Authentication Services

New...

Delete...

Name	Servers	Used By

Cloud Subscriptions

New...

Delete...

Note SD-WAN Orchestrator does not allow you to configure Cloud VPN Hubs from the **Services** screen, but it provides a summary of all configured SD-WAN Edge. The summary information includes edge type, profile where the edge is used, segment, whether the edge is a VPN Hub or/and a Backhaul Hub.

Read the following topics next:

- [About Edge Clustering](#)
- [Configure a Non SD-WAN Destination](#)
- [VMware SD-WAN in Azure Virtual WAN Hub Deployment](#)
- [Cloud Security Services](#)
- [Configure DNS Services](#)
- [Configure Netflow Settings](#)
- [Private Network Names](#)
- [Configure Authentication Services](#)
- [Configure Cloud Subscriptions](#)

About Edge Clustering

The size of a single VMware VPN Network with a VMware SD-WAN 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.

SD-WAN Edge Clustering addresses the issue of SD-WAN 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 SD-WAN Edge would provide. A cluster is functionally treated as an individual Hub from the perspective of other Edges.

The Hubs in a VMware 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 SD-WAN 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.

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 SD-WAN 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 VMware 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 VMware SD-WAN 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.

How are Edge Clusters tracked by the VMware SD-WAN Gateway ?

Once a Hub is added to a VMware 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.



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 SD-WAN 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.




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



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 distributed 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 VMware SD-WAN 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 VMware SD-WAN 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 neighborhood 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.

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.

Configure Edge Clustering

You can configure Edge clusters by following the steps in this section.

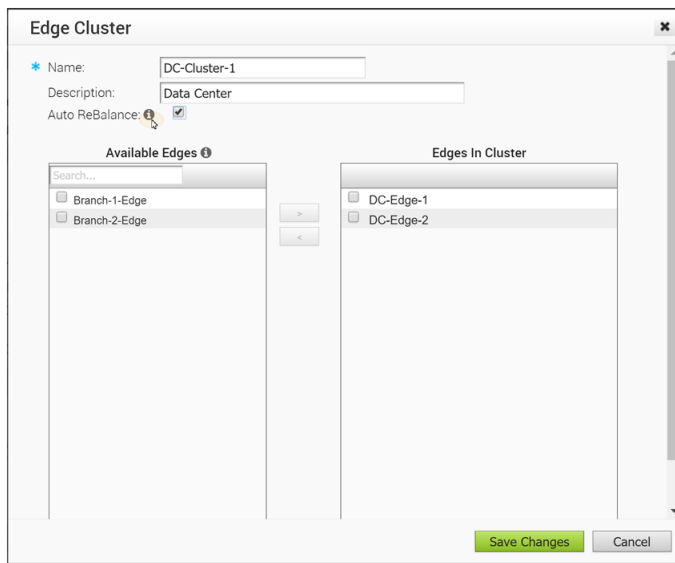
- 1 To access the **Edge Cluster** area, go to **Configure > Network Services**.

Edge Cluster			New Cluster	Delete Cluster
Name	Location	Used in Profiles		
<input checked="" type="checkbox"/> East Coast DC Cluster [3 Edges]	n.a.	1 Profile 1 Edge		

- 2 To add new Cluster:
 - a From the **Edge Cluster** area, click the **New Cluster** button.

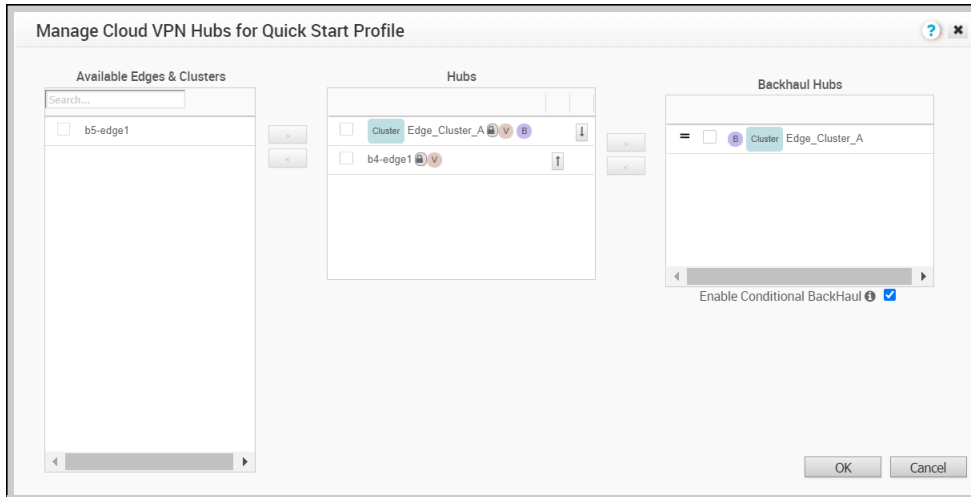
- b In the **Edge Cluster** dialog box, enter the name and description in the appropriate text boxes.
- c Enable **Auto Rebalance** if needed, as this feature is not enabled by default.

Note If this option is enabled, when an individual Edge in a Hub Cluster exceeds a Cluster Score of 70, Spokes will Rebalance at the rate of one Spoke 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 will disconnect 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 will be performed. For more information, see [How Edge Clustering Works](#).



- d In the **Available Edges** section, select an Edge and move it to the **Edges In Cluster** section, by using the > button.
- e Click **Save Changes**. The configured Edge Cluster will appear under **Available Edges & Clusters** area of the **Manage Cloud VPN Hubs** screen for the selected profile.

Note Edges used as a Hub or in Hub Clusters, or configured as an Active/Standby HA pair are not displayed in the **Available Edges** list area.



- 3 From the **Manage Cloud VPN Hubs** screen, you can configure an Edge Cluster and an individual Edge simultaneously as Hubs in a branch profile. Once Edges are assigned to a Cluster, they cannot be assigned as individual Hubs. Choose an Edge Cluster as a Hub in the Branch Profile.
- 4 In order to configure Branch to Branch VPN using Hubs that are also Edge Clusters, you would first select a Hub from the **Hubs** area, and then move it to the **Branch to Branch VPN Hubs** area.
- 5 Hub Clusters can also be configured as Internet Backhaul Hubs in the Business Policy configuration by first selecting a Hub from the **Hubs** area and then moving it to the **Backhaul Hubs** area.
- 6 To enable Conditional Backhaul, select the **Enable Conditional BackHaul** checkbox. With Conditional Backhaul (CBH) enabled, the Edge will be able to failover Internet-bound traffic (Direct Internet traffic, Internet via SD-WAN Gateway and Cloud Security Traffic via IPsec) to MPLS links whenever there is no Public Internet links available. When Conditional Backhaul is enabled, 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 more information, see [Conditional Backhaul](#).

Note It is mandatory to run a dynamic routing protocol, like eBGP, on the LAN side of the clusters.

Troubleshooting Edge Clustering

This section describes the troubleshooting enhancements for Edge Clustering.

Overview

Edge Clustering includes a troubleshooting feature to rebalance VMware 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 VMware SD-WAN Orchestrator

An administrator may rebalance Spokes in a Cluster via **Remote Diagnostics** on the VMware SD-WAN Orchestrator. When an SD-WAN 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 will attempt to evenly re-distribute Spoke Edges among all Hub Edges in the Cluster.

Redistribute Spokes excluding this Hub

- This option will attempt 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.

Shown below is an image of the **Remote Diagnostics** section of the Hub.

Note Rebalancing Spokes will cause 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 VMware support for manually rebalancing the Spoke from its Primary Gateway.

Configure a Non SD-WAN Destination

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

VMware allows the Enterprise users to define and configure a datacenter 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** - Enables a SD-WAN Gateway to establish an IPsec tunnel directly to a Non SD-WAN Destination. VMware supports the following Non SD-WAN Destination configurations through SD-WAN Gateway:

- AWS VPN Gateway

Note The AWS VPN Gateway type is new from the 4.3 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 VMware supports both Generic Route-based and Policy-based Non SD-WAN Destination from Gateway.

- **Non SD-WAN Destinations via Edge** - Enables a SD-WAN Edge to establish an IPsec tunnel directly to a Non SD-WAN Destination (AWS and Azure Datacenter).

Note VMware supports only Generic IKEv2 Router (Route Based VPN) and Generic IKEv1 Router (Route Based VPN) Non SD-WAN Destination from the SD-WAN 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

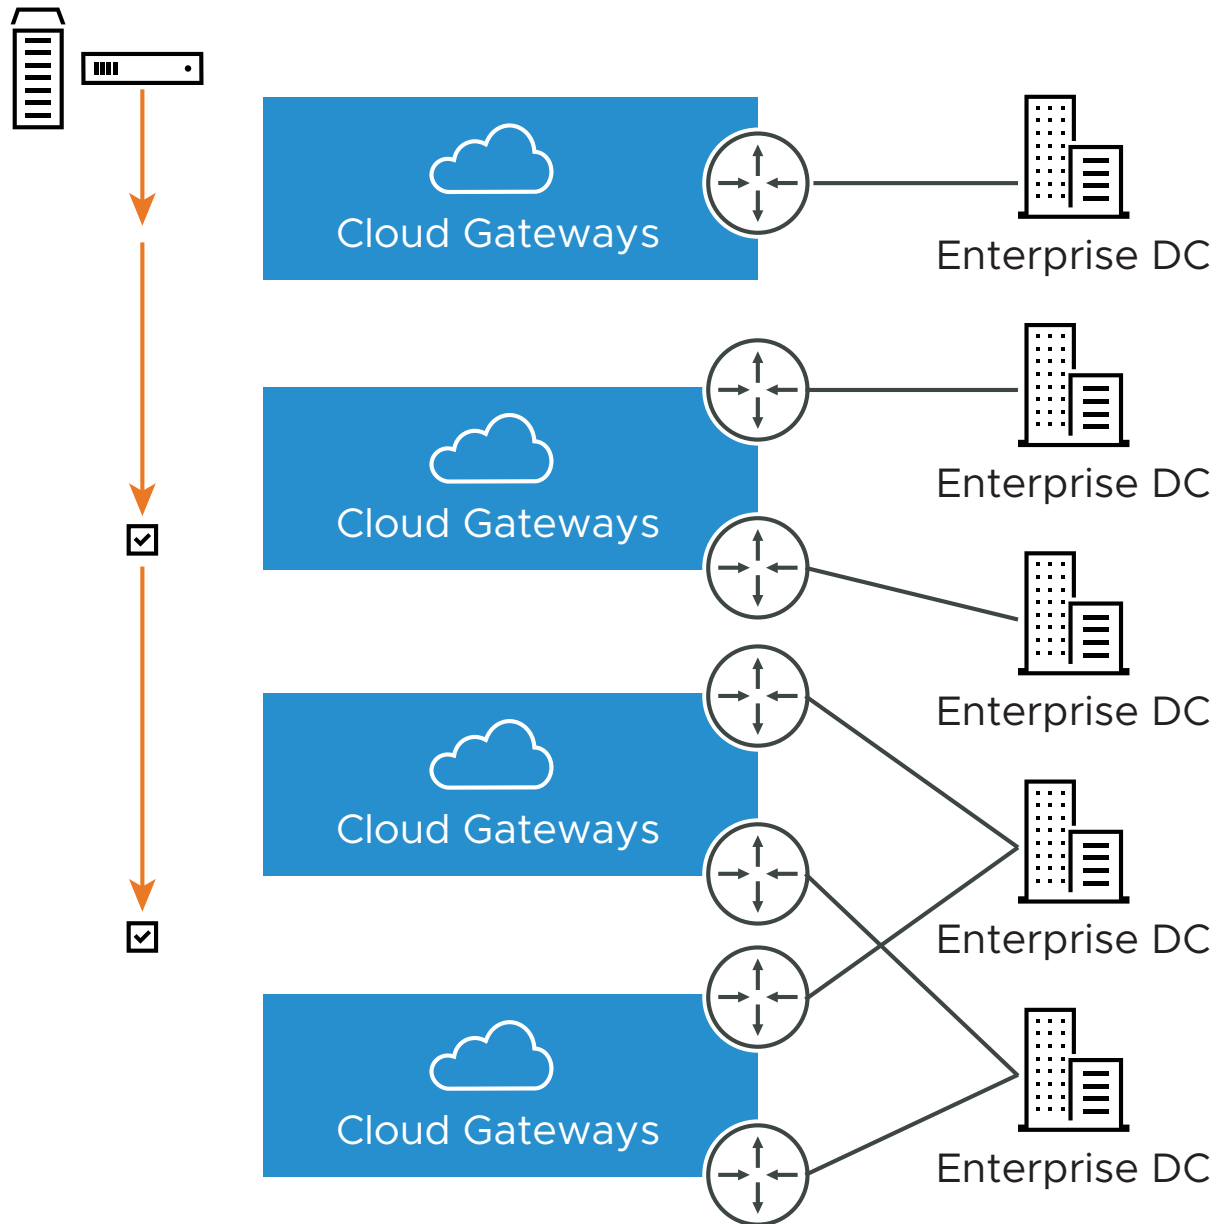
VPN Workflow

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

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 SD-WAN 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 SD-WAN Gateway and the Secondary VPN Gateway. Using Advanced Settings, Redundant VPN Tunnels can be specified for any VPN tunnels you create.

Add Non SD-WAN Destination VPN Gateway

Enter a Name and choose a Gateway Type. Specify the IP address for the Primary VPN Gateway and, optionally, specify an IP address for a Secondary VPN Gateway.

New Non SD-WAN Destination via Gateway...

* Name NVS AWS VPN Gateway Site02

* Type AWS VPN Gateway

VPN Gateways

* Primary VPN Gateway 10.10.10.1

Secondary VPN Gateway 10.1.0.1

Next

Configure Non SD-WAN Destination Subnets

Once you have created a Non SD-WAN Destination configuration, you can add site subnets and configure tunnel settings.

Click the **Advanced** button to configure tunnel settings for VPN Gateways, and to add Redundant VPN tunnel(s).

NSD_AWS_GW

Name

NSD_AWS_GW

Type

AWS VPN Gateway

Enable Tunnel(s)

☒

Tunnel mode

Active/Hot-Standby

Location

Lat,Lng: 37.402889, -122.116859

[Update Location...](#)

Local Auth Id

FQDN

velo.com

Primary VPN Gateway

Public IP

54.183.9.183

Tunnel Settings

PSK

.....

Encryption

AES 128

DH Group

2

PFS

2

Authentication Algorithm

SHA_1

IKE SA Lifetime(min)

1440

IPsec SA Lifetime(min)

480

DPD Type

onDemand

DPD Timeout(sec)

20

Secondary VPN Gateway

Remove

Public IP

54.183.9.185

Tunnel Settings

PSK

.....

Encryption

AES 128

DH Group

2

PFS

2

Authentication Algorithm

SHA_1

IKE SA Lifetime(min)

1440

IPsec SA Lifetime(min)

480

DPD Type

onDemand

DPD Timeout(sec)

20

Redundant VeloCloud Cloud VPN

☒

Primary VPN Gateway

Public IP

54.183.9.183

Tunnel Settings

PSK

.....

Encryption

AES 128

DH Group

2

PFS

2

Authentication Algorithm

SHA_1

IKE SA Lifetime(min)

1440

IPsec SA Lifetime(min)

480

DPD Type

onDemand

DPD Timeout(sec)

20

Secondary VPN Gateway

Restore

Site Subnets

Subnet

Description

Advertise

10.1.2.0/24

(optional)

☒

Deactivate Site Subnets

☐

Advanced

View IKE/IPSec Template

Save Changes


Close

View IKE IPsec Configuration, Configure Non SD-WAN Destination Gateway

If you click the View IKE IPsec Configuration button, the information needed to configure the Non SD-WAN Destination Gateway appears. The Gateway administrator should use this information to configure the Gateway VPN tunnel(s).

VMware, Inc.

157



Enable IPsec Tunnel

The Non SD-WAN Destination VPN tunnel is initially deactivated. You must enable the tunnel(s) after the Non SD-WAN Destination Gateway has been configured and before first use of the Edge-to- Non SD-WAN Destination VPN.

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.

Configure Non SD-WAN Destinations via Gateway

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

To configure a Non SD-WAN Destinations via Gateway:

Procedure

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

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

VMware supports the following Non SD-WAN Destination type configurations through SD-WAN Gateway:

- AWS VPN Gateway

Note AWS VPN Gateway is new in the 4.3 release. In addition, Customers can use different primary Public IPs and Secondary Public IPs for NVS Gateways for AWS.

- 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 VMware supports both Generic Route-based and Policy-based Non SD-WAN Destination from Gateway.

- 5 Enter an IP address for the Primary VPN Gateway (and the Secondary VPN Gateway if necessary), and click **Next**.

A Non SD-WAN Destination is created.

Note To support the datacenter type of Non SD-WAN Destination, besides the IPsec connection, you will need to configure Non SD-WAN Destination local subnets into the VMware system.

What to do next

- Configure tunnel settings for your Non SD-WAN Destination. For more information about configuring tunnel settings for various IPsec tunnel types, see the following sections:
 - [Configure a Non VMware SD-WAN Site 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 IKEv2 Router via Gateway](#)
 - [Configure a Microsoft Azure Non SD-WAN Destination via Gateway](#)
 - [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 via Gateway](#)
 - [Configure a Non SD-WAN Destination of Type Generic Firewall \(Policy Based VPN\)](#)
- Associate your Non SD-WAN Destination to a Profile. For more information, see:
 - [Configure a Tunnel Between a Branch and a Non SD-WAN Destinations via Gateway](#)
 - [Configure a Tunnel Between a Branch and a Non SD-WAN Destinations via Edge](#)
- Configure Business Policy. (Configuring Business Policy is not mandatory for this feature, but if you are going to configure it and would like information, see [Create Business Policy Rules](#).)

Configure a Non VMware SD-WAN Site of Type AWS VPN Gateway

Describes how to configure a Non VMware SD-WAN Site of the type AWS VPN Gateway.

About This Task

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

Procedure

- 1 From the navigation panel in the SD-WAN Orchestrator, go to **Configure > Network Services**.
The **Services** screen appears.
- 2 In the **Non SD-WAN Destinations via Gateway** area, click the **New** button.
The **New Non SD-WAN Destinations via Gateway** dialog box appears.
- 3 In the **Name** text box, enter the name for the Non SD-WAN Destination.
- 4 From the **Type** drop-down menu, select **AWS VPN Gateway**.

New Non SD-WAN Destination via Gateway...

* Name

* Type

VPN Gateways

* Primary VPN Gateway

Secondary VPN Gateway

Next

- 5 Enter the IP address for the Primary VPN Gateway, and click **Next**.

A Non SD-WAN Destination of type AWS VPN Gateway is created and a dialog box for your Non SD-WAN Destination appears.

- 6 To configure tunnel settings for the Non SD-WAN Destination's Primary VPN Gateway, click the **Advanced** button.

NSD_AWS_GW

Name

NSD_AWS_GW

Type

AWS VPN Gateway

Enable Tunnel(s)

☒

Tunnel mode

Active/Hot-Standby

Location

Lat,Lng: 37.402889, -122.116859

Update Location...

Local Auth Id

FQDN

velo.com

Primary VPN Gateway

Public IP

54.183.9.183

Tunnel Settings

PSK

.....

Encryption

AES 128

DH Group

2

PFS

2

Authentication Algorithm

SHA_1

IKE SA Lifetime(min)

1440

IPsec SA Lifetime(min)

480

DPD Type

onDemand

DPD Timeout(sec)

20

Secondary VPN Gateway

Remove

Public IP

54.183.9.185

Tunnel Settings

PSK

.....

Encryption

AES 128

DH Group

2

PFS

2

Authentication Algorithm

SHA_1

IKE SA Lifetime(min)

1440

IPsec SA Lifetime(min)

480

DPD Type

onDemand

DPD Timeout(sec)

20

Redundant VeloCloud Cloud VPN

☒

Primary VPN Gateway

Public IP

54.183.9.183

Tunnel Settings

PSK

.....

Encryption

AES 128

DH Group

2

PFS

2

Authentication Algorithm

SHA_1

IKE SA Lifetime(min)

1440

IPsec SA Lifetime(min)

480

DPD Type

onDemand

DPD Timeout(sec)

20

Secondary VPN Gateway

Restore

Site Subnets

Subnet

Description

Advertise

10.1.2.0/24

(optional)

☒

-

+

Deactivate Site Subnets

☐

Advanced

View IKE/IPSec Template

Save Changes

Close

- 7 In the Primary **VPN Gateway** area, you can configure the following tunnel settings:

Field	Description
Tunnel Mode	Active-Hot-Standby is supported on the SD-WAN Gateway. Active/Hot-Standby automatically displays indicating that if the Active tunnel goes down, the Standby (Hot-Standby) tunnel takes over and becomes the Active tunnel.
PSK	The Pre-Shared Key (PSK), which is the security key for authentication across the tunnel. The SD-WAN Orchestrator generates a PSK by default. If you want to use your own PSK or password, enter it in the textbox.
Encryption	Select either AES 128 or AES 256 as the AES algorithms key size to encrypt data. The default value is AES 128.

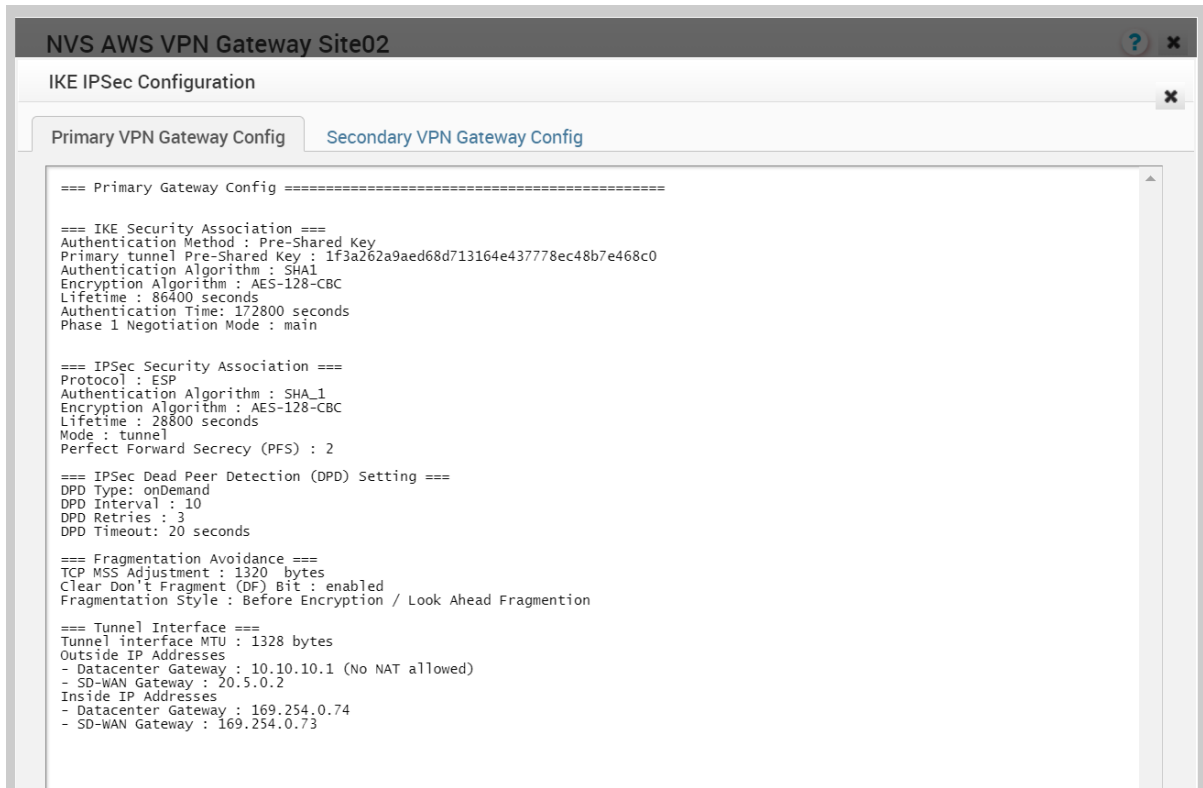
Field	Description
DH Group	Select the Diffie-Hellman (DH) Group algorithm to be used when exchanging a pre-shared key. The DH Group sets the strength of the algorithm in bits. The supported DH Groups are 2, 5, and 14. It is recommended to use DH Group 14.
PFS	Select the Perfect Forward Secrecy (PFS) level for additional security. The supported PFS levels are 2 and 5. The default value is Deactivated.
Authentication Algorithm	The authentication algorithm for the VPN header. Select one of the following supported Secure Hash Algorithm (SHA) function from the drop-down menu list: <ul style="list-style-type: none"> ■ SHA 1 ■ SHA 256 ■ SHA 384 ■ SHA 512 The default value is SHA 1.
IKE SA Lifetime(min)	Time when Internet Key Exchange (IKE) rekeying is initiated for SD-WAN Edges. The minimum IKE life time is 10 minutes and maximum is 1440 minutes. The default value is 1440 minutes.
IPsec SA Lifetime(min)	Time when Internet Security Protocol (IPsec) rekeying is initiated for Edges. The minimum IPsec life time is 3 minutes and maximum is 480 minutes. The default value is 480 minutes.
DPD Type	The Dead Peer Detection (DPD) method is used to detect if the Internet Key Exchange (IKE) peer is alive or dead. If the peer is detected as dead, the device deletes the IPsec and IKE Security Association. Select either Periodic or on Demand from the list. The default value is on Demand.
DPD Timeout(sec)	The maximum time that the device should wait to receive a response to the DPD message before considering the peer to be dead. The default value is 20 seconds. You can deactivate DPD by configuring the DPD timeout timer to 0 second.

- To create a Secondary VPN Gateway for this site, click the **Add** button next to **Secondary VPN Gateway**. In the pop-up window, enter the IP address of the Secondary VPN Gateway and click **Save Changes**.

The Secondary VPN Gateway will be created immediately for this site and will provision a VMware VPN tunnel to this Gateway.

- Select the **Redundant VeloCloud Cloud VPN** checkbox to add redundant tunnels for each VPN Gateway. Any changes made to Encryption, DH Group, or PFS of Primary VPN Gateway will also be applied to the redundant VPN tunnels, if configured.

- 10 After modifying the tunnel settings of the Primary VPN Gateway, save the changes and then click **View IKE/IPsec Template** to view the updated tunnel configuration.



- 11 Click the **Update location** link, located in the top, right corner of the **Non SD-WAN Destination Via Gateway** dialog, to set the location for the configured Non VMware SD-WAN Site. The latitude and longitude details are used to determine the best SD-WAN Edge or SD-WAN Gateway to connect to in the network.
- 12 Under the **Site Subnets** area, you can add subnets for the Non VMware SD-WAN Site by clicking the + button. Use Custom Source Subnets to override the source subnets routed to this VPN device. Normally, source subnets are derived from the SD-WAN Edge LAN subnets routed to this device.

Note Site Subnets should be deactivated for enabling a tunnel if there are no site subnets configured.

- 13 Check the **Enable Tunnel(s)** checkbox once you are ready to initiate the tunnel from the SD-WAN Gateway to the AWS VPN Gateways.
- 14 Click **Save Changes**.
- 15 Assign the newly created Non SD-WAN Site Network Service to a Profile by navigating to **Configure > Profiles** in the SD-WAN Orchestrator. See [Configure a Tunnel Between a Branch and a Non SD-WAN Destinations via Gateway](#).
- 16 Return to the **Non SD-WAN Destinations via Gateway** area in the SD-WAN Orchestrator by going to **Configure > Network Services**.

- 17 In the **Non SD-WAN Destinations via Gateway** area, scroll to the name of your Non SD-WAN Site, and then click the **Edit** link in the **BGP** column.
- 18 Configure the BGP based on the AWS values for the following mandatory fields: Local ASN, Tunnel Type, Neighbor IP, and Local IP (from the Advanced Options section). NOTE: Tunnel type is updated by default. Refer to the AWS documentation if needed. For more information, see [Configure BGP over IPsec from Gateways](#).
- 19 Click the **OK** button to save your changes.
- 20 In the **Non SD-WAN Destinations via Gateway** area, click the **Edit** link in the **BFD** column for a Non SD-WAN Destination, to configure the BFD settings. For more information, see [Configure BFD for Gateways](#).

What to do next

You can check the overall status of the Non SD-WAN Sites in the monitoring tab. See:

- [Monitor Network Services](#)
- [Monitor Non SD-WAN Destinations through Gateway](#)

Configure Check Point

The SD-WAN Gateway connects to the Check Point CloudGuard service using IKEv1/IPsec. There are two steps to configure Check Point: Configuring the Checkpoint CloudGuard service and Configuring Checkpoint on the SD-WAN Orchestrator. You will perform the first step on the Check Point Infinity Portal and the second step on the SD-WAN Orchestrator.

Click the links for the following sections below to complete the instructions to configure Check Point.

Step 1: [Configure the Check Point CloudGuard Connect](#)

Step 2: [Configure a Non SD-WAN Destination of Type Check Point](#)

Prerequisites

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

Configure the Check Point CloudGuard Connect

Instructions on how to configure the Check Point CloudGuard Service.

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

Procedure

- 1 To configure the Check Point CloudGuard service, login to Check Point's Infinity Portal at (<https://portal.checkpoint.com/>).

- 2 Once logged in, create a site at Check Point's Infinity Portal via the following link: <https://sc1.checkpoint.com/documents/integrations/VeloCloud/check-point-VeloCloud-integration.html>

After you create a site at Check Point's Infinity Portal, [Configure a Non SD-WAN Destination of Type Check Point](#)

Configure a Non SD-WAN Destination of Type Check Point

After you create a site at Check Point's Infinity Portal, configure a Non SD-WAN Destination of type Check Point in the SD-WAN Orchestrator.

To configure a Non SD-WAN Destination of type Check Point in the SD-WAN Orchestrator, perform the following steps:

Procedure

- 1 From the navigation panel in the SD-WAN Orchestrator, go to **Configure > Network Services**.
The **Services** screen appears.
- 2 In the **Non SD-WAN Destinations via Gateway** area, click the **New** button.
The **New Non SD-WAN Destinations via Gateway** dialog box appears.
- 3 In the **Name** text box, enter the name for the Non SD-WAN Destination.
- 4 From the **Type** drop-down menu, select **Check Point**.
- 5 Enter the IP address for the Primary VPN Gateway (and the Secondary VPN Gateway if necessary), and click **Next**.

A Non SD-WAN Destination of type Check Point is created and a dialog box for your Non SD-WAN Destination appears.

CPM

Name

CPM

Type

Check Point

Enable Tunnel(s)

☐

Tunnel mode

Active/Hot-Standby

Location

Lat,Lng: 37.402889, -122.116859

[Update Location...](#)

Primary VPN Gateway

Public IP

54.183.9.191

Tunnel Settings

PSK

.....

Redundant Tunnel PSK

.....

Encryption

AES 128

DH Group

2

PFS

2

Local Auth Id

Default

Site Subnets

Subnet	Description	Advertise
Ex: 10.0.2.0/24	(optional)	<input checked="" type="checkbox"/>

Deactivate Site Subnets

☐

Secondary VPN Gateway

Remove

Public IP

54.183.9.192

Tunnel Settings

PSK

.....

Redundant Tunnel PSK

.....

Encryption

AES 128

DH Group

2

PFS

2

Redundant VeloCloud Cloud VPN

☒

Advanced

View IKE/IPSec Template

Save Changes

Close

- To configure tunnel settings for the Non SD-WAN Destination's Primary VPN Gateway, click the **Advanced** button.
- In the **Primary VPN Gateway** area, you can configure the following tunnel settings:

Field	Description
PSK	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 textbox.
Encryption	Select either AES 128 or AES 256 as the AES algorithms key size to encrypt data. The default value is AES 128.
DH Group	Select the Diffie-Hellman (DH) Group algorithm to be used when exchanging a pre-shared key. The DH Group sets the strength of the algorithm in bits. The supported DH Groups are 2, 5, and 14. It is recommended to use DH Group 14.
PFS	Select the Perfect Forward Secrecy (PFS) level for additional security. The supported PFS levels are 2 and 5. The default value is 2.

VMware, Inc.

167

- 8 If you want to create a Secondary VPN Gateway for this site, then click the **Add** button next to **Secondary VPN Gateway**. In the pop-up window, enter the IP address of the Secondary VPN Gateway and click **Save Changes**.

The Secondary VPN Gateway will be created immediately for this site and will provision a VMware VPN tunnel to this Gateway.

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

- 9 Select the **Redundant VeloCloud Cloud VPN** checkbox to add redundant tunnels for each VPN Gateway.

Any changes made to Encryption, DH Group, or PFS of Primary VPN Gateway will also be applied to the redundant VPN tunnels, if configured. After modifying the tunnel settings of the Primary VPN Gateway, save the changes and then click **View IKE/IPSec Template** to view the updated tunnel configuration.
- 10 Click the **Update location** link 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.
- 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 IP address used to communicate with the local gateway.

If you do not specify a value, **Default** is used as the local authentication ID.
- 12 Under **Site Subnets**, you can add subnets for the Non SD-WAN Destination by clicking the **+** button. If you do not need subnets for the site, select the **Deactivate Site Subnets** checkbox.
- 13 Check the **Enable Tunnel(s)** checkbox once you are ready to initiate the tunnel from the SD-WAN Gateway to the Check Point VPN gateways.
- 14 Click **Save Changes**.

Configure a Non SD-WAN Destination of Type Cisco ASA

Describes how to configure a Non SD-WAN Destination of type **Cisco ASA** in SD-WAN Orchestrator.

Procedure

- 1 From the navigation panel in the SD-WAN Orchestrator, go to **Configure > Network Services**.

The **Services** screen appears.

- In the **Non SD-WAN Destinations via Gateway** area, click the **New** button.

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

- In the **Name** text box, enter the name for the Non SD-WAN Destination.
- From the **Type** drop-down menu, select **Cisco ASA**.
- Enter the IP address for the Primary VPN Gateway, and click **Next**.

A Non SD-WAN Destination of type Cisco ASA is created and a dialog box for your Non SD-WAN Destination appears.

NSD_Cisco_ASA

Name: NSD_Cisco_ASA
 Type: Cisco ASA
 Enable Tunnel(s): ☐
 Tunnel mode: Active/Hot-Standby

Location: Lat,Lng: 37.402889, -122.116859
[Update Location...](#)

Primary VPN Gateway

Public IP: 10.10.10.7

Tunnel Settings

PSK:

Redundant Tunnel PSK:

Encryption: AES 128
 DH Group: 2
 PFS: deactivated

Local Auth Id: FQDN
 Local Auth Id: velo.com

Site Subnets

Subnet	Description	Advertise
10.1.2.0/24	(optional)	<input checked="" type="checkbox"/>

Deactivate Site Subnets: ☐

Custom Source Subnets

Subnet	Description	Advertise
Ex: 10.0.2.0/24	(optional)	<input checked="" type="checkbox"/>

Secondary VPN Gateway ✖
 Secondary VPN Gateways are not supported for Cisco ASA. This is a limitation of the Cisco ASA VPN.

Redundant VeloCloud Cloud VPN: ☒

Advanced View IKE/IPSec Template Save Changes Close

- To configure tunnel settings for the Non SD-WAN Destination's Primary VPN Gateway, click the **Advanced** button.
- In the **Primary VPN Gateway** area, you can configure the following tunnel settings:

Field	Description
PSK	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 textbox.
Encryption	Select either AES 128 or AES 256 as the AES algorithms key size to encrypt data. The default value is AES 128.

Field	Description
DH Group	Select the Diffie-Hellman (DH) Group algorithm to be used when exchanging a pre-shared key. The DH Group sets the strength of the algorithm in bits. The supported DH Groups are 2, 5, and 14. It is recommended to use DH Group 14.
PFS	Select the Perfect Forward Secrecy (PFS) level for additional security. The supported PFS levels are 2 and 5. The default value is deactivated .

Note

- The Secondary VPN Gateway is not supported for the Cisco ASA network service type.
- For Cisco ASA Non SD-WAN Destination, by default, the local authentication ID value used is the Local IP address of the SD-WAN Gateway.

- 8 Select the **Redundant VeloCloud Cloud VPN** checkbox to add redundant tunnels for each VPN Gateway.

Any changes made to Encryption, DH Group, or PFS of Primary VPN Gateway will also be applied to the redundant VPN tunnels, if configured. After modifying the tunnel settings of the Primary VPN Gateway, save the changes and then click **View IKE/IPSec Template** to view the updated tunnel configuration.

- 9 Click the **Update location** link 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.
- 10 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 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 will be the SD-WAN Gateway Interface Public IP.

- 11 Under **Site Subnets**, you can add subnets for the Non SD-WAN Destination by clicking the **+** button. If you do not need subnets for the site, select the **Deactivate Site Subnets** check box.
- 12 Use **Custom Source Subnets** to override the source subnets routed to this VPN device. Normally, source subnets are derived from the edge LAN subnets routed to this device.
- 13 Check the **Enable Tunnel(s)** checkbox once you are ready to initiate the tunnel from the SD-WAN Gateway to the Cisco ASA VPN gateways.
- 14 Click **Save Changes**.

Configure a Non SD-WAN Destination of Type Cisco ISR

Describes how to configure a Non SD-WAN Destination of type **Cisco ISR** in SD-WAN Orchestrator.

Procedure

- 1 From the navigation panel in the SD-WAN Orchestrator, go to **Configure > Network Services**.
The **Services** screen appears.
- 2 In the **Non SD-WAN Destinations via Gateway** area, click the **New** button.
The **New Non SD-WAN Destinations via Gateway** dialog box appears.
- 3 In the **Name** text box, enter the name for the Non SD-WAN Destination.
- 4 From the **Type** drop-down menu, select **Cisco ISR**.
- 5 Enter the IP address for the Primary VPN Gateway, and click **Next**.

A Non SD-WAN Destination of type Cisco ISR is created and a dialog box for your Non SD-WAN Destination appears.

NSD_ISR

Name: NSD_ISR
 Type: Cisco ISR
 Enable Tunnel(s): ☒
 Tunnel mode: Active/Hot-Standby

Location: Lat,Lng: 37.402889, -122.116859
[Update Location...](#)

Primary VPN Gateway

Public IP: 54.183.9.185
 Tunnel Settings: PSK:
 Redundant Tunnel PSK:
 Encryption: AES 128
 DH Group: 2
 PFS: deactivated

Secondary VPN Gateway [Remove](#)

Public IP: 54.183.9.187
 Tunnel Settings: PSK:
 Redundant Tunnel PSK:
 Encryption: AES 128
 DH Group: 2
 PFS: deactivated

Site Subnets

Subnet	Description	Advertise
Ex: 10.0.2.0/24	(optional)	<input checked="" type="checkbox"/>

Redundant VeloCloud Cloud VPN ☒

[Advanced](#) [View IKE/IPSec Template](#) [Save Changes](#) [Close](#)

- 6 To configure tunnel settings for the Non SD-WAN Destination's Primary VPN Gateway, click the **Advanced** button.

- 7 In the **Primary VPN Gateway** area, you can configure the following tunnel settings:

Field	Description
Tunnel Mode	Active-Hot-Standby is supported on the SD-WAN Gateway. Active/Hot-Standby automatically displays indicating that if the Active tunnel goes down, the Standby (Hot-Standby) tunnel takes over and becomes the Active tunnel.
PSK	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.
Encryption	Select either AES 128 or AES 256 as the AES algorithms key size to encrypt data. The default value is AES 128.
DH Group	Select the Diffie-Hellman (DH) Group algorithm to be used when exchanging a pre-shared key. The DH Group sets the strength of the algorithm in bits. The supported DH Groups are 2, 5, and 14. It is recommended to use DH Group 14.
PFS	Select the Perfect Forward Secrecy (PFS) level for additional security. The supported PFS levels are 2 and 5. The default value is deactivated .

- 8 If you want to create a Secondary VPN Gateway for this site, then click the **Add** button next to **Secondary VPN Gateway**. In the pop-up window, enter the IP address of the Secondary VPN Gateway and click **Save Changes**.

The Secondary VPN Gateway will be created immediately for this site and will provision a VMware SD-WAN VPN tunnel to this Gateway.

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

- 9 Select the **Redundant VeloCloud Cloud VPN** check box to add redundant tunnels for each VPN Gateway.

Any changes made to Encryption, DH Group, or PFS of Primary VPN Gateway will also be applied to the redundant VPN tunnels, if configured. After modifying the tunnel settings of the Primary VPN Gateway, save the changes and then click **View IKE/IPsec Template** to view the updated tunnel configuration.

- 10 Click the **Update location** link to set the location for the configured Non SD-WAN Destination. The latitude and longitude details are used to determine the best SD-WAN Edge or SD-WAN Gateway to connect to in the network.
- 11 Under **Site Subnets**, add subnets for the Non SD-WAN Destination by clicking the **+** button.

Note For Cisco ISR, Site Subnets are mandatory to be configured.

- 12 Check the **Enable Tunnel(s)** check box once you are ready to initiate the tunnel from the SD-WAN Gateway to the Cisco ISR VPN Gateways.
- 13 Click **Save Changes**.
- 14 Assign the newly created Non SD-WAN Site Network Service to a Profile by navigating to **Configure > Profiles** in the SD-WAN Orchestrator. See [Configure a Tunnel Between a Branch and a Non SD-WAN Destinations via Gateway](#).
- 15 Return to the **Non SD-WAN Destinations via Gateway** area in the SD-WAN Orchestrator by going to **Configure > Network Services**.
- 16 In the **Non SD-WAN Destinations via Gateway** area, scroll to the name of your Non SD-WAN Site, and then click the **Edit** link in the **BGP** column.
- 17 Configure the BGP based on the Cisco ISR values for the following mandatory fields: Local ASN, Tunnel Type, Neighbor IP, and Local IP (from the Advanced Options section). Refer to the Cisco documentation if needed. For more information, see [Configure BGP over IPsec from Gateways](#).

Note The VTI IP (private IP) assigned by the SD-WAN Orchestrator can be used for peer ship in Single-Hop BGP.

- 18 Click **OK** to save your changes.
- 19 In the **Non SD-WAN Destinations via Gateway** area, click the **Edit** link in the **BFD** column for a Non SD-WAN Destination, to configure the BFD settings. For more information, see [Configure BFD for Gateways](#).

What to do next

You can check the overall status of the Non SD-WAN Sites in the monitoring tab. See:

- [Monitor Network Services](#)
- [Monitor Non SD-WAN Destinations through Gateway](#)

Configure a Non SD-WAN Destination of Type Generic IKEv2 Router via Gateway

Describes how to configure a Non SD-WAN Destination of type **Generic IKEv2 Router (Route Based VPN)** in SD-WAN Orchestrator.

Note To configure a **Generic IKEv2 Router (Route Based VPN)** via Edge, see [Configure a Non-VMware SD-WAN Site of Type Generic IKEv2 Router via Edge](#).

Procedure

- 1 From the navigation panel in the SD-WAN Orchestrator, go to **Configure > Network Services**.
The **Services** screen appears.

- 2 In the **Non SD-WAN Destinations via Gateway** area, click the **New** button.

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

- 3 In the **Name** text box, enter the name for the Non SD-WAN Destination.
- 4 From the **Type** drop-down menu, select **Generic IKEv2 Router (Route Based VPN)**.
- 5 Enter the IP address for the Primary VPN Gateway (and the Secondary VPN Gateway if necessary), and click **Next**.

A route-based Non SD-WAN Destination of type IKEv2 is created and a dialog box for your Non SD-WAN Destination appears.

NSD_IKEv2

Name: NSD_IKEv2

Type: Generic IKEv2 Router (Route Based VPN)

Enable Tunnel(s): ☒

Tunnel mode: Active/Hot-Standby

Location: Lat,Lng: 37.402889,-122.116859 [Update Location...](#)

Local Auth Id: FQDN

Local Auth Id: velo.com

Site Subnets

Subnet	Description	Advertise
10.1.2.0/24	(optional)	<input checked="" type="checkbox"/>

Deactivate Site Subnets: ☐

Primary VPN Gateway

Public IP: 54.183.9.180

Tunnel Settings

PSK: [Redacted]

Redundant Tunnel PSK: [Redacted]

Encryption: AES 128

DH Group: 2

PFS: 2

Authentication Algorithm: SHA 1

IKE SA Lifetime(min): 1440

IPsec SA Lifetime(min): 480

DPD Type: onDemand

DPD Timeout(sec): 20

Secondary VPN Gateway [Remove](#)

Public IP: 54.183.9.181

Tunnel Settings

PSK: [Redacted]

Redundant Tunnel PSK: [Redacted]

Encryption: AES 128

DH Group: 2

PFS: 2

Authentication Algorithm: SHA 1

IKE SA Lifetime(min): 1440

IPsec SA Lifetime(min): 480

DPD Type: onDemand

DPD Timeout(sec): 20

Redundant VeloCloud Cloud VPN: ☒

[Advanced](#) [View IKE/IPSec Template](#) [Save Changes](#) [Close](#)

- 6 To configure tunnel settings for the Non SD-WAN Destination's Primary VPN Gateway, click the **Advanced** button.

7 In the **Primary VPN Gateway** area, you can configure the following tunnel settings:

Field	Description
Tunnel Mode	Active-Hot-Standby is supported on the SD-WAN Gateway. Active/Hot-Standby automatically displays indicating that if the Active tunnel goes down, the Standby (Hot-Standby) tunnel takes over and becomes the Active tunnel.
PSK	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 textbox.
Encryption	Select either AES 128 or AES 256 as the AES algorithms key size to encrypt data. The default value is AES 128.
DH Group	Select the Diffie-Hellman (DH) Group algorithm to be used when exchanging a pre-shared key. The DH Group sets the strength of the algorithm in bits. The supported DH Groups are 2, 5, and 14. It is recommended to use DH Group 14.
PFS	Select the Perfect Forward Secrecy (PFS) level for additional security. The supported PFS levels are 2 and 5. The default value is 2.
Authentication Algorithm	<p>The authentication algorithm for the VPN header. Select one of the following supported Secure Hash Algorithm (SHA) function from the list:</p> <ul style="list-style-type: none"> ■ SHA 1 ■ SHA 256 ■ SHA 384 ■ SHA 512 <p>The default value is SHA 1.</p>
IKE SA Lifetime(min)	Time when Internet Key Exchange (IKE) rekeying is initiated for Edges. The minimum IKE life time is 10 minutes and maximum is 1440 minutes. The default value is 1440 minutes.
IPsec SA Lifetime(min)	Time when Internet Security Protocol (IPsec) rekeying is initiated for Edges. The minimum IPsec life time is 3 minutes and maximum is 480 minutes. The default value is 480 minutes.

Field	Description
DPD Type	The Dead Peer Detection (DPD) method is used to detect if the Internet Key Exchange (IKE) peer is alive or dead. If the peer is detected as dead, the device deletes the IPsec and IKE Security Association. Select either Periodic or on Demand from the list. The default value is on Demand.
DPD Timeout(sec)	The maximum time that the device should wait to receive a response to the DPD message before considering the peer to be dead. The default value is 20 seconds. You can deactivate DPD by configuring the DPD timeout timer to 0 second.

Note When AWS initiates the rekey tunnel with a VMware SD-WAN 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 SD-WAN 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.

- 8 If you want to create a Secondary VPN Gateway for this site, then click the **Add** button next to **Secondary VPN Gateway**. In the pop-up window, enter the IP address of the Secondary VPN Gateway and click **Save Changes**.

The Secondary VPN Gateway will be created immediately for this site and will provision a VMware VPN tunnel to this Gateway.

- 9 Select the **Redundant VeloCloud Cloud VPN** check box to add redundant tunnels for each VPN Gateway.

Any changes made to Encryption, DH Group, or PFS of Primary VPN Gateway will also be applied to the redundant VPN tunnels, if configured. After modifying the tunnel settings of the Primary VPN Gateway, save the changes and then click **View IKE/IPsec Template** to view the updated tunnel configuration.

- 10 Click the **Update location** link 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.
- 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 IP address used to communicate with the local gateway.

Note For Generic route based VPN, if the user does not specify a value, **Default** is used as the local authentication ID. The default local authentication ID value will be the SD-WAN Gateway Interface Public IP.

- Under **Site Subnets**, you can add subnets for the Non SD-WAN Destination by clicking the **+** button. If you do not need subnets for the site, select the **Deactivate Site Subnets** check box.
- Check the **Enable Tunnel(s)** check box once you are ready to initiate the tunnel from the SD-WAN Gateway to the Generic IKEv2 VPN gateways.
- Click **Save Changes**.
- Assign the newly created Non SD-WAN Site Network Service to a Profile by navigating to **Configure > Profiles** in the SD-WAN Orchestrator. See [Configure a Tunnel Between a Branch and a Non SD-WAN Destinations via Gateway](#).
- Return to the **Non SD-WAN Destinations via Gateway** area in the SD-WAN Orchestrator by going to **Configure > Network Services**.
- In the **Non SD-WAN Destinations via Gateway** area, scroll to the name of your Non SD-WAN Site, and then click the **Edit** link in the **BGP** column.
- Configure the BGP values for the following mandatory fields: Local ASN, Tunnel Type, Neighbor IP, and Local IP (from the Advanced Options section). For more information, see [Configure BGP over IPsec from Gateways](#).

Note The VTI IP (private IP) assigned by the SD-WAN Orchestrator can be used for peer ship in Single-Hop BGP.

- Click **OK** to save your changes.
- In the **Non SD-WAN Destinations via Gateway** area, click the **Edit** link in the **BFD** column for a Non SD-WAN Destination, to configure the BFD settings. For more information, see [Configure BFD for Gateways](#).

What to do next

You can check the overall status of the Non SD-WAN Sites in the monitoring tab. See:

- [Monitor Network Services](#)
- [Monitor Non SD-WAN Destinations through Gateway](#)

Configure a Microsoft Azure Non SD-WAN Destination

Describes how to configure a Non SD-WAN Destination of type **Microsoft Azure Virtual Hub** in SD-WAN Orchestrator.

To configure a Non SD-WAN Destination of type **Microsoft Azure Virtual Hub** from SD-WAN Gateway:

Prerequisites

- Ensure you have configured a Cloud subscription. For steps, see [Configure a Cloud Subscription Network Service](#).
- Ensure you have created Virtual WAN and Hubs in Azure. For steps, see [Configure Azure Virtual WAN for Branch-to-Azure VPN Connectivity](#).

Procedure

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

New Non SD-WAN Destination via Gateway...

* Name:

* Type:

Virtual Hub Configuration

Subscription:

Virtual WAN:

Resource Group:

Virtual Hub:

Azure Region:

Enable Tunnel(s): ☒

Next

- 3 In the **Name** text box, enter the name for the Non SD-WAN Destination.
- 4 From the **Type** drop-down menu, select **Microsoft Azure Virtual Hub**.
- 5 From the **Subscription** drop-down menu, select a subscription.
The application fetches all the available Virtual WANs dynamically from Azure.
- 6 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.
- 7 From the **Virtual Hub** drop-down menu, select a Virtual Hub.
The application auto-populates the Azure region corresponding to the Hub

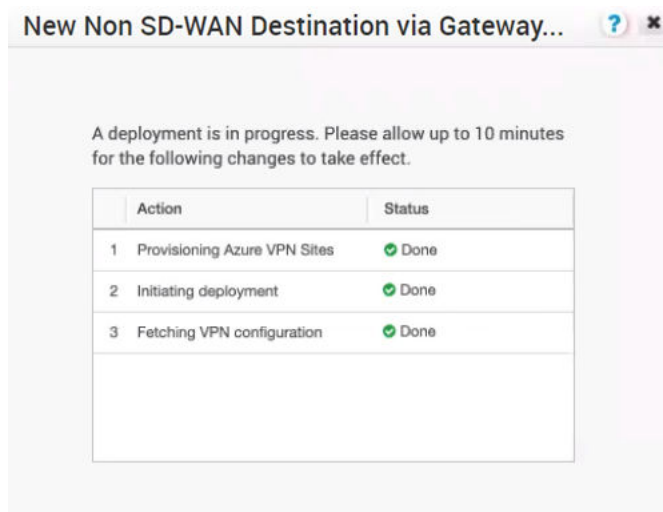
- 8 Select the **Enable Tunnel(s)** checkbox to enable VMware VPN Gateways initiate VPN connections to the target Virtual Hub as soon as the site is successfully provisioned.

Note VMware VPN Gateways will not initiate IKE negotiation until this Non SD-WAN Destination is configured on at least one profile.

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

- 9 Click **Next**.

The SD-WAN Orchestrator automatically initiates deployment, provisions Azure VPN Sites, and downloads the VPN Site Configuration for the newly configured sites and stores the configuration in the SD-WAN Orchestrator's Non SD-WAN Destination configuration database.



Results

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

What to do next

- Associate the Microsoft Azure Non SD-WAN Destination to a Profile to establish a tunnel between a branch and Azure Virtual Hub. For more information, see [Associate a Microsoft Azure Non SD-WAN Destination to a Profile](#).
- You must add SD-WAN routes into Azure network manually. For more 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 configure the BGP settings for the Non SD-WAN Destination. Scroll to the name of your Non SD-WAN Destination, and then click the **Edit** link in the **BGP** column. For more information, see [Configure BGP over IPsec from Gateways](#).
- In the **Non SD-WAN Destinations via Gateway** area, click the **Edit** link in the **BFD** column for a Non SD-WAN Destination, to configure the BFD settings. For more information, see [Configure BFD for Gateways](#).

For information about Azure Virtual WAN Gateway Automation, see [Configure SD-WAN Orchestrator for Azure Virtual WAN IPsec Automation from SD-WAN Gateway](#).

Configure a Non SD-WAN Destination of Type Palo Alto

Describes how to configure a Non SD-WAN Destination of type **Palo Alto** in SD-WAN Orchestrator.

Procedure

- 1 From the navigation panel in the SD-WAN Orchestrator, go to **Configure > Network Services**.
The **Services** screen appears.
- 2 In the **Non SD-WAN Destinations via Gateway** area, click the **New** button.
The **New Non SD-WAN Destinations via Gateway** dialog box appears.
- 3 In the **Name** text box, enter the name for the Non SD-WAN Destination.
- 4 From the **Type** drop-down menu, select **Palo Alto**.
- 5 Enter the IP address for the Primary VPN Gateway, and click **Next**.

A Non SD-WAN Destination of type Palo Alto is created and a dialog box for your Non SD-WAN Destination appears.

NVS Palo Alto Site01

Name

NVS Palo Alto Site01

Type

Palo Alto

Enable Tunnel(s)

☒

Location

Lat,Lng: 37.402889, -122.116859

[Update Location...](#)

Primary VPN Gateway

Public IP

10.10.10.8

Tunnel Settings

PSK

.....

Encryption

AES 128

DH Group

2

PFS

5

Site Subnets

Subnet	Description	Advertise
10.0.2.0/24	(optional)	<input checked="" type="checkbox"/> - +

Secondary VPN Gateway

Add

Redundant VeloCloud Cloud VPN

☐

Advanced

View IKE/IPSec Template

Save Changes

Close

6 To configure tunnel settings for the Non SD-WAN Destination's Primary VPN Gateway, click the **Advanced** button.

7 In the **Primary VPN Gateway** area, you can configure the following tunnel settings:

Field	Description
PSK	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 textbox.
Encryption	Select either AES 128 or AES 256 as the AES algorithms key size to encrypt data. The default value is AES 128.
DH Group	Select the Diffie-Hellman (DH) Group algorithm to be used when exchanging a pre-shared key. The DH Group sets the strength of the algorithm in bits. The supported DH Groups are 2, 5, and 14. It is recommended to use DH Group 14.
PFS	Select the Perfect Forward Secrecy (PFS) level for additional security. The supported PFS levels are 2 and 5. The default value is 2.

- 8 If you want to create a Secondary VPN Gateway for this site, then click the **Add** button next to **Secondary VPN Gateway**. In the pop-up window, enter the IP address of the Secondary VPN Gateway and click **Save Changes**.

The Secondary VPN Gateway will be created immediately for this site and will provision a VMware VPN tunnel to this Gateway.

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

- 9 Select the **Redundant VeloCloud Cloud VPN** checkbox to add redundant tunnels for each VPN Gateway.

Any changes made to Encryption, DH Group, or PFS of Primary VPN Gateway will also be applied to the redundant VPN tunnels, if configured. After modifying the tunnel settings of the Primary VPN Gateway, save the changes and then click **View IKE/IPSec Template** to view the updated tunnel configuration.
- 10 Click the **Update location** link 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.
- 11 Under **Site Subnets**, you can add subnets for the Non SD-WAN Destination by clicking the **+** button.
- 12 Check the **Enable Tunnel(s)** checkbox once you are ready to initiate the tunnel from the SD-WAN Gateway to the Palo Alto VPN gateways.
- 13 Click **Save Changes**.

Configure a Non SD-WAN Destination of Type SonicWALL

Describes how to configure a Non SD-WAN Destination of type **SonicWALL** in SD-WAN Orchestrator.

Procedure

- 1 From the navigation panel in the SD-WAN Orchestrator, go to **Configure > Network Services**.
The **Services** screen appears.
- 2 In the **Non SD-WAN Destinations via Gateway** area, click the **New** button.
The **New Non SD-WAN Destinations via Gateway** dialog box appears.
- 3 In the **Name** text box, enter the name for the Non SD-WAN Destination.
- 4 From the **Type** drop-down menu, select **SonicWALL**.
- 5 Enter the IP address for the Primary VPN Gateway, and click **Next**.

A Non SD-WAN Destination of type SonicWALL is created and a dialog box for your Non SD-WAN Destination appears.

NVS Sonicwall site1 ? ✕

* Name: NVS Sonicwall site1
 Type: SonicWALL
 Enable Tunnel(s): ☒

Location ? Lat,Lng: 37.402889, -122.116859
[Update Location...](#)

Primary VPN Gateway
 * Public IP: 10.10.10.5
 Tunnel Settings ?
 PSK: ?
 Encryption: AES 128
 DH Group: 2
 PFS: 2

Secondary VPN Gateway Add

Redundant VeloCloud Cloud VPN ? ☐

Advanced View IKE/IPSec Template Save Changes Close

Site Subnets ?

Subnet	Description	Advertise
10.0.2.0/24	(optional)	<input checked="" type="checkbox"/> - +

- 6 To configure tunnel settings for the Non SD-WAN Destination's Primary VPN Gateway, click the **Advanced** button.
- 7 In the **Primary VPN Gateway** area, you can configure the following tunnel settings:

Field	Description
PSK	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 textbox.
Encryption	Select either AES 128 or AES 256 as the AES algorithms key size to encrypt data. The default value is AES 128.
DH Group	Select the Diffie-Hellman (DH) Group algorithm to be used when exchanging a pre-shared key. The DH Group sets the strength of the algorithm in bits. The supported DH Groups are 2, 5, and 14. It is recommended to use DH Group 14.
PFS	Select the Perfect Forward Secrecy (PFS) level for additional security. The supported PFS levels are 2 and 5. The default value is 2.

- 8 If you want to create a Secondary VPN Gateway for this site, then click the **Add** button next to **Secondary VPN Gateway**. In the pop-up window, enter the IP address of the Secondary VPN Gateway and click **Save Changes**.

The Secondary VPN Gateway will be created immediately for this site and will provision a VMware VPN tunnel to this Gateway.

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

- 9 Select the **Redundant VeloCloud Cloud VPN** checkbox to add redundant tunnels for each VPN Gateway.

Any changes made to Encryption, DH Group, or PFS of Primary VPN Gateway will also be applied to the redundant VPN tunnels, if configured. After modifying the tunnel settings of the Primary VPN Gateway, save the changes and then click **View IKE/IPSec Template** to view the updated tunnel configuration.
- 10 Click the **Update location** link 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.
- 11 Under **Site Subnets**, you can add subnets for the Non SD-WAN Destination by clicking the **+** button.
- 12 Check the **Enable Tunnel(s)** checkbox once you are ready to initiate the tunnel from the SD-WAN Gateway to the SonicWALL VPN gateways.
- 13 Click **Save Changes**.

Zscaler and VMware SD-WAN Integration

Enterprises can take advantage of secure local Internet breakout by using VMware SD-WAN integrated with Zscaler. Using VMware 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 VMware SD-WAN are:

- Zscaler Internet Access (ZIA)
 - A working instance of ZIA (any cloud)
 - Administrator login credentials
- VMware SD-WAN Orchestrator
 - Enterprise account access to VMware SD-WAN Orchestrator
 - Administrator login credentials
 - One or more VMware SD-WAN Edge appliances with “Online” status in VMware SD-WAN Orchestrator

Zscaler SD-WAN Gateway Selection and Routing Behavior

The VMware SD-WAN Orchestrator configuration process for building tunnels to Zscaler does not require the manual selecting of specific VMware SD-WAN Gateways. Using a geo-IP lookup process, the VMware SD-WAN Gateways are dynamically chosen based on proximity to the provided Zscaler IP endpoint. Operator and Partner Administrators with sufficient permissions can manually override the SD-WAN Orchestrator-default Gateway selections. Normally, this is not necessary, and the recommended best-practice is to accept the SD-WAN Gateways as chosen by the system. After the Zscaler configuration has been completed on the SD-WAN Orchestrator and the tunnels are up and active, Operator and Partner Administrators (with sufficient permissions) can verify which SD-WAN Gateways were chosen. To verify which SD-WAN Gateways were selected, login to the Orchestrator and go to Operator > Gateways. Click on a specific SD-WAN 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 SD-WAN Gateway will be denoted with the *Zscaler_Name* and the redundant SD-WAN Gateway will be denoted as *Zscaler_Name*[redundant].

Primary SD-WAN Gateway

Customer Usage			
	Customer	Pool	Gateway Type
1	SEC	Production	On Premise Gateway ☑ 1 Edge
2	SEC	Production	Secure VPN Gateway ☑ Zscaler
3	SEC	Production	Secure VPN Gateway ☑ Zscaler

Redundant SD-WAN Gateway

Customer Usage			
	Customer	Pool	Gateway Type
1	SEC	Production	Super Gateway ☑
2	SEC	Production	On Premise Gateway ☑ 1 Edge
3	SEC	Production	Secure VPN Gateway ☑ Zscaler [redundant]
4	SEC	Production	Secure VPN Gateway ☑ Zscaler [redundant]

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

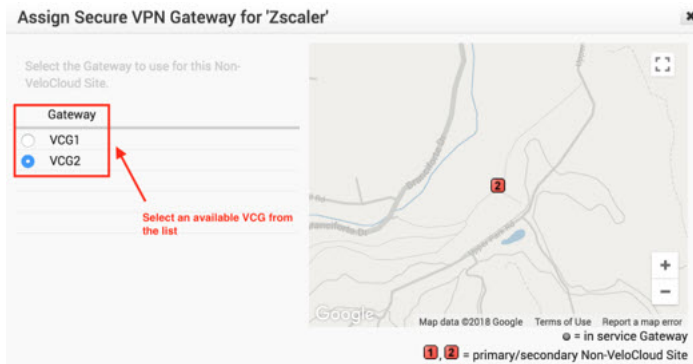
- 1 Locate current tunnel location.

Customer Usage			
	Customer	Pool	Gateway Type
1	SEC	Production	On Premise Gateway ☑ 1 Edge
2	SEC	Production	Secure VPN Gateway ☑ Zscaler
3	SEC	Production	Secure VPN Gateway ☑ Zscaler

- 2 Click on Secure VPN Gateway.

Customer Usage			
	Customer	Pool	Gateway Type
1	SEC	Production	On Premise Gateway ☑ 1 Edge
2	SEC	Production	Secure VPN Gateway ☑ Zscaler
3	SEC	Production	Secure VPN Gateway ☑ Zscaler

3 Select a SD-WAN Gateway.

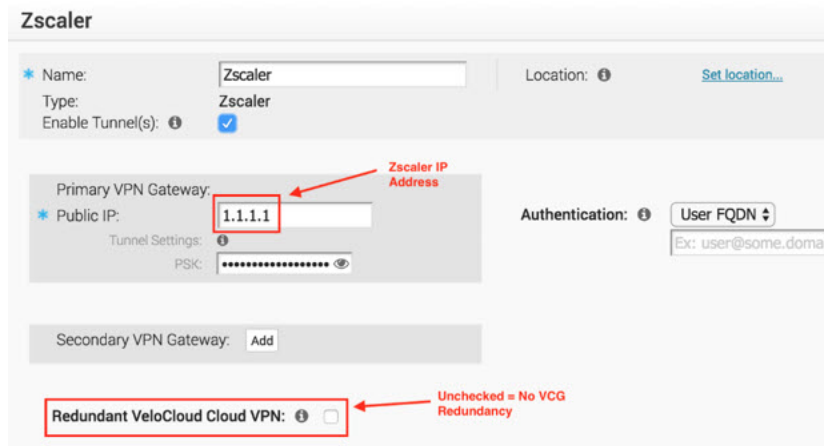



Note Assigning/Moving a tunnel to a different SD-WAN Gateway is service affecting. The existing tunnel connection will terminate and a new tunnel from the newly assigned SD-WAN Gateway will be established.

During the VMware SD-WAN Edge configuration/activation process, each Edge is assigned a pair of cloud SD-WAN Gateways or a set of Partner SD-WAN Gateways, in accordance with the device configuration. If the SD-WAN Gateways used by the Edge are not the same SD-WAN Gateways which contain the Zscaler tunnels, the Edge will automatically build VCMPT tunnels to the SD-WAN Gateways that connect to Zscaler in addition to the SD-WAN 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






In this example, only one Zscaler VPN tunnel is created, and the Redundant Velocloud Cloud VPN checkbox is not selected. A single Gateway (Primary SD-WAN 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 SD-WAN Edge, to the Primary SD-WAN Gateway and then on to Zscaler. Even though the SD-WAN Edge always has VCMP tunnels to at least two SD-WAN Gateways, there is no redundancy in this design. Since the Redundant Velocloud Cloud VPN checkbox is not selected, there will not be a backup SD-WAN Gateway tunnel to Zscaler. If either Zscaler or the primary SD-WAN 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

The screenshot shows the Zscaler configuration interface. The 'Name' field is set to 'Zscaler'. The 'Type' is 'Zscaler'. The 'Enable Tunnel(s)' checkbox is checked. The 'Location' field is empty, with a 'Set location...' link. The 'Primary VPN Gateway' section shows the 'Public IP' set to '1.1.1.1', which is highlighted with a red box and labeled 'Zscaler IP Address'. The 'Tunnel Settings' section shows the 'PSK' field with a masked value. The 'Redundant Tunnel PSK' field is also masked. The 'Authentication' section shows the 'User FQDN' dropdown set to 'User FQDN' and the 'Ex: user@some.domain' field. The 'Secondary VPN Gateway' section has an 'Add' button. At the bottom, the 'Redundant VeloCloud Cloud VPN' checkbox is checked, highlighted with a red box, and labeled 'Checked = VCG Redundancy'.



In this example, only one Zscaler VPN tunnel is created, and the Redundant Velocloud Cloud VPN checkbox is selected. Two SD-WAN 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 SD-WAN Gateway. If the Primary SD-WAN Gateway fails traffic will then shift to the Secondary SD-WAN 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

Zscaler

Name: Location: [Set location...](#)

Type: Zscaler

Enable Tunnel(s): ☒

Primary VPN Gateway:


Public IP: Zscaler Primary IP Address

Authentication: [User FQDN](#)

Secondary VPN Gateway: Zscaler Secondary IP Address

Public IP:

Redundant VeloCloud Cloud VPN: ☐ Unchecked = No VCG Redundancy



In this example, redundant IPsec tunnels to Zscaler is configured in the SD-WAN Orchestrator by adding a secondary Zscaler IP address, however Redundant Velocloud Cloud VPN checkbox is not selected. A single SD-WAN 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 SD-WAN 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 SD-WAN Gateway that originated the request. Should the primary Zscaler location go down, traffic from the SD-WAN Gateway will shift to the secondary Zscaler IPsec tunnel. Since the Redundant Velocloud Cloud VPN checkbox is not selected, there are no redundant SD-WAN Gateway connections to Zscaler. If the SD-WAN 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

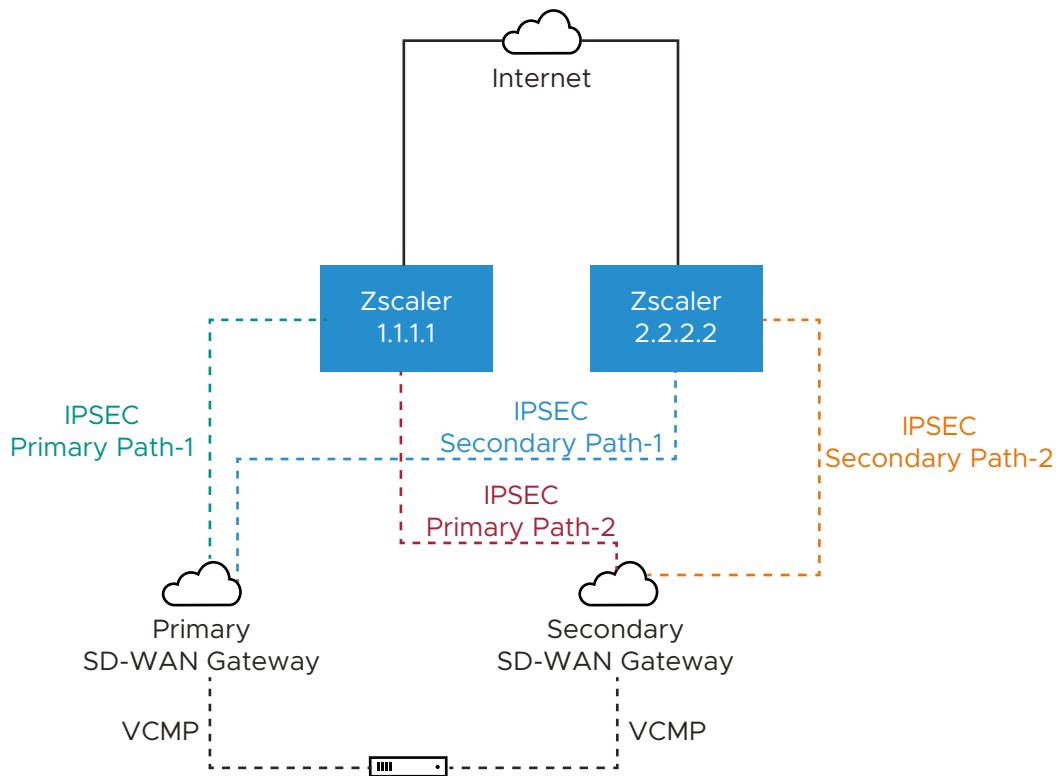
Zscaler

* Name: Location: ⓘ [Set location...](#)
 Type: Zscaler
 Enable Tunnel(s): ⓘ ☒

Primary VPN Gateway:
 * Public IP: Zscaler Primary IP Address
 Authentication: ⓘ
 Ex: user@some.domain

Secondary VPN Gateway: Zscaler Secondary IP Address
 Public IP:

Redundant VeloCloud Cloud VPN: ⓘ ☒ Checked = VCG Redundancy



In this example, redundant IPsec tunnels to Zscaler is configured in the SD-WAN Orchestrator by adding a secondary Zscaler IP address and Redundant Velocloud Cloud VPN checkbox is selected. Two SD-WAN 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 SD-WAN Gateways knows which of the two is the primary SD-WAN Gateway and which is secondary. The SD-WAN 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 SD-WAN Gateway that originated the request. Should the primary Zscaler location go down,

traffic from the primary SD-WAN Gateway will shift to the secondary Zscaler IPsec tunnel. Since the Redundant Velocloud Cloud VPN checkbox is selected, if the primary SD-WAN Gateway fails traffic will shift to the secondary SD-WAN Gateway. The secondary SD-WAN 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 datacenter for Zscaler Internet Access (ZIA), the tunnel is established between the SD-WAN Edge or SD-WAN 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 SD-WAN Edge or SD-WAN Gateway sends probe requests periodically to the HTTP probe URL (in the following format) if probe is enabled.

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

The probe URL is configurable in the SD-WAN Orchestrator, but the probe interval and number of retries are currently not editable in the SD-WAN 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 VMware SD-WAN Deployment Configurations

Describes the configuration steps for integrating Zscaler Internet Access (ZIA) and VMware 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 more information, see <https://www.zscaler.com/resources/solution-briefs/partner-vmware-sdwan-deployment-guide.pdf>. This guide will provide GUI examples for configuring Zscaler Internet Access and VMware SD-WAN Orchestrator.

Layer 7 health check Events

Event	Displayed on Orchestrator UI as	Severity	Notification Configurable	Generated By	Generated When
EDGE_NV_S_TUNNEL_UP	Edge Direct IPsec tunnel up	INFO	N	SD-WAN Orchestrator	A Cloud Security Service tunnel or NSD via Edge tunnel is up.
EDGE_NV_S_TUNNEL_DOWN	Edge Direct IPsec tunnel down	INFO	N	SD-WAN Orchestrator	A Cloud Security Service tunnel or NSD via Edge tunnel is down.
VPN_DATACENTER_STATUS	VPN Tunnel state change	NOTICE	N	SD-WAN Gateway	The VPN Tunnel state is changed.

For information about events related to cloud security services, see [Monitor Cloud Security Services Events](#).

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 SD-WAN Orchestrator, go to **Configure > Network Services**. The **Services** screen appears.
- 2 In the **Non SD-WAN Destinations via Gateway** area, click the **New** button.

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

New Non SD-WAN Destination via Gateway...

* Name: Velo NVS

* Type: Zscaler

VPN Gateways

* Primary VPN Gateway: 10.10.10.1

Secondary VPN Gateway: Ex: 54.183.9.192

Next

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

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

The screenshot shows the 'Velo NVS' configuration window. It includes fields for Name (Velo NVS), Type (Zscaler), Enable Tunnel(s) (checked), and Tunnel mode (Active/Hot-Standby). The Location is set to Lat,Lng: 37.402889, -122.116859. The Primary VPN Gateway Public IP is 10.10.10.1. The Local Auth Id is set to FQDN with the value uddhav.velocloud.net. The Zscaler Login URL is https://admin.zscaler.net, and there is a 'Login to Zscaler' button. The L7 Health Check is checked. The Zscaler Cloud is set to zscaler.net. The HTTP Probe Interval is 5 seconds, the Number of Retries is 3, and the RTT Threshold is 3000 milliseconds. At the bottom, there are buttons for 'Advanced', 'View IKE/IPSec Template', 'Save Changes', and 'Close'.

- To configure tunnel settings for the Non SD-WAN Destination's Primary VPN Gateway, click the **Advanced** button.
- 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 textbox.

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.

- If you want to create a Secondary VPN Gateway for this site, then click the **Add** button next to **Secondary VPN Gateway**. In the pop-up window, enter the IP address of the Secondary VPN Gateway and click **Save Changes**. The Secondary VPN Gateway will be created immediately for this site and will provision a VMware VPN tunnel to this Gateway.
- Select the **Redundant VeloCloud Cloud VPN** checkbox 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 Primary VPN Gateway, save the changes and then click **View IKE/IPSec Template** to view the updated tunnel configuration.

- 10 Click the **Update location** link 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.

- 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 IP 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** checkbox to enable L7 Health check for the Zscaler Cloud Security Service provider, with default probe details (HTTP Probe interval = 5 seconds, Number of Retries = 3, RTT Threshold = 3000 milliseconds). By default, L7 Health Check is 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 textbox.
- 13 To login to Zscaler portal from here, enter the login URL in the **Zscaler Login URL** textbox and then click **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 more information, see [Configure a Cloud Security Service](#).

- 14 Check the **Enable Tunnel(s)** checkbox once you are ready to initiate the tunnel from the SD-WAN Gateway to the Zscaler VPN gateways.
- 15 Click **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 more 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**.

	Name	Public IP	Status	Tunnel Status	Service Status	Used By	Last Contact	Events
1	Craig-GW-ZS Zscaler	199.168.148.132 104.129.194.39				1 Profile	Thu Apr 08, 09:53:31	View
2	Craig-GW2-ZS Zscaler	104.129.194.39			N/A	0	Fri Mar 12, 09:13:27	View

Associate a Non SD-WAN Destination to a Configuration Profile

After configuring a Non SD-WAN Destination of type **Zscaler** in SD-WAN Orchestrator, you have to associate the Non SD-WAN Destination to the desired Profile in order to establish the tunnels between SD-WAN Gateways and Zscaler VPN Gateways. To associate a Non SD-WAN Destination to a configuration profile, perform the following steps:

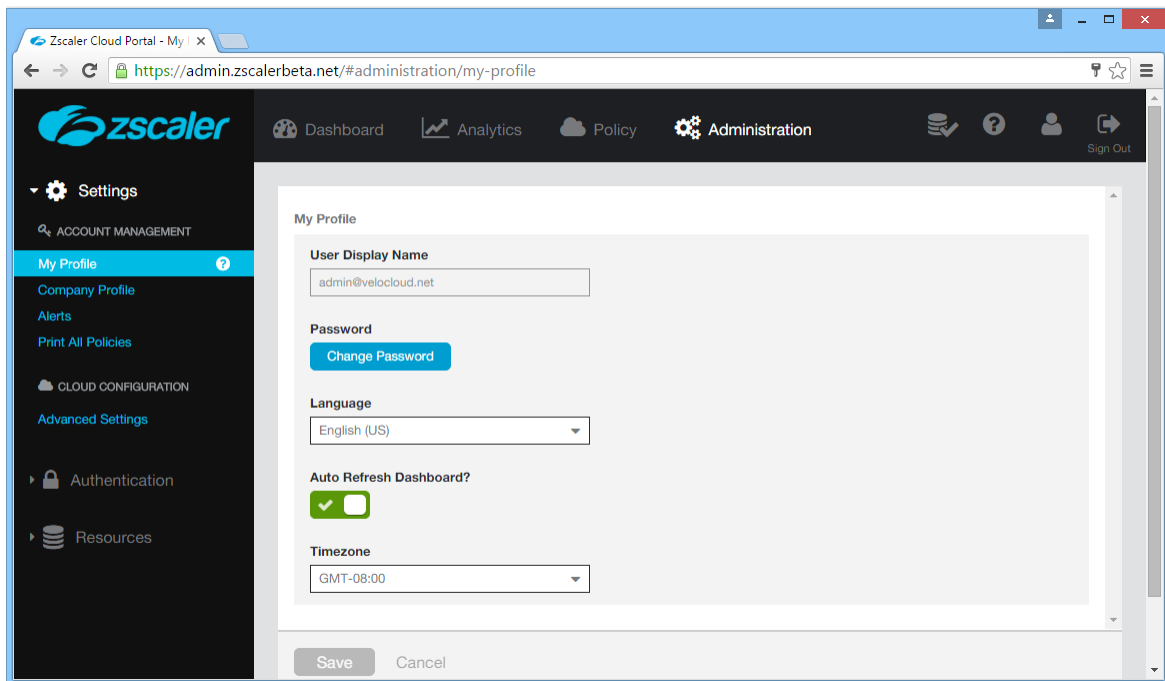
- 1 From the SD-WAN Orchestrator navigation panel, go to **Configure > Profiles**. The **Configuration Profiles** page appears.
- 2 Select a profile you want to associate your Non SD-WAN Destination of type **Zscaler** and click the icon under the **Device** column. The **Device Settings** page for the selected profile appears.
- 3 Go to **Cloud VPN** area and enable Cloud VPN by turning the toggle button to **On**.
- 4 Under **Branch to Non SD-WAN Destinations via Gateway**, select the **Enable** checkbox.
- 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 Click **Save Changes**.

Configure Zscaler

This section describes Zscaler configuration.

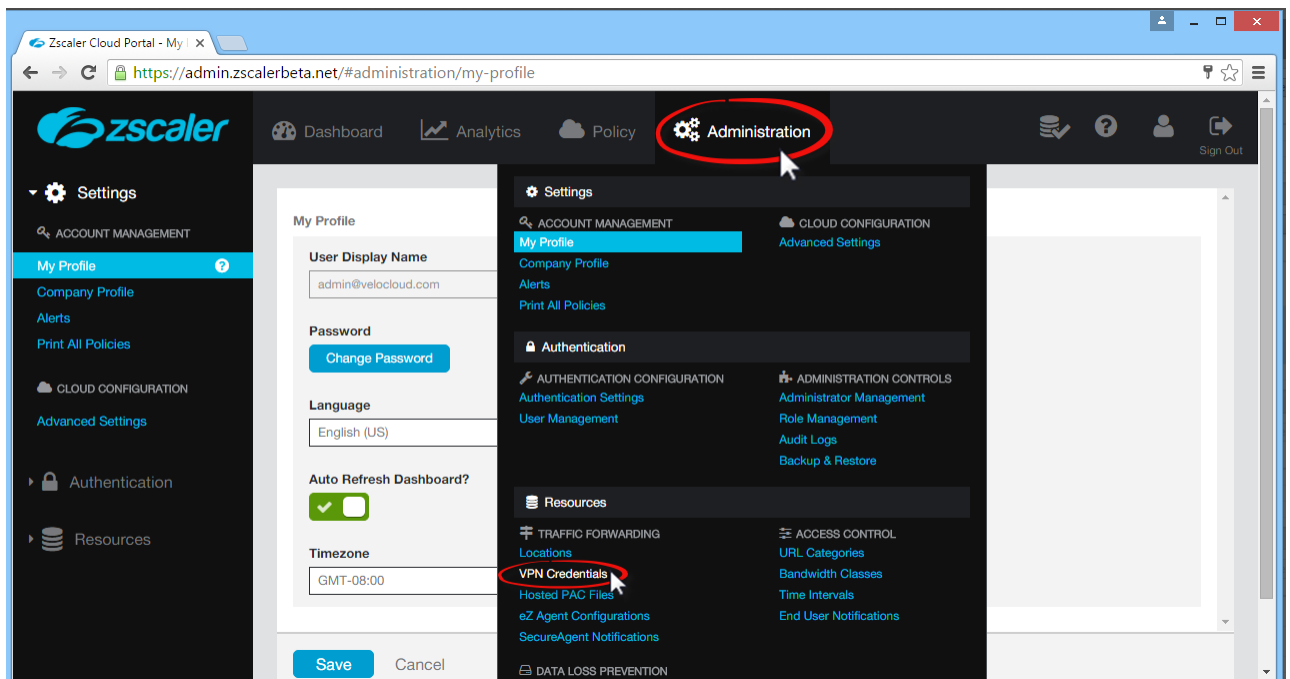
Complete the following these steps on the Zscaler website. From there, you will create a Zscaler account, add VPN credentials, and add a location.

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




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**, click **VPN Credentials**.



- c Click **Add VPN Credentials** at the top left corner.



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 SD-WAN Orchestrator.
- 3 If necessary, type in any comments in the **Comments** section.

 A screenshot of the 'Add VPN Credential' dialog box. The dialog has a title bar 'Add VPN Credential' with a close button. Inside, there is a section 'VPN Credential' with the following fields:

- Authentication Type:** Three buttons: 'FQDN' (selected and highlighted with a red box), 'XAUTH', and 'IP'.
- User ID:** A text box containing 'velocloud01' and a dropdown menu showing 'velocloud.com'.
- New Pre-Shared Key:** A text box with masked characters '*****'.
- Confirm New Pre-Shared Key:** A text box with masked characters '*****'.
- Comments:** A text area containing the text 'The PSK and User ID FQDN was obtained from the VeloCloud portal when the Non-VeloCloud Site was created.' (This text area is circled in red).

 At the bottom of the dialog are two buttons: 'Save' and 'Cancel'.

- 4 Click **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**, click **Locations**.
 - c Click **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 Click **Done**.
 - 5 Click **Save**.

Add Location

Location

Name
VeloCloud Admin

Country
United States

State/Province
San Jose, CA

Time Zone
America/Los Angeles

Addressing

Public IP Addresses
None

VPN Credentials
velocloud01@velocloud.com

Unselected Items	Selected Items (1)
<input checked="" type="checkbox"/> velocloud01@velocloud.com	velocloud01@velocloud.com

Done **Clear Selection**

Save **Cancel**

Configure Business Priority Rules

Define the business policy in your SD-WAN 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.

To create a business policy:

- 1 From the navigation panel in the SD-WAN Orchestrator, go to **Configure > Edges**.
- 2 In the **Edges** screen, click the **Biz. Pol** icon for your Edge.
- 3 Click the **New Rule** button. The **Configure Rule** dialog box appears.

?

x

Configure Rule

Rule Name

BiZ policy Zscaler

Match

Source

Any

Object Group

Define...

Destination

Any

Object Group

Define...

☐ Any
 ☒ Internet
 ☐ Edge
 ☐ Non SD-WAN Destination via Gateway
 ☐ Non SD-WAN Destination via Edge

IP Address

Ex: 10.0.2.0

CIDR prefix

24

Domain Name

Ex: domain.com

Protocol

Ports

Ex: 2224-2226

Application

Any

Define...

Action

Priority

High

Normal

Low

☐ Rate Limit

Network Service

Direct

Multi-Path

Internet Backhaul

☐ Backhaul Hubs
 ☒ Non SD-WAN Destination via Gateway
 ☐ Non SD-WAN Destination via Edge / Cloud Security Service

Zscaler Site1

Link Steering

Auto

Transport Group

Interface

WAN Link

Inner Packet DSCP Tag

Leave as is

Outer Packet DSCP Tag

0 - CS0/DF

NAT

Disabled

Enabled

Service Class

Real Time

Transactional

Bulk

OK

Cancel

- In the **Rule Name** textbox, enter a name for the rule.
- Under the **Match** area, configure the match conditions for the rule.

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

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

- d Click **OK**.

For more information about how to create a business policy rule, see [Create Business Policy Rules](#).

Configure a Non SD-WAN Destination of Type Generic IKEv1 Router via Gateway

Describes how to configure a Non SD-WAN Destination of type **Generic IKEv1 Router (Route Based VPN)** through SD-WAN Gateway in SD-WAN Orchestrator.

Note To configure a **Generic IKEv1 Router (Route Based VPN)** via Edge, see [Configure a Non-VMware SD-WAN Site of Type Generic IKEv1 Router via Edge](#).

Procedure

- 1 From the navigation panel in the SD-WAN Orchestrator, go to **Configure > Network Services**.
The **Services** screen appears.
- 2 In the **Non SD-WAN Destinations via Gateway** area, click the **New** button.
The **New Non SD-WAN Destinations via Gateway** dialog box appears.
- 3 In the **Name** text box, enter the name for the Non SD-WAN Destination.
- 4 From the **Type** drop-down menu, select **Generic IKEv1 Router (Route Based VPN)**.
- 5 Enter the IP address for the Primary VPN Gateway (and the Secondary VPN Gateway if necessary), and click **Next**.

A route-based Non SD-WAN Destination of type IKEv1 is created and a dialog box for your Non SD-WAN Destination appears.

NSD_IKEv1

NameNSD_IKEv1

TypeGeneric IKEv1 Router (Route Based VPN)

Enable Tunnel(s)☒

Tunnel modeActive/Hot-Standby

Location ⓘ

Lat,Lng: 37.402889, -122.116859

[Update Location...](#)

Primary VPN Gateway

Public IP54.183.9.198

Tunnel Settings ⓘ

PSK

Redundant Tunnel PSK

EncryptionAES 128

DH Group2

PFS2

Local Auth Id ⓘ

Default

Site Subnets ⓘ

Subnet	Description	Advertise
Ex: 10.0.2.0/24	(optional)	<input checked="" type="checkbox"/>

Deactivate Site Subnets ⓘ ☐

Secondary VPN Gateway Remove

Public IP54.183.9.199

Tunnel Settings ⓘ

PSK

Redundant Tunnel PSK

EncryptionAES 128

DH Group2

PFS2

Redundant VeloCloud Cloud VPN ⓘ ☒

Advanced

View IKE/IPSec Template

Save Changes

Close

- 6 To configure tunnel settings for the Non SD-WAN Destination's Primary VPN Gateway, click the **Advanced** button.
- 7 In the **Primary VPN Gateway** area, you can configure the following tunnel settings:

Field	Description
Tunnel Mode	Active-Hot-Standby is supported on the SD-WAN Gateway. Active/Hot-Standby automatically displays indicating that if the Active tunnel goes down, the Standby (Hot-Standby) tunnel takes over and becomes the Active tunnel.
PSK	<p>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 textbox.</p> <p>Note Starting from the 4.5 release, the use of the special character "<" in the password is no longer supported. In cases where users have already used "<" in their passwords in previous releases, they must remove it to save any changes on the page.</p>

Field	Description
Encryption	Select either AES 128 or AES 256 as the AES algorithms key size to encrypt data. The default value is AES 128.
DH Group	Select the Diffie-Hellman (DH) Group algorithm to be used when exchanging a pre-shared key. The DH Group sets the strength of the algorithm in bits. The supported DH Groups are 2, 5, and 14. It is recommended to use DH Group 14.
PFS	Select the Perfect Forward Secrecy (PFS) level for additional security. The supported PFS levels are 2 and 5. The default value is 2.

- 8 If you want to create a Secondary VPN Gateway for this site, then click the **Add** button next to **Secondary VPN Gateway**. In the pop-up window, enter the IP address of the Secondary VPN Gateway and click **Save Changes**.

The Secondary VPN Gateway will be created immediately for this site and will provision a VMware VPN tunnel to this Gateway.

- 9 Select the **Redundant VeloCloud Cloud VPN** checkbox to add redundant tunnels for each VPN Gateway.

Any changes made to Encryption, DH Group, or PFS of Primary VPN Gateway will also be applied to the redundant VPN tunnels, if configured. After modifying the tunnel settings of the Primary VPN Gateway, save the changes and then click **View IKE/IPsec Template** to view the updated tunnel configuration.

- 10 Click the **Update location** link 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.

- 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 IP address used to communicate with the local gateway.

Note For Generic route based VPN, if the user does not specify a value, **Default** is used as the local authentication ID. The default local authentication ID value will be the SD-WAN Gateway Interface Public IP.

- 12 Under **Site Subnets**, you can add subnets for the Non SD-WAN Destination by clicking the **+** button. If you do not need subnets for the site, select the **Deactivate Site Subnets** checkbox.
- 13 Check the **Enable Tunnel(s)** checkbox once you are ready to initiate the tunnel from the SD-WAN Gateway to the Generic IKEv1 VPN Gateways.

- 14 Click **Save Changes**.
- 15 Assign the newly created Non SD-WAN Site Network Service to a Profile by navigating to **Configure > Profiles** in the SD-WAN Orchestrator. See [Configure a Tunnel Between a Branch and a Non SD-WAN Destinations via Gateway](#).
- 16 Return to the **Non SD-WAN Destinations via Gateway** area in the SD-WAN Orchestrator by going to **Configure > Network Services**.
- 17 In the **Non SD-WAN Destinations via Gateway** area, scroll to the name of your Non SD-WAN Site, and then click the **Edit** link in the **BGP** column.
- 18 Configure the BGP values for the following mandatory fields: Local ASN, Tunnel Type, Neighbor IP, and Local IP (from the Advanced Options section). For more information, see [Configure BGP over IPsec from Gateways](#).

Note The VTI IP (private IP) assigned by the SD-WAN Orchestrator can be used for peer ship in Single-Hop BGP.

- 19 Click **OK** to save your changes.
- 20 In the **Non SD-WAN Destinations via Gateway** area, click the **Edit** link in the **BFD** column for a Non SD-WAN Destination, to configure the BFD settings. For more information, see [Configure BFD for Gateways](#).

What to do next

You can check the overall status of the Non SD-WAN Sites in the monitoring tab. See:

- [Monitor Network Services](#)
- [Monitor Non SD-WAN Destinations through Gateway](#)

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

Describes how to configure a Non SD-WAN Destination of type **Generic Firewall (Policy Based VPN)** in SD-WAN Orchestrator.

Procedure

- 1 From the navigation panel in the SD-WAN Orchestrator, go to **Configure > Network Services**.
The **Services** screen appears.
- 2 In the **Non SD-WAN Destinations via Gateway** area, click the **New** button.
The **New Non SD-WAN Destinations via Gateway** dialog box appears.
- 3 In the **Name** text box, enter the name for the Non SD-WAN Destination.
- 4 From the **Type** drop-down menu, select **Generic Firewall (Policy Based VPN)**.

- 5 Enter the IP address for the Primary VPN Gateway, and click **Next**.

A Non SD-WAN Destination of type Generic Firewall (Policy Based VPN) is created and a dialog box for your Non SD-WAN Destination appears.

- 6 To configure tunnel settings for the Non SD-WAN Destination's Primary VPN Gateway, click the **Advanced** button.
- 7 In the **Primary VPN Gateway** area, you can configure the following tunnel settings:

Field	Description
PSK	<p>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 textbox.</p> <p>Note Starting from the 4.5 release, the use of the special character "<" in the password is no longer supported. In cases where users have already used "<" in their passwords in previous releases, they must remove it to save any changes on the page.</p>
Encryption	<p>Select either AES 128 or AES 256 as the AES algorithms key size to encrypt data. The default value is AES 128.</p>

Field	Description
DH Group	Select the Diffie-Hellman (DH) Group algorithm to be used when exchanging a pre-shared key. The DH Group sets the strength of the algorithm in bits. The supported DH Groups are 2, 5, and 14. It is recommended to use DH Group 14.
PFS	Select the Perfect Forward Secrecy (PFS) level for additional security. The supported PFS levels are 2 and 5. The default value is deactivated .

Note The Secondary VPN Gateway are not supported for the Generic Firewall (Policy Based VPN) network service type.

- 8 Select the **Redundant VeloCloud Cloud VPN** checkbox to add redundant tunnels for each VPN Gateway.

Any changes made to Encryption, DH Group, or PFS of Primary VPN Gateway will also be applied to the redundant VPN tunnels, if configured. After modifying the tunnel settings of the Primary VPN Gateway, save the changes and then click **View IKE/IPSec Template** to view the updated tunnel configuration.

Note Currently, the supported IKE version is **IKEv1**.

- 9 Click the **Update location** link 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.
- 10 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 IP address used to communicate with the local gateway.

Note For Generic Firewall (Policy based VPN), if the user do not specify a value, **Default** is used as the local authentication ID. The default local authentication ID value will be the SD-WAN Gateway Interface Local IP.

- 11 Under **Site Subnets**, you can add subnets for the Non SD-WAN Destination by clicking the **+** button. If you do not need subnets for the site, select the **Deactivate Site Subnets** checkbox.
- 12 Use **Custom Source Subnets** to override the source subnets routed to this VPN device. Normally, source subnets are derived from the edge LAN subnets routed to this device.
- 13 Check the **Enable Tunnel(s)** checkbox once you are ready to initiate the tunnel from the SD-WAN Gateway to the Generic Firewall (Policy Based VPN) VPN gateways.

14 Click *Save Changes*.**Configure Amazon Web Services**

VMware 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-VMware Network Service in the SD-WAN Orchestrator.

Obtain Amazon Web Services Configuration Details

Describes 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 SD-WAN Gateways associated with the enterprise account in the SD-WAN 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 SD-WAN Orchestrator when you create a Non SD-WAN Destination.

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 via Gateway](#)
- [Configure a Non SD-WAN Destination of Type Generic IKEv2 Router via Gateway](#)

To configure a Non SD-WAN Destination via Edge, see:

- [Configure a Non-VMware SD-WAN Site of Type Generic IKEv1 Router via Edge](#)
- [Configure a Non-VMware SD-WAN Site of Type Generic IKEv2 Router via Edge](#)

Configure a Non SD-WAN Destinations via Edge

VMware allows the Enterprise users to define and configure a Non SD-WAN Destination instance and establish a secure IPSec tunnel directly from a SD-WAN Edge to a Non SD-WAN Destination.


Note VMware supports only Generic IKEv2 Router (Route Based VPN) and Generic IKEv1 Router (Route Based VPN) Non SD-WAN Destination from Edge. This will enable the Edge to establish an IPSec tunnel to AWS datacenter or Azure datacenter. Currently, VMware only verifies IPSec tunnel support to AWS and Azure datacenters.

To configure a Non SD-WAN Destinations via Edge:

Procedure

- 1 From the navigation panel in the SD-WAN Orchestrator, go to **Configure > Network Services**.
The **Services** screen appears.

- 2 In the **Non SD-WAN Destinations via Edge** area, click the **New** button.
The **Non SD-WAN Destinations via Edge** dialog box appears.



- 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 either Generic IKEv2 Router (Route Based VPN) or Generic IKEv1 Router (Route Based VPN) as the IPsec tunnel type.
- 5 Click **Next**.

A Non SD-WAN Destination is created.

Note To support the datacenter type of Non SD-WAN Destination, besides the IPsec connection, you will need to configure Non SD-WAN Destination local subnets into the VMware system.

What to do next

- Configure tunnel settings for your Non SD-WAN Destination. For more information, see:
 - [Configure a Non-VMware SD-WAN Site of Type Generic IKEv1 Router via Edge](#)
 - [Configure a Non-VMware SD-WAN Site of Type Generic IKEv2 Router via Edge](#)
- Associate your Non SD-WAN Destination to a profile or Edge. For more information, see [Configure a Tunnel Between a Branch and a Non SD-WAN Destinations via Edge](#).
- Configure Tunnel parameters (WAN link selection and Per tunnel credentials) at the Edge level. For more information, see [Configure Cloud VPN and Tunnel Parameters at the Edge level](#).
- 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 more information, see [Create Business Policy Rules](#).

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

Describes how to configure a Non SD-WAN Destination of type **Generic IKEv1 Router (Route Based VPN)** through SD-WAN Edge in SD-WAN Orchestrator.

Procedure

- 1 From the navigation panel in the SD-WAN Orchestrator, go to **Configure > Network Services**.

The **Services** screen appears.

- 2 In the **Non SD-WAN Destinations via Edge** area, click the **New** button.

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

- 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 Click **Next**.

A route-based Non SD-WAN Destination of type IKEv1 is created and a dialog box for your Non SD-WAN Destination appears.

Non SD-WAN Destinations via Edge

Name: Gen_IKEv1

Type: Generic IKEv1 Router(Router Based VPN)

Tunnel mode: Active/Active

Primary VPN Gateway

Public IP: 10.0.0.25

Encryption: AES 128

DH Group: 14

PFS: Deactivated

Hash: SHA 256

IKE SA Lifetime(min): 1440

IPsec SA Lifetime(min): 480

DPD Timeout Timer(sec): 20

Secondary VPN Gateway: ☒

Keep Tunnel Active: ☒

Public IP: 10.0.0.5

☒ Tunnel settings are same as Primary VPN Gateway

Encryption: AES 128

DH Group: 14

PFS: Deactivated

Hash: SHA 256

IKE SA Lifetime(min): 1440

IPsec SA Lifetime(min): 480

DPD Timeout Timer(sec): 20

Site Subnets

Subnet	Description	Advertise
Ex: 10.0.2.0/24	(optional)	<input checked="" type="checkbox"/>

Advanced Save Changes Close

- 6 Under **Primary VPN Gateway**, in the **Public IP** text box, enter the IP address of the Primary VPN Gateway.
- 7 To configure tunnel settings for the Non SD-WAN Destination's Primary VPN Gateway, click the **Advanced** button.

8 In the **Primary VPN Gateway** area, you can configure the following tunnel settings:

Field	Description
Encryption	Select either AES 128 or AES 256 as the AES algorithms key size to encrypt data. If you do not want to encrypt data, select Null . The default value is AES 128.
DH Group	Select the Diffie-Hellman (DH) Group algorithm to be used when exchanging a pre-shared key. The DH Group sets the strength of the algorithm in bits. The supported DH Groups are 2, 5, 14, 15, and 16. It is recommended to use DH Group 14.
PFS	Select the Perfect Forward Secrecy (PFS) level for additional security. The supported PFS levels are 2, 5, 14, 15, and 16. The default value is Deactivated.
Hash	The authentication algorithm for the VPN header. Select one of the following supported Secure Hash Algorithm (SHA) function from the list: <ul style="list-style-type: none"> ■ SHA 1 ■ SHA 256 ■ SHA 384 ■ SHA 512 The default value is SHA 256.
IKE SA Lifetime(min)	Time when Internet Key Exchange (IKE) rekeying is initiated for Edges. The minimum IKE life time is 10 minutes and maximum is 1440 minutes. The default value is 1440 minutes.
IPsec SA Lifetime(min)	Time when Internet Security Protocol (IPsec) rekeying is initiated for Edges. The minimum IPsec life time is 3 minutes and maximum is 480 minutes. The default value is 480 minutes.
DPD Timeout Timer(sec)	The maximum time that the device should wait to receive a response to the DPD message before considering the peer to be dead. The default value is 20 seconds. You can deactivate DPD by configuring the DPD timeout timer to 0 second.

Note When AWS initiates the rekey tunnel with a VMware SD-WAN 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 SD-WAN 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.

- 9 If you want to create a Secondary VPN Gateway for this site, then select the **Secondary VPN Gateway** checkbox and then enter the IP address of the Secondary VPN Gateway in the **Public IP** text box.

The Secondary VPN Gateway will be created immediately for this site and will provision a VMware VPN tunnel to this Gateway.

- 10 Select the **Keep Tunnel Active** checkbox to keep the Secondary VPN tunnel active for this site.
- 11 Select the **Tunnel settings are same as Primary VPN Gateway** checkbox to apply the same tunnel settings as that of the Primary VPN Gateway.

Any tunnel setting changes made to the Primary VPN Gateway will also be applied to the Secondary VPN tunnels, if configured.

- 12 Under **Site Subnets**, you can add subnets for the Non SD-WAN Destination by clicking the **+** button.

Note To support the datacenter type of Non SD-WAN Destination, besides the IPSec connection, you will need to configure Non SD-WAN Destination local subnets into the VMware system.

- 13 Click **Save Changes**.

What to do next

- [Configure a Tunnel Between a Branch and a Non SD-WAN Destinations via Edge](#)
- [Configure Cloud VPN and Tunnel Parameters at the Edge Level](#)

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

Describes how to configure a Non SD-WAN Destination of type **Generic IKEv2 Router (Route Based VPN)** through SD-WAN Edge in SD-WAN Orchestrator.

Procedure

- 1 From the navigation panel in the SD-WAN Orchestrator, go to **Configure > Network Services**.
The **Services** screen appears.
- 2 In the **Non SD-WAN Destinations via Edge** area, click the **New** button.
The **Non SD-WAN Destinations via Edge** dialog box appears.
- 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 Click **Next**.

A route-based Non SD-WAN Destination of type IKEv2 is created and a dialog box for your Non SD-WAN Destination appears.

Non SD-WAN Destinations via Edge

* Name: Gen_IKEv2
 Type: Generic IKEv2 Router(Router Based VPN)
 Tunnel mode: Active/Hot-Standby

Primary VPN Gateway
 Public IP: 10.0.0.5
 Encryption: AES 128
 DH Group: 14
 PFS: Deactivated
 Hash: SHA 256
 IKE SA Lifetime(min): 1440
 IPsec SA Lifetime(min): 480
 DPD Timeout Timer(sec): 20

Secondary VPN Gateway
☒ Keep Tunnel Active
☒ Tunnel settings are same as Primary VPN Gateway
 Public IP: 10.0.2.5
 Encryption: AES 128
 DH Group: 14
 PFS: Deactivated
 Hash: SHA 256
 IKE SA Lifetime(min): 1440
 IPsec SA Lifetime(min): 480
 DPD Timeout Timer(sec): 20

Site Subnets

Subnet	Description	Advertise
Ex: 10.0.2.0/24	(optional)	<input checked="" type="checkbox"/>

Advanced Save Changes Close

- 6 Under **Primary VPN Gateway**, in the **Public IP** text box, enter the IP address of the Primary VPN Gateway.
- 7 To configure tunnel settings for the Non SD-WAN Destination's Primary VPN Gateway, click the **Advanced** button.
- 8 In the **Primary VPN Gateway** area, you can configure the following tunnel settings:

Field	Description
Encryption	Select either AES 128 or AES 256 as the AES algorithms key size to encrypt data. If you do not want to encrypt data, select Null . The default value is AES 128.
DH Group	Select the Diffie-Hellman (DH) Group algorithm to be used when exchanging a pre-shared key. The DH Group sets the strength of the algorithm in bits. The supported DH Groups are 2, 5, 14, 15, and 16. It is recommended to use DH Group 14.
PFS	Select the Perfect Forward Secrecy (PFS) level for additional security. The supported PFS levels are 2, 5, 14, 15, and 16. The default value is Deactivated.

Field	Description
Hash	<p>The authentication algorithm for the VPN header. Select one of the following supported Secure Hash Algorithm (SHA) function from the list:</p> <ul style="list-style-type: none"> ■ SHA 1 ■ SHA 256 ■ SHA 384 ■ SHA 512 <p>The default value is SHA 256.</p>
IKE SA Lifetime(min)	<p>Time when Internet Key Exchange (IKE) rekeying is initiated for Edges. The minimum IKE life time is 10 minutes and maximum is 1440 minutes. The default value is 1440 minutes.</p>
IPsec SA Lifetime(min)	<p>Time when Internet Security Protocol (IPsec) rekeying is initiated for Edges. The minimum IPsec life time is 3 minutes and maximum is 480 minutes. The default value is 480 minutes.</p>
DPD Timeout Timer(sec)	<p>The maximum time that the device should wait to receive a response to the DPD message before considering the peer to be dead. The default value is 20 seconds. You can deactivate DPD by configuring the DPD timeout timer to 0 second.</p>

Note When AWS initiates the rekey tunnel with a VMware SD-WAN 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 SD-WAN 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.

- 9 If you want to create a Secondary VPN Gateway for this site, then select the **Secondary VPN Gateway** checkbox and then enter the IP address of the Secondary VPN Gateway in the **Public IP** text box.

The Secondary VPN Gateway will be created immediately for this site and will provision a VMware VPN tunnel to this Gateway.

- 10 Select the **Keep Tunnel Active** checkbox to keep the Secondary VPN tunnel active for this site.
- 11 Select the **Tunnel settings are same as Primary VPN Gateway** checkbox to apply the same tunnel settings as that of the Primary VPN Gateway.

Any tunnel setting changes made to the Primary VPN Gateway will also be applied to the Secondary VPN tunnels, if configured.

- 12 Under **Site Subnets**, you can add subnets for the Non SD-WAN Destination by clicking the **+** button.

Note To support the datacenter type of Non SD-WAN Destination, besides the IPSec connection, you will need to configure Non SD-WAN Destination local subnets into the VMware system.

- 13 Click **Save Changes**.

What to do next

- [Configure a Tunnel Between a Branch and a Non SD-WAN Destinations via Edge](#)
- [Configure Cloud VPN and Tunnel Parameters at the Edge Level](#)

Configure a Microsoft Azure Non SD-WAN Destination via Edge

Describes how to configure a Non SD-WAN Destination of type **Microsoft Azure Virtual Hub** via Edge in SD-WAN Orchestrator.

To configure a Non SD-WAN Destination of type **Microsoft Azure Virtual Hub** via Edge in SD-WAN Orchestrator:

Prerequisites

- Ensure you have configured a Cloud subscription. For steps, see [Configure a Cloud Subscription Network Service](#).
- Ensure you have created Virtual WAN and Hubs in Azure. For steps, see [Configure Azure Virtual WAN for Branch-to-Azure VPN Connectivity](#).

Procedure

- 1 From the navigation panel in the SD-WAN Orchestrator, go to **Configure > Network Services**.
The **Services** screen appears.
- 2 In the **Non SD-WAN Destinations via Edge** area, click the **New** button.
The **New Non SD-WAN Destinations via Edge** dialog box appears.

Non SD-WAN Destinations via Edge

* Service Name: Azure auto

* Service Type: Microsoft Azure Virtual Wan

Virtual Hub Configuration

* Subscription: VeloCloud-ManagementPlane-[...]

* Virtual Wan: azure_automation_test_vwan

Resource Group: azure_vwan_automation

* Virtual Hub: azure_test_vhub_west_us_2

Azure Regions: West US 2

Enable Tunnels: ☐

Next

3 In the **Service Name** text box, enter the name for the Non SD-WAN Destination.

4 From the **Service Type** drop-down menu, select **Microsoft Azure Virtual Hub**.

5 From the **Subscription** drop-down menu, select a cloud subscription.

The application fetches all the available Virtual WANs dynamically from Azure.

6 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.

7 From the **Virtual Hub** drop-down menu, select a Virtual Hub.

The application auto-populates the Azure region corresponding to the Hub

8 Click **Next**.

The Microsoft Azure Non SD-WAN Destination is created and a dialog box for your Non SD-WAN Destination appears.

Non SD-WAN Destinations via Edge

Name Azure auto
Type Microsoft Azure Virtual Wan

Primary VPN Gateway
Public IP Ex: 10.0.2.5

Secondary VPN Gateway ☒
Keep Tunnel Active ☒
Public IP Ex: 10.0.2.5
☒ Tunnel settings are same as Primary VPN Gateway

Site Subnets

Subnet	Description	Advertise
Ex: 10.0.2.0/24	(optional)	<input checked="" type="checkbox"/>

Subscription VeloCloud-ManagementPlane-Dev
Virtual Wan azure_automation_test_vwan
Resource Group azure_vwan_automation
Virtual Hub azure_test_vhub_west_us_2
Azure Regions West US 2

Advanced **Save Changes** **Close**

- 9 To configure tunnel settings for the Non SD-WAN Destination's Primary VPN Gateway, click the **Advanced** button.

Non SD-WAN Destinations via Edge

Name Azure auto
Type Microsoft Azure Virtual Wan

Primary VPN Gateway
Public IP Ex: 10.0.2.5
Encryption AES 128
DH Group 2
PFS Deactivated
Hash SHA 256
IKE SA Lifetime(min) 1440
IPsec SA Lifetime(min) 480
DPD Timeout Timer(sec) 20

Secondary VPN Gateway ☒
Keep Tunnel Active ☒
Public IP Ex: 10.0.2.5
☒ Tunnel settings are same as Primary VPN Gateway
Encryption AES 128
DH Group 2
PFS Deactivated
Hash SHA 256
IKE SA Lifetime(min) 1440
IPsec SA Lifetime(min) 480
DPD Timeout Timer(sec) 20

Site Subnets

Subnet	Description	Advertise
Ex: 10.0.2.0/24	(optional)	<input checked="" type="checkbox"/>

Subscription VeloCloud-ManagementPlane-Dev
Virtual Wan azure_automation_test_vwan
Resource Group azure_vwan_automation
Virtual Hub azure_test_vhub_west_us_2
Azure Regions West US 2

Advanced **Save Changes** **Close**

10 In the **Primary VPN Gateway** area, you can configure the following tunnel settings:

Field	Description
Encryption	Select either AES 128 or AES 256 as the AES algorithms key size to encrypt data. If you do not want to encrypt data, select NONE . The default value is AES 128.
DH Group	The Diffie-Hellman (DH) Group algorithm to be used when exchanging a pre-shared key. The DH Group sets the strength of the algorithm in bits. The supported DH Group is 2.
PFS	Select the Perfect Forward Secrecy (PFS) level for additional security. The supported PFS levels are 2, 5, 14, 15, and 16. The default value is Deactivated.
Hash	The authentication algorithm for the VPN header. Select one of the following supported Secure Hash Algorithm (SHA) function from the list: <ul style="list-style-type: none"> ■ SHA 1 ■ SHA 256 The default value is SHA 256.
IKE SA Lifetime(min)	Time when Internet Key Exchange (IKE) rekeying is initiated for Edges. The minimum IKE lifetime is 10 minutes and maximum IKE lifetime is 1440 minutes. The default value is 1440 minutes.
IPsec SA Lifetime(min)	Time when Internet Security Protocol (IPsec) rekeying is initiated for Edges. The minimum IPsec lifetime is 3 minutes and maximum IPsec lifetime is 480 minutes. The default value is 480 minutes.
DPD Timeout Timer(sec)	The maximum time that the device should wait to receive a response to the DPD message before considering the peer to be dead. The default value is 20 seconds. You can deactivate DPD by configuring the DPD timeout timer to 0 second.

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 SD-WAN Edge act as an Initiator and Azure act as a Responder during an IPsec tunnel setup.

11 Click **Save Changes**.**What to do next**

- [Enable Cloud VPN at the Profile Level](#)
- 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 more information, see [Associate a Microsoft Azure Non SD-WAN Destination to a SD-WAN Edge and Add Tunnels](#).

For information about Azure Virtual WAN Edge Automation, see [Configure SD-WAN Orchestrator for Azure Virtual WAN IPsec Automation from SD-WAN Edge](#).

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

After configuring a Non SD-WAN Destination via Edge in SD-WAN Orchestrator, you have to associate the Non SD-WAN Destination to the desired Profile in order to establish the tunnels between SD-WAN 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.

Procedure

- 1 From the SD-WAN Orchestrator, go to **Configure > Profiles**.
The **Configuration Profiles** page appears.
- 2 Select a profile you want to configure Cloud VPN and click the icon under the **Device** column.
The **Device Settings** page for the selected profile appears.
- 3 Go to **Cloud VPN** area and activate Cloud VPN by turning the toggle button to **On**.
- 4 To establish a VPN connection directly from a SD-WAN Edge to a Non SD-WAN Destination (VPN gateway of Cloud provider such as Azure, AWS), select the **Enable** check box under **Branch to Non SD-WAN Destinations via Edge**.
- 5 From the list of configured Services, select a Non SD-WAN Destination to establish VPN connection. Click the + (plus) 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 more information about configuring a Non SD-WAN Destination Network Service through Edge, see [Configure a Non SD-WAN Destinations via Edge](#).

- 6 To deactivate a particular service, uncheck the respective **Enable Service** check box.
- 7 Click **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.

Configure Cloud VPN and Tunnel Parameters at the Edge Level

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 the Branch to Non SD-WAN Destination via Edge settings inherited from a Profile and configure Tunnel parameters (WAN link selection and Per tunnel credentials).

- 1 From the SD-WAN Orchestrator, go to **Configure > Edges**.

- 2 Select an Edge you want to override Non SD-WAN Destination settings and click the icon under the **Device** column. The Device Setting page for the selected Edge appears.
- 3 Go to the **Branch to Non SD-WAN Destination via Edge** area and select the **Enable Edge Override** checkbox.

Cloud VPN On

Branch to Non SD-WAN Destination via Gateway

Enable: ✖

Branch to Hubs

Enable: ✖

Branch to Branch VPN

Enable: ☒

Isolate Profile: ✖

Use Cloud Gateways ☒ Use Hubs for VPN ✖

Dynamic Branch To Branch VPN ✖

To All Edges ☒

To Edges Within Profile ✖

Branch to Non SD-WAN Destination via Edge ☒ **Enable Edge Override**

Enable ☒

Service				Link			
Action	Name	Automation for all public WAN Links	Enable Service	Enable tunnel	Destination Primary Public IP	Destination Secondary Public IP	Action
+ -	Azure_New ▼ 🔒	N/A	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/> 104.208.31.2... <input checked="" type="checkbox"/> 104.208.26.99	52.185.65.128	52.185.67.187	Edit Add Del Edit Add Del

- 4 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.

- 5 Under **Action**, click **Add** to add tunnels. The **Add Tunnel** pop-up window appears.

Add Tunnel ✖

Public Wan Link

Local Identification Type

Local Identification

PSK

Destination Primary Public IP

Destination Secondary Public IP

104.208.31.249 ▼

FQDN/Hostname ▼

a.com

••••••••
👁

52.185.65.128

52.185.67.187

Save Changes

Cancel

- 6 Enter the following details for configuring a tunnel to the Non SD-WAN Destination and click **Save Changes**.

Field	Description
Public WAN Link	
Local Identification Type	<p>Select any one of the Local authentication type from the drop-down menu:</p> <ul style="list-style-type: none"> ■ 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 IP address used to communicate with the local gateway.
Local Identification	<p>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 address.</p>
PSK	<p>Enter the Pre-Shared Key (PSK), which is the security key for authentication across the tunnel in the textbox.</p>
Destination Primary Public IP	<p>Enter the Public IP address of the destination Primary VPN Gateway.</p>
Destination Secondary Public IP	<p>Enter the Public IP address of the destination Secondary VPN Gateway.</p>

- 7 Click **Save Changes**.

Azure Virtual WAN IPsec Tunnel Automation

VMware SD-WAN Orchestrator supports integration and automation of Azure Virtual WAN from VMware SD-WAN Gateway and VMware SD-WAN Edge to enable Branch-to-Azure VPN Connectivity.

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.

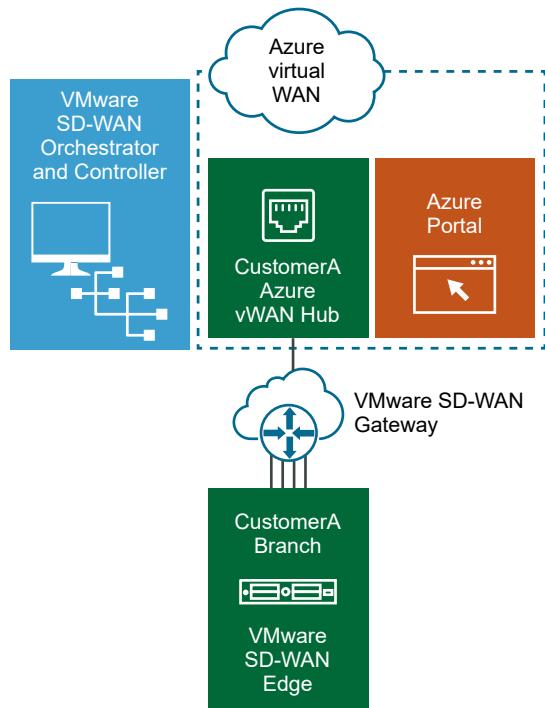
To establish branch-to-Azure VPN connectivity, SD-WAN Orchestrator supports Azure Virtual WAN and VMware SD-WAN integration and automation by leveraging the Azure backbone. Currently, the following Azure deployment options are supported from the VMware SD-WAN perspective:

- IPsec from SD-WAN Gateway to Azure virtual WAN hub with automation.

- Direct IPsec from SD-WAN Edge to Azure virtual WAN hub with automation.

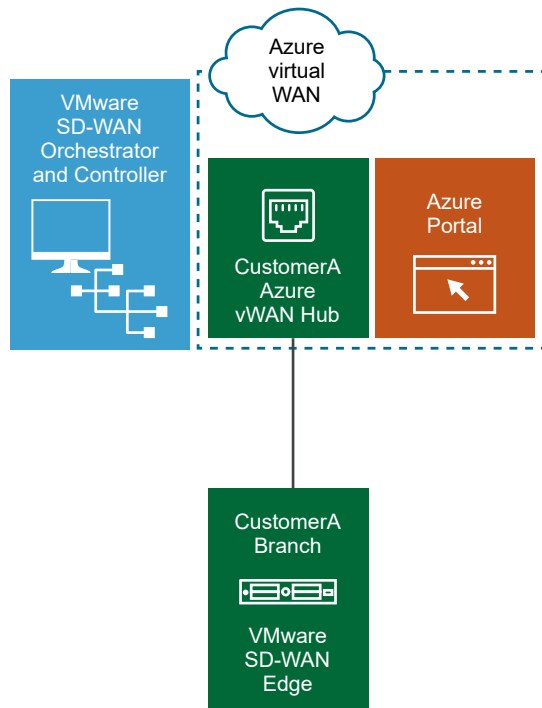
Azure Virtual WAN SD-WAN Gateway automation

The following diagram illustrates the IPsec tunnel from SD-WAN Gateway to Azure virtual WAN hub.



Azure Virtual WAN SD-WAN Edge automation

The following diagram illustrates the IPsec tunnel directly from SD-WAN Edge to Azure virtual WAN hub.



The following topics provide instructions for configuring the SD-WAN Orchestrator and Azure to enable branch-to-Azure VPN connectivity through the SD-WAN Gateway and SD-WAN Edge:

- [Prerequisite Azure Configuration](#)
- [Configure Azure Virtual WAN for Branch-to-Azure VPN Connectivity](#)
- [Configure SD-WAN Orchestrator for Azure Virtual WAN IPsec Automation from SD-WAN Gateway](#)
- [Configure SD-WAN Orchestrator for Azure Virtual WAN IPsec Automation from SD-WAN Edge](#)

Prerequisite Azure Configuration

Enterprise network administrators must complete the following prerequisite configuration tasks at the Azure portal to ensure that the SD-WAN Orchestrator application can function as the Service Principal (identity for the application) for the purposes of Azure Virtual WAN and SD-WAN Gateway integration.

- [Register SD-WAN Orchestrator Application](#)
- [Assign the SD-WAN Orchestrator Application to Contributor Role](#)
- [Register a Resource Provider](#)
- [Create a Client Secret](#)

Register SD-WAN Orchestrator Application

Describes how to register a new application in Azure Active Directory (AD).

To register a new application in Azure AD:

Prerequisites

- Ensure you have an Azure subscription. If not, create a [free account](#).

Procedure

- 1 Log in to your [Microsoft Azure](#) account.

The **Microsoft Azure** home screen appears.

- 2 Click **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.

Register an application

* Name

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

vcd 

Supported account types

Who can use this application or access this API?

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

[Help me choose...](#)

Redirect URI (optional)

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

Web  e.g. <https://myapp.com/auth>

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

Register

- 4 In the **Name** field, enter the name for your SD-WAN Orchestrator application.
- 5 Select a supported account type, which determines who can use the application.
- 6 Click **Register**.

Results

Your SD-WAN 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 SD-WAN Orchestrator configuration for Cloud Subscription.

What to do next

- [Assign the SD-WAN Orchestrator Application to Contributor Role](#)
- [Create a Client Secret](#)

Assign the SD-WAN 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.

To assign a Contributor role at the subscription scope:

Prerequisites

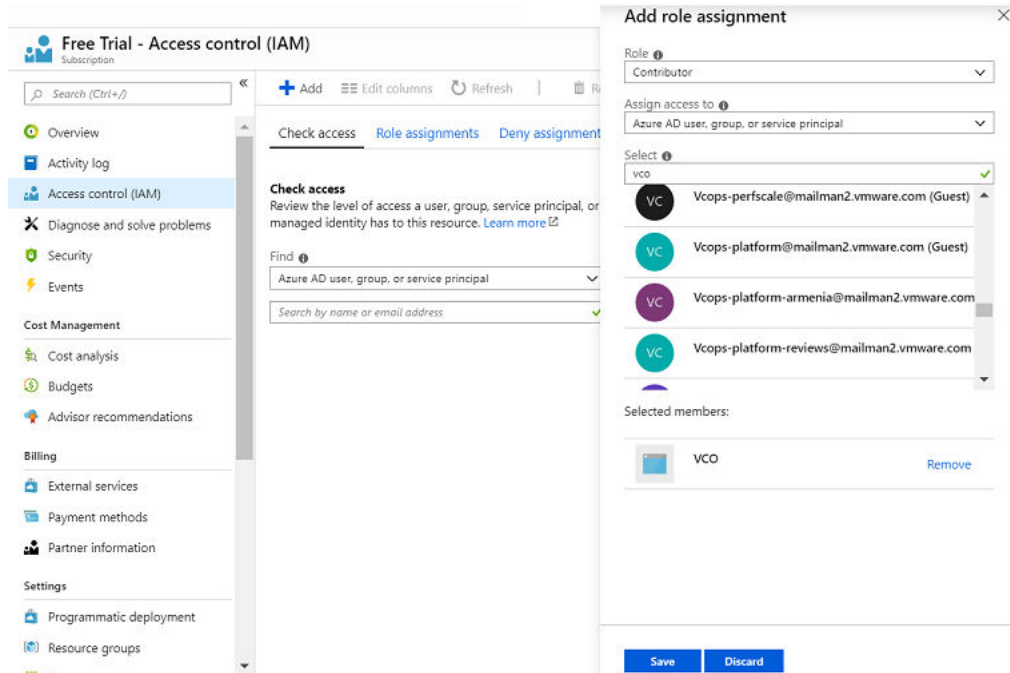
- Ensure you have an Azure subscription. If not, create a [free account](#).

Procedure

- 1 Click **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 Click **Access control (IAM)**.

4 Click **+Add > Add role assignment**.

The **Add role assignment** dialog box appears.



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**.

Results

The application is assigned to the Contributor role and it appears in the list of users assigned to a role for that scope.

What to do next

- [Create a Client Secret](#)
- [Configure Azure Virtual WAN for Branch-to-Azure VPN Connectivity](#)

Register a Resource Provider

To download Virtual WAN Virtual Private Network (VPN) configurations, the SD-WAN Orchestrator requires a Blob Storage Account that acts as an intermediary data store from where the configurations can be downloaded. The SD-WAN 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.

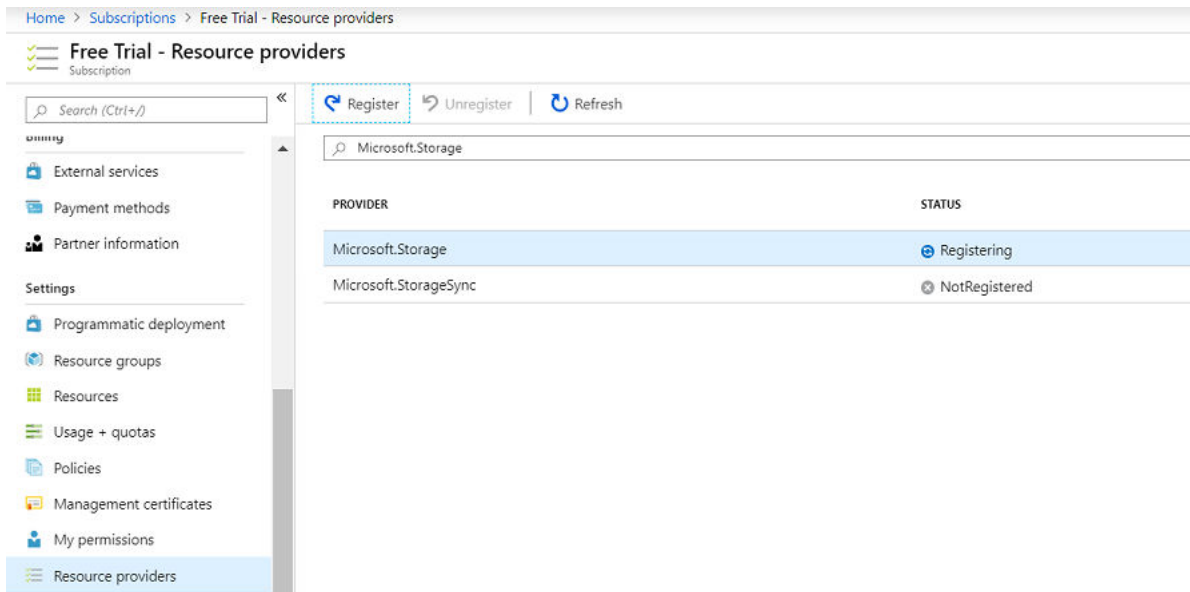
To register a resource provider for your subscription:

Prerequisites

- Ensure you have an Azure subscription. If not, create a [free account](#).
- Ensure you have the Contributor or Owner roles permission.

Procedure

- 1 Log in to your [Microsoft Azure](#) account.
- 2 Click **All Services** and search for **Subscriptions**.
- 3 From the list of subscriptions, select your subscription.
- 4 Under the **Settings** tab, select **Resource providers**.



- 5 From the list of available resource providers, select **Microsoft.Storage**. and click **Register**.

Results

The resource provider is registered and configures your subscription to work with the resource provider.

What to do next

You can create the resources in Azure, for steps, see [Configure Azure Virtual WAN for Branch-to-Azure VPN Connectivity](#).

Create a Client Secret

Describes how to create a new client secret in Azure AD for the purpose of authentication.

To create a new client secret in Azure AD:

Prerequisites

- Ensure you have an Azure subscription. If not, create a [free account](#).

Procedure

- 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, click on your registered SD-WAN Orchestrator application.
- 4 Go to **Certificates & secrets** > **New client secret**.

The **Add a client secret** screen appears.

- 5 Provide details such as description and expiry value for the secret and click **Add**.

Results

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 SD-WAN Orchestrator.

What to do next

- [Configure Azure Virtual WAN for Branch-to-Azure VPN Connectivity](#)

- [Configure SD-WAN Orchestrator for Azure Virtual WAN IPsec Automation from SD-WAN Gateway](#)

Configure Azure Virtual WAN for Branch-to-Azure VPN Connectivity

This section describes the procedures to configure Azure for integrating Azure Virtual WAN and SD-WAN 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 SD-WAN 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](#)

Create a Resource Group

Describes how to create a resource group in Azure.

To create a resource group in Azure:

Prerequisites

- Ensure you have an Azure subscription. If not, create a [free account](#).

Procedure

- 1 Log in to your [Microsoft Azure](#) account.
The **Microsoft Azure** home screen appears.
- 2 Click **All Services** and search for **Resource groups**.

3 Select **Resource groups** and click **+Add**.

The **Create a resource group** screen appears.

[Home](#) > [Resource groups](#) > Create a resource group

Create a resource group

[Basics](#)
[Tags](#)
[Review + create](#)

Resource group - A container that holds related resources for an Azure solution. The resource group can include all the resources for the solution, or only those resources that you want to manage as a group. You decide how you want to allocate resources to resource groups based on what makes the most sense for your organization. [Learn more](#)

Project details

* Subscription ⓘ

Free Trial ▼

* Resource group ⓘ

Sasi_RG1 ✓

Resource details

* Region ⓘ

(US) Central US ▼

Review + create

< Previous

Next : Tags >

4 From the **Subscription** drop-down menu, select your Microsoft Azure subscription.

5 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.

6 From the **Region** drop-down menu, select the location for your resource group, where the majority of your resources will reside.

7 Click **Review+create** and then click **Create**.

Results

A resource group is created and appears on the Azure portal dashboard.

What to do next

Create an Azure Virtual WAN. For steps, see [Create a Virtual WAN](#).

Create a Virtual WAN

Describes how to create a Virtual WAN in Azure.

To create a Virtual WAN in Azure:

Prerequisites

- Ensure you have an Azure subscription. If not, create a [free account](#).
- Ensure you have a resource group created to add the Virtual WAN.

Procedure

- 1 Log in to your [Microsoft Azure](#) account.
The **Microsoft Azure** home screen appears.
- 2 Click **All Services** and search for **Virtual WANs**.
- 3 Select **Virtual WANs** and click **+Add**.
The **Create WAN** screen appears.

Create 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 metadata 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 Click **Create**.

Results

A Virtual WAN is created and appears on the Azure portal dashboard.

What to do next

Create Virtual Hubs. For steps, see [Create a Virtual Hub](#).

Create a Virtual Hub

Describes how to create a Virtual Hub in Azure.

To create a Virtual Hub in Azure:

Prerequisites

- 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.

Procedure

- 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, click **Hubs**.

4 Click **+New Hub**.

The **Create virtual hub** screen appears.

Create virtual hub

[Basics](#) [Site to site](#) [Point to site](#) [ExpressRoute](#) [Routing](#) [Tags](#) [Review + create](#)

A virtual hub is a Microsoft-managed virtual network. The hub contains various service endpoints to enable connectivity from your on-premises network (vpnsite). The hub is the core of your network in a region. There can only be one hub per Azure region. When you create a hub using Azure portal, it creates a virtual hub VNet and a virtual hub vpngateway. [Learn more](#)

Project details

The hub will be created under the same subscription and resource group as the vWAN.

* Subscription

* Resource group

Virtual Hub Details

* Region

* Name

* Hub private address space

Creating a hub with a gateway will take 30 minutes.

[Review + create](#) [Previous](#) [Next : Site to site >](#)

5 In the **Basics** tab, enter the following Virtual Hub details.

- From the **Region** drop-down menu, select the location where the Virtual Hub resides.
- In the **Name** text box, enter the unique name for your Hub.
- In the **Hub private address space** text box, enter the address range for the Hub in Classless inter-domain routing (CIDR) notation.

6 Click **Next: Site to site >** and enable Site to site (VPN gateway) before connecting to VPN sites by selecting **Yes**.


Note A VPN Gateway is required for tunnel automation to work, otherwise it is not possible to create VPN connections.

Create virtual hub


[Basics](#)
[Site to site](#)
[Point to site](#)
[ExpressRoute](#)
[Routing](#)
[Tags](#)
[Review + create](#)

You will need to enable Site to site (VPN gateway) before connecting to VPN sites. You can do this after hub creation, but doing it now will save time and reduce the risk of service interruptions later. [Learn more](#)

Do you want to create a Site to site (VPN gateway)? ☒ Yes ☐ No

AS Number 

* Gateway scale units

 Creating a hub with a gateway will take 30 minutes.

[Review + create](#)
[Previous](#)
[Next : Point to site >](#)

a From the **Gateway scale units** drop-down menu, select a scaling value.

7 Click **Review + Create**.

Results

A Virtual Hub is created and appears on the Azure portal dashboard.

What to do next

- 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](#).

Create a Virtual Network

Describes how to create a Virtual Network in Azure.

To create a Virtual Network in Azure:

Prerequisites

- Ensure you have an Azure subscription. If not, create a [free account](#).

Procedure

- 1 Log in to your [Microsoft Azure](#) account.
The **Microsoft Azure** home screen appears.
- 2 Click **All Services** and search for **Virtual networks**.

- 3 Select **Virtual networks** and click **+Add**.

The **Create virtual network** screen appears.

Create virtual network ☐

* **Name**
 ✓

* **Address space** ⓘ
 ✓
 10.0.0.0 - 10.0.0.255 (256 addresses)

* **Subscription**
 ▼

* **Resource group**
 ▼
[Create new](#)

* **Location**
 ▼

Subnet

* **Name**
 ✓

* **Address range** ⓘ
 ✓
 10.0.0.0 - 10.0.0.255 (256 addresses)

DDoS protection ⓘ
☒ Basic ☐ Standard

Service endpoints ⓘ

[Automation options](#)

- 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 Click **Create**.

Results

A Virtual network is created and appears on the Azure portal dashboard.

What to do next

Create Virtual Connection between Hubs and Virtual Networks (VNETs). For steps, see [Create a Virtual Connection between VNet and Hub](#).

Create a Virtual Connection between VNet and Hub

Describes how to create a virtual connection between Virtual Networks (VNETs) and the Virtual Hub in a particular Azure region.

To create a virtual network connection between a VNet and a Virtual Hub in a particular Azure region:

Prerequisites

- Ensure you have an Azure subscription. If not, create a [free account](#).
- Ensure you have Virtual Hubs and Virtual Networks created.

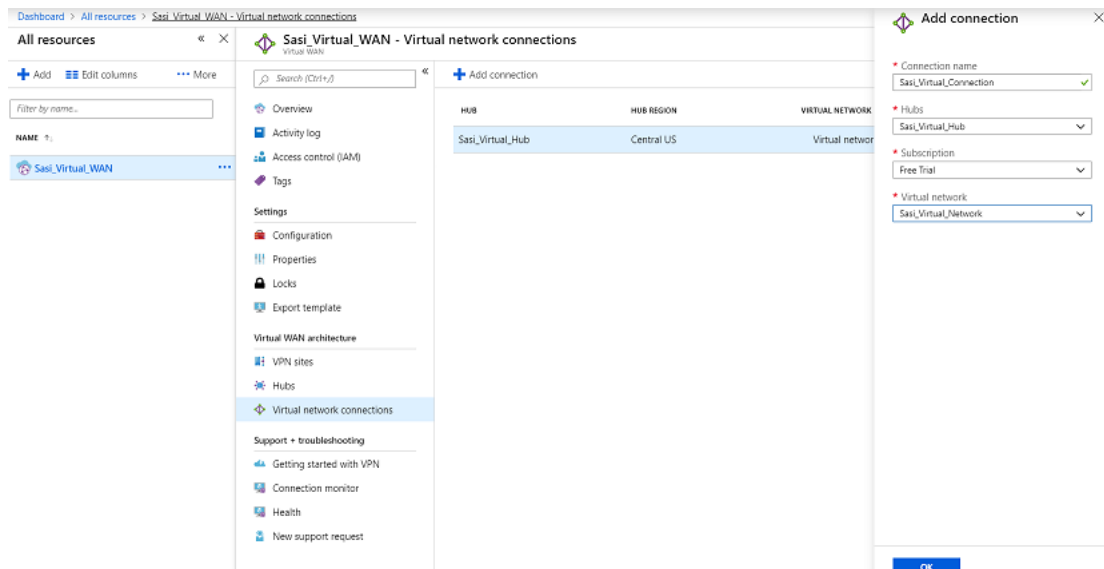
Procedure

- 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, click **Virtual network connections**.
- 4 Click **+Add connection**.

The **Add connection** screen appears.



- 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 Click **OK**.

Results

A peering connection is established between the selected VNet and the Hub.

What to do next

- [Configure SD-WAN Orchestrator for Azure Virtual WAN IPsec Automation from SD-WAN Gateway](#)

Configure SD-WAN Orchestrator for Azure Virtual WAN IPsec Automation from SD-WAN Gateway

You can configure SD-WAN Orchestrator for integrating Azure Virtual WAN and SD-WAN 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 SD-WAN 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 SD-WAN Orchestrator configuration for Azure Virtual WAN - SD-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 SD-WAN Orchestrator side for integrating Azure Virtual WAN and SD-WAN Gateway, see:

- [Configure a Cloud Subscription Network Service](#)
- [Configure a Microsoft Azure Non SD-WAN Destination via Gateway](#)
- [Synchronize VPN Configuration](#)

To view the details of Non SD-WAN Destinations network services configured for an enterprise, see [Monitor Non SD-WAN Destinations](#).

Configure a Cloud Subscription Network Service

Describes how to configure a Cloud subscription in SD-WAN Orchestrator.

To configure a Cloud subscription in SD-WAN Orchestrator:

Prerequisites

Ensure you have registered the SD-WAN Orchestrator application and created Client secret in the Azure portal. For steps, see [Prerequisite Azure Configuration](#).

Procedure

- 1 From the navigation panel in the SD-WAN Orchestrator, go to **Configure > Network Services**.

The **Services** screen appears.

- 2 In the **Cloud Subscriptions** area, click the **New** button.

The **Configure Cloud Subscription** dialog box appears.

Configure Cloud Subscription

* Subscription Type: Microsoft Azure Subscription

* Active Directory Tenant ID: 22eb73a3-5c68-47b6-8098-08952150a401

* Client ID: 5188a0f1-8215-49d0-9085-ea3043a12721

* Client Secret:

* Subscription: Pay-As-You-Go(Converted to EA)

Save Changes Cancel

- 3 From the **Subscription Type** drop-down-menu, select **Microsoft Azure Subscription**.
- 4 Enter the Active Directory Tenant ID, Client ID, and Client Secret corresponding to your SD-WAN Orchestrator Application Registration.
- 5 Click the **Get Subscriptions** button to retrieve the list of Azure Subscriptions for which the App Registration has been allocated an IAM role.
- 6 Click **Save Changes**.

What to do next

Configure a Non SD-WAN Destination of type Microsoft Azure Virtual Hub.

- To configure a Microsoft Azure Non SD-WAN Destination from SD-WAN Gateway, see [Configure a Microsoft Azure Non SD-WAN Destination via Gateway](#).
- To configure a Microsoft Azure Non SD-WAN Destination from SD-WAN Edge, see [Configure a Microsoft Azure Non SD-WAN Destination via Edge](#).

Configure a Microsoft Azure Non SD-WAN Destination via Gateway

Describes how to configure a Non SD-WAN Destination of type **Microsoft Azure Virtual Hub** from SD-WAN Gateway.

To configure a Non SD-WAN Destination of type **Microsoft Azure Virtual Hub** from SD-WAN Gateway:

Prerequisites

- Ensure you have configured a Cloud subscription. For steps, see [Configure a Cloud Subscription Network Service](#).
- Ensure you have created Virtual WAN and Hubs in Azure. For steps, see [Configure Azure Virtual WAN for Branch-to-Azure VPN Connectivity](#).

Procedure

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

New Non SD-WAN Destination via Gateway...

* Name:

* Type:

Virtual Hub Configuration

Subscription:

Virtual WAN:

Resource Group:

Virtual Hub:

Azure Region:

Enable Tunnel(s): ☒

Next

- 3 In the **Name** text box, enter the name for the Non SD-WAN Destination.
- 4 From the **Type** drop-down menu, select **Microsoft Azure Virtual Hub**.
- 5 From the **Subscription** drop-down menu, select a subscription.
The application fetches all the available Virtual WANs dynamically from Azure.
- 6 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.
- 7 From the **Virtual Hub** drop-down menu, select a Virtual Hub.
The application auto-populates the Azure region corresponding to the Hub


- 8 Select the **Enable Tunnel(s)** checkbox to enable VMware VPN Gateways initiate VPN connections to the target Virtual Hub as soon as the site is successfully provisioned.

Note VMware VPN Gateways will not initiate IKE negotiation until this Non SD-WAN Destination is configured on at least one profile.

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

- 9 Click **Next**.

The SD-WAN Orchestrator automatically initiates deployment, provisions Azure VPN Sites, and downloads the VPN Site Configuration for the newly configured sites and stores the configuration in the SD-WAN Orchestrator's Non SD-WAN Destination configuration database.



Results

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

What to do next

- Associate the Microsoft Azure Non SD-WAN Destination to a Profile to establish a tunnel between a branch and Azure Virtual Hub. For more information, see [Associate a Microsoft Azure Non SD-WAN Destination to a Profile](#).
- You must add SD-WAN routes into Azure network manually. For more 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 configure the BGP settings for the Non SD-WAN Destination. Scroll to the name of your Non SD-WAN Destination, and then click the **Edit** link in the **BGP** column. For more information, see [Configure BGP over IPsec from Gateways](#).
- In the **Non SD-WAN Destinations via Gateway** area, click the **Edit** link in the **BFD** column for a Non SD-WAN Destination, to configure the BFD settings. For more information, see [Configure BFD for Gateways](#).

For information about Azure Virtual WAN Gateway Automation, see [Configure SD-WAN Orchestrator for Azure Virtual WAN IPsec Automation from SD-WAN Gateway](#).

Associate a Microsoft Azure Non SD-WAN Destination to a Profile

After configuring a Non SD-WAN Destination of type **Microsoft Azure Virtual Hub** in SD-WAN Orchestrator, you must associate the Non SD-WAN Destination to the desired Profile to establish the tunnels between SD-WAN Gateways and Microsoft Azure Virtual Hub.

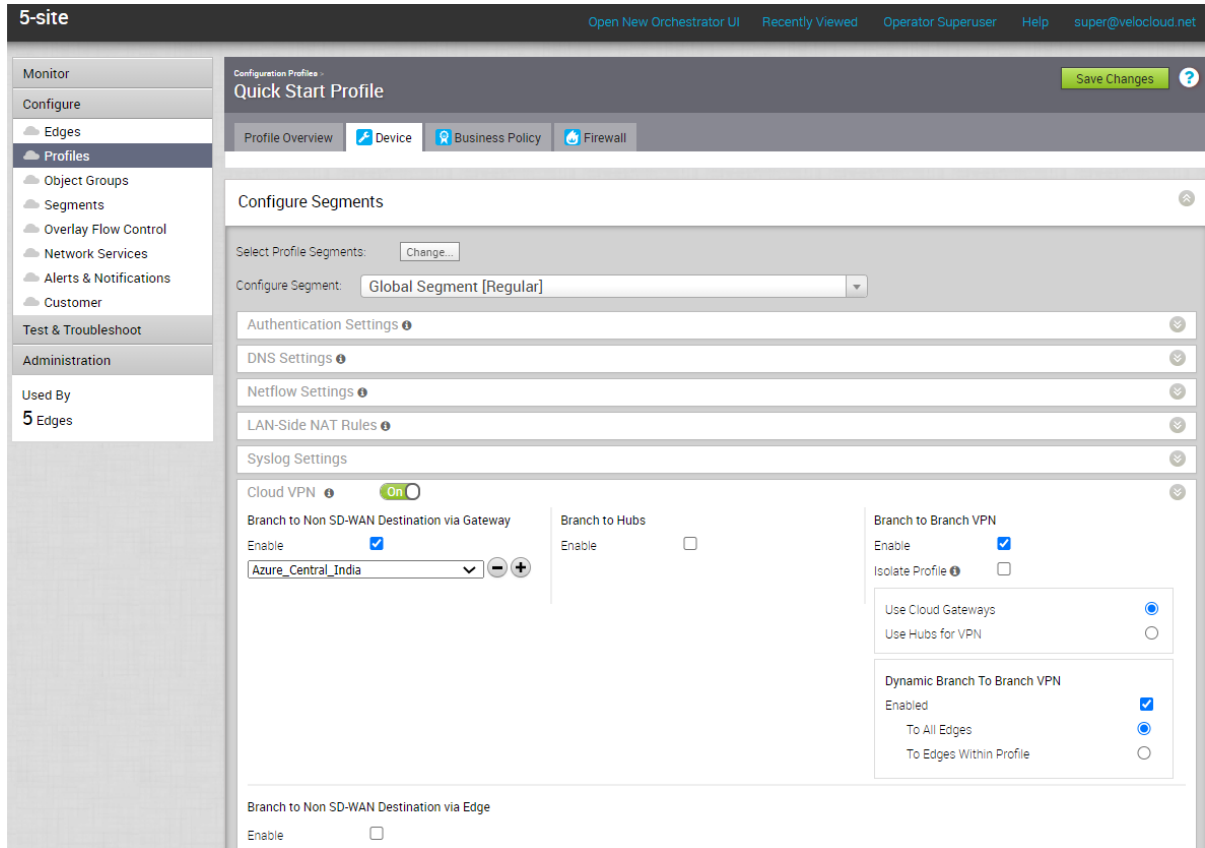
To associate a Non SD-WAN Destination to a Profile, perform the following steps:

Procedure

- 1 From the SD-WAN Orchestrator navigation panel, go to **Configure > Profiles**.

The **Configuration Profiles** page appears.

- 2 Select a profile you want to associate your Non SD-WAN Destination of type **Microsoft Azure Virtual Hub** and click the icon under the **Device** column.



The **Device Settings** page for the selected profile appears.

- 3 Go to **Cloud VPN** area and enable Cloud VPN by turning the toggle button to **On**.
- 4 Under **Branch to Non SD-WAN Destinations via Gateway**, select the **Enable** checkbox.
- 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 Click **Save Changes**.

Results

A tunnel is established between the branch and the Microsoft Azure Non SD-WAN Destination.

Edit a VPN Site

Describes how to add SD-WAN routes into the Azure network manually.

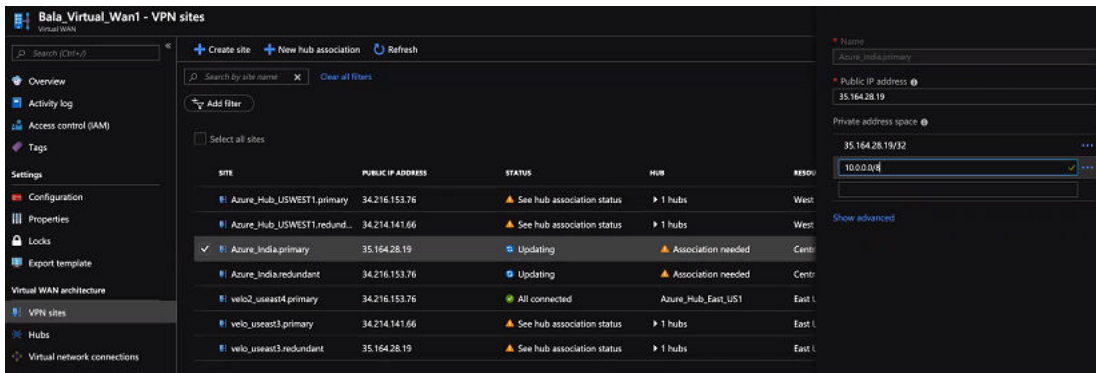
To add SD-WAN routes manually into the Azure network:

Prerequisites

Ensure you have completed provisioning the Azure VPN sites at the SD-WAN Orchestrator side.

Procedure

- 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, click **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 SD-WAN Orchestrator.
- 5 Click on the name of the selected VPN site and from the top of the next screen, select **Edit site**.



- 6 In the **Private address space** text box, enter the address range for the SD-WAN routes.
- 7 Click **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 doesn't 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 SD-WAN 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 be handled with caution.

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. Clicking the **Resync configuration** button in the **Non-VeloCloud Sites** area will automatically fetch the VPN configuration details from the Azure portal and will update the SD-WAN Orchestrator local configuration.

Delete a Non SD-WAN Destination

Describes the steps to delete Non SD-WAN Destination corresponding to the Azure's Virtual Hub and thereby ensure Virtual WAN deployment state is consistent between the SD-WAN Orchestrator and Azure following the deletion.

Procedure

- 1 Delete the Azure VPN Connections associated to the VPN Sites targeted for deletion.
- 2 Delete the Azure VPN Sites provisioned on behalf of the Non SD-WAN Destination SD-WAN Gateways selected for that Virtual Hub by using an Azure API.

Note Deletion of the Azure VPN Sites will fail if the VPN connections associated to the VPN Sites (targeted for deletion) are not removed.

Configure SD-WAN Orchestrator for Azure Virtual WAN IPsec Automation from SD-WAN Edge

You can configure SD-WAN Orchestrator for integrating Azure Virtual WAN and SD-WAN Edge to enable the branch-to-Azure VPN connectivity directly from SD-WAN Edge.

Note When using the Azure Virtual WAN Automation from SD-WAN 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 SD-WAN Orchestrator configuration for Azure Virtual WAN - SD-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 SD-WAN Orchestrator side for integrating Azure Virtual WAN and SD-WAN Edge, see:

- [Configure a Cloud Subscription Network Service](#)
- [Configure a Microsoft Azure Non SD-WAN Destination via Edge](#)
- [Enable Cloud VPN at the Profile Level](#)
- [Associate a Microsoft Azure Non SD-WAN Destination to a SD-WAN Edge and Add Tunnels](#)

Configure a Microsoft Azure Non SD-WAN Destination via Edge

Describes how to configure a Non SD-WAN Destination of type **Microsoft Azure Virtual Hub** via Edge in SD-WAN Orchestrator.

To configure a Non SD-WAN Destination of type **Microsoft Azure Virtual Hub** via Edge in SD-WAN Orchestrator:

Prerequisites

- Ensure you have configured a Cloud subscription. For steps, see [Configure a Cloud Subscription Network Service](#).

- Ensure you have created Virtual WAN and Hubs in Azure. For steps, see [Configure Azure Virtual WAN for Branch-to-Azure VPN Connectivity](#).

Procedure

- 1 From the navigation panel in the SD-WAN Orchestrator, go to **Configure > Network Services**.

The **Services** screen appears.

- 2 In the **Non SD-WAN Destinations via Edge** area, click the **New** button.

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

Non SD-WAN Destinations via Edge

* Service Name: Azure auto

* Service Type: Microsoft Azure Virtual Wan

Virtual Hub Configuration

* Subscription: VeloCloud-ManagementPlane- [v]

* Virtual Wan: azure_automation_test_vwan [v]

Resource Group: azure_vwan_automation

* Virtual Hub: azure_test_vhub_west_us_2 [v]

Azure Regions: West US 2

Enable Tunnels: ☐

Next

- 3 In the **Service Name** text box, enter the name for the Non SD-WAN Destination.
- 4 From the **Service Type** drop-down menu, select **Microsoft Azure Virtual Hub**.
- 5 From the **Subscription** drop-down menu, select a cloud subscription.
The application fetches all the available Virtual WANs dynamically from Azure.
- 6 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.
- 7 From the **Virtual Hub** drop-down menu, select a Virtual Hub.
The application auto-populates the Azure region corresponding to the Hub
- 8 Click **Next**.
The Microsoft Azure Non SD-WAN Destination is created and a dialog box for your Non SD-WAN Destination appears.

Non SD-WAN Destinations via Edge

Name Azure auto
Type Microsoft Azure Virtual Wan

Primary VPN Gateway
Public IP Ex: 10.0.2.5

Secondary VPN Gateway ☒
Keep Tunnel Active ☒
Public IP Ex: 10.0.2.5
☒ Tunnel settings are same as Primary VPN Gateway

Site Subnets

Subnet	Description	Advertise
Ex: 10.0.2.0/24	(optional)	<input checked="" type="checkbox"/>

Subscription VeloCloud-ManagementPlane-Dev
Virtual Wan azure_automation_test_vwan
Resource Group azure_vwan_automation
Virtual Hub azure_test_vhub_west_us_2
Azure Regions West US 2

Advanced **Save Changes** **Close**

- 9 To configure tunnel settings for the Non SD-WAN Destination's Primary VPN Gateway, click the **Advanced** button.

Non SD-WAN Destinations via Edge

Name Azure auto
Type Microsoft Azure Virtual Wan

Primary VPN Gateway
Public IP Ex: 10.0.2.5
Encryption AES 128
DH Group 2
PFS Deactivated
Hash SHA 256
IKE SA Lifetime(min) 1440
IPsec SA Lifetime(min) 480
DPD Timeout Timer(sec) 20

Secondary VPN Gateway ☒
Keep Tunnel Active ☒
Public IP Ex: 10.0.2.5
☒ Tunnel settings are same as Primary VPN Gateway
Encryption AES 128
DH Group 2
PFS Deactivated
Hash SHA 256
IKE SA Lifetime(min) 1440
IPsec SA Lifetime(min) 480
DPD Timeout Timer(sec) 20

Site Subnets

Subnet	Description	Advertise
Ex: 10.0.2.0/24	(optional)	<input checked="" type="checkbox"/>

Subscription VeloCloud-ManagementPlane-Dev
Virtual Wan azure_automation_test_vwan
Resource Group azure_vwan_automation
Virtual Hub azure_test_vhub_west_us_2
Azure Regions West US 2

Advanced **Save Changes** **Close**

10 In the **Primary VPN Gateway** area, you can configure the following tunnel settings:

Field	Description
Encryption	Select either AES 128 or AES 256 as the AES algorithms key size to encrypt data. If you do not want to encrypt data, select NONE . The default value is AES 128.
DH Group	The Diffie-Hellman (DH) Group algorithm to be used when exchanging a pre-shared key. The DH Group sets the strength of the algorithm in bits. The supported DH Group is 2.
PFS	Select the Perfect Forward Secrecy (PFS) level for additional security. The supported PFS levels are 2, 5, 14, 15, and 16. The default value is Deactivated.
Hash	The authentication algorithm for the VPN header. Select one of the following supported Secure Hash Algorithm (SHA) function from the list: <ul style="list-style-type: none"> ■ SHA 1 ■ SHA 256 The default value is SHA 256.
IKE SA Lifetime(min)	Time when Internet Key Exchange (IKE) rekeying is initiated for Edges. The minimum IKE lifetime is 10 minutes and maximum IKE lifetime is 1440 minutes. The default value is 1440 minutes.
IPsec SA Lifetime(min)	Time when Internet Security Protocol (IPsec) rekeying is initiated for Edges. The minimum IPsec lifetime is 3 minutes and maximum IPsec lifetime is 480 minutes. The default value is 480 minutes.
DPD Timeout Timer(sec)	The maximum time that the device should wait to receive a response to the DPD message before considering the peer to be dead. The default value is 20 seconds. You can deactivate DPD by configuring the DPD timeout timer to 0 second.

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 SD-WAN Edge act as an Initiator and Azure act as a Responder during an IPsec tunnel setup.

11 Click **Save Changes**.**What to do next**

- [Enable Cloud VPN at the Profile Level](#)
- 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 more information, see [Associate a Microsoft Azure Non SD-WAN Destination to a SD-WAN Edge and Add Tunnels](#).

For information about Azure Virtual WAN Edge Automation, see [Configure SD-WAN Orchestrator for Azure Virtual WAN IPsec Automation from SD-WAN Edge](#).

Enable Cloud VPN at the Profile Level

To enable Cloud VPN at the Profile level, perform the following steps:

Procedure

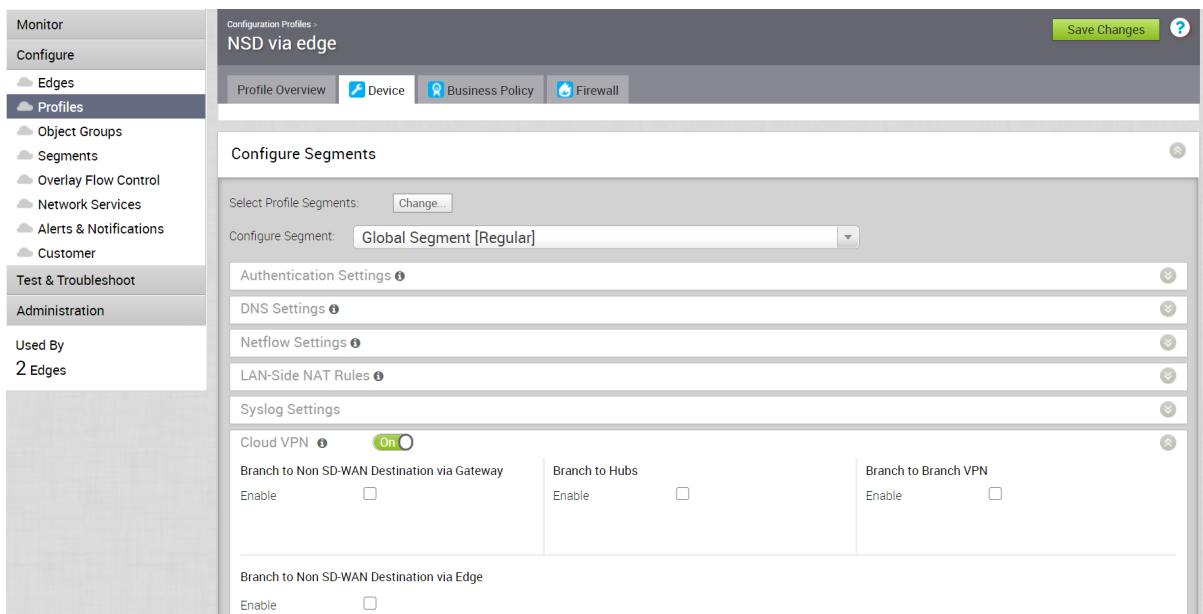
- 1 From the SD-WAN Orchestrator navigation panel, go to **Configure > Profiles**.

The **Configuration Profiles** page appears.

- 2 Select a profile you want to associate your Non SD-WAN Destination of type **Microsoft Azure Virtual Hub** and click the icon under the **Device** column.

The **Device Settings** page for the selected profile appears.

- 3 Go to **Cloud VPN** area and enable Cloud VPN by turning the toggle button to **On**.



Note Automation of all Public WAN links through SD-WAN Edge is not supported at the Profile level.

- 4 Click **Save Changes**.

What to do next

[Associate a Microsoft Azure Non SD-WAN Destination to a SD-WAN Edge and Add Tunnels](#)

Associate a Microsoft Azure Non SD-WAN Destination to a SD-WAN Edge and Add Tunnels

After configuring a Non SD-WAN Destination of type **Microsoft Azure Virtual Hub** from SD-WAN 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 a SD-WAN Edge, perform the following steps:

Procedure

1 Go to **Configure > Edges**

The **Edges** page appears.

2 Select an Edge you want to associate your Microsoft Azure Non SD-WAN Destination and click the icon under the **Device** column.

3 In the **Device Settings** page, under **Branch to Non SD-WAN Destinations via Edge**, select the **Enable Edge Override** checkbox.

Cloud VPN ⓘ On

Branch to Non SD-WAN Destination via Gateway

Enable: ✕

Branch to Hubs

Enable: ✕

Branch to Branch VPN

Enable: ✕

Branch to Non SD-WAN Destination via Edge Enable Edge Override ⓘ

Enable ☑

Service				Link			
Action	Name	Automation for all public WAN Links	Enable Service	Enable tunnel	Destination Primary Public IP	Destination Secondary Public IP	Action
+ -	Azure auto ▼	N/A	☑	No Sites Added			Add

4 Select the **Enable** checkbox.

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**, click the **Add** link. The **Add Tunnel** dialog box appears.

Add Tunnel

Public Wan Link ⓘ 34.213.104.253 ▾

Local Identification Type IP Address ▾

Local Identification ⓘ

PSK ⓘ

Destination Primary Public IP

Destination Secondary Public IP


Save Changes Cancel

- a From the **Public WAN Link** drop-down menu, select a WAN link to establish IPsec tunnel and click **Save Changes**.

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 will be used as the Local Identification value of the tunnel. You will be able to 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 will send 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 will be able to view the details of configured tunnel and Public WAN link as shown in the following screenshot.

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 Enterprise portal. See [Monitor Non SD-WAN Destinations](#).



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**, click the **Edit** link to view the tunnel settings such as PSK.
- Delete a tunnel - Under **Action**, click the **Del** link to delete a specific tunnel.
- Deactivate a tunnel - Under **Enable tunnel**, unselect a tunnel to deactivate the specific tunnel.
- Delete a network service - Under **Action**, click the **-** icon to delete a specific network service.
- Deactivate a network service - Under **Enable Service**, unselect a network service checkbox to deactivate a specific network service.

7 Click **Save Changes**.

What to do next

Once the automation is complete and tunnel is created, you can monitor the tunnel status from the **Monitor > Edges** page.

Search...	Cols	Reset View	Refresh	CSV	Display 1 items				
Edge	Status	HA	Links	VM Status	VNF	Edge Tunnels	Gateways	Profile	
1 edge-1	●		↔ 2			↔ 1 ↔ 1	View	NSD via edge	

NVS Via Edge Up Tunnels
 Azure auto
 Microsoft Azure Virtual Wan
 GE2 Global Segment 51.143.61.143 ●Up

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 Enterprise portal.

In 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 Enterprise portal, click **Monitor > Network Services**.

The **Network Services** page appears.

The screenshot shows the NSD Dev interface. On the left is a navigation menu with sections: Monitor (Network Overview, Edges, Network Services, Routing, Alerts, Events, Reports), Configure, Test & Troubleshoot, and Administration. The main area is titled 'Network Services' and features a world map. Below the map are two sections: 'Non SD-WAN Destinations via Gateway' (showing 'No Items') and 'Non SD-WAN Destinations via Edge'. The latter contains a table with columns: Name, Public IP, Tunnel Status, Used By, Last Contact, and Deployment Status. One entry is visible: 'Azure auto' (Microsoft Azure Virtual Wan) with a tunnel status of 2 and a deployment status of 'View'.

- Under **Non SD-WAN Destinations via Edge**, click the link in the **Deployment Status** column to view the deployment status of the Non SD-WAN Destinations.

The screenshot shows the 'Non SD-WAN Destinations Deployment Status for Service: Azure auto' page. It includes a search bar, a status filter set to 'COMPLETE', and a refresh button. A summary row shows counts for various states: Enqueued (0), Pending (0), Notified (0), Completed (2), Errored (0), Timed Out (0), and Pending Delete (0). Below this is a table with columns: Edge, Link, Segment, Action, Status, and API Tracking Info. Two rows are shown, both with 'edge-1' as the edge, 'Global Segment' as the segment, and 'COMPLETE' as the status.

Edge	Link	Segment	Action	Status	API Tracking Info
edge-1	34.213.104.253	Global Segment	createNvsFromEdgeSite	COMPLETE	Details
edge-1	52.38.235.212	Global Segment	createNvsFromEdgeSite	COMPLETE	Details

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.

- 3 Under **API Tracking Info**, click **Details** to view the Event details.

VMware SD-WAN in Azure Virtual WAN Hub Deployment

About VMware SD-WAN in Azure Virtual WAN Hub Deployment

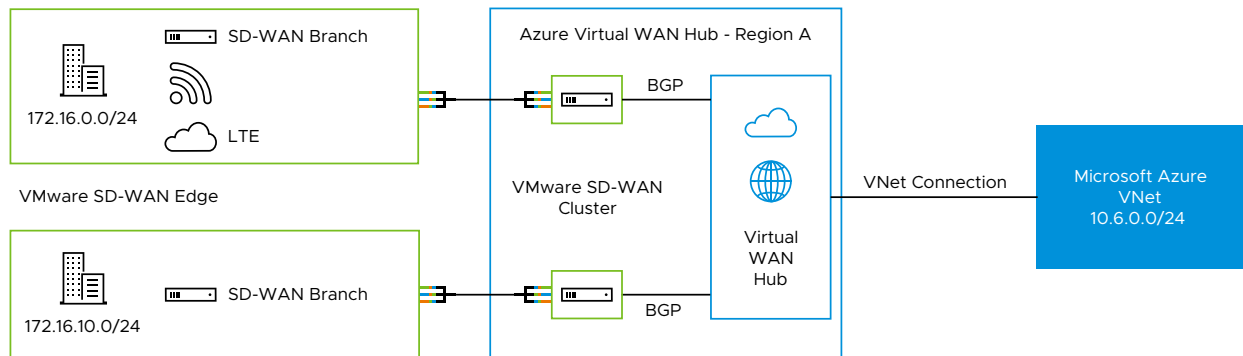
The VMware SD-WAN in Azure Virtual WAN (vWAN) Hub deployment describes the configurations that are required to manually deploy a Virtual SD-WAN Edge as a Network Virtual Appliance (NVA) in Azure vWAN Hub network.

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. VMware SD-WAN addresses these problems by leveraging Dynamic Multipath Optimization™ (DMPO) technologies and distributed cloud gateway coverage across the globe. VMware 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, VMware SD-WAN have been progressively adding more capabilities to the solution. With this new integration, customers can now deploy VMware SD-WAN Edges directly inside Azure Virtual WAN hubs manually, resulting in an offering that natively integrates Azure Virtual WAN's customizable routing intelligence with VMware SD-WAN's optimized last-mile connectivity.

The following diagram illustrates the VMware SD-WAN and Azure vWAN NVA Manual Deployment scenario.



Deploy VMware SD-WAN in Azure Virtual WAN Hub

To deploy VMware SD-WAN Edges in Virtual Hub manually, you must have already created Resource Group, virtual WAN (vWAN), and virtual Hub (vHUB) on the Azure side.

Configuration Steps:

Prerequisites

Once vWAN Hub is up and running and routing status is completed, you must ensure the following prerequisites are met before proceeding with the Manual deployment of Azure vWAN NVA via VMware SD-WAN Orchestrator:

- Obtain Enterprise account access to VMware SD-WAN Orchestrator.
- Obtain access to the Microsoft Azure portal with the appropriate IAM roles.
- Software image requirements for this deployment are as follows:
 - VMware SD-WAN Orchestrator: 4.2.1 and above.
 - VMware SD-WAN Gateway: 4.2.1 and above.
 - VMware SD-WAN Edges: 4.2.1 and above.


Procedure

- 1 In the Orchestrator, create virtual Edges 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 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.



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 more details regarding this IP configuration, see <https://docs.microsoft.com/en-us/azure/virtual-network/what-is-ip-address-168-63-129-16>

- 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 "VMware SD-WAN in vWAN" application in the Azure Market place. The **VMware SD-WAN in vWAN** managed application page appears. You can use this application to automate the deployment of Virtual Edges in Virtual WAN Hub.

[Home](#) >

VMware SD-WAN in vWAN ...

VeloCloud



VMware SD-WAN in vWAN

VeloCloud

Create

- 6 Click **Create** on the managed application and enter the following basic details:

[Home](#) > [VMware SD-WAN in vWAN \(preview\)](#) >

Create VMware SD-WAN in vWAN

Basics VMware SD-WAN in Virtual WAN Review + create

Project details

Select the subscription to manage deployed resources and costs. Use resource groups like folders to organize and manage all your resources.

Subscription * ⓘ ▼

Resource group * ⓘ ▼

[Create new](#)

Instance details

Region * ⓘ ▼

Managed Application Details

Provide a name for your managed application, and its managed resource group. Your application's managed resource group holds all the resources that are required by the managed application which the consumer has limited access to.

Application Name *

Managed Resource Group * ⓘ ✓

- **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 In the **VMware SD-WAN in Virtual WAN** tab, select Virtual WAN Hub in the selected region. The Virtual Edges will be deployed in this Hub.

[Home](#) > [VMware SD-WAN in vWAN \(preview\)](#) >

Create VMware SD-WAN in vWAN

Basics	VMware SD-WAN in Virtual WAN	Review + create
Virtual WAN Hub	<div>WestUSHub</div>	
Scale unit * ⓘ	<div>2 Scale Units - 1.0 Gbps</div>	
VMware SD-WAN Orchestrator * ⓘ	<div></div>	
IgnoreCertErrors * ⓘ	<div>False</div>	
ActivationKey for VMware SD-WAN Edge1 * ⓘ	<div></div>	
ActivationKey for VMware SD-WAN Edge2 * ⓘ	<div></div>	
BGP ASN * ⓘ	<div></div>	
ClusterName * ⓘ	<div></div>	

Once the customer selects a Virtual WAN hub, the following information appears which lists the BGP neighbor IP Addresses and the ASN of the Virtual WAN Hub. Make a note of this information as it will be needed to configure BGP neighborships on the Orchestrator.

i BGP neighbor IPs ["10.101.32.4","10.101.32.5"]Virtual WAN Hub BGP ASN 65515

- **Scale unit:** Select the scale as required.
- **VMware 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 VMware SD-WAN 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 an unique name for the deployment and ensure not to use special characters such as #, @, _, -, and so on.
- 8 After entering all the required fields, click **Review + create**.
- 9 The deployment process will start and it will take 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 will 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. 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 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.

Interface Status

View the MAC address and connection status of physical interfaces.

Routed Interfaces										
Name	MAC Address	Link Detected	IP Address	Netmask	IPv6 Address	Speed	Autonegotiation	RX errors	TX errors	Collisions
GE1	00:22:48:06:81:2B	true	10.101.112.133	255.255.255.128		40000 Mbps, full duplex	on	0	0	0
GE2	00:22:48:06:98:45	true	10.101.112.6	255.255.255.128		40000 Mbps, full duplex	on	0	0	0

Two static routes are configured on the Edge to reach BGP neighbors as shown in the following screenshot.

Static Route Settings

IPv4 IPv6

Local Routes

+ ADD - REMOVE CLONE

<input type="checkbox"/>	Subnet	Source IP	Next Hop IP	Interface	VLAN	Cost	Preferred	Advertise	ICMP Probe	Description
<input type="checkbox"/>	10.101.32.5/32	N/A	10.101.112.1	GE2		0	<input checked="" type="checkbox"/> Yes	<input checked="" type="checkbox"/> Yes		+ NEW Enter Descrip...
<input type="checkbox"/>	10.101.32.4/32	N/A	10.101.112.1	GE2		0	<input checked="" type="checkbox"/> Yes	<input checked="" type="checkbox"/> Yes		+ NEW Enter Descrip...

2 items

- 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.

Neighbors

IPv4 IPv6


+ ADD - DELETE CLONE

<input type="checkbox"/>	Neighbor IP *	ASN *	Inbound Filter	Outbound Filter	Additional Options
<input type="checkbox"/>	10.101.32.5	65515	[None]	[None]	<p>VIEW LESS</p> <p>Max-Hop 2</p> <p>Local IP IP Address</p> <p>Source Interface Auto</p> <p>Uplink <input type="checkbox"/></p> <p>Allow AS <input type="checkbox"/></p> <p>Default Route <input type="checkbox"/></p> <p>Enable BFD <input type="checkbox"/></p> <p>Keep Alive Example: 10</p> <p>Hold Timer Example: 10</p> <p>Connect <input type="checkbox"/> Example: 10</p> <p>MDS Auth <input type="checkbox"/></p> <p>MDS Password Password</p>
<input type="checkbox"/>	10.101.32.4	65515	[None]	[None]	<p>VIEW ALL</p>

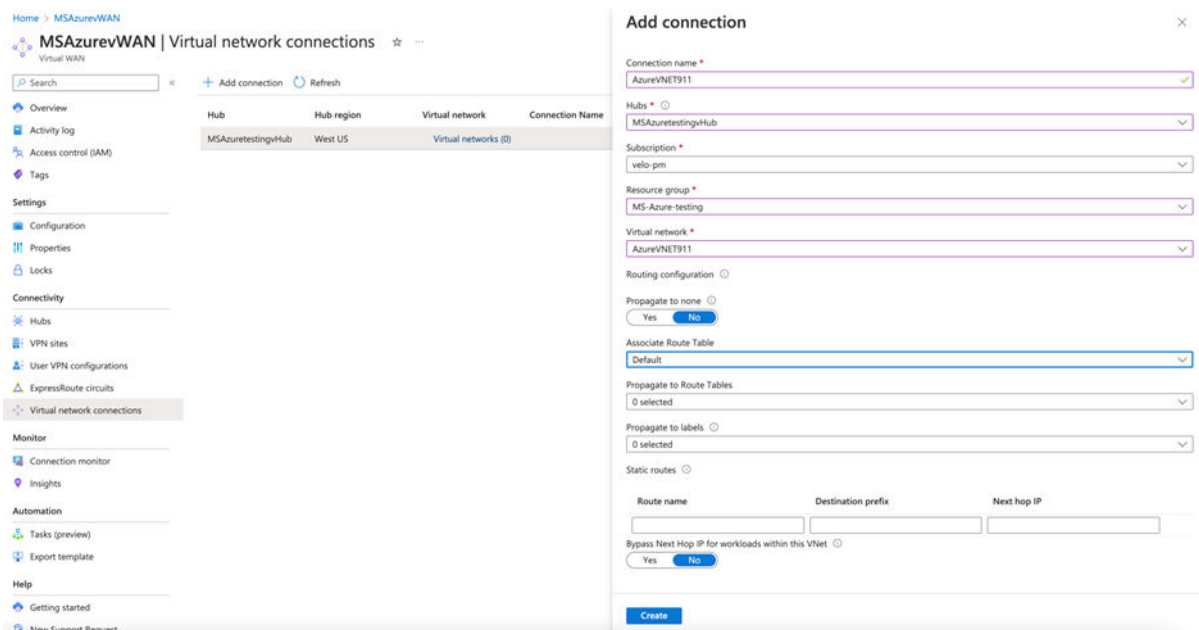
Once static routes and BGP neighborships are configured, the Virtual Edges should begin learning routes from the Azure Virtual WAN Hub. BGP neighborhood status can be verified under **Monitor > Network Services**.

- 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**.



Click on **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.



For the vWAN NVA Edge, the image is a 2 NIC Deployment i.e. the GE1 interface is not used for 'management' interface. This is unique to the vWAN NVA image. In the cloud_init, the 'management_interface' flag needs to be set to 'False'.

```
#cloud-config
password: Velocloud123
chpasswd: { expire: False }
ssh_pwauth: True
velocloud:
```

```
vce:
  management_interface: false
  vco: $vco
  activation_code: $velo2_token
  vco_ignore_cert_errors: $velo_ignore_cert_errors
```

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 VMware SD-WAN Edge Deployed as Azure vWAN vWAN NVA](#) .

Hub Upgrade Instructions for VMware SD-WAN Edge Deployed as Azure vWAN NVA

This document is intended for customers who use VMware SD-WAN Edges in Azure and deploy them as Network Virtual Appliances (NVAs) in the Azure Virtual WAN (vWAN) Hub.

For more information,

see <https://learn.microsoft.com/en-us/azure/virtual-wan/virtual-wan-faq#why-am-i-seeing-a-message-and-button-called-update-router-to-latest-software-version-in-portal>.

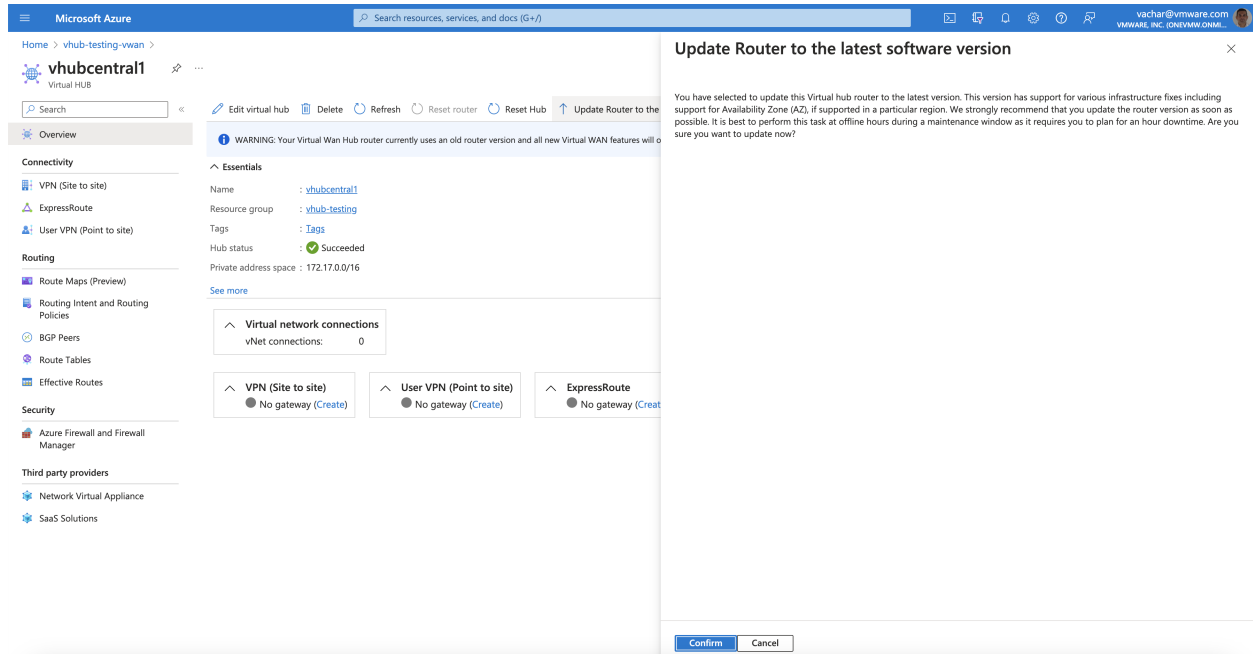
Upgrade Instructions

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, click "Update router to latest software version" button to initiate router upgrade.

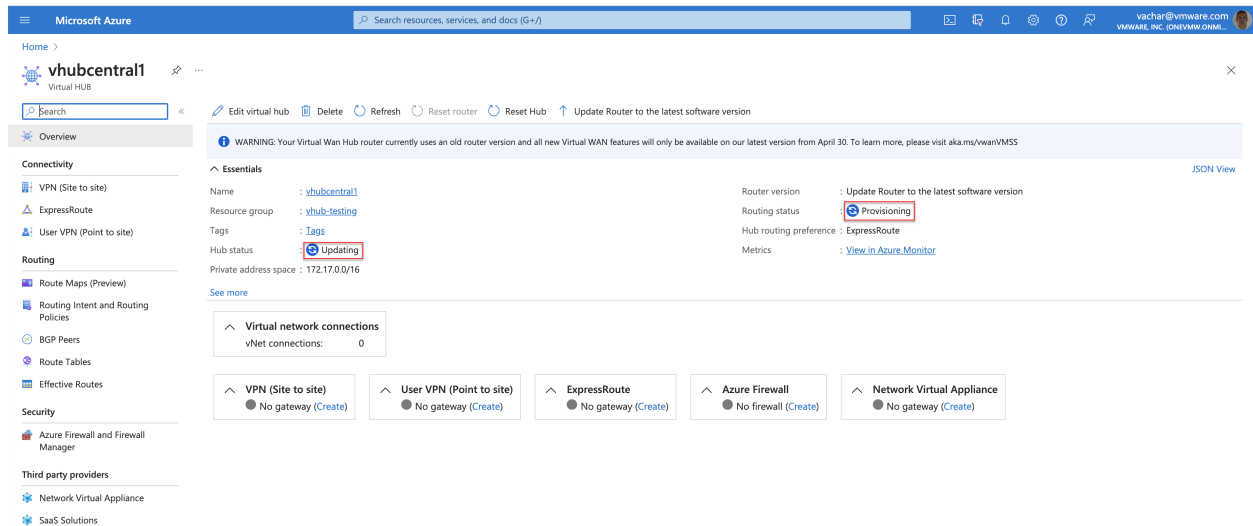
Note All newly created Virtual Hubs will be automatically deployed on the latest Virtual Machine Scale Sets-based infrastructure and do not require this upgrade.

Microsoft Azure portal interface showing the 'vhubcentral1' Virtual WAN Hub resource. The 'Update Router to the latest software version' button is highlighted. A warning message indicates that the current router version is outdated and features will be unavailable after April 30. The Essentials section shows the router version as 'Update Router to the latest software version' and the routing status as 'Provisioned'. Below, there are five expandable sections for connectivity options: VPN (Site to site), User VPN (Point to site), ExpressRoute, Azure Firewall, and Network Virtual Appliance, each with a 'No gateway (Create)' button.

After clicking **"Upgrade Router to the latest software version"**, a message will indicate that this operation must be performed during a maintenance window.



The Hub Status would display **"Updating"** and the Routing State as **"Provisioning"**. This process will take approximately 30 to 60 minutes to complete.



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.

The screenshot shows the Microsoft Azure portal interface for a Virtual Hub named 'vhubcentral1'. The left sidebar contains navigation options: Overview, Connectivity, Routing, Security, and Third party providers. The main content area displays the 'Essentials' section with the following details:

- Name: vhubcentral1
- Resource group: vhub-testing
- Tags: Tags
- Hub status: Succeeded
- Private address space: 172.16.0.0/16
- Routing status: Provisioned
- Hub routing preference: ExpressRoute
- Metrics: View in Azure Monitor

Below the Essentials section, there are five cards representing different connection types, each with a 'No gateway (Create)' button:

- Virtual network connections (vNet connections: 0)
- VPN (Site to site)
- User VPN (Point to site)
- ExpressRoute
- Azure Firewall
- Network Virtual Appliance

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.

The screenshot shows the Microsoft Azure portal interface for the 'BGP Peers' configuration of the Virtual Hub 'vhubcentral1'. The left sidebar is the same as the previous screenshot, but the 'BGP Peers' option is selected. The main content area displays the 'Essentials' section with the following details:

- Name: vhubcentral1
- Location: West Central US
- ASN: 65515
- IP: 172.16.32.9, 172.16.32.8

Below the Essentials section, there is a table with the following columns: Name, ASN, IPv4 Address, and Virtual Network connection. The table currently shows 'No results'.

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 (VMware 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. To configure BGP neighbors for Virtual Edges, see [BGP Neighbor Configuration](#).

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. See [Static Routes Configuration](#).

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.

The screenshot shows the VMware Orchestrator interface. The top navigation bar includes 'vmw Orchestrator', 'TbanPM tbanPLM', 'SD-WAN', and user icons. The left sidebar has 'Monitor', 'Configure', 'Diagnostics', and 'Service Settings'. Under 'Diagnostics', 'Remote Diagnostics' is selected. The main panel shows the 'Edges / nva1-azurevwan' section with a 'Connected' status. Below this, there are two sections: 'IPv6 Route Table Dump' and 'Interface Status', each with a 'RUN' button. The 'Interface Status' section shows a table of 'Routed Interfaces' with columns: Name, MAC Address, Link Detected, IP Address, Netmask, IPv6 Address, Speed, Autonegotiation, RX errors, TX errors, and Collisions. The table lists interfaces GE1 through GE6. The row for GE2 is highlighted with a red border, showing an IP address of 172.16.112.5 and a netmask of 255.255.255.128. The test duration is noted as 3.002 seconds.

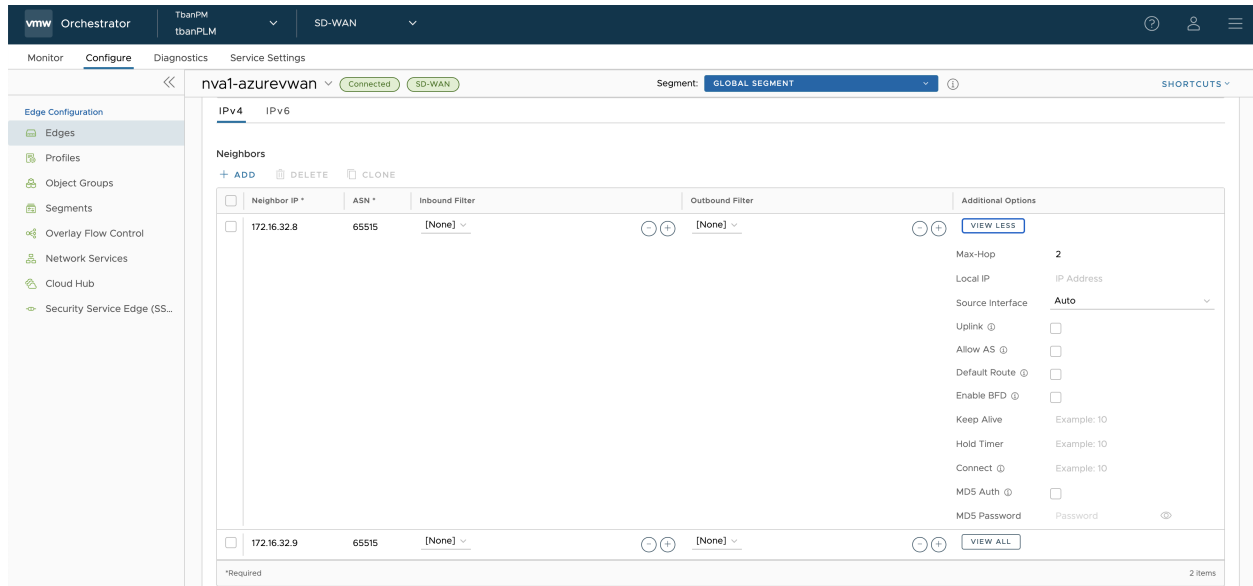
Name	MAC Address	Link Detected	IP Address	Netmask	IPv6 Address	Speed	Autonegotiation	RX errors	TX errors	Collisions
GE1	00:22:48:5E:82:84	true	172.16.112.132	255.255.255.128		50000 Mbps, full duplex	on	0	0	0
GE2	00:22:48:5E:82:84	true	172.16.112.5	255.255.255.128		50000 Mbps, full duplex	on	0	0	0
GE3		false	N/A	N/A	N/A	N/A	N/A	-1	-1	-1
GE4		false	N/A	N/A	N/A	N/A	N/A	-1	-1	-1
GE5		false	N/A	N/A	N/A	N/A	N/A	-1	-1	-1
GE6		false	N/A	N/A	N/A	N/A	N/A	-1	-1	-1

Two static routes are configured on the Edge to reach BGP neighbors, as illustrated in the following screenshot.



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.



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**.

vmw

Orchestrator

TbaniPM

tbaniPLM

SD-WAN

Monitor

Configure

Diagnostics

Service Settings

Monitor

Network Overview

Edges

Network Services

Routing

Alerts

Events

Firewall Logs

Reports

Edge Network Intelligence

Routing

Multicast Groups

PIM Neighbors

BGP Edge Neighbor State

BFD

BGP Gateway Neighbor State

Gateway Route Table

Q Search

1

<div></div>	nva1-azurevwan	Global Segment	172.16.32.5	<div>Removed</div>	Sep 27, 2023, 7:25:58 AM 12 days ago	165	149	4	0
<div></div>	nva2azurevwan	Global Segment	172.16.32.8	<div>Established</div>	Sep 27, 2023, 8:18:27 AM 12 days ago	20,325	17,785	0	01w5d08h 3
<div></div>	nva1-azurevwan	Global Segment	172.16.32.8	<div>Established</div>	Sep 27, 2023, 7:25:58 AM 12 days ago	20,305	17,782	0	01w5d07h 3
<div></div>	nva1-azurevwan	Global Segment	172.16.32.9	<div>Established</div>	Sep 27, 2023, 8:23:43 AM 12 days ago	20,309	17,779	2	01w5d07h 3
<div></div>	nva2azurevwan	Global Segment	172.16.32.9	<div>Established</div>	Sep 27, 2023, 8:18:27 AM 12 days ago	20,285	17,785	0	01w5d08h 3

COLUMNS

REFRESH

28 Items

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

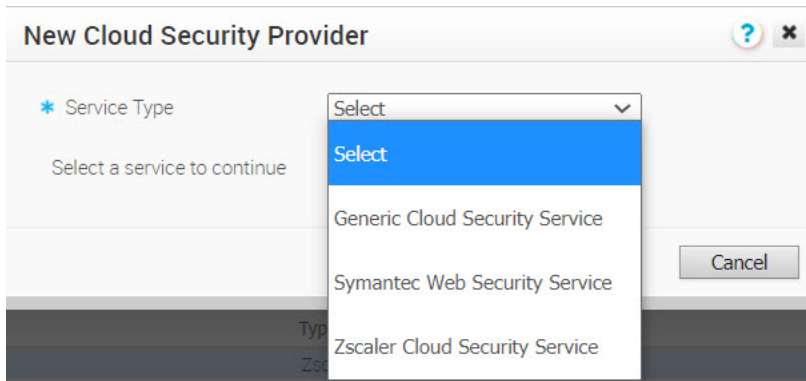
The cloud security service establishes a secure tunnel from an Edge to the cloud security service sites. This ensures secured traffic flow to the cloud security services.

To configure a Cloud Security Service, perform the following steps.

Procedure

- 1 In the Enterprise portal, click **Configure > Network Services**.

- 2 In the **Cloud Security Service** section, click **New**.



- 3 In the **New Cloud Security Provider** window, select a service type from the drop-down menu.
- a If you have selected either "Generic" or "Symantec Web" Cloud Security Service as the Service Type, then configure the following required details and click **Add**.

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 automation deployment by selecting the **Automate Cloud Service Deployment** checkbox. 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 an SD-WAN Edge to Zscaler

This section describes how to automatically create a GRE or IPsec tunnel from an SD-WAN Edge to a Zscaler service provider.

New Cloud Security Provider

* Service Name:

* Service Type:

* Automate Cloud Service Deployment: ☒

Tunneling Protocol: ☐ IPsec ☒ GRE

Domestic Preference: ☒

* Zscaler Cloud:

* Partner Admin Username:

* Partner Admin Password:

* Partner Key:

* Domain:

L7 Health Check: ☒

HTTP Probe Interval: sec

Number of Retries:

RTT Threshold: msec

Zscaler Login URL:

- a In the **New Cloud Security Provider** window, enter a service name.
- b Select the **Automate Cloud Service Deployment** checkbox.
- c 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.

- d Configure additional details such as Domestic Preference, Zscaler Cloud, Partner Admin Username, Password, API Key, and Domain, as described in the following table.

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 This option is configurable only if GRE is selected for establishing tunnels.
Zscaler Cloud	Select a Zscaler cloud service from the drop-down menu or enter the Zscaler cloud service name in the textbox.
Partner Admin Username	Enter the provisioned username of the partner admin.

Option	Description
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.

- e Click **Validate Credentials**. If the validation is successful, the **Add** button will be enabled.

Note You must validate the credentials to add a new CSS Provider.

- f Optional: Configure the following **L7 Health Check** details to monitor the health of the Zscaler Server.

Note The **L7 Health Check** feature tests HTTP reachability to the Zscaler backend server. Upon enabling L7 Health Check, the Edge sends HTTP L7 probes to a Zscaler destination (Example: `http://<zscaler cloud>/vpntest`) which is Zscaler's backend server for the HTTP health check. This method is an improvement over using network level keep-alive (GRE or IPsec) as that method only tests for network reachability to the frontend of a Zscaler server.

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.

Option	Description
L7 Health Check	<p>Select the checkbox to enable L7 Health Check for the Zscaler Cloud Security Service provider, with default probe details (HTTP Probe interval = 5 seconds, Number of Retries = 3, RTT Threshold = 3000 milliseconds). By default, L7 Health Check is not enabled.</p> <p>Note Configuration of health check probe details is not supported.</p> <p>Note For a given Edge/Profile, a user cannot override the L7 health check parameters configured in the Network Service.</p>
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.

Option	Description
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	<p>Enter the login URL and then click Login to Zscaler. This will redirect you to the Zscaler Admin portal of the selected Zscaler cloud.</p> <p>Note The Login to Zscaler button will be enabled if you have entered the Zscaler login URL.</p>

- g If you want to login to the Zscaler Admin portal from the Orchestrator, enter the Zscaler login URL and then click **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 more information about Zscaler CSS automation, see [Zscaler and VMware 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 describes how to manually create a GRE or IPsec tunnel from SD-WAN Edge to Zscaler service provider. Unlike automatic tunnels, configuring manual tunnels requires you to specify a tunnel destination to bring up the tunnels.

New Cloud Security Provider

* Service Name: Zscaler Manual

* Service Type: Zscaler Cloud Security Service

* Automate Cloud Service Deployment: ☐

* Primary Server: 199.168.148.131

Secondary Server: 10.64.4.40

* Zscaler Cloud: zscalerbeta.net

L7 Health Check: ☒

HTTP Probe Interval: 5 sec

Number of Retries: 3

RTT Threshold: 3000 msec

Zscaler Login URL: https://admin.zscaler.net

Login to Zscaler

Add Cancel

- In the **New Cloud Security Provider** window, enter a service name.
- Enter the IP address or hostname for the Primary server.
- Optionally, you can enter the IP address or hostname for the Secondary server.
- Select a Zscaler cloud service from the drop-down menu or enter the Zscaler cloud service name in the textbox.
- Configure other parameters as desired, and then click **Add**.

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.

Results

The configured cloud security services are displayed under the **Cloud Security Service** area in the **Network Services** window.

Cloud Security Service

[New...](#)
[Delete...](#)

Name	Type	Used By
<input type="checkbox"/> zscaler_gre_auto	Zscaler Cloud Security Service	2 Edges
<input type="checkbox"/> zscaler_ipsec_auto	Zscaler Cloud Security Service	1 Edge

What to do next

Associate the cloud security service with a Profile or an Edge:

- [Configure Cloud Security Services for Profiles](#)
- [Configure Cloud Security Services for Edges](#)

Configure Cloud Security Services for Profiles

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, VMware SD-WAN and Zscaler integration supports automation of IPsec and GRE tunnels.

Note Only one CSS with GRE is allowed per Profile.

Before you begin:

- Ensure that you have access permission to configure network services.
 - Ensure that your SD-WAN 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.
- 1 In the Enterprise portal, click **Configure > Profiles**.
 - 2 Click the Device Icon next to a profile, or click the link to the profile, and then click 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:

Cloud Security Service 

Cloud Security Service zscalerbeta ▼

Hash SHA 1 ▼

Encryption None ▼

Key Exchange Protocol ☐ IKEv1 ☒ IKEv2

Login to Zscaler

Option	Description
Cloud Security Service	<p>Select a cloud security service from the drop-down menu to associate with the profile. You can also click New Cloud Security Service from the drop-down to create a new service type. For more information about how to create a new CSS, see Configure a Cloud Security Service.</p> <hr/> <p>Note For cloud security services with Zscaler login URL configured, Login to Zscaler button appears in the Cloud Security Service area. Clicking the Login to Zscaler button will redirect you to the Zscaler Admin portal of the selected Zscaler cloud.</p>
Tunneling Protocol	<p>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.</p> <hr/> <p>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.</p>
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	<p>Select the key exchange method as IKEv1 or IKEv2. By default, IKEv2 is selected.</p> <p>This option is not available for Symantec cloud security service.</p>
Login to Zscaler	Click Login to Zscaler to login to the Zscaler Admin portal of the selected Zscaler cloud.

5 Click **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](#).

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 Enterprise portal, click **Configure > Edges**.
- 2 Select an Edge you want to override CSS settings and click the icon under the **Device** column. The **Device Settings** page for the selected Edge appears.
- 3 In the **Cloud Security Service** area, the CSS parameters of the associated profile are displayed. Select **Enable Edge Override** to select a different CSS or to modify the attributes inherited from the profile associated with the Edge. For more information on the attributes, see [Configure Cloud Security Services for Profiles](#).
- 4 Click **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.

Cloud Security Service On

Cloud Security Service manual

Tunneling Protocol ☒ IPsec ☐ GRE

Hash SHA 1

Encryption None

Key Exchange Protocol ☐ IKEv1 ☒ IKEv2

Credentials

FQDN	PSK
S15.L2B13.E1.V05d7@velocloud.net
S15.L8F3D.E1.Vfefd@velocloud.net
S15.LBB35.E1.V4380@velocloud.net

Login to Zscaler

Enable Edge Override

Note For cloud security services with Zscaler login URL configured, **Login to Zscaler** button appears in the **Cloud Security Service** area. Clicking 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 Click **Add Tunnel**.

Cloud Security Service On

Cloud Security Service Zscaler Manual

Tunneling Protocol ☐ IPsec ☒ GRE

GRE Tunnels Add Tunnel

Action Pending Tunnel Creation

Login to Zscaler

Enable Edge Override

- 2 In the **Add Tunnel** window appears, configure the following GRE tunnel parameters, and click **OK**.

Add Tunnel

WAN Links: 169.254.7.10

Tunnel Source Public IP: Custom WAN IP 216.66.5.49

Tunnel Addressing	Point-of-Presence	Router IP/Mask	Internal ZEN IP/Mask
Primary Address	10.1.1.1	172.18.58.121/30	172.18.58.122/30
Secondary Address	10.2.2.2	172.18.58.125/30	172.18.58.126/30

OK Cancel

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 Datacenter.
Secondary Point-of-Presence	Enter the secondary Public IP address of the Zscaler Datacenter.
Primary Router IP/Mask	Enter the primary IP address of Router.
Secondary Router IP/Mask	Enter the secondary IP address of Router.

Option	Description
Primary ZEN IP/Mask	Enter the primary IP address of Internal Zscaler Public Service Edge.
Secondary 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, VMware SD-WAN and Zscaler integration supports:

- [IPsec/GRE Tunnel Automation](#)
- [Zscaler Location/Sub-Location Configuration](#)

For a selected automated Zscaler CSS provider at the Edge level, you can override the CSS settings inherited from the profile, establish automatic IPsec/GRE tunnels for each Edge Segment, create Sub-locations, and configure Gateway options and Bandwidth controls for Location and Sub-locations.

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 Enterprise portal, click **Configure > Edges**.
- 2 Select an Edge you want to establish automatic tunnels.
- 3 Click the icon under the **Device** column. The **Device Settings** page for the selected Edge appears.
- 4 In the **Cloud Security Service** section, select **Enable Edge Override**.

- From the **Cloud Security Service** drop-down menu, select an automated CSS provider and click **Save Changes**.

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/about-sub-locations>.
- 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:

- In the Enterprise portal, click **Configure > Edges**.

- 2 Select an Edge and click 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.

Zscaler On Enable Edge Override

Cloud zscalerbeta.net

Location

Name	Gateway Options	Action Details
edge_6762b07b-789e-4f00-9f1c-2986191d99a4	Edit	View

Sub-Locations


Sub-Location Name	LAN Networks	Subnets	Gateway Options	Action Details	Action
1 Other			Edit		+

- 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.

Note 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, clicking **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, click the **Edit** button under **Gateway Options**. For more information, see [Configure Zscaler Gateway Options and Bandwidth Control](#).

- 5 To create a Sub-location, in the **Sub-Locations** table, click the  icon under the **Action** column.
 - a In the **Sub-Location Name** textbox, enter a unique name for the Sub-location. The Sub-location name should be unique across all 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 Click **Save Changes**.

Zscaler On Enable Edge Override

Cloud:


Location:

Name	Gateway Options	Action Details
edge_6762b07b-789e-4f00-9f1c-2986191d99a4	Edit	View

Sub-Locations:

Sub-Location Name	LAN Networks	Subnets	Gateway Options	Action Details	Action
1 Other			Edit	View	+
2 corporate1	<input type="text" value="x 1 - Corporate"/>	<input type="text" value="x 172.21.11.0/24"/>	Edit	View	+ -

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 clicking the **Edit** button under **Gateway Options** in the **Sub-Locations** table. For more information, see [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 click **Save Changes**, the Sub-location configurations on the Zscaler side will be updated automatically.
- e To delete a Sub-location, click the  icon under the **Action** column.

Note When the last Sub-location is deleted from the table, the "other" Sub-location will also be deleted automatically.

Configure Zscaler Gateway Options and Bandwidth Control

To configure Gateway options and Bandwidth controls for the Location and Sub-location, click the **Edit** button under **Gateway Options**, in the respective table.

The **Zscaler Gateway Options and Bandwidth Control** window appears.

Zscaler Gateway Options and Bandwidth Control

Gateway Options

Use XFF from Client Request ☐ Off

Enable Caution ☐ Off

Enable AUP ☐ Off

Enforce Firewall Control ☐ Off

Authentication ☐ Off

Bandwidth Control

Bandwidth Control ☐ Off

[Save Changes](#)
[Cancel](#)

Configure the Gateway options and Bandwidth controls for the Location and Sub-location, as needed, and click **Save Changes**.

Note The Zscaler Gateway Options and Bandwidth Control parameters that can be configured for 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 more information about Zscaler Gateway Options and Bandwidth Control parameters, see <https://help.zscaler.com/zia/configuring-locations>

Option	Description
Gateway Options for Location/Sub-Location	
Use XFF from Client Request	<p>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.</p> <p>Note This Gateway option is only configurable for Parent location.</p>
Enable Caution	If you have not enabled Authentication , you can enable this feature to display a caution notification to unauthenticated users.
Enable AUP	<p>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:</p> <ul style="list-style-type: none"> ■ 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: <ul style="list-style-type: none"> ■ 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.

Option	Description
Enforce Firewall Control	<p>Select to enable the service's firewall control.</p> <hr/> <p>Note Before enabling this option, user must ensure if its Zscaler account has subscription for "Firewall Basic".</p>
Enable IPS Control	<p>If you have enabled Enforce Firewall Control, select this to enable the service's IPS controls.</p> <hr/> <p>Note Before enabling this option, user must ensure if its Zscaler account has subscription for "Firewall Basic" and "Firewall Cloud IPS".</p>
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	<p>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.</p> <ul style="list-style-type: none"> ■ 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.
Refresh Time for re-validation of Surrogacy	<p>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.</p> <ul style="list-style-type: none"> ■ 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	<p>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.</p>

Option	Description
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)	
<div> <div>Zscaler Gateway Options and Bandwidth Control</div> <div> <div> Gateway Options <div> <div>Enable Caution</div> <div>Off</div> </div> <div> <div>Enable AUP</div> <div>Off</div> </div> <div> <div>Enforce Firewall Control</div> <div>Off</div> </div> <div> <div>Authentication</div> <div>Off</div> </div> </div> <div> Bandwidth Control <div> <div>Use Location Bandwidth</div> <div>Override</div> <div>Disabled</div> </div> </div> <div> <div>Save Changes</div> <div>Cancel</div> </div> </div> </div>	
<p>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).</p>	
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 **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.

Related links

- [Monitor Cloud Security Services](#)
- [Monitor Cloud Security Services Events](#)
- [Monitor Network Services](#)

Configure Business Policies with Cloud Security Services

You can create business policies to redirect the traffic to a Cloud Security Service.

For more information on business policies, see [Create Business Policy Rules](#).

Procedure

- 1 In the Enterprise portal, click **Configure > Profiles**.
- 2 Select a profile from the list and click the **Business Policy** tab.
- 3 Click **New Rule** or **Actions > New Rule**.
- 4 Enter a name for the business rule.
- 5 Choose the **Match** options to match the traffic.

- 6 In the **Action** area, click the **Internet Backhaul** button and choose a **Cloud Security Service** from the drop-down list. You must have already associated the cloud security service to the profile.

Configure Rule

Rule Name:

Match

Source: **Any** | Object Group | Define...

Destination: **Any** | Object Group | Define...

☐ Any
☒ Internet
☐ Edge
☐ Non SD-WAN Destination via Gateway
☐ Non SD-WAN Destination via Edge

IP Address:

CIDR prefix:

Domain Name:

Protocol:

Ports:

Application: **Any** | Define...

Action

Priority: **High** | **Normal** | Low

☐ Rate Limit

Network Service: **Direct** | Multi-Path | **Internet Backhaul**

☐ Backhaul Hubs
☐ Non SD-WAN Destination via Gateway
☒ Non SD-WAN Destination via Edge / Cloud Security Service

Link Steering: **Auto** | Transport Group | Interface | WAN Link

Inner Packet DSCP Tag:

Outer Packet DSCP Tag:

NAT: **Disabled** | Enabled

Service Class: **Real Time** | **Transactional** | Bulk

OK Cancel

- 7 Choose the other actions as required and click **OK**.

Results

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 click the **Business Policy** tab.
- 2 Click **New Rule** or **Actions > New Rule**.
- 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.

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 Enterprise portal, click **Monitor > Network Services**. The **Network Services** page appears.
- 2 The **Cloud Security Services Sites** section displays all the CSS configured for the Enterprise along with the following configuration details.

	Name	Public IP	Status	Tunnel Status	Service Status	State Changed Time	Events	DeploymentStatus
1	zscaler_gre_auto	199.168.14... 104.129.19... 104.129.19...		 	10	Mon Jul 26, 12:43:08 a few sec...	56 View	View
2	zscaler_ipsec_auto	199.168.14... 104.129.19...		 	 	Mon Jul 26, 12:42:56 a few sec...	54 View	View

Field	Description
Name	The name of the CSS provider.
Public IP	The Public IP address of the CSS provider.
Status	<p>The overall status of the CSS provider:</p> <ul style="list-style-type: none"> ■ White - Specifies two possible states: <ul style="list-style-type: none"> ■ 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.

Field	Description
Tunnel Status	<p>The status of tunnels created from the CSS provider from different Edges:</p> <ul style="list-style-type: none"> ■ White - Specifies two possible states: <ul style="list-style-type: none"> ■ 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: <ul style="list-style-type: none"> ■ 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: <ul style="list-style-type: none"> ■ ALL_DOWN - All the tunnels associated with the CSS provider are DOWN. ■ DOWN - A specific tunnel associated with the CSS provider is DOWN. <p>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.</p>
Service Status	<p>The status of the external service as recorded by each Edge:</p> <ul style="list-style-type: none"> ■ 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: <ul style="list-style-type: none"> ■ 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.
Events	The number of related state change events.
DeploymentStatus	The deployment status of the CSS provider.

3 Click the link in the **Events** column to view the related state change Events.

Events

Provider: Zscaler automated -1
Type: Zscaler Web Security Service

	Edge	Identifier	Public IP	State	State Changed Time
1	b1-edge1	Link 00000004-f47d-4d4b-935a-58144ff27e27	104.129.194.39	STANDBY	Tue Dec 15, 15:39
2	b1-edge1	Link 00000005-f47d-4d4b-935a-58144ff27e27	104.129.194.39	STANDBY	Tue Dec 15, 15:39
3	b1-edge1	Link 00000003-f47d-4d4b-935a-58144ff27e27	104.129.194.39	STANDBY	Tue Dec 15, 15:39
4	b1-edge1	S15.LFF67.E1.V2f46@velocloud.net	199.168.148.132	UP	Tue Dec 15, 15:39
5	b1-edge1	S15.L8C73.E1.Vcb90@velocloud.net	199.168.148.132	UP	Tue Dec 15, 15:39
6	b1-edge1	S15.L6E80.E1.Vd099@velocloud.net	199.168.148.132	UP	Tue Dec 15, 15:39

Close

- 4 Click the link in the **Deployment Status** column to view the deployment status of the CSS provider.

Cloud Security Service Automated Deployment Status for Zscaler-Beta-GRE

Search Status in COMPLETE Refresh

Display 3 items

Pending Location 0 Pending 0 Notified 0 **Completed 3** Errored 0 Timed Out 0 Pending Delete 0

Edge	Segment	Action	Status	Zscaler Object Name	Details
b1-edge1		createCloudServiceGRESite	COMPLETE		Details
b1-edge1		createCloudServiceGRESite	COMPLETE		Details
b1-edge1		createCloudServiceGRESite	COMPLETE		Details

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 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.
- 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.
- **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 Click **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 as shown in the following sample screenshot.

Edge	Status	HA	Links	VM Status	VNF	Edge Tunnels	Gateways	Profile
1 b1-edge1	●		↔ 3			↔ 3	View	Quick Start Proc
2 b2-edge1	●		↔			↔	View	Quick Start Proc
3 b3-edge1	●		↔			↔	View	Quick Start Proc

Interface	Segment	Tunnel IP	Tunnel State	Location	L7 Health	L7 Health RTT
GE5	Global Segment	104.129.194.43	● Up	Fremont, CA	● N/A	N/A
GE3	Global Segment	104.129.194.45	● Up	Fremont, CA	● N/A	N/A
GE4	Global Segment	199.168.148.131	● Up	Fremont, CA	● N/A	N/A

Monitor Cloud Security Services Events

You can view the events related to cloud security services from the **Monitor > Events** page.

In the enterprise portal, click **Monitor > Events**.

To view the events related to cloud security service sites, 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.

Monitor

Network Overview

Edges

Network Services

Routing

Alerts

Events

Reports

Configure

Test & Troubleshoot

Administration

Events

Tue Jul 27, 7:12

Past 12 Hours

now

<

>

Search

Cols

Reset View

Refresh

CSV

Display 281 items

	Time	Event	Segment	Edge	User	Severity
i	Tue Jul 27, 12:05:26	Zscaler Location object created		b1-edge1		Info
i	Tue Jul 27, 12:05:26	Network Service created				Info
i	Tue Jul 27, 12:05:07	Call made to external API	Global Segment	b1-edge1		Info
i	Tue Jul 27, 12:04:57	Zscaler Location object deleted				Info
i	Tue Jul 27, 12:04:42	Call made to external API				Info
i	Tue Jul 27, 12:04:22	Edge Non SD-WAN Destination tunnel down	Global Segment	b2-edge1		Info
i	Tue Jul 27, 12:04:21	Cloud Security Service site creation enqueued	Global Segment	b1-edge1		Info
i	Tue Jul 27, 12:04:21	Cloud Security Service site creation enqueued	Global Segment	b1-edge1		Info

The following table includes the Enterprise events which help track various Edge actions related to CSS deployment, Location and Sub-location automation.

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.

Events	Description
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.

You can also view the events in the new Orchestrator UI.

Click **Launch New Orchestrator UI** in the pop-up window. The UI opens in a new tab displaying the monitoring options.

Click **Events**. Click the Filter Icon in the **Search** option to filter the events.

Configure DNS Services

This is an optional service that allows you to create a configuration for DNS.

The DNS Service can be for a public DNS service or a private DNS service provided by your company. A **Primary Server** and **Backup Server** can be specified. The service is preconfigured to use Google and Open DNS servers.

The following figure shows a sample configuration for a Public DNS.

For a private service, you can also specify one or more **Private Domains**.

Configure Netflow Settings

In an Enterprise network, Netflow monitors traffic flowing through SD-WAN Edge and exports Internet Protocol Flow Information Export (IPFIX) information directly from SD-WAN 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. VMware 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 SD-WAN 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 SD-WAN 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.

Procedure

- 1 From the SD-WAN Orchestrator, go to **Configure > Network Services**.
The **Services** page appears.
- 2 To configure a collector, go to the **Netflow Settings** area and click the **New** button at the right side of the Collector table. The **Add New Collector** dialog box appears.
 - 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 Click **Save Changes**.

Under **Network Services**, the newly added collector appears in the Collector table.

- 3 SD-WAN Orchestrator allows filtering of traffic flow records by source IP, destination IP, and application ID associated with the flow. To configure a filter, go to the **Netflow Settings** area and click the **New** button at the right side of the Filter table. The **Add New Filter** dialog box appears.

The screenshot shows the 'Add New Filter' dialog box. At the top, the 'Filter Name' is 'Allow_ICMP'. Below this is the 'Match' section with three rows: 'Source', 'Destination', and 'Application'. Each row has 'Any' and 'Define...' buttons. 'Source' is currently set to 'IP Address' with the value '0.0.0.0/1'. Below the 'Match' section is a list of applications: Internet Instant Messaging, Management, Media, Network Service, and Other TCP/UDP. 'Network Service' is expanded, showing a sub-list with IATP, IBM RPL, ICMP (selected), ICMPv6, and IDPR. At the bottom is the 'Action' section with 'Filter Action' set to 'Allow' and 'Deny' buttons. The dialog has 'OK' and 'Cancel' buttons at the bottom right.

- In the **Filter Name** text box, enter a unique name for the filter.
- Under the **Match** area, click **Define** to define per collector filtering rules to match by source IP or destination IP or application associated with the flow, or click **Any** to use any of the source IP or destination IP or application associated with the flow as the match criteria for Netflow filtering.
- Under the **Action** area, select either **Allow** or **Deny** as the filter action for the traffic flow, and click **OK**.

Under **Network Services**, the newly added filter appears in the Filter table.

Results

At the profile and edge level, the configured collectors and filters appears as a list under the **Netflow Settings** 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 SD-WAN Edge, it periodically sends messages to the configured collector. The contents of these messages are defined using IPFIX templates. For more information on templates, see [IPFIX Templates](#).

IPFIX Templates

After you enable Netflow on the VMware SD-WAN 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.

Non-NAT Template

<https://www.iana.org/assignments/ipfix/ipfix.xhtml>. This is an aggregated flow. Keys for this flow record are: sourceIPv4Address, destinationIPv4Address, destinationTransportPort, ingressVRFID, ApplicationID, protocolIdentifier. Source port is aggregated out.

Template ID: 256

The Non-NAT template is the common Netflow template.

Element ID	Name	Type	Description	Recommended Implementation	Applicable Edge Release
1	octetDeltaCount	unsigned64	The number of octets includes IP header(s) and IP payload.	Used to report on total bytes (aggregate of bytesTX and bytesRX) and BytesRX.	3.3.0
2	packetDeltaCount	unsigned64	The number of incoming packets since the previous report (if any) for this flow at the observation point.	Used to report on 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

Element ID	Name	Type	Description	Recommended Implementation	Applicable Edge Release
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 valueDistribution Method 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 RFC2863 .	This value maps to Interface option template 272 'ingressInterface' value where to map the flow to SD-WAN link interface number.	3.3.0

Element ID	Name	Type	Description	Recommended Implementation	Applicable Edge Release
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 RFC2863 .	Egress interface	3.3.0
15	ipNextHopIPv4Address	ipv4Address	The IPv4 address of the next IPv4 hop. http://www.iana.org/go/rfc2863	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

Element ID	Name	Type	Description	Recommended Implementation	Applicable Edge Release
239	biflowDirection	unsigned8	<p>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.</p> <hr/> <p>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.</p> <hr/>	See IPFIX Information Element Definitions .	3.3.0

Element ID	Name	Type	Description	Recommended Implementation	Applicable Edge Release
95	applicationId	octetArray(8)	Specifies an application ID. RFC6759. 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 timestamp of the first packet of this flow.	Implement as per description.	3.3.0
153	flowEndMilliseconds	dateTimeMilliseconds	The absolute timestamp of the last packet of this flow.	Implement as per description.	3.3.0

Element ID	Name	Type	Description	Recommended Implementation	Applicable Edge Release
136	flowEndReason	unsigned8	<p>The reason for flow termination. The range of values includes the following:</p> <ul style="list-style-type: none"> ■ 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 	Implement as per description.	3.3.0

Element ID	Name	Type	Description	Recommended Implementation	Applicable Edge Release
			<p>example, a shutdown of the metering process initiated by a network management application.</p> <ul style="list-style-type: none"> ■ 0x05: lack of resources - The flow was terminated because of lack of resources available to the metering process and/or the exporting process. 		
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 VMware SD-WAN Orchestrator segments. A segment should be visualized and reported as a separated L3 domain within the Edge.	3.3.0

Enterprise-Specific Fields (ID>32767)

VMware 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	<ul style="list-style-type: none"> ■ 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	<ul style="list-style-type: none"> ■ 0 - Unset ■ 1 - Gateway (using hosted GW svc) ■ 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	<ul style="list-style-type: none"> ■ 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	<ul style="list-style-type: none"> ■ 0 - Realtime ■ 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	<ul style="list-style-type: none"> ■ 0 - Edge2CloudViaGateway (SaaS optimized) ■ 1 - Edge2CloudDirect (SaaS not optimized) ■ 2 - Edge2EdgeViaGateway (spoke2hub2 spoke via VCG) ■ 3 - Edge2EdgeViaHub (spoke2hub2 spoke via PDC Hub) ■ 4 - Edge2EdgeDirect (Edge2Edge dynamic) ■ 5 - Edge2DataCenterDirect (Edge2PDC using underlay routing) ■ 6 - Edge2DataCenterViaGateway (Edge2PDC using NVS) ■ 7 - Edge2Backhaul (Edge2Internet using PDC Hub) ■ 8 - Edge2Proxy 	This identifies the type of 'path' the flow is taking.	3.3.0

Element ID (Enterprise Element ID)	Name	Type	Description	Recommended Implementation	Applicable Edge Release
			<ul style="list-style-type: none"> ■ 9 - Edge2OPG (PGW) ■ 10 – Routed (path using underlay routing) ■ 11 - Edge2CloudV iaSecuritySer vice (path using a CASB service to internet) 		
45009 (12241)	replicatedPacket sRxDeltaCount	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)	replicatedPacket sTxDeltaCount	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)	lostPacketsRxDel taCount	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)	retransmittedPac ketsTxDeltaCoun t	unsigned64	Count of packets retransmitted for the flow	This value provides the number of retransmitted packets due to loss (applies to transactional traffic).	3.3.0

Element ID (Enterprise Element ID)	Name	Type	Description	Recommended Implementation	Applicable Edge Release
45085 (12317)	tcpRttMs	unsigned16	Maximum RTT observed for a TCP flow	The maximum 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
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

NAT Template

Template ID: 259

Common + NAT template

Element ID	Name	Type	Description	Applicable Edge Release
225	postNATSourceIPv4 Address	ipv4Address	The definition of this information element is identical to the definition of information element <i>sourceIPv4Address</i> , 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 <i>destinationIPv4Address</i> , 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. VMware SD-WAN needs to export flow stats as two flow records or implement RFC5103 (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.

Flow Link Stats Template

The Flow Link Stats template captures the flow stats broken down by link.

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 observationpoint. The number of octets includes IP header(s) and IP payload.	3.3.0
2	packetDeltaCount	unsigned64	The number of incoming packets since the previous report (if any) for this flow at the observation point.	3.3.0
32769	octetDeltaCount_rev	unsigned64	Biflow RFC 5103. The number of outgoing bytes.	3.3.0
32770	packetDeltaCount_rev	unsigned64	Biflow RFC 5103. The number of outgoing packets.	3.3.0
14	egressInterface	unsigned32	The index of the IP interface where packets of this flow are being sent. The value matches the value of managed object as defined in [RFC2863] .	3.3.0
45008 (12240)	linkUUID	octetArray(16)	The VMware internal link ID.	3.3.0

Element ID (Enterprise Element ID)	Name	Type	Description	Applicable Edge Release
45009 (12241)	replicatedPacketsRx DeltaCount	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)	replicatedPacketsTx DeltaCount	unsigned64	Count of packets replicated for the flow.	3.3.2 (This field was part of template Id 256 in 3.3.0)
45012 (12244)	retransmittedPackets TxDeltaCount	unsigned64	Count of packets retransmitted for the flow.	3.3.2 (This field was part of template Id 256 in 3.3.0)

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 interfaceIndex 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 = $((packetsLostDeltaRxCount) / (packetsLostDeltaRxCount + packetsLostCompDeltaRxCount)) * 100$

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

Element ID	Name	Type	Description	Applicable Edge Release
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

Application Option Template

<https://tools.ietf.org/html/rfc6759>. The Application Option template is sent every 5 minutes or when changed. Only applications that have been referenced in flows are exported.

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	applicationCategory Name	string	An attribute that provides a first level categorization for each application ID.	3.3.0

Application ID Format

```

0           1           2           3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|      20      |                enterprise ID = 45346                ...|
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|...Ent.ID.contd|                                app ID                                |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

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 VMware SD-WAN PEN
- App ID is internal application ID

Interface Option Template


Interfaces in the VMware 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 VMware devices. On the overlay, there may be several tunnels between a pair of VMware 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 VMware 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.



Template ID: 272

The interface option template is sent every 5 minutes by default. The timer is configurable.

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 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	<ul style="list-style-type: none"> ■ 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

VMware Segment ID to Segment Mapping Template

The template is sent every 10 minutes and utilizes VRF as the nomenclature to define a segment.

Template ID: 273

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

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 .

Template ID: 276

The Link Option template is sent every 5 minutes.

Element ID (Enterprise Element ID)	Name	Type	Description	Applicable Edge Release
45008 (12240)	linkUUID	octetArray(16)	The VMware 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 as defined in [RFC2863] .	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

Element ID (Enterprise Element ID)	Name	Type	Description	Applicable Edge Release
8	sourceIP	unsigned32	The source IP for this link.	3.3.2
15	nextHopIP	unsigned32	The nextHop IP for this link.	3.3.2

Netflow Source Address and Segmentation

Netflow source interface's primary IP address should come from VMware SD-WAN 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 - RFC1213).

```

0..7      0..7      0..16
destination_type    reserved    destination_if_idx

```

destination_type:

- E2E
- E2DC
- CLOUD
- ANY/DIRECT

destination_if_idx:

- E2E, E2DC, CLOUD: map(next_hop_id) -> if_idx
- ANY/DIRECT: map(link_logical_id) -> if_idx

Filtering

Allow Netflow to be filtered by:

- ingressVRFID (or all segments)
- ApplicationID
- sourceIPv4Address (mask)
- destinationIPv4Address (mask)
- protocolIdentifier

IPFIX Information Element Definitions

The following table lists the IPFIX information element definitions.

38	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.	

		<pre> 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. +-----+-----+-----+-----+-----+-----+ </pre>
23 9	biflowDirection	<p>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.</p> <p>Note 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.</p> <p>This field may take the following values:</p> <pre> +-----+-----+-----+-----+-----+-----+ 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 </pre>

			Destination addresses exported in the
			Biflow Records.
	+-----+	+-----+	+-----+

Private Network Names

You can define multiple private networks and assign them to individual private WAN overlays.

Configure Private Networks

To configure private networks:

- 1 From the SD-WAN Orchestrator navigation panel, go to **Configure > Network Services**.
- 2 In the **Private Network Names** area, click the **New** button.
- 3 In **New Private Network Name** dialog box, enter a unique name in the appropriate text box.

New Private Network Name

?

✕

* Private Network Name:

Save Changes

Cancel

- 4 Click **Save Changes**.

The private network name appears in the **Private Network Name** area.

Private Network Names		New...	Delete...
Name	Used By		
<input type="checkbox"/> MPLS A	0		
<input type="checkbox"/> MPLS B	0		

Delete a Private Network Name

Only private network names that are not used by an Edge device can be deleted.

To delete a private network name not used by an Edge device:

- 1 Select the name by clicking the name's checkbox, and then click the **Delete** button.
- 2 In the **Confirm Deletion dialog box**, click **OK**.


You can select private link tags when you define a User Defined Overlay. See section titled, "*Select a Private Network Name.*"

Configure Authentication Services

Authentication Services is an optional configuration. 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.


The following figure shows an example configuration.



New Radius Service

* Service Name: VeloAcme-radius

* Server Address: 4.4.4.4

* Shared Secret: 

* Authentication Port: 1812

Accounting Port: 1813

Save Changes Cancel

Note Source interfaces are configured only at Edge level. For more information, see [Configure Authentication Settings](#).

Configure Cloud Subscriptions

Before you configure a Cloud Subscription, ensure to register the SD-WAN Orchestrator application and create Client secret in the Azure portal.

To configure a Cloud subscription in SD-WAN Orchestrator:

Prerequisites

For more information, see [Prerequisite Azure Configuration](#).

Procedure

- 1 From the navigation panel in the SD-WAN Orchestrator, go to **Configure > Network Services**.
The **Services** screen appears.
- 2 In the **Cloud Subscriptions** area, click the **New** button.
The **Configure Cloud Subscription** dialog box appears.

Configure Cloud Subscription	
* Subscription Type:	Microsoft Azure Subscription
* Active Directory Tenant ID	22eb73a3-5c68-47b6-8098-08952150a401
* Client ID	5188a0f1-8215-49d0-9085-ea3043a12721
* Client Secret
* Subscription	Pay-As-You-Go(Converted to EA)
<div>Save Changes Cancel</div>	

- 3 From the **Subscription Type** drop-down menu, select **Microsoft Azure Subscription**.
- 4 Enter the Active Directory Tenant ID, Client ID, and Client Secret corresponding to your SD-WAN Orchestrator Application Registration.
- 5 Click the **Get Subscriptions** button to retrieve the list of Azure Subscriptions for which the App Registration has been allocated an IAM role.
- 6 Click **Save Changes**.

What to do next

Configure a Non SD-WAN Destination of type Microsoft Azure Virtual Hub.

- To configure a Microsoft Azure Non SD-WAN Destination from SD-WAN Gateway, see [Configure a Microsoft Azure Non SD-WAN Destination via Gateway](#).
- To configure a Microsoft Azure Non SD-WAN Destination from SD-WAN Edge, see [Configure a Microsoft Azure Non SD-WAN Destination via Edge](#).

Configure Profiles

11

Profiles provide a composite of the configurations created in Segments and Network Services. It also adds configuration for Business Policy and Firewall rules.

Note If you are logged in using a user ID that has Customer Support privileges, you will only be able to view SD-WAN Orchestrator objects. You will not be able to create new objects or configure/update existing ones.

Complete the following tasks to configure a new profile:

- 1 [Create a Profile](#)
- 2 [Configure a Device](#)
 - a [Assign Segments in Profile](#)
 - b [Configure Authentication Settings](#)
 - c [Configure DNS Settings](#)
 - d [Configure Interface Settings](#)
- 3 [Configure Cloud VPN for Profiles](#)
- 4 [Configure Business Policy for Profiles](#)
- 5 [Configure Firewall for Profiles](#)

You can review the newly configured profile in the **Profile Overview** page. To access the page, go to **Configure > Profiles**, and then select the required profile. The **Profile Overview** page appears. In this page, you can review all the configurations, such as the profile name and description, local credentials to access the profile, the Edge models that are enabled for the profile, the network services configured for the profile, and the segments assigned to the profile.

Configuration Profiles > **Quick Start Profile** Save Changes ?

Profile Overview Device Business Policy Firewall

* Name Local Credentials ***** Modify

Description

Profile Overview

Enabled Models Edge 500, Edge 510, Edge 510-LTE, Edge 5X0, Edge 6X0, Edge 610-LTE, Edge 840, Edge 1000, Edge 2000, Edge 3X00, Edge 3X10, Virtual Edge

Services

Dynamic Multi-Path Optimization	On
Application Recognition	On
Identity	On
DHCP	On
Wireless	On
802.1x	Off

Segments

Segment	Netfl...	Cloud ...	OSPF	BGP	Multicast	Cloud S...	Auth	Business ...	Firewall
Global Segment	Off	Off	Off	Off	Off	Off	Off	22 rules	1 outbound ...
segment1	Off	Off	-	Off	Off	Off	Off	22 rules	1 outbound ...
segment2	Off	Off	-	Off	Off	Off	Off	22 rules	1 outbound ...
a	Off	Off	-	Off	Off	Off	Off	22 rules	1 outbound ...
1	Off	Off	-	Off	Off	Off	Off	22 rules	1 outbound ...
c	Off	Off	-	Off	Off	Off	Off	22 rules	1 outbound ...

Read the following topics next:

- [Create a Profile](#)
- [Modify a Profile](#)
- [Configure Local Credentials](#)

Create a Profile

After a new installation, the SD-WAN Orchestrator has predefined profiles that are segment-based.

To create a new Profile:

- 1 Go to **Configure > Profiles**, and click **New Profile**.
- 2 In the **New Profile** dialog, enter a Profile Name and Description in the appropriate textboxes.
- 3 Click **Create**.

The newly added profile is listed in the appears in the list of profiles **Configuration Profiles** page.

Modify a Profile

Enterprise Administrators can manually assign a profile to an Edge.

One scenario in which this is necessary is for Edge Staging Profiles. In this case, the Edge gets activated against the staging profile by default due to Zero Touch Provisioning. Enterprise Administrators must manually assign a final production profile to the Edge. For instructions on how to manually assign Profiles, see [Provision a New Edge](#).

Name	Network	Used By	Device	Biz. Pol.	Firewall
Edge Staging Profile <small>Use by edges that are newly assigned, but yet to be configured with a profile.</small>	Segment Based Profile	2 Edges			
Quick Start Profile <small>ACME</small>	Segment Based Profile	1 Edge			

Configure Local Credentials

Use the local credentials to gain access to the Profile through SD-WAN Orchestrator.

After a new installation of SD-WAN Orchestrator, a local credential with a default username as 'admin' and a randomly generated password is assigned to the Profiles. You can change the local credentials at your first login to SD-WAN Orchestrator.

Note Enterprise users with Read-Only and Customer Support access cannot view or change the password.

To change the local credentials at the Profile level:

- 1 From the SD-WAN Orchestrator, go to **Configure > Profiles**.
- 2 Click the Profile for which you want to change the local credentials.

Configuration Profiles - Quick Start Profile Save Changes ?

Profile Overview Device Business Policy Firewall

* Name: Quick Start Profile

Description: 3-site

Local Credentials: ***** Modify

- 3 In the **Profile Overview** tab, click **Modify**. The **Local Configuration Credentials** modal popup appears.
- 4 In the **User** field, enter the required user name.

- 5 Select the **Change Password** check box, and then enter the new password of your choice. Click the Eye icon to view the 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.

- 6 Click **Submit**.

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 [Edge Properties](#).

Manage Profiles with New Orchestrator UI

12

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 SD-WAN Orchestrator.

You can manage the Profiles using the new Orchestrator UI.

- 1 In the Enterprise portal, click the **Open New Orchestrator UI** option available at the top of the Window.
- 2 Click **Launch New Orchestrator UI** in the pop-up window.
- 3 The UI opens in a new tab displaying the monitoring and configuring options.


In the new Orchestrator UI, click the **Configure** tab.

Note The **Configure** tab is available only when the Operator has enabled the option. If the tab is not available for you, contact your Operator.

- 1 Click **Configure > Profiles**.
- 2 The **Profiles** page displays the existing Profiles with the following details:

Option	Description
Name	The name of the Profile. Click the link to the Profile to modify the Profile configurations. See Configure Profiles with new Orchestrator UI .
Network	Displays whether the Profile is Segment based or Network based.
Used By	The number of Edges associated with the Profile.
Device	Click the View link to modify the configurations of the Profile. See Configure Profiles with new Orchestrator UI .

- 3 Select one or more Profiles to perform the following activities:



- **New Profile** – Creates a new Profile. See [Create Profile with new Orchestrator UI](#).
- **Duplicate** – Creates a duplicate of the selected Profile.
- **Modify** – Allows to edit the configurations of the selected Profile. See [Configure Profiles with new Orchestrator UI](#).
- **Download** – Downloads the details of Profiles into an MS Excel file.

Click **More** to perform the following:

- **Delete** – Deletes the selected Profiles. You cannot delete the Profiles that are associated with Edges.

Read the following topics next:

- [Create Profile with new Orchestrator UI](#)
- [Configure Profiles with new Orchestrator UI](#)

Create Profile with new Orchestrator UI

After installing SD-WAN Orchestrator, a default profile is available. If required, you can create additional Profiles.

To create a Profile using the new Orchestrator UI:

Procedure

- 1 In the Enterprise portal, click the **Open New Orchestrator UI** option available at the top of the Window.
- 2 Click **Launch New Orchestrator UI** in the pop-up window. The UI opens in a new tab displaying the monitoring and configuring options.

Note The Configure tab is available only when the Operator has enabled the option. If the tab is not available for you, contact your Operator.

- 3 Click **Configure > Profiles**.

- 4 In the **Profiles** page, click **New Profile**.

New Profile

Profile name* ACME_Profile

Description

CANCEL CREATE

- 5 Enter a Name and Description for the new Profile and click **Create**.
- 6 The **Device** tab opens, which provides options to configure the Profile settings. For more information, see [Configure Profiles with new Orchestrator UI](#).

Configure Profiles with new Orchestrator UI

You can configure the Profile settings using the new Orchestrator UI.

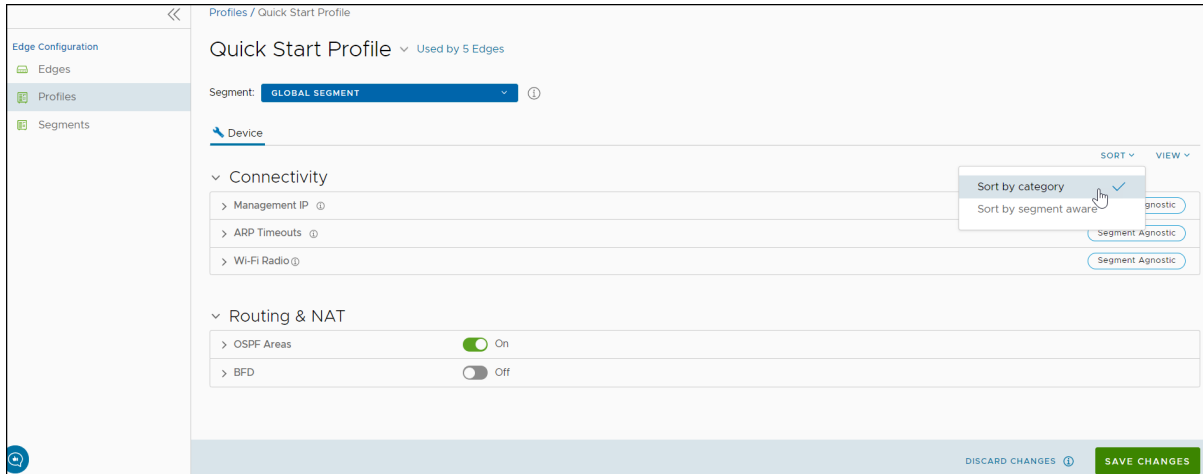
- 1 In the Enterprise portal, click the **Open New Orchestrator UI** option available at the top of the Window.
- 2 Click **Launch New Orchestrator UI** in the pop-up window.
- 3 The UI opens in a new tab displaying the monitoring and configuring options.

In the new Orchestrator UI, click the **Configure** tab.

Note The **Configure** tab is available only when the Operator has enabled the option. If the tab is not available for you, contact your Operator.

To configure a specific Profile:

- 1 Click **Configure > Profiles**.
- 2 The **Profiles** page displays the existing Profiles.
- 3 Click the link to a Profile or click the **View** link in the **Device** column of the Profile. You can also select a Profile and click **Modify** to configure the Profile.
- 4 The configuration options for the selected Profile are displayed in the **Device** tab.



- 5 Click **View** to expand or collapse the view of available settings.
- 6 You can view the configuration settings sorted by category or segmentation. 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.

The following settings are available when you choose to sort by category:

Connectivity

- **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. For more information, see [Configure the Management IP Address](#).
- **ARP Timeouts** – By default the ARP Timeout values are configured. If required, select the **Override default ARP Timeouts** checkbox, to modify the default values. For more information, see [Configure Layer 2 Settings for Profiles](#).
- **Wi-Fi Radio** – Turn on or turn off Wi-Fi Radio and configure the band of radio frequencies. For more information, see [Configure Wi-Fi Radio Settings](#).

Routing & NAT

- **OSPF Areas** – Configure OSPF areas for the selected Profile. For more information, see [Enable OSPF](#).
- **BFD** – Configure BFD settings for the selected Profile. For more information, see [BFD Settings](#).

Note In the new Orchestrator UI, you can configure only IPv4 addresses for BFD. To configure IPv6 addresses for BFD, see [Configure BFD](#).

Configure a Profile Device

13

This section describes 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 SD-WAN Orchestrator objects. You will not be able to create new objects or configure/update existing ones.

VMware provides device settings using the **Device** tab (**Configure > Profiles > Device Tab**) in a profile. The **Device Settings** tab is used to assign segments, create VLANs, configure interfaces, configure DNS settings, Configure Authentication Settings. For more information about Segmentation, see [Chapter 8 Configure Segments](#).

Read the following topics next:


- [Configure a Device](#)

Configure a Device

Device configuration allows you to assign segments to a Profile and configure Interfaces to be associated with a Profile.

For segment aware profiles, there are two sections on the UI:

Configuration Type	Description
Segment-aware configurations	Configure Segments area of the Device tab screen. Customers can choose the segment from the drop-down menu, select the segment, and then the configuration for that segment will display in the Configure Segments area.
Common configurations	The lower part of the Device tab screen. Features and configurations that apply to multiple segments, which include VLAN configs, Device Settings, Wi-Fi and Multi-source QoS.



You can perform the following steps for Device Configuration:

Segment-aware Configurations

- Authentication Settings
- DNS Settings
- Netflow Settings
- Syslog Settings
- Cloud VPN
- OSPF Areas
- BGP Settings
- Multicast Settings
- Cloud Security Service

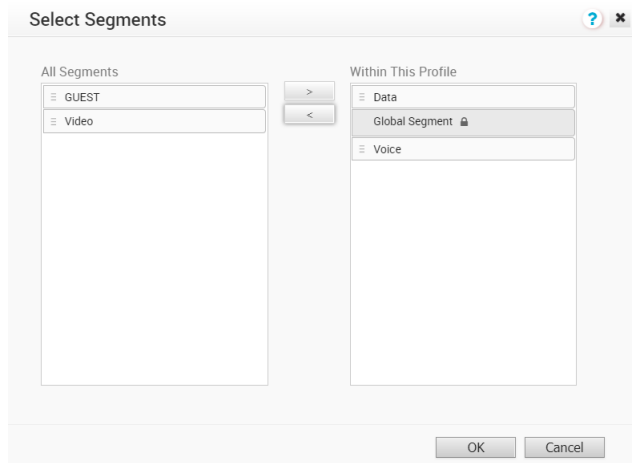
Common Configurations:

- VLAN
- Device Settings
- Wi-Fi Radio Settings
- Multi-Source QoS
- SNMP Settings
- NTP Servers
- Visibility Mode

Assign Segments in Profile

After creating a Profile, you can select Profile Segments by clicking the **Change** button in the image **Configure Segments** window.

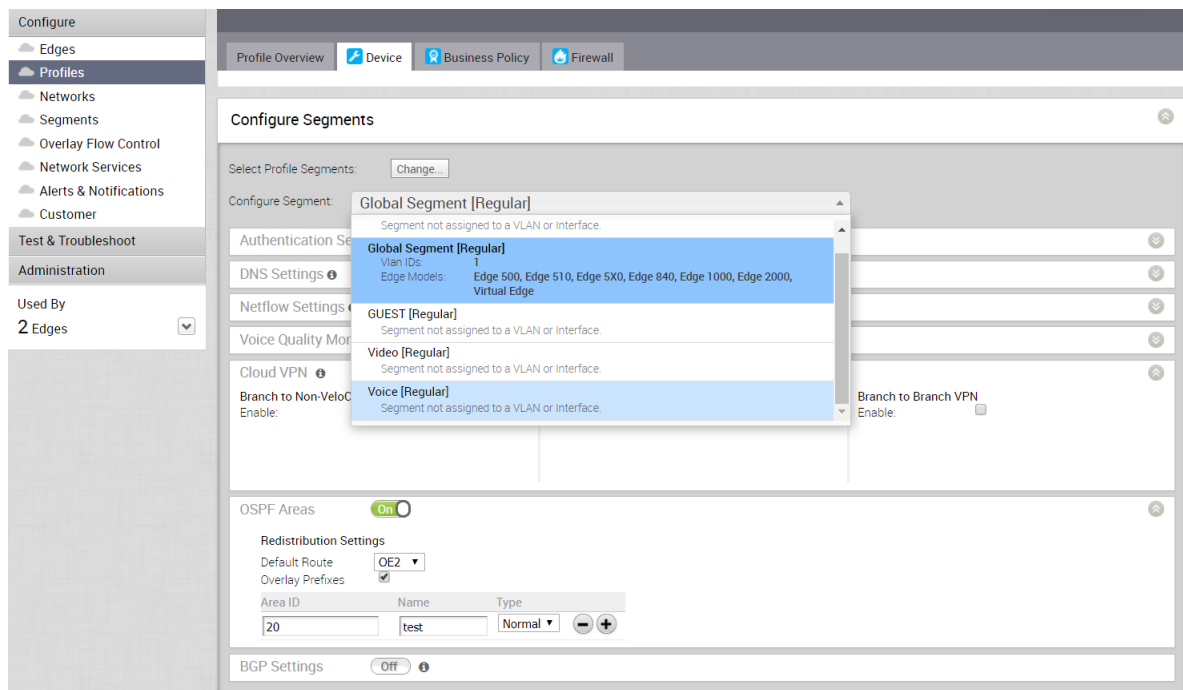
Clicking the **Change** button opens the **Select Segments** dialog box.



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 on the left side of the dialog under **All Segments**.

After you have selected a Segment, you can configure your Segment through the **Configure Segment** drop-down menu. All Segments available for configuration are listed in the **Configure 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 **Configure Segment** drop-down menu, depending upon the Segment's options, the settings associated that Segment display in the **Configure Segments** area.



Configure Authentication Settings

The **Device Authentication Settings** allows you to select a Radius server used to authenticate a user.

To configure the Authentication Settings for a Profile:

- 1 Log in to VMware SD-WAN Orchestrator, and then go to **Configure > Profiles**.
- 2 Either click the Device icon next to the profile for which you want to configure the Authentication Settings, or click the link to the Profile, and then go to the **Device** tab.
- 3 In the **Authentication Settings** area, from the **RADIUS Server** drop-down list, select the Radius server that you want to use for authentication.

Note Configure the Radius server in the **Authentication Services** area in the **Network Services** page. Alternatively, you can configure a new authentication service by selecting the **New Authentication Service...** option from the **RADIUS Server** drop-down list. For instructions about how to configure Authentication Services, see [Configure Authentication Services](#).

At the Edge-level, you can choose to override the Authentication Settings configured for the Profile.

- 1 Log in to VMware SD-WAN Orchestrator, and then go to **Configure > Edges**.
- 2 Either click the Device icon next to the Edge for which you want to override the Authentication Settings, or click the link to the Edge, and then go to the **Device** tab.
- 3 In the **Authentication Settings** area, select the **Enable Edge Override** check box, and then expand the area.
- 4 From the **RADIUS Server** drop-down menu, select the Radius server that you want to use for authentication.
- 5 From the **Source Interface** drop-down list, 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.
-

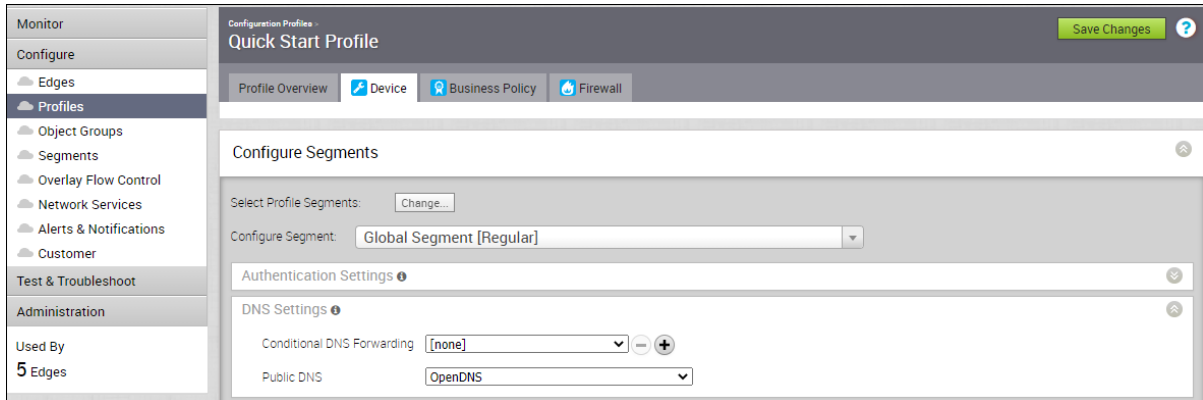
- 6 Click **Save Changes**.

Configure DNS Settings

The **DNS Settings** can be used to configure conditional DNS forwarding through a private DNS service and to specify a public DNS service to be used for querying purpose.

To configure the DNS settings for a Profile:

- 1 Log in to VMware SD-WAN Orchestrator, and then go to **Configure > Profiles**.
- 2 Either click the Device icon next to the profile for which you want to configure the DNS settings, or click the link to the Profile, and then go to the **Device** tab.
- 3 Configure the following settings in the **DNS Settings** area:



- **Conditional DNS Forwarding**—Select a private DNS service from the drop-down list to forward the DNS requests related to the domain name. You can also choose the **New Private DNS Service** to create a new private DNS service.
- **Public DNS**—Select a public DNS service from the drop-down list to be used for querying the domain names. You can also choose the **New DNS Service** to create a new public DNS service.

Note The public DNS service is enabled on a VLAN or a routed interface only if the DHCP service is enabled on that VLAN or routed interface. For instructions, see [Configure DHCP Server on Routed Interfaces](#).

For more information on creating new DNS service, see [Configure DNS Services](#).

- 4 Click **Save Changes**.

Note The global segment configuration for DNS applies to all the customer-created segments.

At the Edge-level, you can choose to override the DNS Settings configured for the Profile.

- 1 Log in to VMware SD-WAN Orchestrator, and then go to **Configure > Edges**.
- 2 Either click the Device icon next to the Edge for which you want to override the DNS settings, or click the link to the Edge, and then go to the **Device** tab.
- 3 In the **DNS Settings** area, select the **Enable Edge Override** check box, and then expand the area.
- 4 From the **Conditional DNS Forwarding** and **Public DNS** drop-down list, select the required private or public DNS service if you choose to override the DNS service at the Edge-level.

- From the **Source Interface** drop-down list, 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 will have the IP address of the selected source interface, whereas the packets can be sent through any interface based on the destination route.

- Click **Save Changes**.

Configure Netflow Settings for Profiles

As an enterprise Administrator, you can configure Netflow settings at the Profile level.

Procedure

- From the SD-WAN Orchestrator, go to **Configure > Profiles**.

The **Configuration Profiles** page appears.

- Select a profile you want to configure Netflow settings and click the icon under the **Device** column.

The Device Setting page for the selected profile appears.

Netflow Settings ⓘ

Netflow Enabled: ☒

Version: ⓘ v10

Collector: ⓘ C_Global_10.2.1.25 Filter: ⓘ Allow_All ⓘ

Allow All ⓘ

Intervals:

* Flow Stats:	61
* FlowLink Stats:	61
* Segment Table:	100
* Application Table:	100
* Interface Table:	100
* Link Table:	100
* Tunnel Stats:	60

- From the **Configure Segment** drop-down menu, select a profile segment to configure Netflow settings.

4 Go to the **Netflow Settings** area and configure the following details.

- a Select the **Netflow Enabled** checkbox.

SD-WAN 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 SD-WAN Edge, or click **New Collector** to configure a new Netflow collector.

For more 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 clicking the + button. When the number of configured collectors reaches the maximum allowable limit, the + button will be deactivated.

- c From the **Filter** drop-down menu, select an existing Netflow filter for the traffic flows from SD-WAN Edge, or click **New Filter** to configure a new Netflow filter.

For more information about how to add a new filter, see [Configure Netflow Settings](#).

Note You can configure a maximum of 16 filters per collector by clicking the + button. However, the 'Allow All' filtering rule is added implicitly at the end of the defined filter list, per collector.

- d Enable the **Allow All** checkbox 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 is 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 is exported every 60 seconds. The allowable export interval range is from 60 seconds to 300 seconds.
 - **VRF Table** - Export interval for VRF 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 more information on various Netflow templates, see IPFIX Templates.

- 5 Click **Save Changes**.

Configure Syslog Settings for Profiles

In an Enterprise network, SD-WAN Orchestrator supports collection of SD-WAN Orchestrator bound events and firewall logs originating from enterprise SD-WAN Edge to one or more centralized remote Syslog collectors (Servers), in the native Syslog format. For the Syslog collector to receive SD-WAN 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 SD-WAN Orchestrator by performing the steps on this procedure.

Prerequisites

- Ensure that Cloud Virtual Private Network (branch-to-branch VPN settings) is configured for the SD-WAN Edge (from where the SD-WAN Orchestrator bound events are originating) to establish a path between the SD-WAN Edge and the Syslog collectors. For more information, see [Configure Cloud VPN for Profiles](#).

Procedure

- 1 From the SD-WAN Orchestrator, go to **Configure > Profiles**.

The **Configuration Profiles** page appears.

- 2 Select a profile you want to configure Syslog settings and click the icon under the **Device** column.

The Device Settings page for the selected profile appears.

- 3 From the **Configure Segment** drop-down menu, select a profile segment to configure syslog settings. By default, **Global Segment [Regular]** is selected.

- 4 Go to the **Syslog Settings** area and configure the following details.

- a From the **Facility Code** 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 SD-WAN Edge. The allowed values are from **local0** through **local7**.

Note The **Facility Code** field is configurable only for the **Global Segment**, even if the Syslog settings is enabled or not for the profile. The other segments will inherit the facility code value from the Global segment.

- b Select the **Syslog Enabled** checkbox.
- c In the **IP** text box, enter the destination IP address of the Syslog collector.
- d From the **Protocol** drop-down menu, select either **TCP** or **UDP** as the Syslog protocol.
- e In the **Port** text box, enter the port number of the Syslog collector. The default value is 514.
- f 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.
- g From the **Roles** drop-down menu, select one of the following:
 - **EDGE EVENT**
 - **FIREWALL EVENT**
 - **EDGE AND FIREWALL EVENT**

- h From the **Syslog Level** drop-down menu, select the Syslog severity level that need to be configured. For example, If **CRITICAL** is configured, the SD-WAN 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**
- i Optionally, in the **Tag** textbox, 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.
 - j When configuring a Syslog collector with **FIREWALL EVENT** or **EDGE AND FIREWALL EVENT** role, select the **All Segments** checkbox if want the Syslog collector to receive firewall logs from all the segments. If the checkbox 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 Click the **+** button to add another Syslog collector or else click **Save Changes**. The remote syslog collector is configured in SD-WAN 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.

Syslog Settings

Facility : local0

Syslog Enabled: ☒

* IP	* Protocol	* Port	* Source Interface	* Roles	* Syslog Level	Tag	All Segments
10.1.1.25	TCP	514	Auto	FIREWALL EVENT	INFO	VMware.SDWAN.FW	<input checked="" type="checkbox"/>
10.1.2.25	TCP	514	Auto	EDGE EVENT	ERROR	VMware.SDWAN.Edge	<input checked="" type="checkbox"/>

Firewall logs are forwarded at INFO level by default

You are at the maximum limit of 2 collectors per segment

Note 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 SD-WAN Orchestrator auto-generated local events to be received at the Syslog collector, you must configure Syslog at the SD-WAN 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](#).

What to do next

SD-WAN Orchestrator allows you to enable Syslog Forwarding feature at the profile and the Edge level. On the **Firewall** page of the Profile configuration, enable the **Syslog Forwarding** button if you want to forward firewall logs originating from enterprise SD-WAN 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 more information about Firewall settings at the profile level, see [Configure Firewall for Profiles](#).

Syslog Message Format for Firewall Logs

Describes the Syslog message format for Firewall logs with an example.

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 Open FW_POLICY_NAME=test
SID=0000012278 SEGMENT_NAME=Global Segment IN="IFNAME" PROTO=ICMP SRC=x.x.x.x DST=x.x.x.x
DEST_NAME=Internet-via-gateway-3
```

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 Open FW_POLICY_NAME=test SID=0000012278
SEGMENT_NAME=Global Segment IN="IFNAME" PROTO=ICMP SRC=x.x.x.x DST=x.x.x.x
DEST_NAME=Internet-via-gateway-3

VMware 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 13-1. Firewall Log Message Fields

Field	Description
FW_POLICY_NAME	The name of the firewall policy applied to the session.
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.
SEGMENT_NAME	The name of the segment to which the session belongs to.
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.

Table 13-1. Firewall Log Message Fields (continued)

Field	Description
SPT	The source port number of the session. This field is applicable only if the underlaying transport is UDP/TCP.
DPT	The destination port number of the session. This field is applicable only if the underlaying transport is UDP/TCP.
DEST_NAME	<p>The name of the remote-end device of the session. The possible values are:</p> <ul style="list-style-type: none"> ■ CSS-Backhaul - For traffic which is destined to Cloud Security Service from edge. ■ Internet-via-<i><egress-iface-name></i> - For Cloud traffic going directly from edge using business policy. ■ Internet-BH-via-<i><backhaul hub name></i> - For Cloud-bound traffic going to Internet through Backhaul hub using business policy. ■ <i><Remote edge name></i>-via-Hub - For VPN traffic flowing through Hub. ■ <i><Remote edge name></i>-via-DE2E - For VPN traffic flowing between the edges through direct VCMP tunnel. ■ <i><Remote edge name></i>-via-Gateway - For VPN traffic flowing through Cloud gateway. ■ NVS-via-<i><gateway name></i> - For Non SD-WAN Destination traffic flowing through Cloud gateway. ■ Internet-via-<i><gateway name></i> - For Internet traffic flowing through Cloud gateway.
NAT_SRC	The source IP address used for source natting 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.

Table 13-1. Firewall Log Message Fields (continued)

Field	Description
DURATION_SECS	The duration for which the session has been active. This field is available only for Close log messages.
REASON	<p>The reason for closure or denial of the session. The possible values are:</p> <ul style="list-style-type: none"> ■ State Violation ■ Reset ■ Purged ■ Aged-out ■ Fin-Received ■ RST-Received ■ Error <p>This field is available for Close and Deny log messages.</p>

Configure Cloud VPN for Profiles

At the profile level, SD-WAN Orchestrator allows you to configure Cloud Virtual Private Network (VPN). To initiate and respond to VPN connection requests, you must enable Cloud VPN. You can configure the Cloud VPN from the **Configure > Profiles > Device** page.

Cloud VPN On

Branch to Non SD-WAN Destination via Gateway

Enable ☒

NVS Check Point Site01 ▼ ⊖ ⊕

Branch to Hubs

Enable ☒

Select Hubs...

Hubs	E2E	Backhaul	Order
b1-edge1	x	<input checked="" type="checkbox"/>	1

Conditional BackHaul Enabled ☒

Branch to Branch VPN

Enable ☒

Isolate Profile ☐

Use Cloud Gateways ☒

Use Hubs for VPN ☐

Dynamic Branch To Branch VPN

Enabled ☒

To All Edges ☒

To Edges Within Profile ☐

Branch to Non SD-WAN Destination via Edge

Enable ☒

Service			
Action	Name	Automation for all public WAN Links	Enable Service
⊕ ⊖	IKEv1 Site01 ▼	N/A	<input checked="" type="checkbox"/>

On enabling Cloud VPN for a profile, you can configure the following Cloud VPN types:

- [Configure a Tunnel Between a Branch and a Non SD-WAN Destinations via Gateway](#)
- [Configure a Tunnel Between a Branch and a SD-WAN Hubs VPN](#)
- [Configure a Tunnel Between a Branch and a Branch VPN](#)

- 1 Non SD-WAN Destination
- 2 Branch to SD-WAN 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 VMware Gateway and the data center firewall (any VPN router) provides connectivity between branches (with SD-WAN 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:



1

Primary tunnel

2

Redundant tunnel

3

Secondary VPN Gateway

VMware supports the following Non SD-WAN Destination configurations through SD-WAN 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 VMware 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 SD-WAN 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

VMware Cloud VPN is simple to set up (global networks of SD-WAN 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.

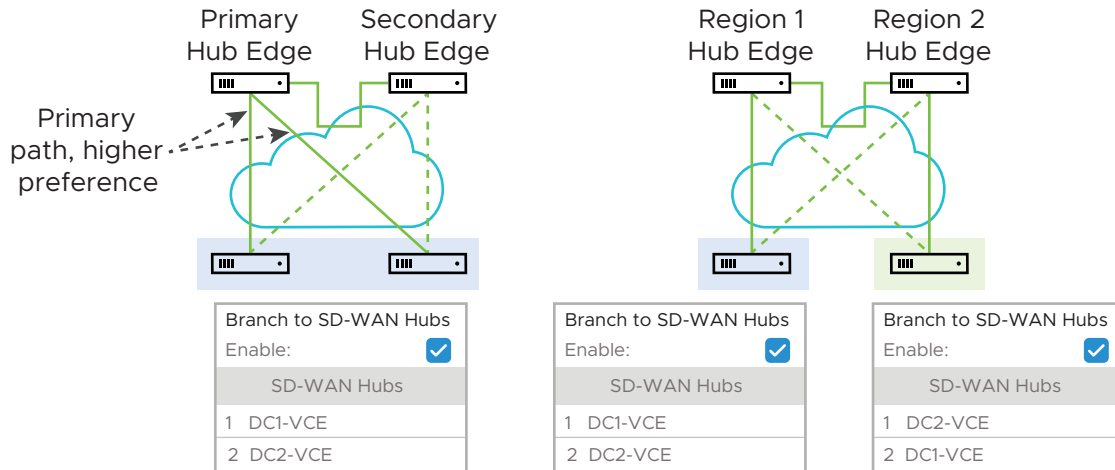
Configuring using Zscaler provides the following benefits:

- Performance: Direct to Zscaler (Zscaler via Gateway)
- Managing proxy is complex: Allows simple click policy aware Zscaler

Branch to SD-WAN Hub

The SD-WAN Hub is an Edge deployed in Data Centers for branches to access Data Center resources. You must set up your SD-WAN Hub in the SD-WAN Orchestrator. The SD-WAN Orchestrator notifies all the SD-WAN Edges about the Hubs, and the SD-WAN Edges build secure overlay multi-path tunnel to the Hubs.

The following figure shows how both Active-Standby and Active-Active are supported.



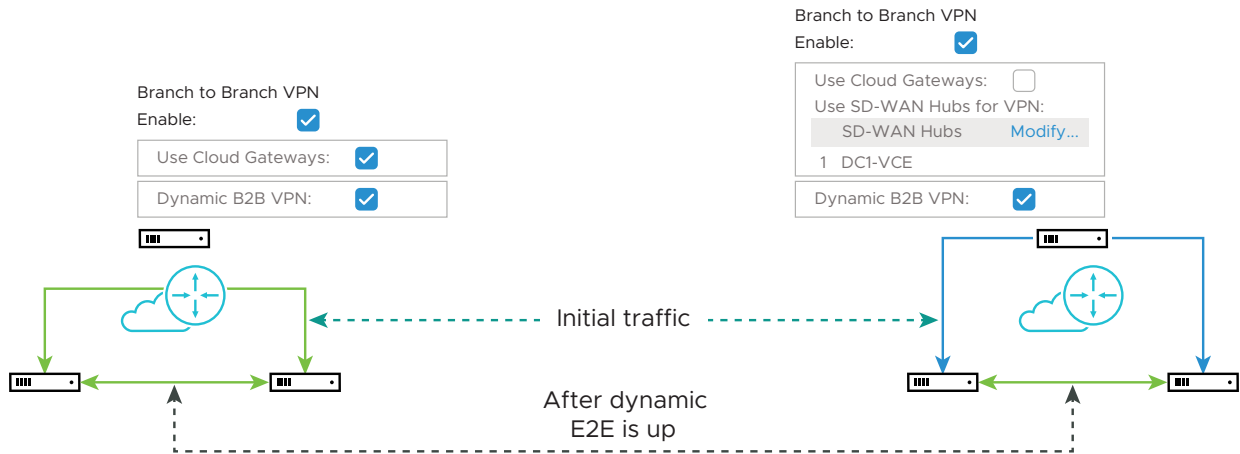
Branch to Branch VPN

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
- SD-WAN Hubs for VPN

The following figure shows Branch to Branch traffic flows for both Cloud Gateway and a SD-WAN Hub.



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 SD-WAN 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 VMware supports only Route-based Non SD-WAN Destination configurations through Edge.

For more information, see [Configure a Non SD-WAN Destinations via Edge](#).

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 SD-WAN Gateway by enabling **Branch to Non SD-WAN Destinations via Gateway** under **Cloud VPN**.

Procedure

- 1 From the SD-WAN Orchestrator, go to **Configure > Profiles**.
The **Configuration Profiles** page appears.
- 2 Select a profile you want to configure Cloud VPN and click the icon under the **Device** column.
The **Device Settings** page for the selected profile appears.

- 3 Go to **Cloud VPN** 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 SD-WAN Gateway, select the **Enable** check box under **Branch to Non SD-WAN Destinations via Gateway**.
- 5 From the drop-down menu, select a Non SD-WAN Destination to establish VPN connection. Click the **+** (plus) button to add additional Non SD-WAN Destinations.
- 6 You can also create VPN connections by selecting the **New Non SD-WAN Destinations via Gateway** option from the drop-down menu. The **New Non SD-WAN Destinations via Gateway** dialog appears.
 - a In the **Name** textbox, enter the name for the Non SD-WAN Destination.
 - b From the **Type** drop-down menu, select a Non SD-WAN Destination.
 - c In the **Primary VPN Gateway** textbox, enter the IP address that you want to configure as the primary VPN gateway for the selected Non SD-WAN Destination.
 - d Click **Next**. A new Non SD-WAN Destination will be created and gets added to the Non SD-WAN Destination drop-down menu.

For more information about configuring a Non SD-WAN Destination Network Service through Gateway, see [Configure Non SD-WAN Destinations via Gateway](#).
- 7 Click **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.

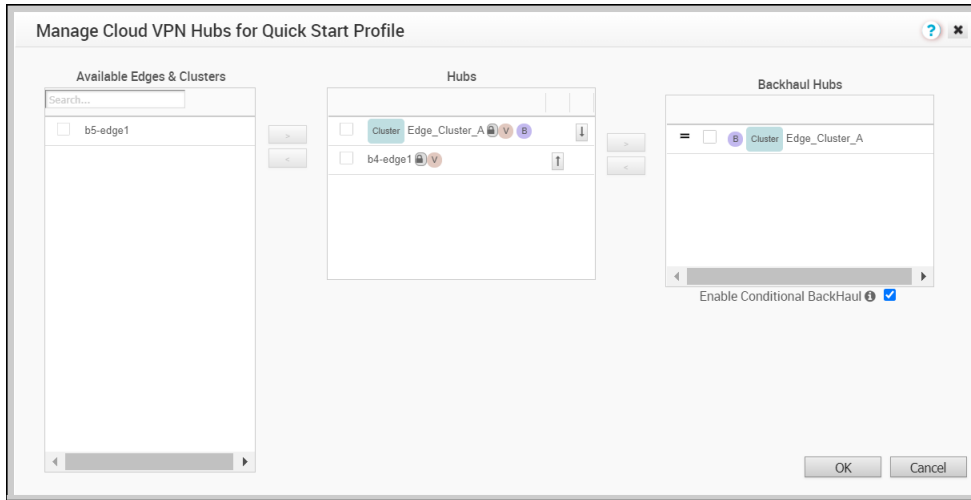
Configure a Tunnel Between a Branch and a SD-WAN Hubs VPN

Configure Branch to SD-WAN Hubs VPN to establish VPN connection between branch and hubs.

Procedure

- 1 From the SD-WAN Orchestrator, go to **Configure > Profiles**.
The **Configuration Profiles** page appears.
- 2 Select a profile you want to configure Cloud VPN and click the icon under the **Device** column.
The **Device Settings** page for the selected profile appears.
- 3 Go to **Cloud VPN** area and enable Cloud VPN by turning the toggle button to **On**.
- 4 To configure Branch to SD-WAN Hubs, under **Branch to Hubs**, select the **Enable** checkbox.

- 5 Click the **Select Hubs** link. The **Manage Cloud VPN Hubs** page for the selected profile appears.



- 6 From **Available Edges & Clusters**, you can select and configure the Edges to act as SD-WAN 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 To enable Conditional Backhaul, select the **Enable Conditional BackHaul** checkbox.

With Conditional Backhaul (CBH) enabled, the Edge will be able to failover Internet-bound traffic (Direct Internet traffic, Internet via SD-WAN Gateway and Cloud Security Traffic via IPsec) to MPLS links whenever there is no Public Internet links available. When Conditional Backhaul is enabled, 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 more information, see [Conditional Backhaul](#).

- 8 Click **Save Changes**.

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 VMware SD-WAN Edge, tunnels to VMware SD-WAN Gateway, Cloud Security Service (CSS), and Direct breakout to Internet are not established. In this scenario, the Conditional Backhaul feature, if enabled, will make use of the connectivity through Private links to designated Backhaul Hubs, giving the SD-WAN Edge the ability to failover Internet-bound traffic over Private overlays to the Hub and provide reachability to Internet destinations.

Whenever Public Internet link fails and Conditional Backhaul is enabled, the Edge can failover the following Internet-bound traffic types:

- 1 Direct to Internet
- 2 Internet via SD-WAN 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 SD-WAN Gateway as per the Business Policies configured.




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



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.



Use case 2: Cloud Security Service (CSS) Link Failure


Whenever there is a CSS (Zscaler) link failure on a SD-WAN 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 enabled, will allow the business policy to perform conditional backhaul and route the traffic to the Hub.

The Policy-based Conditional Backhaul provides the SD-WAN 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 backhailed, 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 backhailed to the Hub, irrespective of the status of the public link.



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.



Behavioral Characteristics of Conditional Backhaul

- When Conditional Backhaul is enabled, 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 enable Cloud VPN and then establish VPN connection between Branch and SD-WAN Hubs by performing the following steps:

- 1 From the SD-WAN Orchestrator, go to **Configure > Profiles**. The **Configuration Profiles** page appears.
- 2 Select a profile you want to configure Cloud VPN and click the icon under the Device column. The Device Settings page for the selected profile appears.
- 3 From the **Configure 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 enabled at each Segment where it is intended to work.

- 4 Go to **Cloud VPN** area and enable Cloud VPN by turning the toggle button to **On**.
- 5 To configure Branch to SD-WAN Hubs, under **Branch to Hubs**, select the **Enable** checkbox.
- 6 Click the **Select Hubs** link. The **Manage Cloud VPN Hubs** page for the selected profile appears.



From **Hubs** area, select the Hubs to act as Backhaul Hubs and move them to **Backhaul Hubs** area by using the > arrow.

- 7 To enable Conditional Backhaul, select the **Enable Conditional BackHaul** checkbox.

With Conditional Backhaul enabled, the SD-WAN Edge will be able to failover:

- Internet-bound traffic (Direct Internet traffic, Internet via SD-WAN 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 SD-WAN Edge, while the Public Internet link is still up.

Conditional Backhaul when enabled 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** checkbox in the **Action** area of the **Configure Rule** screen for the selected business policy. For more information, see [Configure Network Service for Business Policy 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 enabled 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 more information about SD-WAN reachability, see [SD-WAN Service Reachability via MPLS](#).
- When there are multiple candidate hubs, Conditional Backhaul will use the first hub in the list unless the Hub has lost connectivity to Gateway.

8 Click **Save Changes**.

Troubleshooting Conditional Backhaul

Consider a user with the following two Business Policy rules created at the Branch level.

Business Policy		Match			Action			
		Source	Destination	Application	Network Service	Link	Priority	Service Class
Edge Overrides	<input type="checkbox"/> Rule							
	<input type="checkbox"/> 1 TEST_MULTIPATH	IP: 10.0.5.25	Internet IP: 8.8.4.4	Any	Multi-Path	auto	Normal	Transactional
	<input type="checkbox"/> 2 TEST_DIRECT	IP: 10.0.5.25	Internet IP: 1.1.1.1	Any	Direct	auto	Normal	Transactional

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.

List Active Flows

List active flows in the system. Use source and destination IP address filters to view the exact flows you want to see. This output is limited to a maximum of 1000 flows.

Run

Segment: all

Max Flows: 100

Source IP/Port: 10.0.5.25

Destination IP/Port:

Test Duration: 5.002 seconds

Src IP	Dst IP	Segment	Protocol	Src Port	Dst Port	Application	Link Policy	Route	Business Policy
10.0.5.25	8.8.4.4	Global Segment	ICMP	N/A	N/A	icmp	Loadbalance	Cloud via Gateway	TEST_MULTIPATH
10.0.5.25	1.1.1.1	Global Segment	ICMP	N/A	N/A	icmp	Loadbalance	Direct to Cloud	TEST_DIRECT

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.

List Active Flows

List active flows in the system. Use source and destination IP address filters to view the exact flows you want to see. This output is limited to a maximum of 1000 flows.

Run

Segment: all

Max Flows: 100

Source IP/Port: 10.0.5.25

Destination IP/Port:

Test Duration: 5.008 seconds

Src IP	Dst IP	Segment	Protocol	Src Port	Dst Port	Application	Link Policy	Route	Business Policy
10.0.5.25	8.8.4.4	Global Segment	ICMP	N/A	N/A	icmp	Loadbalance	Internet Backhaul	TEST_MULTIPATH
10.0.5.25	1.1.1.1	Global Segment	ICMP	N/A	N/A	icmp	Loadbalance	Internet Backhaul	TEST_DIRECT

Note that 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 of CBH becomes operational.

List Active Flows

List active flows in the system. Use source and destination IP address filters to view the exact flows you want to see. This output is limited to a maximum of 1000 flows.

Run

Segment:

Max Flows:

Source IP/Port:

Destination IP/Port:

Test Duration: 5.002 seconds

Src IP	Dst IP	Segment	Protocol	Src Port	Dst Port	Application	Link Policy	Route	Business Policy
10.0.5.25	8.8.4.4	Global Segment	ICMP	N/A	N/A	icmp	Loadbalance	Internet Backhaul	User Default
10.0.5.25	1.1.1.1	Global Segment	ICMP	N/A	N/A	icmp	Loadbalance	Internet Backhaul	User Default

Configure a Tunnel Between a Branch and a Branch VPN

Configure Branch to Branch VPN to establish a VPN connection between branches.

Procedure

- 1 In the Enterprise portal, click **Configure > Profiles**.
The **Configuration Profiles** page appears.
- 2 Select a profile you want to configure Cloud VPN and click the icon under the **Device** column.
The **Device Settings** page for the selected profile appears.
- 3 Go to **Cloud VPN** area and activate Cloud VPN by turning the toggle button to **On**.
- 4 To configure a Branch to Branch VPN, under **Branch to Branch VPN**, select the **Enable** check box.

Branch to Branch VPN supports two configurations for establishing a VPN connection between branches:

Configuration	Description
Using SD-WAN Gateway	In this option, Edges establish VPN tunnel with the closest gateway and connections between Edges go through this gateway. The SD-WAN Gateway may have traffic from other customers.
Using SD-WAN Hub	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.

- 5 To activate profile isolation, select the **Isolate Profile** check box.

If profile isolation is activated, then the edges within the profile will not learn routes from other edges outside the profile via the SD-WAN Overlay.

You can activate **Dynamic Branch To Branch VPN** to all Edges or to Edges within a Profile. On selecting the **Enabled** check box, by default the dynamic branch to branch VPN is configured for all edges. To configure dynamic Branch to Branch VPN by profile, make sure the **Isolate Profile** check box is unselected.

Note When Profile Isolation is activated, Dynamic Branch To Branch VPN can only be activated to Edges within Profile.

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.

6 Click **Save Changes**.

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

After configuring a Non SD-WAN Destination via Edge in SD-WAN Orchestrator, you have to associate the Non SD-WAN Destination to the desired Profile in order to establish the tunnels between SD-WAN 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.

Procedure

- 1 From the SD-WAN Orchestrator, go to **Configure > Profiles**.
The **Configuration Profiles** page appears.
- 2 Select a profile you want to configure Cloud VPN and click the icon under the **Device** column.
The **Device Settings** page for the selected profile appears.
- 3 Go to **Cloud VPN** area and activate Cloud VPN by turning the toggle button to **On**.
- 4 To establish a VPN connection directly from a SD-WAN Edge to a Non SD-WAN Destination (VPN gateway of Cloud provider such as Azure, AWS), select the **Enable** check box under **Branch to Non SD-WAN Destinations via Edge**.
- 5 From the list of configured Services, select a Non SD-WAN Destination to establish VPN connection. Click the + (plus) 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 more information about configuring a Non SD-WAN Destination Network Service through Edge, see [Configure a Non SD-WAN Destinations via Edge](#).


- 6 To deactivate a particular service, uncheck the respective **Enable Service** check box.
- 7 Click **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.

Configure Multicast Settings

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.



Multicast support includes:

- Multicast support on both overlay and underlay
- Protocol-Independent Multicast - Sparse Mode (PIM-SM) on SD-WAN Edge
- Internet Group Management Protocol (IGMP) version 2 on SD-WAN 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.

Configure Multicast Globally

To configure Multicast globally:

- 1 Log in to VMware SD-WAN Orchestrator, and then go to **Configure > Profiles**.
- 2 Either click the Device icon next to the Profile for which you want to configure Multicast Settings, or click the link to the Profile, and then go to the **Device** tab.
- 3 Scroll down to the **Multicast Settings** area.
- 4 If the **Multicast Settings** button is in the **Off** position, click the **Off** button to turn **On** Multicast Settings.

The RP Selection is set to **Static** by default.

Multicast Settings **On**

RP Selection: **Static**

	RP Address	Multicast Group	
1.	10.1.1.1	230.0.0.1/32 231.0.0.0/8	Clone
2.	10.2.2.2	240.0.0.1/32 231.0.0.0/8	Clone

Enable PIM on Overlay ☒

Source IP Address: 172.16.3.3

Advanced Settings

PIM Timers

Join Prune Send Interval: 30

Keep Alive Timer: 60

- In the appropriate textboxes for the RP Selection, type in the RP Address and Multicast Group. (See the table below for a description of **RP Address** and **Multicast Group**.)
- If applicable, select the **Enable PIM on Overlay** check box and enter the IP Source Address.
- Set **Advanced Settings**, if necessary. Refer to the table that follows for a description of each setting. In the appropriate text boxes, enter PIM Timers for **Join Prune Send Interval** (default 60 seconds) and **Keep Alive Timer** (default 60 seconds).

Multicast Settings

The following table describes Multicast settings.

Multicast Setting	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 SD-WAN Edge and hub SD-WAN 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	
Join Prune Send Interval	The Join Prune Interval Timer. Default value is 60 seconds.
Keep Alive Timer	PIM keep alive timer. Default value is 60 seconds.

Configure Multicast Settings at the Edge-Level

To override the Multicast settings for an Edge:

- Log in to VMware SD-WAN Orchestrator, and then go to **Configure > Edges**.

- 2 Either click the Device icon next to the Edge for which you want to override the Multicast Settings, or click the link to the Edge, and then go to the **Device** tab.
- 3 Scroll down to the **Multicast Settings** area, and then select the **Enable Edge Override** check box.
- 4 If the **Multicast Settings** button is in the **Off** position, click the **Off** button to turn **On** Multicast Settings, and then configure the required settings. For details, see the [Multicast Settings](#) table above.

To configure the multicast settings at the Interface level, see: [Configure Interface Settings](#).

To monitor the multicast information, see

- [Monitor Routing](#)
- [Monitor Multicast Groups](#)

Configure VLAN for Profiles

As an Enterprise Administrator, you can configure a VLAN for a profile.

- 1 In the Enterprise portal, click **Configure > Profiles**.
- 2 Select a Profile to configure a VLAN and either click the Device icon or click the Profile and click the **Device** tab. In the **Device** page, scroll down to the **Configure VLAN** section.

Configure VLAN + Add VLAN									
Action	VLAN	Network	IP Address	DHCP	Segment	IGMP	PIM	VNF Insertion	
Edit Del	1 - Corporate			Enabled (242)	Global Segment			✖	
Edit Del	100 - VLAN-100			Enabled (242)	segment1			✖	
Edit Del	101 - VLAN-101			Enabled (242)	segment2			✖	

- 3 Click **Add VLAN**.

VLAN

Segment

Global Segment

VLAN Name

Corporate

VLAN Id

1

Assign Overlapping Subnets

☒

Edge LAN IP Address

10.0.1.1

Cidr Prefix

24

Network

10.0.1.0

Advertise

☒

ICMP Echo Response

☒

VNF Insertion

✖ VNF Insertion requires that the selected segment have a Service VLAN

Multicast

Multicast is not enabled for the selected segment

Fixed IPs

Applicable at the edge level.

LAN Interfaces

Applicable at the edge level.

SSID

Applicable at the edge level.

DHCP

Type

Enabled **Relay** Not Enabled

Relay Agent IP(s)

☒ Source from Secondary IP(s)

Ext: 8.8.8.8

OSPF

Enabled

✖ OSPF not enabled.

Add VLAN

Cancel

4 In the **VLAN** window, configure the following details:**Table 13-2.**

Option	Description
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.
Assign Overlapping Subnets	<p>Select the checkbox 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 checkbox and configure the subnets on each Edge individually.</p> <hr/> <p>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.</p>
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 checkbox to advertise the VLAN to other branches in the network.
ICMP Echo Response	Select the checkbox to enable the VLAN to respond to ICMP echo messages.
VNF Insertion	Select the checkbox 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. For more information about VNF, see Security VNFs
Multicast	<p>This option is enabled only when you have configured multicast settings for the Edge. You can configure the following multicast settings for the VLAN.</p> <ul style="list-style-type: none"> ■ IGMP ■ PIM <p>Click toggle advanced multicast settings to set the following timers:</p> <ul style="list-style-type: none"> ■ PIM Hello Timer ■ IGMP Host Query Interval ■ IGMP Max Query Response Value
Fixed IPs	You can configure the fixed IP only at the Edge level.
LAN Interfaces	You can configure the LAN Interfaces only at the Edge level.

Table 13-2. (continued)

Option	Description
SSID	You can configure the Wi-Fi SSID details for the VLAN only at the Edge level.
DHCP Type	<p>Choose one of the following DHCP settings:</p> <p>Enabled – Enables DHCP with the Edge as the DHCP server. Configure the following details:</p> <ul style="list-style-type: none"> ■ 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. <p>Relay – Enables DHCP with the DHCP Relay Agent installed at a remote location. If you choose this option, configure the following:</p> <ul style="list-style-type: none"> ■ Source from Secondary IP(s) – When you select this checkbox, 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. <p>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.</p> <ul style="list-style-type: none"> ■ Relay Agent IP(s) – Specify the IP address of Relay Agent. Click the Plus(+) Icon to add more IP addresses. <p>Not Enabled – Deactivates DHCP.</p>
OSPF	This option is enabled only when you have configured OSPF for the Edge. Select the checkbox and choose an OSPF from the drop-down list.

- 5 Click **Add VLAN**. The VLAN is configured for the Profile. You can change the VLAN settings by clicking the **Edit** link under **Actions** column.

To configure VLANs for Edges, see [Configure VLAN for Edges](#).

Configure the Management IP Address

The Management IP address is deprecated and is replaced with Loopback Interfaces.

For more information about Loopback Interfaces, see [Loopback Interfaces Configuration](#).


IPv6 Settings

VMware 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

An IPv4 only Interface can establish overlay only with either IPv4 or dual stack regardless of the overlay initiator and the preference value is ignored. The same rule applies to IPv6 only Interface as well. You cannot establish overlay between an IPv4 only and IPv6 only Interfaces.



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 (unpreferred tunnel) and IPv6 VCMP to the IPv6 only Interface on Edge B3 (preferred tunnel).

Dual Stack Environment


When all the Edges and Gateways are on dual stack, the tunnel preference is selected as follows:

- **Edge to Gateway** – The initiator, Edge, always chooses the tunnel type based on the tunnel preference.
- **Edge to Hub** – The initiator, Spoke Edge, always chooses the tunnel type based on the tunnel preference.
- **Dynamic Branch to Branch** – When 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.



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 more information, see [Configure Edge WAN Overlay Settings](#).

Limitations of IPv6 Address Configuration

- SD-WAN 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 SD-WAN Edge would initiate the tunnel using the preferred address family configured on the routed Interface.
- If all the WAN Interfaces are migrated to IPv6 only, the Edge loses its direct path to Orchestrator communication as fallback. In this environment, the Orchestrator services require at least one routed interface with IPv4 address and a default Gateway to forward the Orchestrator communication through multi-path routes.
- 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.

You can configure IPv6 addresses for the following:

- [Configure Static Route Settings](#)
- [Configure Interface Settings](#)
- [Configure Edge WAN Overlay Settings](#)
- [Configure BGP from Edge to Underlay Neighbors](#)
- [Configure BGP Over IPsec from Edge to Non SD-WAN Neighbors](#)
- [Configure BFD](#)
- [Configure a Loopback Interface for an Edge](#)

Configure Device Settings

Device Settings allows you configure the Interface Settings for one or more Edge models in a profile.

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 SD-WAN Orchestrator. This auto-discovered WAN link can then be modified via the SD-WAN 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 Datacenter 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 SD-WAN Gateway, which will handle the dissemination of information needed for peer connectivity.

Private (MPLS) WAN Links

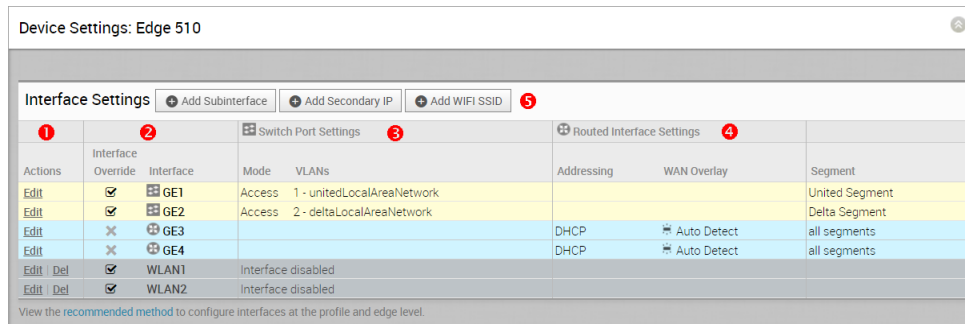
Private WAN links belong to a private network and can only connect to other WAN links within the same private network. Because there can be multiple MPLS networks, within a single enterprise, for example, the user must identify which links belong to which network. The SD-WAN Gateway will use this information 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

The following screen captures illustrate the top-level user interface for the SD-WAN Edge 500, SD-WAN Edge 1000, and introducing SD-WAN Edge 610 for the 3.4 release. The following table describes the major features.



- 1 Actions you can perform on the network interface, such as Edit or **Delete**.
 - 2 The Interface name. This name matches the Edge port label on the Edge device or is predetermined for wireless LANs.
 - 3 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.
 - 4 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.
 - 5 The list of Wireless Interfaces (if available on the Edge device). You can add additional wireless networks by clicking the **Add Wi-Fi SSID** button. Wireless Interfaces are highlighted with a light, gray background.
- 5
 - You can add additional wireless networks by clicking the **Add Wi-Fi SSID** button. Wireless Interfaces are highlighted with a light gray background.
 - You can add sub interfaces by clicking the **Add Sub Interfaces** button. Sub interfaces are displayed with "SIF" next to the interface. Sub interface for PPPoE interfaces is not supported.
 - You can add secondary IPs by clicking the **Add Secondary IP** button. Secondary IPs are displayed with "SIP" next to the interface.

Edges Without Wifi Modules

VMware supports Edge models 510, 610, 620, 640, and 680 without WiFi modules for the following releases: 3.4.6, 4.2.2, 4.3.0, 4.3.1, and 4.5.0. For specific model names, see table below. The Edge 6X0 series device and 510 Edge device are shipped with default images, but the working image is typically downloaded from the SD-WAN Orchestrator upon activation.

Note When an Edge without Wi-Fi is activated, the Wi-Fi settings in the SD-WAN Orchestrator will not be visible, as shown in the image below.

Device Settings: Edge 620

Interface Settings [Add Subinterface](#) [Add Secondary IP](#) [Add Wi-Fi SSID](#)

Actions	Interface Override	Interface	Switch Port Settings		Routed Interface Settings		Segment	Multicast		VNF Insertion
			Mode	VLANs	Addressing	WAN Overlay		IGMP	PIM	
Edit	<input checked="" type="checkbox"/>	GE1	Access	1 - Corporate			Global Segment			
Edit	<input checked="" type="checkbox"/>	GE2	Access	1 - Corporate			Global Segment			
Edit	<input checked="" type="checkbox"/>	GE3			IPv4 - DHCP	Auto Detect	all segments			<input checked="" type="checkbox"/>
Edit	<input checked="" type="checkbox"/>	GE4			IPv4 - DHCP	Auto Detect	all segments			<input checked="" type="checkbox"/>
Edit	<input checked="" type="checkbox"/>	GE5			IPv4 - DHCP	Auto Detect	all segments			<input checked="" type="checkbox"/>
Edit	<input checked="" type="checkbox"/>	GE6			IPv4 - DHCP	Auto Detect	all segments			<input checked="" type="checkbox"/>
Edit	<input checked="" type="checkbox"/>	SFP1			IPv4 - DHCP	Auto Detect	all segments			<input checked="" type="checkbox"/>
Edit	<input checked="" type="checkbox"/>	SFP2			IPv4 - DHCP	Auto Detect	all segments			<input checked="" type="checkbox"/>
Edit	<input checked="" type="checkbox"/>	WLAN1	Interface not enabled							
Edit	<input checked="" type="checkbox"/>	WLAN2	Interface not enabled							

View the [recommended method](#) to configure interfaces at the profile and edge level.

WAN Settings [Add User Defined WAN Overlay](#)

Actions	Type	Name	Address		Link Type	Public IP	Pre-Notifications	Alerts
			Type	Interfaces				
Edit Del	Auto Detect	2406:7400:bf:9254:956c:1f30:8768:6f21	IPv6	INTERNET1	Public Wired	2406:7400:bf:9254:956c:1f30:8768:6f21	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Edit Del	Auto Detect	AT&T Wireless	IPv4	USB3	Public Wireless	166.177.250.247	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Edit Del	Auto Detect	2406:7400:bf:9254:956c:1f30:8768:9f21	IPv6	INTERNET2	Public Wired	2406:7400:bf:9254:956c:1f30:8768:9f21	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

Table 13-3. Model Names: Edges Without WiFi Modules

Marketing Name	Hardware Model	Hardware Part Number
Edge 510N	Edge 510	Edge 510-NW
Edge 610N	E42W	Edge 610N
Edge 610-LTE	E42W	Edge 610LTE-RW, Edge 610LTE-AM
Edge 620N	E42W	Edge 620N
Edge 640N	E42W	Edge 640N
Edge 680N	E42W	Edge 680N

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 more information, see [Configure Interface Settings](#).

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 section titled, [Remote Diagnostics](#)
- 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 SD-WAN 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 SD-WAN 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.



Edge 3810

Edge 3810 is an evolution of the Edge 3800 platform, which includes 6 GE ports and 8 SFP ports. Otherwise, the functionality is identical to the Edge 3800.

Edge 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

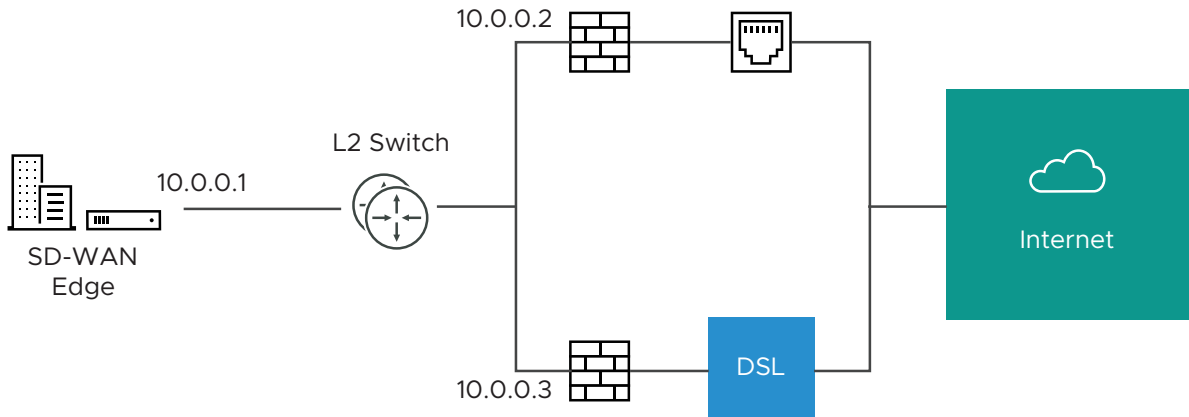
For the Edge 510-LTE model, a new routed interface (CELL1) is displayed in the **Interface Settings**. To edit the Cell Settings, see [Configure Interface Settings](#).

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 more information, see [Remote Diagnostics](#)

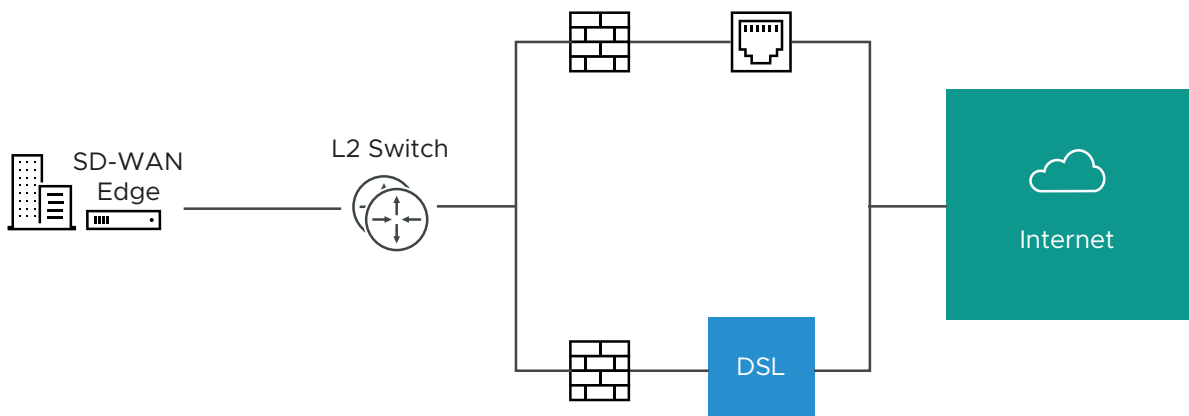
User-defined WAN Overlay Use Cases

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 SD-WAN Edge is connected to an L2 switch in the DMZ that is connected to multiple firewalls, each connected to a different upstream WAN link.



In this topology, the VMware 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:




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 SD-WAN 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 SD-WAN Edge.



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.

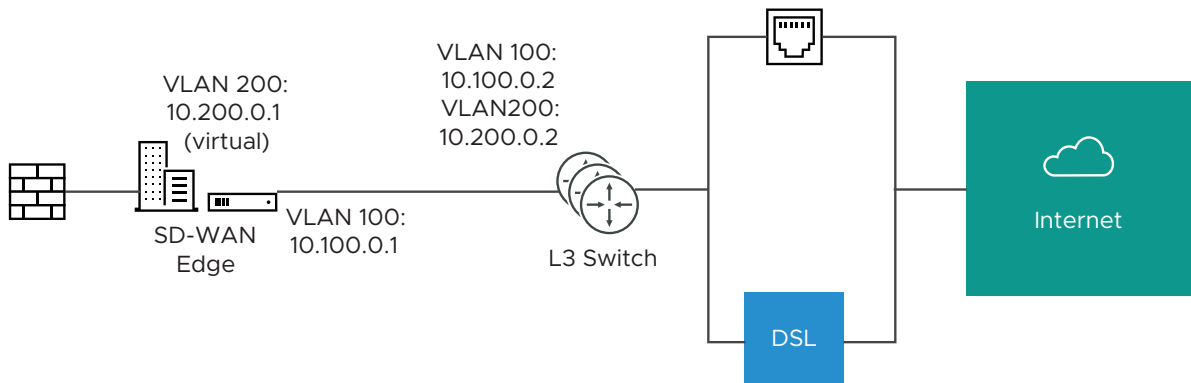
Steering by IP: If the L3 device is capable of policy-based routing by source IP address, then both devices may reside on the same VLAN. In this case, the only configuration required is a custom source IP to differentiate the devices.



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 SD-WAN 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.



- 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 SD-WAN 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.



Again, the SD-WAN 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 VMware 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.



- 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.


Configure Interface Settings

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.

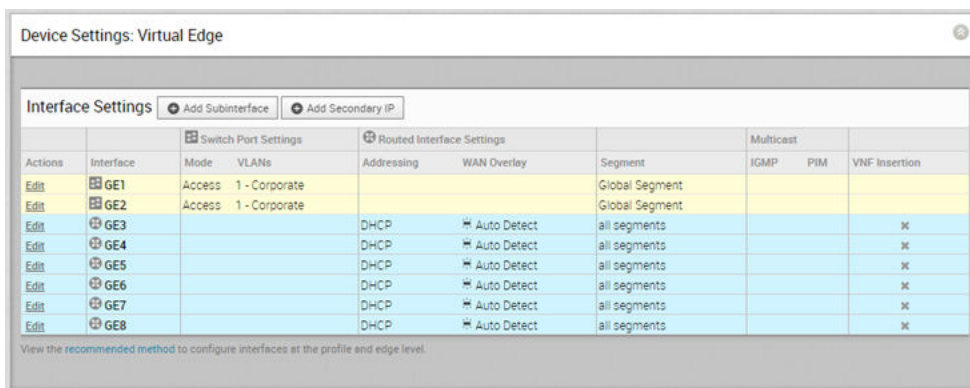
The Interface Settings options vary based on the Edge model. For more information on different Edge models and deployments, see [Configure Device Settings](#).

Procedure

- 1 In the Enterprise portal, click **Configure > Profiles**.
- 2 Click the Device Icon next to a profile, or click the link to the profile, and then click the **Device** tab.
- 3 Scroll down to the **Device Settings** section, which displays the existing Edge models in the Enterprise.




- 4 Click the DOWN arrow next to an Edge model to view the **Interface Settings** for the Edge.



The **Interface Settings** section displays the existing interfaces available in the selected Edge model.

- 5 Click the **Edit** option for an Interface to view and modify the settings.
- 6 The following image shows the Switch Port settings of an Interface.



You can modify the existing settings as follows:

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.
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 enabled by default. When enabled, 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.

Click **Update** to save the settings.

- 7 The following image shows the Routed Interface settings.

Virtual Edge

Interface GE3

Interface Enabled ☒

Capability **Routed**

Segments **All Segments**

RADIUS Authentication ☒ Require User Authentication to access WAN ✖ WAN Overlay must be disabled to configure RADIUS Authentication.

ICMP Echo Response ☒

Underlay Accounting ☒

Enable WAN Overlay ☒

VLAN **101**

IP Preference ☐ IPv4 ☒ IPv6

IPv4 Settings ☒ Active

Addressing Type **Static** Static/PPPoE addressing details must be configured individually per edge.

WAN Overlay **Auto-Detect Overlay**

OSPF ☒ ✖ OSPF not enabled for the selected Segment.

Multicast ☒ Multicast is not enabled for the selected segment

VNF Insertion ☒ ✖ VNF insertion is disallowed when an interface is configured for WAN overlays

Advertise ☒

NAT Direct Traffic ☒

Trusted Source ☒

Reverse Path Forwarding **Specific**

IPv6 Settings ☒ Active

Addressing Type **Static** Static addressing details must be configured individually per edge.

WAN Overlay **Auto-Detect Overlay**

Trusted Source ☒

Reverse Path Forwarding **Specific**

L2 Settings

Autonegotiate ☒

MTU **1500**

Update GE3 **Cancel**

You can modify the existing settings as follows:

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.
Capability	For a Routed Interface, the option Routed is selected by default. You can choose to convert the Interface to a Switch Port by selecting the option Switched from the drop-down list.
Segments	By default, the configuration settings are applicable to all the segments.

Option	Description
RADIUS Authentication	You must turn off WAN Overlay to configure RADIUS Authentication. Select the checkbox to enable RADIUS Authentication on the Interface and add the MAC addresses that should not be forwarded to RADIUS for re-authentication. For more information, see Enable RADIUS on a Routed Interface .
ICMP Echo Response	Select the checkbox to enable the Interface to respond to ICMP echo messages. You can turn off this option for the Interface, for security purposes.
Underlay Accounting	<p>This option is enabled by default. If a private WAN overlay is defined on the Interface, all underlay traffic traversing the interface will be counted against the measured rate of the WAN link to prevent over-subscription. If you do not want this behavior (for example, while using one-arm deployments), turn off the option.</p> <hr/> <p>Note Underlay Accounting is supported for both the IPv4 and IPv6 addresses.</p>
Enable WAN Overlay	Select the checkbox to enable WAN overlay for the Interface.
VLAN	Enter a VLAN ID for the Interface to support VLAN tagging over the port.
Overlay Preference	Choose whether the WAN Overlay link should be using IPv4 or IPv6 address when initiating tunnels.
IPv4 Settings – Select the Active checkbox to enable IPv4 Settings.	
Addressing Type	By default, DHCP is selected, which assigns an IPv4 address dynamically. If you select Static or PPPoE, you should configure the addressing details for each Edge.
WAN Overlay	<p>By default, this option is enabled with Auto-Detect Overlay. You can choose the User Defined Overlay and configure the Overlay settings. For more information, see Configure Edge WAN Overlay Settings.</p> <hr/> <p>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.</p>
OSPF	This option is enabled only when you have configured OSPF for the Profile. Select the checkbox and choose an OSPF from the drop-down list. Click toggle advance ospf settings to configure the Interface settings for the selected OSPF. For more information on OSPF settings, see Enable OSPF .

Option	Description
Multicast	<p>This option is enabled only when you have configured multicast settings for the Profile. You can configure the following multicast settings for the selected Interface.</p> <ul style="list-style-type: none"> ■ IGMP - Select the checkbox to enable Internet Group Management Protocol (IGMP) and only IGMP v2 is supported. ■ PIM - Select the checkbox to enable Protocol Independent Multicast and only PIM Sparse Mode (PIM-SM) is supported. <p>Click toggle advanced multicast settings to configure the following timers:</p> <ul style="list-style-type: none"> ■ 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.
VNF Insertion	<p>You must turn off WAN Overlay and enable Trusted Source to allow VNF insertion. When you insert the VNF into Layer 3 interfaces or sub-interfaces, the system redirects traffic from the Layer 3 interfaces or subinterfaces to the VNF.</p>
Advertise	<p>Select the checkbox to advertise the Interface to other branches in the network.</p>
NAT Direct Traffic	<p>Select the checkbox to apply NAT to the network traffic sent from the Interface.</p>
Trusted Source	<p>Select the checkbox to set the Interface as a trusted source.</p>

Option	Description
Reverse Path Forwarding	<p>You can choose an option for Reverse Path Forwarding (RPF) only when you have enabled Trusted Source. This option allows traffic on the interface only if return traffic can be forwarded on the same interface. This helps to prevent traffic from unknown sources 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:</p> <ul style="list-style-type: none"> ■ 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 not enabled. 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 Active checkbox to enable IPv6 Settings.	
Addressing Type	<p>Choose one of the options from the following to assign an IPv6 address dynamically.</p> <ul style="list-style-type: none"> ■ DHCP Stateless – Allows the Interface to self-configure the IPv6 address. It is not necessary to have a DHCPv6 server available at the ISP and an ICMPv6 discover message will originate from the Edge and is used for auto-configuration. <hr/> <p>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.</p> <ul style="list-style-type: none"> ■ 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. <hr/> <p>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.</p> <ul style="list-style-type: none"> ■ Static – If you select this option, you should configure the addressing details for each Edge. <hr/> <p>Note For Cell Interfaces, the Addressing Type would be Static by default.</p>

Option	Description
WAN Overlay	By default, this option is enabled with Auto-Detect Overlay. You can choose the User Defined Overlay and configure the Overlay settings. For more information, see Configure Edge WAN Overlay Settings .
Trusted Source	Select the checkbox to set the Interface as a trusted source.
Reverse Path Forwarding	<p>You can choose an option for Reverse Path Forwarding (RPF) only when you have enabled Trusted Source. This option allows traffic on the interface only if return traffic can be forwarded on the same interface. This helps to prevent traffic from unknown sources 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:</p> <ul style="list-style-type: none"> ■ 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 not enabled. 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.
L2 Settings	
Autonegotiate	This option is enabled by default. When enabled, 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 enabled. 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 enabled. 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
Enable LoS Detection	<p>This option is available only for a routed Interface of an Edge. Select the checkbox to enable Loss of Signal (LoS) detection by using ARP monitoring. For more information, see HA LoS Detection on Routed Interfaces.</p> <hr/> <p>Note You can select the checkbox only when you have enabled High Availability on the Edge.</p>
ARP Polling Interval	<p>This option is available only when Enable LoS Detection is enabled. Select the ARP Interval. 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.</p>
Cell Settings – This cellular related configuration option is available only for Edge models that support cellular connectivity, such as Edge 510-LTE and Edge 610-LTE.	
SIM PIN	Enter the PIN number used to unlock the SIM card.
Network	Select the Network of the Cell from the drop-down list. The following options are available: AT&T, Sprint, Verizon, or Other network.
APN	Name of optional carrier specific Access point.
IP Type	Select the type of IP address to be assigned to the Interface, as IPv4 or IPv6.
Username	Optional username provided by the carrier.
Password	Optional password provided by carrier.
SFP Settings – This option is available only for Edge models that support SFP ports.	
SFP Module	By default, Standard is selected. You can select DSL or GPON 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 .	

Option	Description
Mode	<p>Choose the DSL mode from the following options:</p> <ul style="list-style-type: none"> ■ 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. <p>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.</p> <ul style="list-style-type: none"> ■ 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. <p>If you choose ADSL2/2+, configure the following settings:</p> <ul style="list-style-type: none"> ■ 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.
GPON Settings – The option to configure Gigabit Passive Optical Network (GPON) settings is available when you select the SFP module as GPON .	
GPON Settings	<p>Configure the GPON mode settings:</p> <ul style="list-style-type: none"> ■ Subscriber Location ID Mode – Choose the mode of the Subscriber Location ID from the following options: <ul style="list-style-type: none"> ■ ASCII – Allows up to 10 ASCII characters. ■ HEX – Allows up to 20 Hexadecimal characters. ■ Subscriber Location ID – Enter the location ID according to the selected mode.

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'
```

Note For a list of modems supported for use on a SD-WAN Edge, see the [Supported Modems](#) page.

- 8 Some of the Edge models support Wireless LAN. The following image shows WLAN Interface settings.

Edge 500

Interface WLAN1

Interface Enabled ☒

VLAN 1 - Corporate

SSID vc-wifi

☒ Broadcast

Security WPA2 / Personal


Passphrase

Update WLAN1 Cancel

You can modify the settings as follows:

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 checkbox to broadcast the SSID name to the surrounding devices.
Security	<p>Select the type of security for the Wi-Fi connection, from the drop-down list. The following options are available:</p> <ul style="list-style-type: none"> ■ Open – No security is enforced. ■ WPA2 / Personal – A password is required for authentication. Enter the password in the Passphrase field. <hr/> <p>Note Starting from the 4.5 release, the use of the special character "<" in the password is no longer supported. In cases where users have already used "<" in their passwords in previous releases, they must remove it to save any changes on the page.</p> <ul style="list-style-type: none"> ■ 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. <p>To configure a RADIUS server, see Configure Authentication Services.</p> <p>To select the RADIUS server for a Profile, see Configure Authentication Settings.</p>

- 9 You can add Sub Interfaces to an existing Interface.
 - a In the **Interface Settings** section, click **Add Sub Interface**.
 - b In the **Select Interface** window, select the Interface for which you want to add a Sub Interface.




The **Select Interface** dialog box contains the following fields:

- Select Interface:** A dropdown menu with **GE3** selected.
- Subinterface ID:** A text input field containing **101**.
- Subinterface Type:** A dropdown menu with **Sub Interface** selected.
- Next:** A button at the bottom right.

Enter the **Subinterface ID** and click **Next**.

- c In the **Sub Interface** window, configure the Interface settings.



The **Virtual Edge** window for **Sub Interface GE3:101** shows the following configuration options:

- Segments:** A dropdown menu with **segment2** selected.
- RADIUS Authentication:** A checkbox labeled **Require User Authentication to access WAN** is unchecked.
- ICMP Echo Response:** A checkbox is checked.
- Enable VLAN Tagging:** A checkbox is checked, with a text input field containing **101**.
- IPv4 Settings:**
 - Addressing Type:** A dropdown menu with **DHCP** selected.
 - IP Address:** N/A
 - CIDR prefix:** N/A
 - Gateway:** N/A
 - VNF Insertion:** A checkbox is checked.
 - Multicast:** A checkbox is unchecked, with a note: "Multicast is not enabled for the selected segment".
 - Advertise:** A checkbox is checked.
 - NAT Direct Traffic:** A checkbox is checked.
 - Trusted Source:** A checkbox is checked.
 - Reverse Path Forwarding:** A dropdown menu with **Specific** selected.
- IPv6 Settings:**
 - Addressing Type:** A dropdown menu with **DHCP Stateless** selected.
 - IP Address:** N/A
 - CIDR prefix:** N/A
 - Gateway:** N/A
 - Trusted Source:** A checkbox is checked.
 - Reverse Path Forwarding:** A dropdown menu with **Specific** selected.

At the bottom right, there are buttons for **Update GE3:101** and **Cancel**.

For more information on the configuration options, refer to the Routed Interface settings.

10 You can add Secondary IP addresses to an existing Interface.

- a In the **Interface Settings** section, click **Add Secondary IP**.
- b In the **Select Interface** window, select the Interface for which you want to add a secondary IP address.




The **Select Interface** dialog box contains the following fields:

- Select Interface:** A dropdown menu with **GE4** selected.
- Subinterface ID:** A text input field containing **201**.
- Subinterface Type:** A label indicating the type is **Secondary IP**.
- Next:** A button at the bottom right to proceed.

Enter the **Subinterface ID** and click **Next**.

- c In the **Secondary IP** window, configure the Interface settings.



The **Virtual Edge** window for **Secondary IP GE4:201** shows the following configuration options:

- Segments:** A dropdown menu set to **Global Segment**.
- RADIUS Authentication:** A section with the text "Require User Authentication to access WAN" and a close icon (X).
- ICMP Echo Response:** A checkbox that is checked.
- IPv4 Settings:**
 - Addressing Type:** A dropdown menu set to **Static**. Below it, a note states: "Static addressing details must be configured individually per edge."
 - Advertise:** A checkbox that is checked.
 - NAT Direct Traffic:** A checkbox that is checked.
- Buttons:** **Update GE4:201** and **Cancel** buttons at the bottom right.

For more information on the configuration options, refer to the Routed Interface settings.

11 In the **Devices** tab, click **Save Changes**.

What to do next

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 as follows:

- 1 In the Enterprise portal, click **Configure > Edges**.
- 2 Click the Device Icon next to an Edge, or click the link to an Edge and then click the **Device** tab.
- 3 In the **Device** tab, scroll down to the **Interface Settings** section, which displays the interfaces available in the selected Edge.
- 4 Click 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.

Configure Wi-Fi Radio Settings

At the profile level, you can activate or deactivate WI-FI Radio and configure the band of radio frequencies.

Procedure

- 1 In the Enterprise portal, click **Configure > Profiles**.
The **Configuration Profiles** page appears.
- 2 Select a profile you want to configure WI-FI Radio settings and click the icon under the **Device** column.
The **Device Settings** page for the selected profile appears.
- 3 In the **WI-FI Radio Settings** area, by default, the **Radio Enabled** checkbox is selected and **Channel** is set to **Automatic**.
- 4 Select the radio band. It can be **2.4 GHz** or **5 GHz**.
- 5 Click **Save Changes**.

Wi-Fi Radio Settings

Radio Enabled: ☒

Band: 2.4 GHz 5 GHz

Channel: Automatic

At the Edge level, you can override the WI-FI Radio settings specified in the Profile by selecting the **Enable Edge Override** checkbox. For more information, see [Configure Wi-Fi Radio Overrides](#).

Activate Multi-Source QOS

Multi-Source QOS intelligently assigns the bandwidth to remote sources such as Gateways, Hubs, and other Edges based on the local availability and traffic priority.

To activate Multi-Source QOS for a Profile:

Procedure

- 1 Log in to VMware SD-WAN Orchestrator, and then go to **Configure > Profiles**.
- 2 Either click the Device icon next to the Profile for which you want to activate Multi-Source QOS or click the Profile link, and then go to the **Device** tab.
- 3 Scroll down to the **Multi-Source QOS** area, and then turn on the **Multi-Source QOS** toggle button.
- 4 Click **Save Changes**.

Results

The Multi-Source QOS setting is activated for the Profile.

To override this setting at the Edge-Level, go to **Configure > Edges > Device**, and then scroll down to the **Multi-Source QOS** area. Select the **Enable Edge Override** check box, and then either turn off or turn on the **Multi-Source QOS** toggle button depending on your requirement.

Configure Layer 2 Settings for Profiles

VMware SD-WAN Orchestrator supports Address Resolution Protocol (ARP) timeout configuration to allow the user to override the default timeout values of the ARP table entries. VMware SD-WAN 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:

Procedure

- 1 From the SD-WAN Orchestrator, go to **Configure > Profiles**.
The **Configuration Profiles** page appears.
- 2 Select a profile you want to configure L2 settings and click the icon under the **Device** column.
The Device Settings page for the selected profile appears.
- 3 Go to the **L2 Settings** area and select the **Override default ARP Timeouts** checkbox.

L2 Settings

Override default ARP Timeouts ☒

ARP Stale Timeout:	<input type="text" value="0"/>	Hours	<input type="text" value="2"/>	Minutes
ARP Dead Timeout:	<input type="text" value="0"/>	Hours	<input type="text" value="25"/>	Minutes
ARP Cleanup Timeout:	<input type="text" value="2"/>	Hours	<input type="text" value="3"/>	Minutes

4 Configure the various ARP timeouts in hours and minutes as follows:

Field	Description
ARP Stale Timeout	<p>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.</p> <p>The allowable value ranges from 1 minute to 23 hours and 58 minutes.</p>
ARP Dead Timeout	<p>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.</p> <p>The allowable value ranges from 2 minutes to 23 hours and 59 minutes.</p>
ARP Cleanup Timeout	<p>When an ARP's age exceeds the Cleanup time, the entry will be completely removed from ARP table.</p> <p>The allowable value ranges from 3 minutes to 24 hours.</p>

Note The ARP timeout values can only be in increasing order of minutes.

5 Click **Save Changes**.

What to do next

At the edge-level, you can override the L2 settings for specific edges. For more information, see [Configure Layer 2 Settings for Edges](#).

Configure SNMP Settings for Profiles

SNMP is a commonly used protocol for network monitoring and MIB is a database associated with SNMP to manage entities. SNMP can be enabled by selecting the desired SNMP version as described in the steps below.

Before you begin:

- To download the SD-WAN Edge MIB: go to the **Remote Diagnostic** screen (**Test & Troubleshooting > Remote Diagnostics**) and run MIB for SD-WAN Edge. Copy and paste 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 the above-mentioned MIBs are available on the Remote Diagnostics page.

Supported MIBs

- SNMP MIB-2 System
- SNMP MIB-2 Interfaces
- VELOCLOUD-EDGE-MIB

Procedure to Configure SNMP Settings at Profile Level:

- 1 Obtain the VELOCLOUD-EDGE-MIB from **Remote Diagnostic**.
- 2 Install all MIBs required by VELOCLOUD-EDGE-MIB. (See "Before you begin" for more information.
- 3 From the SD-WAN Orchestrator, go to **Configure > Profiles**.

The **Configuration Profiles** screen appears.

- 4 Select a profile you want to configure SNMP settings for, and click the **Device** icon under the Device column.

The **Configuration Profiles** screen for the selected Profile appears.

- 5 Scroll down to the **SNMP Settings** area. You can choose between two versions, v2c or v3.
- 6 For a SNMP v2c Config follow the steps below:
 - a Check the **v2c** check box.
 - b Type in a Port in the **Port** textbox. The default setting is 161.
 - c In the **Community** textbox, type in a word or sequence of numbers that will act as a 'password' that will allow you access to the SNMP agent.

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.

- d For Allowed IPs:
 - Check the **Any** check box to allow any IP to access the SNMP agent.
 - To restrict access to the SNMP agent, uncheck the **Any** check box and enter the IP address(es) that will be allowed access to the SNMP agent.

The screenshot shows a 'SNMP Settings' form. It has four main sections: 'SNMP Version' with a dropdown menu currently showing 'v2c'; 'Port' with a text input field containing '161'; 'Community' with an empty text input field; and 'Allowed IPs' with a text input field containing 'Allowed IP' and a button with a plus sign to add more IP addresses.

7 For a SNMP v3 Config, which provides added security support follow the steps below:

- a Type in a port in the **Port** text box. 161 is the default setting.
- b Type in a username and password in the appropriate text boxes. Ensure that the Privacy Password is same as the Authentication Password configured on the Edge.

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.

- c Check the **Privacy** check box if you want your packet transfer encrypted.
- d If you have checked the **Privacy** check box, choose **DES** or **AES** from the **Algorithm** drop-down menu.

Note Algorithm **AES** indicates **AES-128**.

8 Configure Firewall Settings. After you have configured SNMP Settings, go to Firewall settings (**Configure > Profiles > Firewall**) to configure the Firewall settings that will enable your SNMP settings.

Note SNMP interface monitoring is supported on DPDK enabled interfaces for 3.3.0 and later releases.

Configure NTP Settings for Profiles

The Network Time Protocol (NTP) provides the mechanisms to synchronize time and coordinate time distribution in a large, diverse network. VMware 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 SD-WAN Edge to set its own time accurately by configuring a set of upstream NTP Servers to get its time. While 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 SD-WAN Edge to act as an NTP Server to its own clients.

Prerequisites

NTP has the following prerequisites:

- To configure an SD-WAN 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.

Procedure

- 1 From the SD-WAN Orchestrator, go to **Configure > Profiles**.

The **Configuration Profiles** page appears.

- 2 Select a profile for which you want to configure NTP and click the icon under the **Device** column.

The Device Settings page for the selected profile appears.

- 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 Go to the **NTP** area and select the **Private NTP Servers Enabled** check box.
 - b In the **Servers** textbox, 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, click the **+** button.

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 SD-WAN 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.

NTP

Edge as NTP Client

Source Interface **Auto** ⓘ

Private NTP Servers Enabled ⓘ ☒

Servers

IP Address or DNS Name	
10.1.1.1	<input type="button" value="-"/> <input type="button" value="+"/>

Edge as NTP Server

Enabled ⓘ ☒

Authentication **None** **MD5**

Keys

Trusted Key #	Key Value	
1	1	<input type="button" value="-"/> <input type="button" value="+"/>

- 4 Once you have defined Private NTP Servers, Orchestrator allows you to configure the SD-WAN Edge to act as an NTP Server for its clients:
 - a Under **Edge as NTP Server**, select the **Enabled** 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.
- 5 Click **Save Changes**. The NTP configuration settings are applied to the selected profile.

What to do next

At the Edge-level, you can override the NTP settings for specific Edges. For more information, see [Configure NTP Settings for Edges](#).

Configure Visibility Mode

This section describes how to configure Visibility mode.

About Visibility Mode

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.

Visibility Mode	<input type="checkbox"/> Enable Edge Override ⓘ
<input type="radio"/> Visibility by MAC address <input checked="" type="radio"/> Visibility by IP address	

Choosing Visibility Mode

To choose a **Visibility Mode**, go to **Configure > Profile > Devices** tab. Select one of the following:

- **Visibility by MAC address**
- **Visibility by IP address**

Considerations for Using Visibility Mode

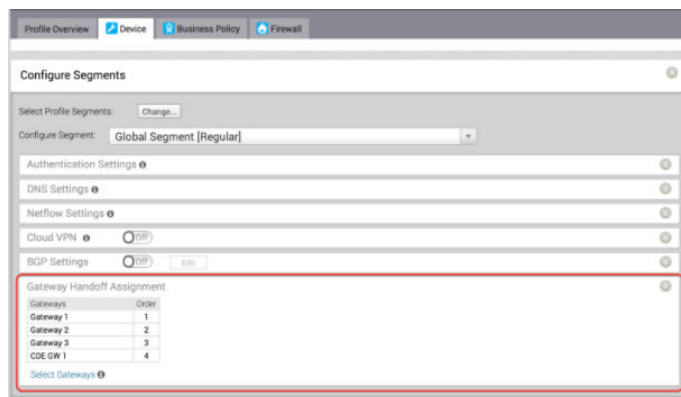
Keep in mind the following when choosing a Visibility mode:

- 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

Assign Partner Gateways

In order for customers to be able to use partner gateways, your Operator must select the **Enable Partner Handoff** check box for the Gateway to activate this feature. If this feature is available to you, will see the **Partner Gateway Assignment** area in the **Configure > Profiles > Device** tab screen.

Note The Partner Gateway 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.

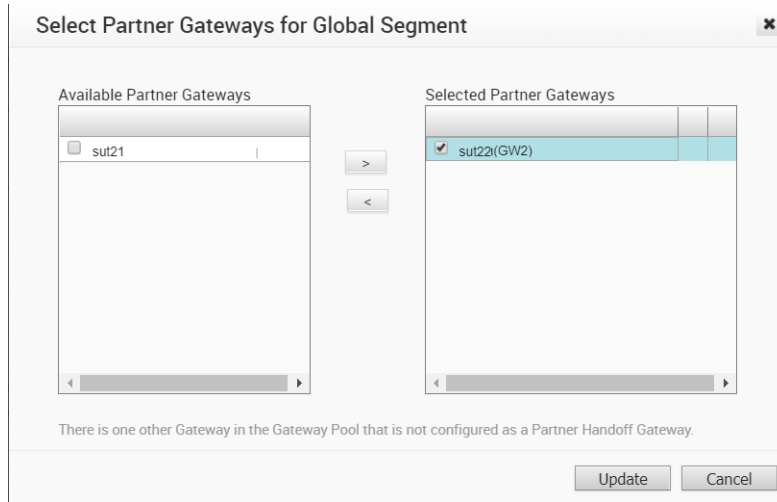


Select Gateways

To complete this section, you must have this feature activated. See your Operator for more information.

If there are no Gateways listed in the **Gateway Handoff Assignment** area:

- 1 Click the **Select Gateways** link to select Partner Gateways.
- 2 In the **Select Partner Gateways for Global Segment** dialog box, select an available Partner Gateway from the **Available Partner Gateway** area and move it (using the appropriate arrow) to the **Selected Partner Gateway** area.



Note that only Gateways configured as a Partner Handoff Gateway will be visible in the **Available Partner Gateways** area. If there are other Gateways not configured as a Partner Handoff Gateway, the following 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.**

Selecting CDE Gateways

In normal scenarios, the PCI traffic runs between customer branch and the 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 unchecking 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).

To complete this section, you must have this feature activated. See your Operator for more information.

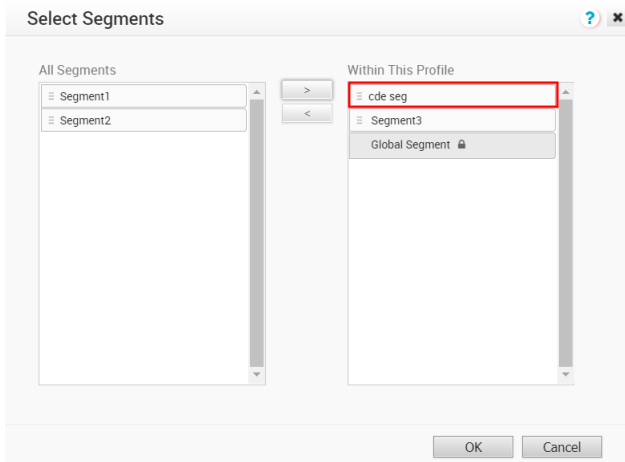
Assign a CDE Gateway

To assign a CDE Gateway:

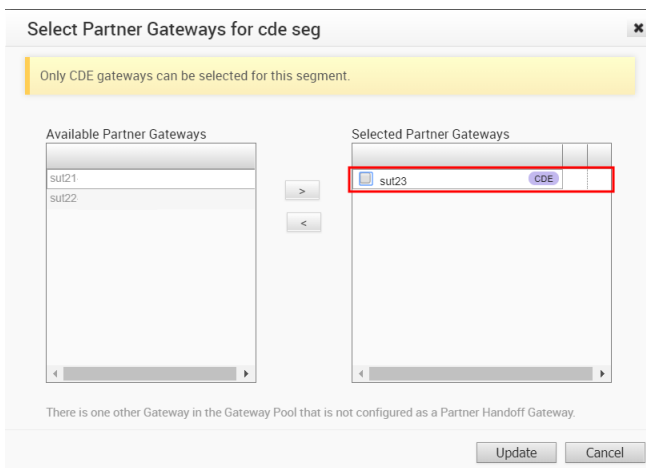
- 1 From the **Configure Segments** window, click the **Select Profile Segments Change** button.



- 2 In the **Select Segments** dialog box, move the available CDE segment from the **Available Segments** area (using the appropriate arrow) to the **Within This Profile** area.



- 3 In the **Gateway Handoff Assignment** area, click the **Select Gateways** link.
- 4 In the **Select Partner Gateways for cde seg** dialog box, select an available CDE Partner Gateway (from the **Available Partner Gateways** area) and move it to the **Selected Partner Gateways** area.



- 5 Click the **Update** button.

The **Gateway Handoff Assignment** area refreshes with the selected Gateways.

Note As indicated in the **Select Partner Gateways for cde seg** dialog box, only CDE gateways can be selected for the segment.

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 **Configure Segments** window, contact your Operator to activate this feature.

Assign Controllers

The SD-WAN Gateway is activated for supporting both the data and control plane. In the 3.2 release, VMware introduces a Controller-only feature (Controller Gateway Assignment).

There are multiple use cases which require the SD-WAN 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.




Configuring a Gateway as a Controller

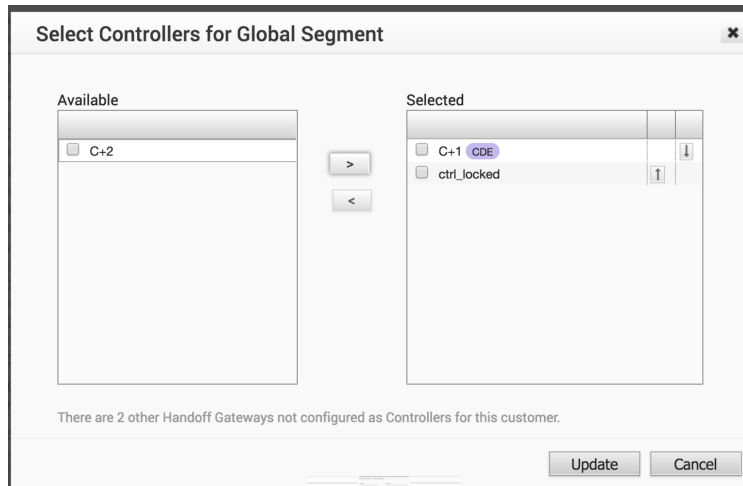
In order for customers to be able to use partner gateways, your Operator must select the **Enable Partner Handoff** check box for the Gateway to activate this feature. If this feature is available to you, you will see the **Controller Assignment** area in the **Configure > Profiles > Device** tab screen.

Note At least one Gateway in the Gateway Pool should be a "Controller Only" Gateway.

- 1 Go to **Configure > Profiles > Device** tab.
- 2 Scroll down to the **Controller Assignment** area.



- 3 In the **Controller Assignment** area, click the **Select Gateways** link.
- 4 In the **Select Controllers for Global Segment** dialog, move controllers from the **Available** area



to the **Selected** area.

- 5 Click **Update**.

The **Controller Assignment** area refreshes.



Configure Business Policy

14

VMware provides an enhanced Quality of Service feature called Business Policy. SD-WAN 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.

Read the following topics next:

- [Configure Business Policy for Profiles](#)
- [Configure Business Policy for Edges](#)
- [Create Business Policy Rules](#)

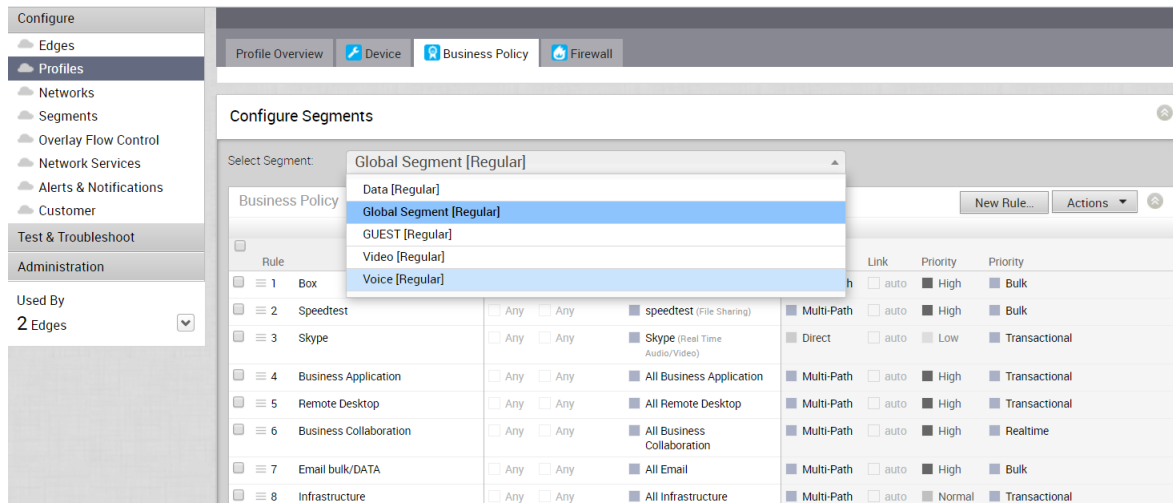
Configure Business Policy for Profiles

You can configure Business Policy rules using the **Business Policy** tab in the **Profile Configuration** dialog. Optionally, at the edge-level, you can also override the Profile Business Policy rules.

Note If you are logged in using a user ID that has Customer Support privileges, you will only be able to view SD-WAN Orchestrator objects. You will not be able to create new objects or configure/update existing ones.

Based on the business policy configuration, VMware 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.

The following screenshot shows 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 clicking the **New Rule** button. For steps to create a new business policy rule, see [Create Business Policy Rules](#).



Business Policy Rules are now Segment aware. All Segments available for configuration are listed in the **Configure Segment** drop-down menu.

When you choose a Segment to configure from the **Configure Segment** drop-down menu, the settings and options associated with that Segment appear in the **Configure Segments** area. **Global Segment [Regular]** is the default segment.

For more information about Segmentation, see [Chapter 8 Configure Segments](#) and [Chapter 13 Configure a Profile Device](#).

Note You can move your configured rules up or down in the list of rules to establish precedence by hovering over the numeric value at the left side of the rule and moving the rule up or down. If you hover over the right side of a rule, click the **– (minus) sign** next to the rule to remove it from the list or the **+ (plus) sign** to add a new rule.

Related Information: To override the Profile Business Policy rules at the edge level, see [Configure Business Policy for Edges](#).

Configure Business Policy for Edges

All the edges inherit the Business Policy rules from the associated Profile. Under the **Business Policy** tab of the **Edge Configuration** dialog, you can view all the inherited Business Policy rules in the **Rule From Profile** area. Overriding Profile Business Policy rules at the Edge is an optional step.

At the Edge level, Business Policy Rules from the assigned Profile can be overridden using the Edge Business Policy dialog shown below. Any Business Policy override match value that is the same as any Profile Business Policy rule, will override that Profile rule. You can create override rules in the same way as you create Profile rules (see [Configure Business Policy for Profiles](#)).

As shown in the image below, Business Policy is Segment aware. All Segments available for configuration are listed in the **Configure Segment** drop-down menu.

When you choose a Segment to configure from the **Configure Segment** drop down, the settings and options associated with that Segment display in the **Configure Segments** area. **Global Segment [Regular]** is the default Segment.

For more information about Segmentation, see [Chapter 8 Configure Segments](#) and *Configure Edge Device*.

Create Business Policy Rules

SD-WAN Orchestrator allows you to configure business policy rules at the Profile and Edge levels. 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.

Before you begin: Know the IP Addresses of your devices and understand the implications of setting a wildcard mask.

To create a business policy:

- 1 From the SD-WAN Orchestrator, go to **Configure > Profiles > Business Policy**.
- 2 Under **Business Policy** area, click **New Rule**. The **Configure Rule** dialog box appears.

Configure Rule

Rule Name:

Match

Source:

Destination:

☐ Any
☒ Internet
☐ Edge
☐ Non SD-WAN Destination via Gateway
☐ Non SD-WAN Destination via Edge

IP Address:

CIDR prefix:

Domain Name:

Protocol:

Ports:

Application:

Action

Priority:

☐ Rate Limit

Network Service:

☐ Backhaul Hubs
☐ Non SD-WAN Destination via Gateway
☒ Non SD-WAN Destination via Edge / Cloud Security Service
☐ VMware Cloud Web Security Gateway

Gen IKEv1:

☒ Turn off Conditional Backhaul

Link Steering:

Inner Packet DSCP Tag:

Outer Packet DSCP Tag:

NAT:

Service Class:

- 3 In the **Rule Name** box, enter a unique name for the rule.

- 4 Under the **Match** area, configure the match conditions for the traffic flow. The option you choose may change the fields in the dialog box:

Settings	Description
Source	<p>Allows to specify match criteria for the source traffic. Select any of the following options:</p> <ul style="list-style-type: none"> ■ Any - Matches all source traffic, by default. ■ Object Group - Allows you to select a combination of address group and port group to be matched for the source. For more information, see Chapter 21 Object Groups and Configure Business Policies with Object Groups. <hr/> <p>Note If the selected address group contains any domain names, then they would be ignored when matching for the source.</p> <hr/> <ul style="list-style-type: none"> ■ Define - Allows you to define the matching criteria for the source traffic from a specific VLAN, Interface, IP Address, Port, or Operating System. Select one of the following options, by default, None is selected: <ul style="list-style-type: none"> ■ 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. <hr/> <p>Note If an interface cannot be selected, then the interface is either not activated or not assigned to this segment.</p> <hr/> ■ IP Address - Matches traffic from the specified IP address. Along with the IP address, you can specify one of the following options to match the source traffic: <ul style="list-style-type: none"> ■ 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

Settings	Description
	<p>= 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.</p> <ul style="list-style-type: none"> ■ Port - 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.
Destination	<p>Allows to specify match criteria for the destination traffic. Select any of the following options:</p> <ul style="list-style-type: none"> ■ Any - Matches all destination traffic, by default. ■ Object Group - Allows you to select a combination of address group and port group to be matched for the destination. For more information, see Chapter 21 Object Groups and Configure Business Policies with Object Groups. ■ 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: <ul style="list-style-type: none"> ■ 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. <p>Protocol - Matches traffic for the specified protocol, selected from the drop-down menu. The supported protocols are: GRE, ICMP, TCP, and UDP.</p> <p>Domain - 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\".</p>
Application	<p>Select any one of the following options:</p> <ul style="list-style-type: none"> ■ Any - Applies the business policy rule to any application by default.

Settings	Description
	<ul style="list-style-type: none"> ■ 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. <p>Note</p> <ul style="list-style-type: none"> ■ 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 is 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.

Depending on your **Match** choices, some Actions may not be available.

5 Under the **Action** area, configure the actions for the rule:

Settings	Description
Priority	<p>Designate the priority of the rule as one of the following:</p> <ul style="list-style-type: none"> ■ High ■ Normal ■ Low <p>Select the Rate Limit checkbox to set limits for inbound and outbound traffic directions.</p> <hr/> <p>Note 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.</p>
Network Service	<p>Set the Network Service to one of the following options:</p> <ul style="list-style-type: none"> ■ Direct - Sends the traffic out of the WAN circuit directly to the destination, bypassing the SD-WAN Gateway. <hr/> <p>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.</p> <p>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.</p> <hr/> <ul style="list-style-type: none"> ■ Multi-Path - Sends the traffic from one SD-WAN Edge to another SD-WAN Edge. ■ Internet Backhaul - This network service is activated only if the Destination is set as Internet. <hr/> <p>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).</p> <hr/> <p>For information about these options, see Configure Network Service for Business Policy Rule.</p>

Settings	Description
	<p>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 checkbox.</p> <p>For more information about how to activate and troubleshoot the Conditional Backhaul feature, see Conditional Backhaul.</p>
Link Steering	<p>Select one of the following link steering modes:</p> <ul style="list-style-type: none"> ■ 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. Enter an Inner Packet DSCP Tag from the drop-down menu and an Outer Packet DSCP Tag from the drop-down menu. ■ 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: <ul style="list-style-type: none"> ■ Public Wired ■ Public Wireless ■ Private Wired ■ Interface - Link steering is tied to a physical interface and will be used primarily for routing purposes. <div style="border: 1px solid #ccc; padding: 5px; margin-top: 5px;"> <p>Note This option is only allowed at the Edge override level.</p> </div> ■ 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. <div style="border: 1px solid #ccc; padding: 5px; margin-top: 5px;"> <p>Note This option is only allowed at the Edge override level.</p> </div> <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;"> <p>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.</p> </div> <p>For more information about the link steering modes and DSCP, DSCP marking for both Underlay and Overlay traffic, see Configure Link Steering Modes.</p>

Settings	Description
NAT	Activate or deactivate NAT. For more information, see Configure Policy-based NAT .
Service Class	Select one of the following Service Class options: <ul style="list-style-type: none"> ■ Real-time ■ Transactional ■ Bulk <hr/> Note This option is only for a custom application. <hr/> VMware Apps/Categories fall in one of these categories.

- 6 Click **OK**. The business policy rule is created for the selected profile and it appears under the **Business Policy** area of the **Profile Business Policy** page.

Related Information: [Overlay QoS CoS Mapping](#)

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 SD-WAN Gateway. NAT is applied to the traffic if the **NAT Direct Traffic** checkbox 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 SD-WAN Edge to another SD-WAN Edge, and from a SD-WAN Edge to a SD-WAN 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 SD-WAN 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
- VMware Cloud Web Security Gateway

Note The VMware Cloud Web Security Gateway option is available only if a user has subscribed to use the VMware Cloud Web Security service.

For more information, see VMware SD-WAN Cloud Web Security Configuration Guide published at <https://docs.vmware.com/en/VMware-Cloud-Web-Security/index.html>.

You should be able to configure multiple VMware SD-WAN 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.

Configure Rule

Rule Name: CBH Rule

Match

Source: Any | Object Group | Define...

Destination: Any | Object Group | Define...

☐ Any
☒ Internet
☐ Edge
☐ Non SD-WAN Destination via Gateway
☐ Non SD-WAN Destination via Edge

IP Address: Ex: 10.0.2.0

CIDR prefix: 24

Domain Name: Ex: domain.com

Protocol: [dropdown]

Ports: Ex: 2224-2226

Application: Any | Define...

Action

Priority: High | Normal | Low

☐ Rate Limit

Network Service: Direct | Multi-Path | Internet Backhaul

☐ Backhaul Hubs
☐ Non SD-WAN Destination via Gateway
☒ Non SD-WAN Destination via Edge / Cloud Security Service
☐ VMware Cloud Web Security Gateway

Gen_IKEv1: [dropdown] ☒ Turn off Conditional Backhaul

Link Steering: Auto | Transport Group | Interface | WAN Link

Inner Packet DSCP Tag: Leave as is

Outer Packet DSCP Tag: 0 - CS0/DF

NAT: Not Enabled | Enabled

Service Class: Real Time | Transactional | Bulk

OK Cancel

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** checkbox in the **Action** area of the **Configure Rule** screen for the selected business policy.

For more information about how to enable and troubleshoot the Conditional Backhaul feature, see [Conditional Backhaul](#).

Configure Link Steering Modes

In the Business Policy, there are four link steering modes: **Auto**, **Transport Group**, **WAN Link**, and **Interface**.

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. As mentioned earlier, traffic within the Enterprise (VPN) always goes through the DMPO tunnels, hence it always receives the benefits of on-demand remediation.

Link Steering: **Auto** Transport Group Interface WAN Link

Inner Packet DSCP Tag: **Leave as is** ▼

Outer Packet DSCP Tag: **0 - CS0/DF** ▼

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 un-usable and steer the flow to the next best available link.
Both Jitter and Loss on both links.	<p>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.</p> <p>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.</p>

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** checkbox displays.

If you select the **Error Correct Before Steering** checkbox, the Loss% variable textbox 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** checkbox 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 checkbox, the application will steer before Error Correction occurs.

The screenshot shows the VMware SD-WAN configuration interface. Under 'Network Service', 'Multi-Path' is selected. Under 'Link Steering', 'Transport Group' is selected. The 'Transport Group' dropdown is set to 'Public Wired'. Below it, the 'Preferred' radio button is selected. The 'Error Correct Before Steering' checkbox is checked and highlighted with a red box. Below this checkbox, the 'Loss (%)' field is set to '4.00'. At the bottom, 'Inner Packet DSCP Tag' is set to 'Leave as is' and 'Outer Packet DSCP Tag' is set to '0 - CS0/DF'.

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 VMware SD-WAN 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** checkbox displays. If you select the checkbox, 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 checkbox, 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 SD-WAN Edge hardware model.

The screenshot shows the 'Link Steering' configuration window with the 'Interface' tab selected. The configuration includes:

- Interface:** A dropdown menu set to 'INTERNET1'.
- VLAN:** A text input field.
- ICMP Probe:** A dropdown menu set to '[none]'.
- Inner Packet DSCP Tag:** A dropdown menu set to '46 - EF'.
- Outer Packet DSCP Tag:** A dropdown menu set to '0 - CS0/DF'.
- Policy Options:** Three radio buttons: 'Mandatory' (selected), 'Preferred', and 'Available'.

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 [Private Network Names](#) and *Selecting a Private Name Link*.

If you choose the **Preferred** option, the **Error Correct Before Steering** checkbox displays. If you do not select this checkbox, the application will steer before Error Correction occurs.

Note This option is only allowed at the Edge override level.

The screenshot shows the 'Link Steering' configuration window with the 'WAN Link' tab selected. The configuration includes:

- WAN Link:** A dropdown menu set to 'e-commerce'.
- Policy Options:** Three radio buttons: 'Mandatory' (selected), 'Preferred', and 'Available'.
- Inner Packet DSCP Tag:** A dropdown menu set to 'Leave as is'.
- Outer Packet DSCP Tag:** A dropdown menu set to '0 - CS0/DF'.

For the **Interface** and **WAN Link** choices, you must select one of the following options:

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

VMware supports DSCP remarking of packets forwarded by the Edge to the Underlay. The SD-WAN 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 Rules](#) . 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**.

Rule Name:

Match

Source:

Destination:

☒ Any ☐ Internet ☐ VeloCloud Edge ☐ Non-VeloCloud Site ⓘ

IP Address:

Hostname: ⓘ

Protocol:

Ports:

Application:

Action

Priority:

☐ Rate Limit

Network Service: ⓘ

Link Steering:

Inner Packet DSCP Tag: ⓘ

Outer 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 enabled for WAN Overlay by default in the SD-WAN Orchestrator (**Configure > Edge Devices > Device Settings** area).

- 2 From the SD-WAN Orchestrator, go to **Configure > Edges > Business Policy**.
- 3 From the **Business Policy** screen, click an existing rule or click the **New Rule** button to create a new rule.
- 4 In the **Action** section, go to the **Link Steering** area.
- 5 Click one of the following as applicable: Auto, Transport Group, Interface, or WAN Link.
- 6 Configure **Match** criteria for the underlay traffic and configure **Inner Packet DSCP Tag**.

Linking Steering: Overlay DSCP Configuration

- 1 Verify that **Underlay Accounting** is enabled for WAN Overlay by default in the SD-WAN Orchestrator (**Configure > Edge Devices > Device Settings** area).
- 2 From the SD-WAN Orchestrator, go to **Configure > Edges > Business Policy**.
- 3 From the **Business Policy** screen, click an existing rule or click the **New Rule** button to create a new rule.
- 4 In the **Action** section, go to the **Link Steering** area.
- 5 Click one of the following as applicable: **Auto**, **Transport Group**, **Interface**, or **WAN Link**.

- 6 Configure **Match** criteria for the Overlay traffic and configure **Inner Packet DSCP Tag** and **Outer Packet DSCP Tag**.

Link Steering: **Auto** Transport Group Interface WAN Link ⓘ

Inner Packet DSCP Tag: Leave as is ⬆⬆

Outer Packet DSCP Tag: 0 - CS0/DF ⬆⬆

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 click the **New Rule** button. The NAT feature is located under the **Action** area.

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 100 - Corporate 2** (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.

Many to One NAT

Source NAT all traffic coming thru Vlan100 to 72.4.3.1

Match

Source: **Any** **Define...**

☐ None

☒ **VLAN:** 100 - Corporate 2

☐ **IP Address:** Ex: 10.0.2.0/24

☐ **Ports:** Ex: 2224-2226

☐ **Operating System:**

NAT: **NAT Enabled** **Enabled**

Source NAT IP: 72.4.3.1

Destination NAT IP:

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 SD-WAN Gateway will forward that traffic to the local IP address of the host or server in the Branch.

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 more information, see [Configure Class of Service](#).

	HIGH	NORMAL	LOW
REAL TIME	Business Collaboration	Audio/Video	
TRANSACTIONAL	Remote Desktop, Business App	Infrastructure, Advertisement, Management, Network Services, Streaming	IM, Web, Photos, Games, Media, Social
BULK	Email	File Sharing	Storage/Backup, FTP

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.

	HIGH	NORMAL	LOW
REAL TIME	35	15	1
TRANSACTIONAL	20	7	1
BULK	15	5	1

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 more information.

To activate Overlay QoS CoS Mapping:

- 1 Go to **Configure > Profiles**.
- 2 Click the link of the appropriate configuration Profile.
- 3 Click 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** checkbox for a Service Class, if necessary.

SD-WAN Traffic Class and Weight Mapping

Service Class / Priority	High	Policing	Normal	Policing	Low	Policing
Real Time	35	<input checked="" type="checkbox"/>	15	<input type="checkbox"/>	1	<input type="checkbox"/>
Transactional	20	<input type="checkbox"/>	7	<input type="checkbox"/>	1	<input type="checkbox"/>
Bulk	15	<input type="checkbox"/>	5	<input type="checkbox"/>	1	<input type="checkbox"/>

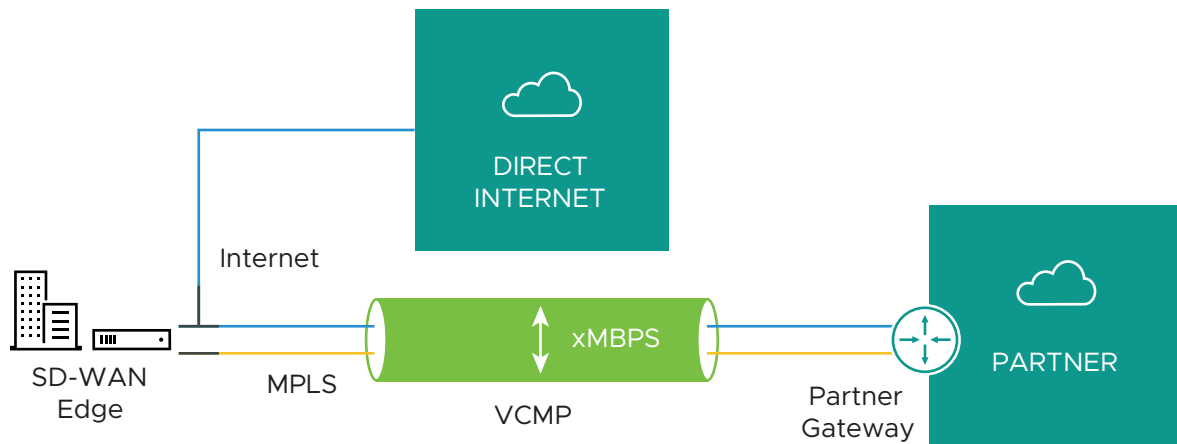
Note The QoS CoS Mapping feature is not supported for Underlay direct traffic.

Tunnel Shaper for Service Providers with Partner Gateway

This section describes 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 VMware 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



Consider a SD-WAN 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 enabled by your Operator. To gain access to this feature, see your Operator for more information.

To enable Rate-Limit Tunnel Traffic:

- 1 Go to **Configure > Profiles** from the navigation panel.
- 2 Click the link of the appropriate configuration profile.
- 3 Click the **Business Policy** tab.
- 4 In the **SD-WAN Overlay Rate Limit** area, check the **Rate-Limit Tunnel Traffic** check box. (See image below).
- 5 Select either the **Percent** or **Rate (Mbps)** radial buttons.
- 6 In the **Limit** text box, type in a numerical limit to the Tunnel Traffic.
- 7 Click **Save Changes**.

SD-WAN Overlay Rate Limit

Rate-Limit Tunnel Traffic: ☒

Percent (%): ☐

Rate (Mbps): ☒

Limit:

Configure Firewall

15

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. SD-WAN Orchestrator supports configuration of stateless and stateful firewalls for profiles and edges.

A Stateful firewall monitors and tracks the operating state and characteristics of every network connections 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.

The Stateful firewall feature provides the following benefits:

- Prevent attacks such as denial of service (DoS) and spoofing
- More robust logging
- Improved network security

The main differences between a Stateful firewall and a Stateless firewall are:

- Matching is directional. For example, 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 VMware 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 VMware SD-WAN Edge must be using Release 3.4.0 or later.
- By default, the **Stateful Firewall** feature is activated for new customers on an SD-WAN Orchestrator using 3.4.0 or later releases. Customers created on a 3.x Orchestrator will need assistance from a Partner or VMware SD-WAN Support to activate this feature.
- The SD-WAN 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 enabled edges.

To configure firewall settings at the profile and edge level, see:

- [Configure Firewall for Profiles](#)
- [Configure Firewall for Edges](#)

Stateful Firewall Logs

With the Stateful Firewall enabled, more 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, Rule, Bytes Received/Sent, and Duration.

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.

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 sending the logs originating from enterprise SD-WAN Edge to one or more centralized remote Syslog collectors (Servers). By default, the **Syslog Forwarding** feature is deactivated for an enterprise. To forward the logs to remote Syslog collectors, you must:

- 1 Enable **Syslog Forwarding** feature under **Configure > Edge/Profile > Firewall** tab.

- 2 Configure a Syslog collector under **Configure > Edges > Device > Syslog Settings**. For steps on how to configure Syslog collector details per segment in the SD-WAN Orchestrator, see [Configure Syslog Settings for Profiles](#).

Read the following topics next:

- [Configure Firewall for Profiles](#)
- [Configure Firewall for Edges](#)
- [Configure Firewall Rules](#)
- [Configure Stateful Firewall Settings](#)
- [Configure Network and Flood Protection Settings](#)
- [Configure Edge Access](#)
- [Troubleshooting Firewall](#)

Configure Firewall for Profiles

As an enterprise administrator, you can configure firewall rules, stateful firewall settings, network and flood protection settings, edge access information, and activate or deactivate firewall status and logs, using the **Firewall** tab in the **Profile Configuration** dialog.

Firewall Profiles are Segment aware. All Segments available for the configuration are listed in the **Configure Segment** drop-down menu. When you select a Segment to configure from the **Configure Segment** drop-down menu, the settings and options associated with that Segment appear in the **Configure Segments** area. **Global Segment [Regular]** is the default Segment.

For more information about Segmentation, see [Chapter 8 Configure Segments](#).

Monitor

Configure

Edges

Profiles

Object Groups

Segments

Overlay Flow Control

Network Services

Alerts & Notifications

Customer

Test & Troubleshoot

Administration

Used By

3 Edges

Configuration Profiles

Quick Start Profile

Save Changes ?

Profile Overview Device Business Policy Firewall

Firewall Status: On

Syslog Forwarding: Off

Stateful Firewall: On

Configure Segments

Configure Segment: Global Segment [Regular]

Firewall Rules

Rule	Match	Source	Destination	Application	Action
1 AllowAny		Any	Any	Any	Allow

New Rule... Actions

Inbound ACLs

Port forwarding and 1:1 NAT rules can be configured individually for each edge.

Stateful Firewall Settings

Established TCP Flow Timeout (seconds) 7440

Non Established TCP Flow Timeout (seconds) 240

UDP Flow Timeout (seconds) 300

Other Flow Timeout (seconds) 60

Network & Flood Protection Settings

New Connection Threshold (connections per second) 25 %

Denylist

TCP Based Attacks

Invalid TCP Flags

TCP Land

TCP SYN Fragment

ICMP Based Attacks

ICMP Ping of Death

ICMP Fragment

IP Based Attacks

IP Unknown Protocol

IP Options

Edge Access

Support Access

Deny All

Allow the following IPs

Ex: 54.183.9.192, 2001:9098:85a3:0000:0000:8a2e:03

Separate each IP with a comma (,)

Console Access

Deny

Allow

SNMP Access

Deny All

Allow All LAN

Allow the following IPs

Ex: 54.183.9.192, 46.2.142.142

Separate each IP with a comma (,)

Local Web UI Access

Deny All

Allow All LAN

Allow the following IPs

Ex: 54.183.9.192, 46.2.142.142

Separate each IP with a comma (,)

Local Web UI Port Number

80

The firewall configuration at the profile level includes:

- Enabling Syslog Forwarding. By default, the Syslog Forwarding feature is deactivated for an enterprise. To collect SD-WAN Orchestrator bound events and firewall logs originating from enterprise SD-WAN Edge to one or more centralized remote Syslog collectors (Servers), an enterprise user must enable this feature at the enterprise level. For steps on how to configure Syslog collector details per segment in the SD-WAN Orchestrator, see [Configure Syslog Settings for Profiles](#).
- Enabling Stateful Firewall at the Profile and Edge level. By default, the Stateful Firewall feature is enabled for an enterprise. To deactivate the Stateful Firewall feature for an enterprise, contact an Operator with Super User permission.
- [Configure Firewall Rules](#)
- [Configure Stateful Firewall Settings](#)
- [Configure Network and Flood Protection Settings](#)
- [Configure Edge Access](#)

Note You can deactivate the Firewall function for profiles by turning the **Firewall Status** to OFF.

Related Links

- [Configure Firewall for Edges](#)
- [Troubleshooting Firewall](#)

Configure Firewall for Edges

All the Edges inherit the firewall rules 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 Profile Firewall rules and Edge access configuration.

Test & Troubleshoot
Administration

- Monitor this Edge
- Events from this Edge
- Remote Actions
- Generate Diagnostic Bundle
- Remote Diagnostics

Configure Segments

Configure Segment: Global Segment [Regular]

Firewall Rules

☐ Rule

Match	Match	Match	Match	Action	
Source	Destination	Application		Action	
There are no Edge specific overrides.					
1	AllowAny	<input type="checkbox"/> Any	<input type="checkbox"/> Any	<input type="checkbox"/> Any	Allow

* Firewall rules applied from the assigned Profile of this Edge. [Quick Start Profile](#)

Inbound ACLs

Port Forwarding Rules

Name	Protocol	Interface	Outside IP	WAN Port(s)	LAN IP	LAN Port	Segment	Remote IP/Subnet
(optional)	Select...	Select...	Ex: 10.0.2.5	Ex: 27015	Ex: 10.0.2.5	Ex: 27030	Select...	Ex: 46.2.142.142/24

1:1 NAT Rules

Name	Outside IP	Interface	Inside (LAN) IP	Segment	Outbound Traffic	Allowed Traffic Source	Remote IP/Subnet
(optional)	Ex: 54.183.9.192	Select...	Ex: 10.0.2.5	Select...	<input type="checkbox"/>	<input type="checkbox"/> [all] <input type="checkbox"/> Ex: 27015	Ex: 46.2.142.142/24

Stateful Firewall Settings

☒ Enable Edge Override

Established TCP Flow Timeout (seconds)	<input type="text" value="7440"/>	Non Established TCP Flow Timeout (seconds)	<input type="text" value="240"/>
UDP Flow Timeout (seconds)	<input type="text" value="300"/>	Other Flow Timeout (seconds)	<input type="text" value="60"/>

Network & Flood Protection Settings

☒ Enable Edge Override

New Connection Threshold (connections per second) %

Denylist ☐

TCP Based Attacks Invalid TCP Flags <input type="checkbox"/> TCP Land <input type="checkbox"/> TCP SYN Fragment <input type="checkbox"/>	ICMP Based Attacks ICMP Ping of Death <input checked="" type="checkbox"/> ICMP Fragment <input type="checkbox"/>	IP Based Attacks IP Unknown Protocol <input type="checkbox"/> IP Options <input type="checkbox"/>
--	---	--

Edge Access

☒ Enable Edge Override

Support Access

☐ Deny All

☒ Allow the following IPs

Separate each IP with a comma (,)

Console Access

☐ Deny

☒ Allow

SNMP Access

☒ Deny All

☐ Allow All LAN

☐ Allow the following IPs

Separate each IP with a comma (,)

Local Web UI Access

☒ Deny All

☐ Allow All LAN

☐ Allow the following IPs

Separate each IP with a comma (,)

Local Web UI Port Number

As an Enterprise Administrator, you can configure Port Forwarding and 1:1 NAT firewall rules individually for each Edge by following the instructions on this page.

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.

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 enable 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 SD-WAN Edge.

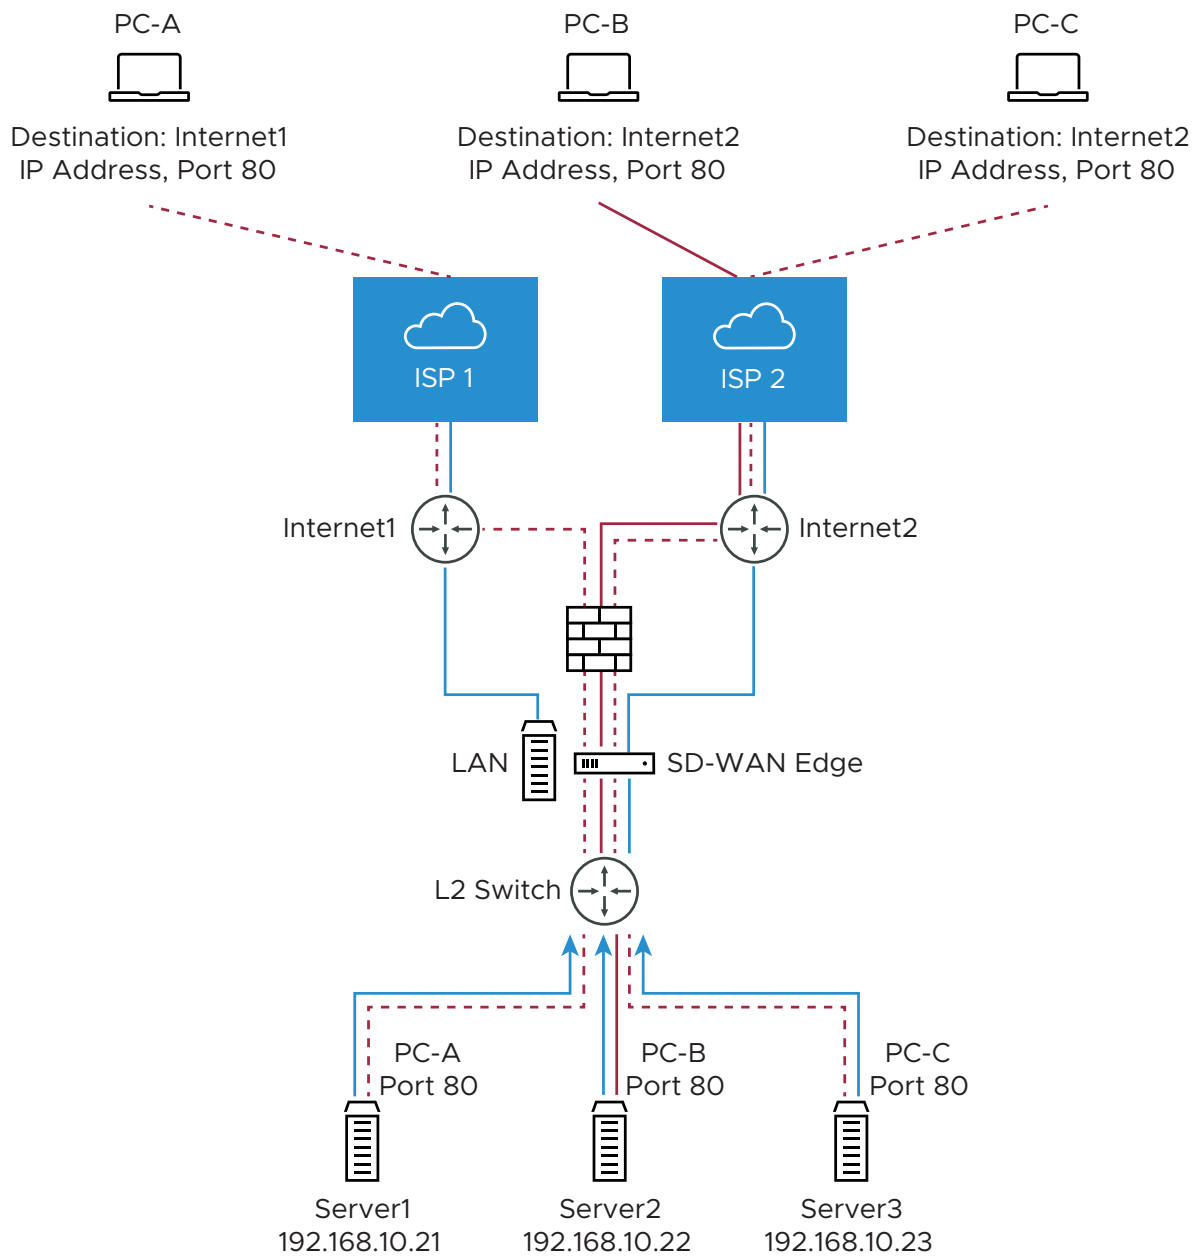
To configure a port forwarding rule, provide the following details.

- 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 IP address using which the host (application) can be accessed from the outside network.
- 5 In the WAN Ports text box, enter one WAN port or range of ports separated with a dash (-), for example 20-25.
- 6 In the **LAN IP** and **LAN Port** text boxes, enter the IP 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.

Port Forwarding Rules ⓘ

Port Forward Rule								Allowed Traffic Source	
Name	Protocol ⓘ	Interface ⓘ	Outside IP	WAN Port(s) ⓘ	LAN IP ⓘ	LAN Port ⓘ	Segment ⓘ	Remote IP/Subnet ⓘ	
Server1	TCP	GE4	30.0.1.2	80	192.168.10.21	80	Global Segment	Ex: 48.2.142.143/24	
Server2	TCP	GE5	30.0.2.2	80	192.168.10.22	80	Global Segment	Ex: 48.2.142.143/24	
Server3	TCP	GE5	30.0.2.3	80	192.168.10.23	80	Global Segment	Ex: 48.2.142.143/24	

The following figure illustrates the port forwarding configuration.



1:1 NAT Settings

These are used to map an Outside IP address supported by the SD-WAN 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 SD-WAN 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.

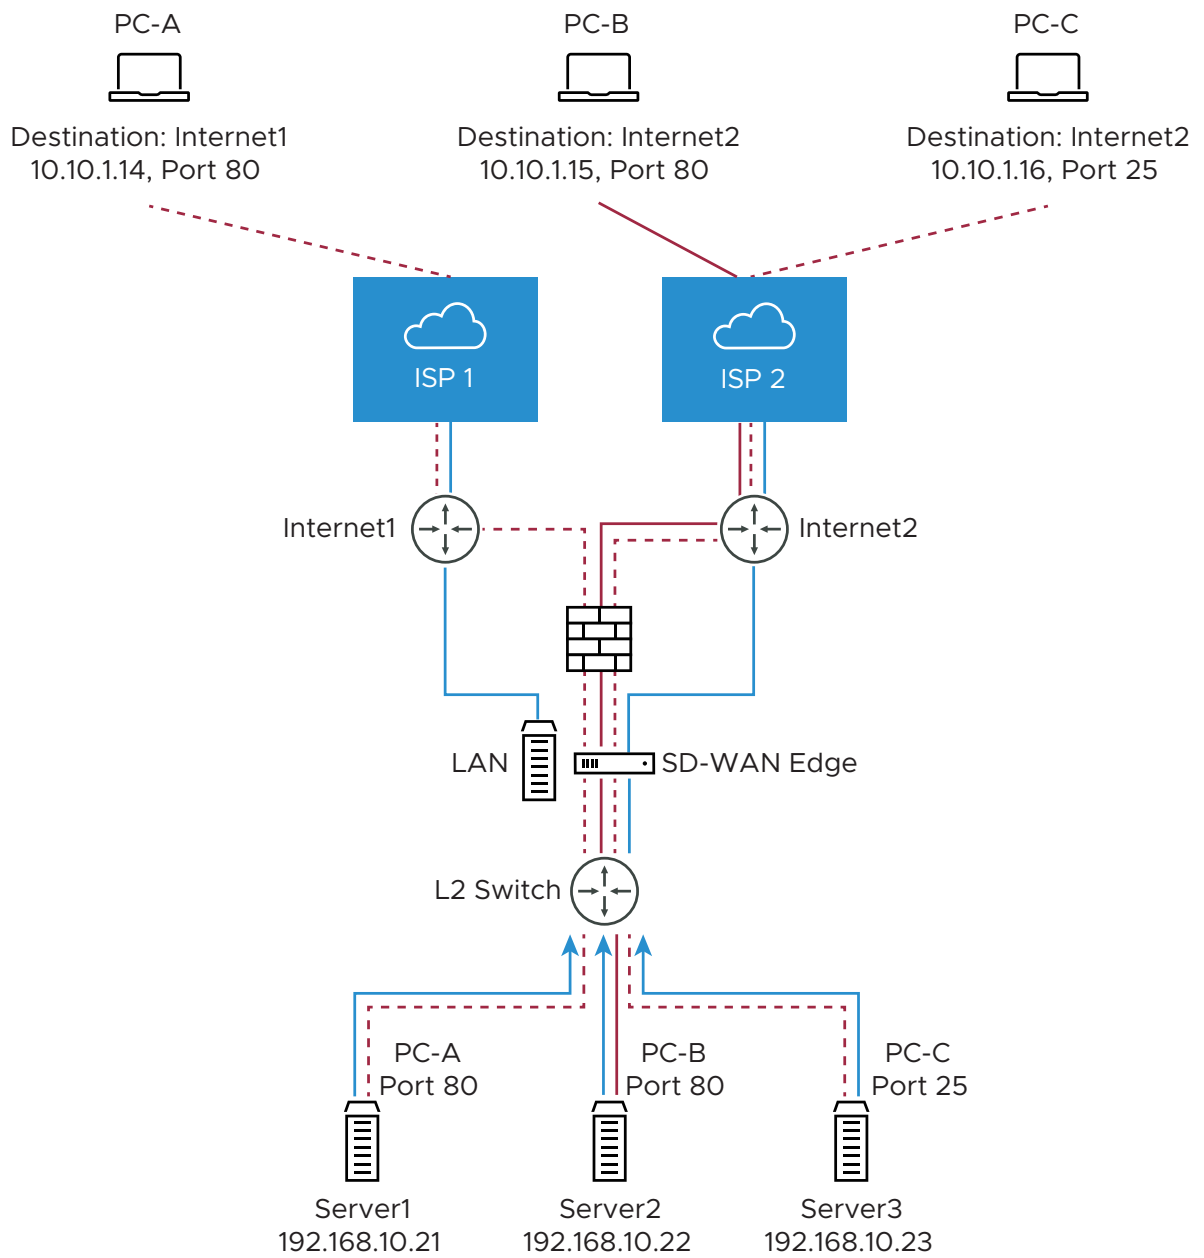
To configure a 1:1 NAT rule, provide the following details.

- 1 In the **Name** text box, enter a name for the rule.
- 2 In the **Outside IP** text box, enter the IP 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 IP (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 the mapping in the respective fields.

1:1 NAT Rules

1:1 NAT Rule						Allowed Traffic Source				
Name	★ Outside IP	★ Interface	★ Inside (LAN) IP	★ Segment ⓘ	Outbound Traffic ⓘ	Protocol	Port(s) ⓘ	Remote IP/Subnet ⓘ		Log ⓘ
Server1	10.10.1.14	GE4	192.168.10.21	Global Segment	<input checked="" type="checkbox"/>	TCP	80	Ex: 46.2.142.142/24	<input type="checkbox"/>	
Server2	10.10.1.15	GE5	192.168.10.22	Global Segment	<input checked="" type="checkbox"/>	TCP	80	Ex: 46.2.142.142/24	<input type="checkbox"/>	
Server3	10.10.1.16	GE5	192.168.10.23	Global Segment	<input checked="" type="checkbox"/>	TCP	25	Ex: 46.2.142.142/24	<input type="checkbox"/>	

The following figure illustrates the 1:1 NAT configuration.



Configure Edge Overrides

Optionally, at the Edge level, you can override the inherited profile firewall rules. To override firewall rules at the Edge level, click **New Rule** under **Firewall Rules**, and follow the steps in [Configure Firewall Rules](#). The override rules will appear in the **Edge Overrides** area. 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.

Override Stateful Firewall Settings

Optionally, at the Edge level, you can override the Stateful Firewall settings by selecting the **Enable Edge Override** check-box in the **Stateful Firewall Settings** area. For more information about Stateful Firewall settings, see [Configure Stateful Firewall Settings](#).

Override Network and Flood Protection Settings

Optionally, at the Edge level, you can override the network and flood protection settings by selecting the **Enable Edge Override** check-box in the **Network and Flood Protection Settings** area. For more information about network and flood protection settings, see [Configure Network and Flood Protection Settings](#).

Override Edge Access Configuration Settings

Optionally, at the Edge level, you can also override the Edge access configuration by selecting the **Enable Edge Override** check-box in the **Edge Access** area. For more information about Edge access configuration, see [Configure Edge Access](#).

Related Links

- [Configure Firewall for Profiles](#)
- [Configure Syslog Settings for Edges](#)
- [Troubleshooting Firewall](#)

Configure Firewall Rules

SD-WAN Orchestrator allows you to configure Firewall rules at the Profile and Edge levels to allow, drop, reject, or skip inbound and outbound traffic. If stateful firewall feature is enabled, the firewall rule will be validated to filter both inbound and outbound traffic. With stateless firewall, you can control to filter only outbound traffic. The firewall rule matches parameters such as IP addresses, ports, VLAN IDs, interfaces, MAC addresses, domain names, protocols, 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.

To configure a firewall rule with stateful firewall-enabled at the profile level, perform the steps on this procedure.

Procedure

- 1 From the SD-WAN Orchestrator, go to **Configure > Profiles > Firewall**.
- 2 Enable **Stateful Firewall** for the selected profile.
- 3 Under **Firewall Rules** area, click **New Rule**. The **Configure Rule** dialog box appears.

Configure Rule

Rule Name:

Match

Source: **Any** **Object Group** **Define...**

Address Group: **?**

Port Group:

Destination: **Any** **Object Group** **Define...**

☐ None
☐ VLAN
☐ Interface
☒ IP Address

Subnet mask:

Domain Name: **?**

Protocol:

Ports:

Application: **Any** **Define...**

Action

Firewall: **Allow** **Drop**

Log: ☐

Audit Comment

[Audit History](#)

OK **Cancel**

- 4 In the **Rule Name** box, enter a unique name for the rule.

5 Under the **Match** area, configure the match conditions for the rule:

Settings	Description
Source	<p>Allows to specify the source for packets. Select any of the following options:</p> <ul style="list-style-type: none"> ■ Any - Allows all source addresses by default. ■ Object Group - Allows you to select a combination of address group and port group. For more information, see Chapter 21 Object Groups and Configure Firewall Rules with Object Groups. <hr/> <p>Note If the selected address group contains any domain names, then they would be ignored when matching for the source.</p> <hr/> <ul style="list-style-type: none"> ■ Define - Allows you to define the source traffic to a specific VLAN, Interface, IP Address, MAC Address, or Port. <p>For IP Address, choose one of the three options:</p> <ul style="list-style-type: none"> ■ 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. <hr/> <p>Note If an interface cannot be selected, then the interface is either not enabled or not assigned to this segment.</p>
Destination	<p>Allows to specify the destination for packets. Select any of the following options:</p> <ul style="list-style-type: none"> ■ Any - Allows all destination addresses by default. ■ Object Group - Allows you to select a combination of address group and port group. For more information, see Chapter 21 Object Groups and Configure Firewall Rules with Object Groups.

Settings	Description
	<ul style="list-style-type: none"> ■ Define - Allows you to define the destination traffic to a specific VLAN, Interface, IP Address, Domain Name, Protocol, or Port. For IP address, choose one of the three options: CIDR prefix, Subnet mask, or Wildcard mask. <p>If an interface cannot be selected, then the interface is either not enabled or not assigned to this segment.</p> <p>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 \"www.salesforce.com\".</p>
Application	<p>Select any of the following options:</p> <ul style="list-style-type: none"> ■ 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. <p>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. Generally 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.</p>

6 Under the **Action** area, configure the actions for the rule:

Settings	Description
Firewall	<p>Select any of the following action the firewall should perform on packets, when the conditions of the rule are met:</p> <ul style="list-style-type: none"> ■ 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.
Log	<p>Select this checkbox if you want a log entry to be created when this rule is triggered.</p>

- 7 While creating or updating a Firewall rule, you can add audit comments using the **Audit Comment** textbox. A maximum of 50 characters is allowed and you can add any number of comments for the same rule.
- 8 Click the **Audit History** button to view all the audit comments added for the rule. You can search for a specific comment by entering the search text in the **Search** field.

Timestamp	Audit Comment	Administrator
2021-04-08 16:32:37	firewall version2	super@velocloud.net
2021-04-08 16:24:13	firewall version1	super@velocloud.net

- 9 Click **OK**.

Results

A firewall rule is created for the selected profile and it appears under the **Firewall Rules** area of the **Profile Firewall** page.

Configure Stateful Firewall Settings

SD-WAN 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.

To configure Stateful Firewall settings at the profile level, perform the following steps.

Procedure

- 1 From the SD-WAN Orchestrator, go to **Configure > Profiles > Firewall**.
- 2 Enable **Stateful Firewall** for the selected profile.
- 3 Under **Stateful Firewall Settings** area, configure the following settings:

Stateful Firewall Settings

Established TCP Flow Timeout (seconds):	<input type="text" value="7440"/>	Non Established TCP Flow Timeout (seconds):	<input type="text" value="300"/>
UDP Flow Timeout (seconds):	<input type="text" value="300"/>	Other Flow Timeout (seconds):	<input type="text" value="60"/>

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 through 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 through 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 through 15999999 seconds. The default value is 300 seconds.
Other Flow Timeout (seconds)	Sets the inactivity timeout period (in seconds) for other flows such as ICMP, after which they are no longer valid. The allowable value ranges from 60 seconds through 15999999 seconds. The default value is 60 seconds.

Note The configured timeout values apply only when the memory usage 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 and Flood Protection Settings

VMware SD-WAN provides detection and protection against various attacks to combat exploits at all stages of their execution.

To secure all connection attempts in an Enterprise network, VMware SD-WAN Orchestrator allows you to configure Network and Flood Protection settings at the Profile and Edge levels, to protect against the following types of attacks:

- 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 and IP Insecure Options

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 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:

- Packet that has no flags set in its TCP header such as SYN, FIN, ACK, etc.,
- TCP header that has SYN and FIN flags combined, which are mutually-exclusive flags 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 generally small used for checking 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 system crash or freeze or reboot, as they cannot handle such huge packets.

ICMP Fragment

An ICMP Fragmentation attack is a common DoS attack which involves the flooding of fraudulent ICMP fragments that cannot be defragmented on the target server. As defragmentation can only take place 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

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 crash if not handled properly on the end device. A cautious stance would be to block such IP packets from entering the protected network.

IP Insecure Options


Attackers sometimes configure IP option fields within an IP packet incorrectly, producing either incomplete or malformed fields. Attackers use these malformed packets to compromise vulnerable hosts on the network. Exploitation of the vulnerability may potentially allow for arbitrary code execution. The vulnerability may be exploited after processing a packet containing a specific crafted IP option in the packet's IP header. Enabling IP Insecure Options protection blocks transit IP packets with incorrectly formatted IP option field in the IP packet header.

Configure Network and Flood Protection Settings

To configure Network and Flood Protection settings at the profile level, perform the following steps.

Procedure

- 1 From the SD-WAN Orchestrator, go to **Configure > Profiles > Firewall**.
- 2 Enable **Stateful Firewall** for the selected profile.
- 3 Under **Network & Flood Protection Settings** area, configure the following settings:



Network & Flood Protection Settings

New Connection Threshold (connections per second) %

Denylist ☒

Detect Duration (seconds)

Denylist Duration (seconds)

TCP Based Attacks		ICMP Based Attacks		IP Based Attacks	
Invalid TCP Flags	<input type="checkbox"/>	ICMP Ping of Death	<input checked="" type="checkbox"/>	IP Unknown Protocol	<input type="checkbox"/>
TCP Land	<input type="checkbox"/>	ICMP Fragment	<input type="checkbox"/>	IP Options	<input type="checkbox"/>
TCP SYN Fragment	<input type="checkbox"/>				

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 percentage through 100 percentage. The default value is 25 percentage.
Denylist	<p>Enable the checkbox to block a source IP address, which is violating the new connection threshold by sending flood traffic either due to misconfiguration of network or malicious user attacks.</p> <p>Note The New Connection Threshold (connections per second) settings will not work unless Denylist is enabled.</p>
Detect Duration (seconds)	<p>Before blocking a Source IP address, it is the grace time duration for which the violating source IP is allowed to send traffic flows.</p> <p>If a host sends flood traffic of new connection requests (port scan, TCP SYN flood, etc..) exceeding the maximum allowed connection per second (CPS) for this duration, it will be considered as eligible for denylisting instead of immediately denylisting it as soon as it exceeds the CPS per source once. For example, consider that the maximum allowed CPS is 10 with detect duration of 10 seconds, if the host floods new connection requests greater than 100 requests for 10 seconds, then the host will be denylisted.</p> <p>The allowable value ranges from 10 seconds through 100 seconds. The default value is 10 seconds.</p>
Denylist Duration (seconds)	The time duration for which the violated source IP is blocked from sending any packets. The allowable value ranges from 10 seconds through 86400 seconds. The default value is 10 seconds.
TCP Based Attacks	<p>Supports protection from the following TCP-based attacks by enabling the respective checkboxes:</p> <ul style="list-style-type: none"> ■ Invalid TCP Flags ■ TCP Land ■ TCP SYN Fragment
ICMP Based Attacks	<p>Supports protection from the following ICMP-based attacks by enabling the respective checkboxes:</p> <ul style="list-style-type: none"> ■ ICMP Ping of Death ■ ICMP Fragment
IP Based Attacks	<p>Supports protection from the following IP-based attacks by enabling the respective checkboxes:</p> <ul style="list-style-type: none"> ■ IP Unknown Protocol ■ IP Insecure Options

Optionally, you can also override the Network and Flood Protection settings at the Edge level. For more information, see [Configure Netflow Settings for Edges](#).

Configure Edge Access

When configuring 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.

To configure Edge access for profiles, perform the following steps:

Procedure

- 1 From the SD-WAN Orchestrator, go to **Configure > Profiles > Firewall**. The **Firewall** page appears.

Edge Access

Support Access

- ☒ Deny All
☐ Allow the following IPs

Ex: 54.183.9.192,
2001:0db8:85a3:0000:0000:8a2e:0370:7334

Separate each IP with a comma (,)

Console Access ⓘ

- ☒ Deny
☐ Allow

USB Port Access ⓘ

- ☐ Deny (Only applicable for edge models 510 and 6X0)
☒ Allow

SNMP Access ⓘ

- ☒ Deny All
☐ Allow All LAN
☐ Allow the following IPs

Ex: 54.183.9.192, 46.2.142.142

Separate each IP with a comma (,)

Local Web UI Access ⓘ

- ☒ Deny All
☐ Allow All LAN
☐ Allow the following IPs

Ex: 54.183.9.192, 46.2.142.142

Separate each IP with a comma (,)

Local Web UI Port Number

80

2 Under **Edge Access** area, you can configure device access using the following options:

Field	Description
Support Access	<p>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 (,).</p> <p>By default, Deny All is selected.</p>
Console Access	<p>Select Allow to enable 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.</p> <p>Note Whenever the console access setting is changed from Allow to Deny or vice-versa, the Edge must be rebooted manually.</p>
USB Port Access	<p>Select Allow to enable and select Deny to deactivate the USB port access on Edges.</p> <p>This option is available only for Edge models 510 and 6x0.</p> <p>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 SD-WAN Orchestrator. For instructions, refer to Remote Actions.</p>
SNMP Access	<p>Allows Edge access from routed interfaces/WAN through SNMP. Select one of the following options:</p> <ul style="list-style-type: none"> ■ 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. The IP addresses must be separated by comma (,).
Local Web UI Access	<p>Allows Edge access from routed interfaces/WAN through a Local Web UI. Select one of the following options:</p> <ul style="list-style-type: none"> ■ 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. The IP addresses must be separated by comma (,).
Local Web UI Port Number	<p>Enter the port number of the local Web UI from where you can access the Edge.</p>

3 Click **Save Changes**.

What to do next

If you want to override the Edge access settings for a specific Edge, use **Enable Edge Override** option available on the **Edge Firewall** page. For related information, see [Configure Firewall for Edges](#)

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 enabled, you can use the following remote diagnostic tests to obtain firewall diagnostic information:

- **Flush Firewall Sessions** - Run this test to reset established sessions from the 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 limit the number of sessions returned by using filters: source and destination IP address, source and destination port, and Segment.

List Active Firewall Sessions

List active sessions in the firewall. Use source and destination IP address filters to view the exact sessions you want to see. This output is limited to a maximum of 1000 sessions.

Run

Segment:

Max Flows:

Source IP/Port:

Destination IP/Port:

Test Duration: 5.002 seconds

Segment	Src IP	Dst IP	Protocol	Src Port	Dst Port	Application	Firewall Policy	TCP State	Bytes Sent	Bytes
Global Segment	10.5.1.25	10.2.1.25	ICMP	N/A	N/A	icmp	AllowAny	N/A	672	672
Global Segment	10.5.1.25	10.5.1.25	TCP	36720	22	ssh	AllowAny	ESTABLISHED	3441	4153

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.

The Remote Diagnostics output displays the following information: Segment name, Source IP, Source Port, Destination IP, Destination Port, Protocol, Application, Firewall Policy, current TCP state of any flows, Bytes Received/Sent, and Duration. There are 11 distinct TCP states as defined in RFC 793:

- **LISTEN** - represents waiting for a connection request from any remote TCP and port. (This state is not shown in a Remote Diagnostic output).
- **SYN-SENT** - represents waiting for a matching connection request after having sent a connection request.
- **SYN-RECEIVED** - represents waiting for a confirming connection request acknowledgment after having both received and sent a connection request.

- ESTABLISHED - represents an open connection, data received can be delivered to the user. The normal state for the data transfer phase of the connection.
- FIN-WAIT-1 - represents waiting for a connection termination request from the remote TCP, or an acknowledgment of the connection termination request previously sent.
- FIN-WAIT-2 - represents waiting for a connection termination request from the remote TCP.
- CLOSE-WAIT - represents waiting for a connection termination request from the local user.
- CLOSING - represents waiting for a connection termination request acknowledgment from the remote TCP.
- LAST-ACK - represents waiting for an acknowledgment of the connection termination request previously sent to the remote TCP (which includes an acknowledgment of its connection termination request).
- TIME-WAIT - represents waiting for enough time to pass to be sure the remote TCP received the acknowledgment of its connection termination request.
- CLOSED - represents no connection state at all.

For more information about how to run remote diagnostics on an Edge, see [Remote Diagnostics](#).

Provision an Edge

16

This section describes how to provision an Edge.

Read the following topics next:

- [Provision a New Edge](#)
- [Provision a New Edge with Analytics](#)
- [Activate SD-WAN Edges](#)
- [Manage Edges](#)
- [Manage Edges with New Orchestrator UI](#)

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.

Procedure

- 1 In the Enterprise portal, click **Configure > Edges**.

- In the **Edges** screen, click **New Edge** at the top-right corner of the screen.

The **Provision New Edge** dialog box appears.



The **Provision New Edge** dialog box contains the following fields and options:

- Name:** Textbox with value "veloedge".
- Model:** Drop-down menu with value "Edge 510".
- Profile:** Drop-down menu with value "Quick Start Profile".
- Authentication:** Drop-down menu with value "Certificate Acquire".
- Edge License:** Drop-down menu with value "ENTERPRISE | 1 Gbps | Asia Paci".
- Custom Info:** Empty textbox.
- High Availability:** Checkbox (unchecked).
- Serial Number:** Textbox with value "Ex: VC00000490". Below it, a note states: "When specified, the Edge must present this serial number on activation."
- Contact Name:** Textbox with value "Super User".
- Contact Email:** Textbox with value "super@velocloud.net".
- Location:** Information icon and a link "Set Location..."

At the bottom right, there are **Create** and **Cancel** buttons.

- In the **Name** textbox, enter a unique name for the Edge.
- From the **Model** drop-down menu, select an Edge model.
- From the **Profile** drop-down menu, select a profile to be assigned to the Edge.

Note If an Edge Staging Profile is displayed as an option due to Zero Touch Provisioning, this profile is used by a newly assigned Edge, but has not been configured with a production Profile.

For information about how to create a new profile, see [Create a Profile](#).

- From the **Authentication** drop-down menu, you can select one of the following certificates-based authentication options:
 - **Certificate Not Required** - Edge uses a pre-shared key mode of authentication.
 - **Certificate Acquire** - This option is selected by default and instructs the Edge to acquire a certificate from the certificate authority of the 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**.

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.

- Certificate Required - Edge uses the PKI certificate.
- 7 From the **Edge License** drop-down menu, select an Edge License from the available list. The list displays the licenses assigned to the Enterprise, by the Operator.
 - 8 In the **Custom Info** textbox, enter custom information associated with the edge
Customer information must not exceed 255 characters.

Note Super User and Standard Admin users of Enterprise/MSP/Operator roles (with UPDATE_EDGE privilege) can add or update the Custom Info for an edge.

- 9 To apply High Availability (HA), select the **High Availability** checkbox. (Edges can be installed as a single standalone device or paired with another Edge to provide High Availability (HA) support. For more information about HA, see the [High Availability Deployment Models](#) section).
- 10 In the **Serial Number** textbox, enter the serial number of the Edge . If specified, the serial number must match the serial number of the Edge that will be activated.
- 11 In the **Contact Name** and **Contact Email** textboxes, enter the name and email address of the site contact for the Edge.
- 12 Click the **Set Location** link to set the location of the Edge.
- 13 Click **Create**.

Results

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. The **LTE Modem Information** diagnostic test will retrieve diagnostic information, such as signal strength, connection information, and so on. For information on how to run a diagnostic test, see sections titled, [Remote Diagnostics](#) and [Performing Remote Diagnostics Tests](#).

Note For Enterprise customers with Analytics enabled, you can provision an Analytics Edge by following the steps in [Provision a New Edge with Analytics](#).

What to do next

- To manage the provisioned edges, see [Manage Edges](#).
- To view Edge details or to make any changes to edge, see [Chapter 17 View or Modify Edge Information](#).

- To configure an Edge, see [Chapter 18 Edge Device Configurations](#).

Provision a New Edge with Analytics

Analytics functionality is built natively into the VMware SD-WAN Edge and enables collecting data inline. However, by default, Analytics is deactivated for Edges. For Enterprise customers with Analytics enabled, allows the Enterprise Administrators to create Edges with Analytics enabled.

To create a new SD-WAN Edge with Analytics, perform the following steps.

Prerequisites

- Ensure that all the necessary system properties to enable Analytics are properly set in the SD-WAN Orchestrator. For more information, contact your Operator Super User.
- Ensure that the Analytics functionality is enabled for the Customer before provisioning an Analytics Edge.

Note For more information, see *VMware Edge Network Intelligence Configuration Guide* available at <https://docs.vmware.com/en/VMware-SD-WAN/index.html>.

Procedure

- 1 In the Enterprise portal, navigate to **Manage Customers**.
- 2 Select a customer and then go to **Configure > Edges**.

The **Edges** screen appears.

- 3 Click **New Edge** at the top-right corner of the screen.

The **Provision New Edge** dialog box appears.

Provision New Edge

* Name: Veloedge

* Model: Edge 6X0

* Analytics: Application Analytics
12 out of 12 analytics licences available.

* Profile: Quick Start Profile

Authentication: Certificate Acquire

Edge License: ENTERPRISE | 1 Gbps | Asia Paci

Custom Info:

High Availability: ☐

Serial Number: Ex: VC00000490
When specified, the Edge must present this serial number on activation.

* Contact Name: Super User

* Contact Email: super@velocloud.net

Location:

Create Cancel

- 4 In the **Name** textbox, enter a unique name for the Edge.
- 5 From the **Model** drop-down menu, select an Edge model.
- 6 From the **Analytics** drop-down menu, select one of the following Analytics modes to be configured for the Edge:
 - Application Analytics - Gains access to fault isolation and Application-specific Analytics.
 - Application and Branch Analytics - Gains access to Application-specific Analytics and Branch Analytics.
 - By default, **None** is selected, which implies Analytics is deactivated for the Edge.

Under the **Analytics** drop-down menu, you can find the remaining number of Analytics licenses that is available to be provisioned as an Analytics Edge. As an Administrator, you can also change the Analytics mode for a specific Edge from the **Edge Overview** screen.

- 7 From the **Profile** drop-down menu, select a profile to be assigned to the Edge.
- 8 From the **Edge License** drop-down menu, select an Edge License from the available list. The list displays the licenses assigned to the Enterprise, by the Operator.
- 9 From the **Authentication** drop-down menu, you can select one of the following certificate-based authentication options:
 - Certificate Not Required - Edge uses a pre-shared key mode of authentication.

- **Certificate Acquire** - This option is selected by default, and instructs the Edge to acquire a certificate from the certificate authority of the SD-WAN 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 SD-WAN Orchestrator and for establishment of VCMP tunnels.

Note After acquiring the certificate, the option can be updated to **Certificate Required**.

- **Certificate Required** - Edge uses the PKI certificate.

- 10 In the **Custom Info** textbox, enter custom information associated with the Edge, if needed. Customer information should not exceed 255 characters.

Note Super User and Standard Admin users of Enterprise/MSP/Operator roles (with UPDATE_EDGE privilege) can add or update the Custom Info for an Edge.

- 11 To apply High Availability (HA), select the **High Availability** checkbox.
- 12 In the **Serial Number** textbox, enter the serial number of the Edge, which is optional. If specified, the serial number must match the serial number of the Edge when activated.
- 13 In the **Contact Name** and **Contact Email** textboxes, enter the name and email address of the site contact for the Edge.
- 14 Click the **Set Location** link to set the location of the Edge.
- 15 Click **Create**.

Results

An Analytic Edge is provisioned for the selected customer. Once the Edge is provisioned, the Analytics functionality collects data, performs deep packet inspection of all traffic, identifies network application and correlates traffic with user information.

What to do next

To send the collected analytics data to the Cloud Analytics Engine, you must configure an Analytics interface on which the Edge transmits Analytics data. For more information, see [Configure an Analytics Interface on an Edge](#).

Enable Analytics for an Existing Edge

VMware SD-WAN Orchestrator allows the Administrator (Enterprise or Partner) to enable Analytics on an existing SD-WAN Edge.

To enable Analytics on an existing SD-WAN Edge, perform the following steps.

Prerequisites

- Ensure that all the necessary system properties to enable Analytics are properly set in the SD-WAN Orchestrator. For more information, contact your Operator Super User.

- Ensure that the Analytics functionality is enabled for the Customer associated with the Edge.

Procedure

- 1 In the Enterprise portal, navigate to **Manage Customers**.
- 2 Select a customer and then go to **Configure > Edges**.

The **Edges** screen appears.

- 3 Click the Edge name to enable Analytics.

- 4 In **Edge Overview** tab, from the **Analytics** drop-down menu, select one of the following Analytics modes for the Edge:
 - Application Analytics - Gains access to fault isolation and Application-specific Analytics.
 - Application and Branch Analytics - Gains access to Application-specific Analytics and Branch Analytics.
 - By default, **None** is selected.
- 5 Click **Save Changes**.

Results

An Analytic Edge is provisioned for the selected customer. Once the Edge is provisioned, the Analytics functionality collects data, performs deep packet inspection of all traffic, identifies network application and correlates traffic with user information.

What to do next

To send the collected analytics data to the Cloud Analytics Engine, you must configure an Analytics interface on which the Edge transmits Analytics data.

- [Configure an Analytics Interface on an Edge](#)
- [Configure Analytics Endpoint Settings](#)

Configure an Analytics Interface on an Edge

Analytics Interface specifies the interface and interface IP that an Edge uses for SNMP polling, receiving AMON, traps, and so on. Once you have provisioned an Analytics Edge, you can override the default Analytics interface on the Global segment for the Edge to ingest data such as SNMP, AMON, traps, and syslog by selecting the **Analytics Enabled** checkbox under **Analytics Interface** in the Device Setting page of the Edge.

To configure an Analytics interface on an SD-WAN Edge, perform the following steps:

Procedure

- 1 From the SD-WAN Orchestrator, go to **Configure > Edges**.

The **Edges** page appears.

- 2 Select an Edge for which you want to configure an Analytics interface and click the icon under the **Device** column.

The Device Setting page for the selected Edge appears.

- 3 From the **Configure Segment** drop-down menu, select Global segment to configure an Analytics interface.

Note Currently, source interface and Analytics-enabled flag are only supported for the Global segment. Settings for non-global segments are ignored even if set.

- 4 Go to the **Analytics Interface** area and select the **Analytics Enabled** checkbox if you want to override the default Analytics interface on the Global segment for the Edge.

Analytics Interface



Analytics Enabled



Source Interface

GE2



- 5 From the **Source Interface** drop-down menu, select an Analytics interface for the Edge to ingest data.

The Edge automatically selects an interface with 'Advertise' field set as the source interface, if the **Analytics Enabled** checkbox is not selected or the **Analytics Enabled** checkbox is selected and Source Interface is set to **None**.

- 6 Click **Save Changes**.

What to do next

- You can change the Analytics Endpoint settings at the Edge-level. For steps, see [Configure Analytics Endpoint Settings](#).
- To view the Analytics data, see [View Analytics Data](#).

Configure Analytics Endpoint Settings

At the Edge-level, an Enterprise or Partner Administrator can configure the Analytics endpoint settings to either Dynamic IP address or Static IP address for a specific Analytics Edge. By default, the Analytics endpoint is set to Dynamic IP address.

For Dynamic IP Analytics endpoint setting, ensure to allow this URL (loupe-m.nyansa.com). If you require Static IP address to open the firewall to allow communication between an Analytics Edge and Cloud Analytics Engine, the Analytics endpoint setting should be set to Static IP address. For Static IP Analytics endpoint setting, ensure to allow this URL (loupe-m2.nyansa.com).

To configure the Analytics endpoint settings to Static IP address, perform the following steps.

Procedure

- 1 From the SD-WAN Orchestrator, go to **Configure > Edges**.

The **Edges** page appears.

- 2 Select an Analytics Edge to configure Analytics endpoint settings and click the icon under the **Device** column.

The Device Setting page for the selected Edge appears.

- 3 From the **Configure Segment** drop-down menu, select a profile segment to configure Analytics settings.
- 4 Go to the **Analytics Settings** area and from the **Analytics Endpoint** drop-down menu, select **Static IP** as the Analytics endpoint for the selected Analytics Edge.

Analytics Settings ⓘ

Analytics Endpoint

Dynamic IP

☒ Enable Edge Override ⓘ

- 5 Click **Save Changes**.

What to do next

- To view the Analytics data, see [View Analytics Data](#).

Activate SD-WAN Edges

You can deploy and activate SD-WAN Edges using the following two methods:

- Zero Touch Provisioning—In this method, you must power-on the Edges and connect them to the internet to activate the Edges. For more information, refer to [Activate SD-WAN Edges Using Zero Touch Provisioning](#).
- 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 more information, refer to [Activate SD-WAN Edges Using Email](#).

Following table shows a comparison of activities that are allowed in each of the activation methods:

Activity	Zero Touch Provisioning (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		✓

Activate SD-WAN Edges Using Zero Touch Provisioning

Zero Touch Provisioning allows you to activate Edges by powering on the Edges and connecting them to the Internet.

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

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

Note

- Zero Touch Provisioning supports Edge models: 510, 510 LTE, 6x0, and 3xx0.
- For Zero Touch Provisioning push 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 Zero Touch Provisioning:

- [Sign-Up for Zero Touch Provisioning](#)
- [Assign Profile and License to Edges](#)
- [Assign Inventory to an Edge](#)

Sign-Up for Zero Touch Provisioning

To sign-up for Zero Touch Provisioning:

Prerequisites

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 [VMware Customer Support](#). Outbound internet connectivity via DHCP is required to complete the push activation successfully.

Procedure

- 1 Log in to SD-WAN Orchestrator, and then go to **Administration > System Settings > General Information**.
- 2 Scroll down to the **Zero Touch Provisioning Sign Up** area, and then in the **SID** field, enter the Subscription Identifier.
- 3 Click **Submit**.

Results

You can view the Edge inventory in the **Pending Assignment** tab only after the successful validation of SID. The validation process may take up to a maximum of 1 week. To view the Edge inventory, go to **Administration > Zero Touch Provisioning > Pending Assignment**.

Note Only the Edges that were shipped to you after the successful completion of the sign-up process appear in the **Pending Assignment** tab. Ensure that the SID assigned to you is used in all your future orders so that the inventory is reflected correctly.

What to do next

You must assign a profile and a license to the Edges. For instructions, refer to [Assign Profile and License to Edges](#).

Assign Profile and License to Edges

To assign profile and license to the Edges:

Prerequisites

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

Procedure

- 1 Log in to SD-WAN Orchestrator, and then go to **Administration > Zero Touch Provisioning > Pending Assignment**. A list of Edge inventory with Serial number and Model appears.

- 2 Select all the Edges for which you want to assign a profile and license, and then click **Actions > Assign....** The **Edge Inventory Assignment** modal popup appears.

Edge Inventory Assignment

Profile: Quick Start Profile

Edge License: ENTERPRISE | 1 Gbps | 12 Months

	Serial Number	Model	Profile	Edge License
1	VCE08400010418	Edge 840	Quick Start Profile	ENTERPRISE 1 Gbps
2	VCE08400010422	Edge 840	Quick Start Profile	ENTERPRISE 1 Gbps
3	VCE08400010421	Edge 840	Quick Start Profile	ENTERPRISE 1 Gbps

OK Cancel

- 3 From the **Profile** and **Edge License** drop-down lists, select the required profile and license that you want 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.

- 4 Click **OK**.

Results

The Edges for which you have assigned a profile and license appears in the **Assigned** tab. The **Inventory State** for the assigned Edges will be **Assigned to Customer** and the **Edge State** will be **Pending**.

What to do next

Power-on the assigned physical Edges and connect them to the internet so that they are redirected to the SD-WAN Orchestrator where they are automatically activated. After an Edge is activated, the **Edge State** in the **Assigned** tab changes from **Pending** to **Activated**.

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, SD-WAN 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 Zero Touch Provisioning feature works and the physical Edge is activated.

To assign inventory to a logical Edge:

Procedure

- 1 Log in to SD-WAN Orchestrator, and then go to **Configure > Edges**.

- 2 Click **New Edge**. The **Provision New Edge** modal popup appears.
- 3 Enter a name for the Edge, and then select the required Edge model, profile, and license.
- 4 Click **Create**. The newly added logical Edge appears in the Edges table.
- 5 Select the logical Edge entry that you just created, and then click **Actions > Assign Inventory to Edge**.
- 6 From the **Assign Inventory** drop-down list, select the inventory that was originally assigned to the Edge, which you had accidentally deleted.
- 7 Click **OK**.

Activate SD-WAN Edges Using Email

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

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

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

Send an Activation Email

The process of activating the SD-WAN Edge begins with the initiation of an activation procedure email that is sent to the Site Contact by the IT Admin.

To send the activation procedure email:

- 1 Go to **Configure > Edges** from the Orchestrator.
- 2 Select the SD-WAN Edge you want to activate. The **Edge Overview Tab** window appears.
- 3 As an optional step, in the **Properties** area, enter the serial number of the SD-WAN Edge that will be activated in the **Serial Number** text field. Serial numbers are case sensitive, so make sure that “VC” is capitalized.

Note This step is optional. However, if specified, the serial number must match the activated SD-WAN Edge.

- 4 Click the **Send Activation Email** button to send the activation email to the Site Contact.

The screenshot shows the 'Properties' window for an SD-WAN Edge. It contains the following fields and controls:

- Name:** ACME-Mountain View 1
- Description:** (empty text field)
- Status:** Pending
- Serial Number:** UNF4-C4HS-LLKS-R4J8
- Activation Key:** UNF4-C4HS-LLKS-R4J8
- Enable Pre-Notifications:** ☒
- Enable Alerts:** ☒
- Authentication Mode:** Certificate Required (dropdown menu)
- Send Activation Email:** (button)

- 5 The **Send Activation Email** pop-up window appears. It describes the steps for the Site Contact to complete to activate the SD-WAN Edge device.

Send Activation Email

Edge: ACME- Mountain View 1

Recipients: Site Contact

* From: support@velocloud.net

* To: jdoe@acme.com

CC:

* Subject: Edge Activation

* Message Body:

Hi,

To activate your VeloCloud Edge, please follow these steps:

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

http://192.168.2.1/?activation_key=UNF4-C4HS-LLKS-R4J8&custom_vco=34.232.58.228

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

Send Close

Note

- For the SD-WAN Edge 510 LTE device, the Activation Email consists of Cellular Settings like SIM PIN, Network, APN, and User name.
- 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.

- 6 Click the **Send** button to send the activation procedure email to the Site Contact.

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 [Remote Diagnostics](#).

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

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

- 3 Click 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 SD-WAN 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.

Edge Activation using an iOS Device and an Ethernet Cable

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

Prerequisites

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.

Procedure

- 1 Complete the Edge configuration on the Orchestrator software. For details, refer to the *Configure Edge Device* section in the *VMware SD-WAN Administration Guide*.
- 2 Navigate to **Configure > Edges > Edge Overview tab**, and then click the **Send Activation Email** button.

Send Activation Email

Edge	VCE iOS Ethernet
Recipients	Site Contact

* From: no-reply@velocloud.net
 * To:
 CC:
 * Subject: Edge Activation
 * Message Body:

Hi,

To activate your Edge, please follow these steps:

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

http://192.168.2.1/?activation_key=K979-QR34-CFQD-JV6V&custom_vco=vco134-usvi1.velocloud.net

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

- 3 Enter the email address of the person activating the Edge, and then click **Send**.

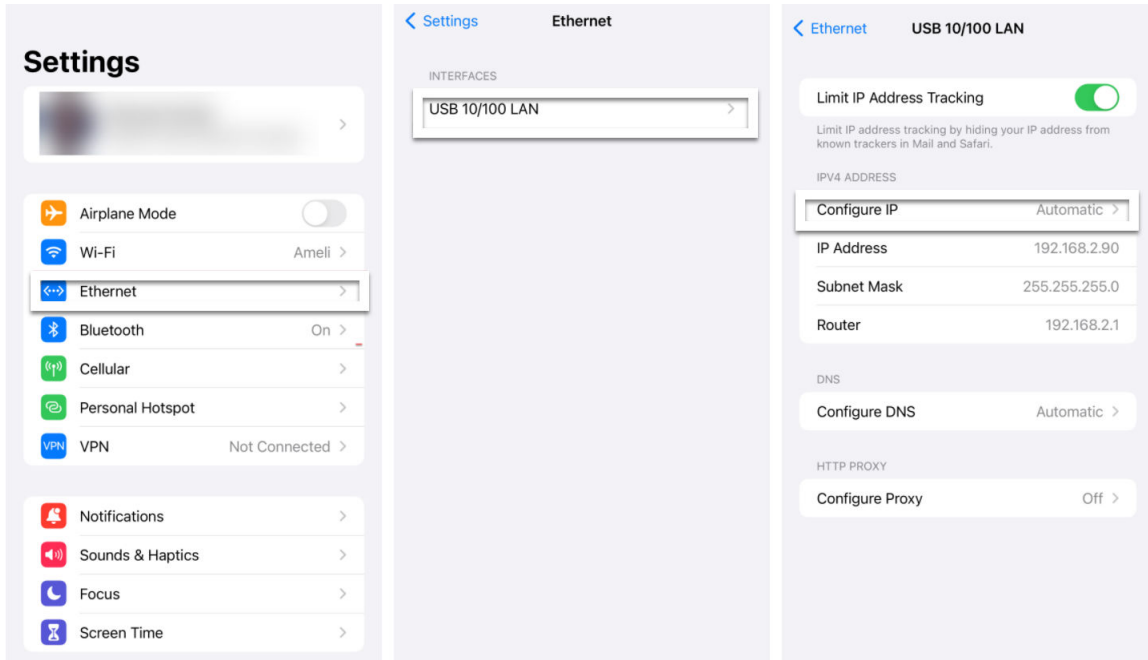
- 4 Power up the Edge, and then connect it to an available internet connection using an Ethernet cable.

Note Refer to [Edge Activation Guides](#) to check details of the model you are installing to determine the correct port.

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

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

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



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

Hi,

To activate your Edge, please follow these steps:

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

http://192.168.2.1/?activation_key=K979-QR34-CFQD-JV6V&custom_vco=vco134-usv1.velocloud.net

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

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

Results

Your Edge device is now activated.

Edge Activation using an Android Device and an Ethernet Cable

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

Prerequisites

The components required for this procedure are:


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

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

Procedure

- 1 Complete the Edge configuration on the Orchestrator software. For details, refer to the *Configure Edge Device* section in the *VMware SD-WAN Administration Guide*.
- 2 Navigate to **Configure > Edges > Edge Overview tab**, and then click the **Send Activation Email** button.

- 3 Enter the email address of the person activating the Edge, and then click **Send**.



The image shows a 'Send Activation Email' dialog box. At the top, it displays 'Edge: Test_VCE' and 'Recipients: Site Contact'. Below this, there are input fields for 'From' (no-reply@sase.vmware.com), 'To' (remote_hands@sp.com), 'CC' (empty), and 'Subject' (Edge Activation). The 'Message Body' section contains a pre-written email template: 'Hi, To activate your Edge, please follow these steps: 1. Connect your device to power and any Internet cables or USB modems. 2. Find and connect to the Wi-Fi network that looks like "velocloud-" followed by 3 more letters/numbers (e.g. "velocloud-01c"), and use "vcsecret" as the password. If your device does not have Wi-Fi, connect to it using an Ethernet cable. 3. Click the following link to activate your edge http://192.168.2.1/?activation_key=8LEC-YTE9-TEXN-UQSQ&custom_vco=vco12-usv1.velocloud.net'. There is a text area below the message body with the text 'If you experience any difficulty, please contact your IT admin.' At the bottom right, there are 'Send' and 'Close' buttons.

- 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 click the activation link displayed at the bottom of the screen to activate your Edge. The following screenshot is an example.

Hi,
 To activate your Edge, please follow these steps:

1. Connect your device to power and any Internet cables or USB modems.
2. Find and connect to the Wi-Fi network that looks like "velocloud-" followed by 3 more letters/numbers (e.g. "velocloud-01c"), and use "vcsecret" as the password. If your device does not have Wi-Fi, connect to it using an Ethernet cable.
3. Click the following link to activate your edge
http://192.168.2.1/?activation_key=9CLJ-GVS4-X3NE-8CMR&custom_vco=vco12-usvi1.velocloud.net

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

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

Results

Your Edge device is now activated.

Request RMA Reactivation

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

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

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

Note RMA reactivation request is allowed only for activated Edges.

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

- [Request RMA Reactivation Using Zero Touch Provisioning](#)
- [Request RMA Reactivation Using Email](#)

Request RMA Reactivation Using Zero Touch Provisioning

To request RMA reactivation using Zero Touch Provisioning:

Procedure

- 1 Log in to SD-WAN Orchestrator, and then go to **Configure > Edges**.

- 2 Click the Edge that you want to replace. The **Edge Overview** page appears.
- 3 Scroll down to the **RMA Reactivation** area, and then click **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, click **Generate New Activation Key**. For details, refer to [RMA Reactivation](#).

- 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 Click **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 **Administration > Zero Touch Provisioning > Assigned**.

- 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.

Results

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

What to do next

Return the old Edge to VMware so that the logical entry for the old Edge with the state **RMA Requested** gets removed from the **Administration > Zero Touch Provisioning > Assigned** page.

Request RMA Reactivation Using Email

To request RMA reactivation using email:

Prerequisites

Procedure

- 1 Log in to SD-WAN Orchestrator, and then go to **Configure > Edges**.
- 2 Click the Edge that you want to replace. The **Edge Overview** page appears.

- 3 Scroll down to the **RMA Reactivation** area, and then click **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, click **Generate New Activation Key**. For details, refer to [RMA Reactivation](#).

- 4 Click **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 SD-WAN 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.

Results

The Edge downloads the configuration and software from the SD-WAN Orchestrator and gets activated.

What to do next

Manage Edges








As an enterprise user, you can manage all the edges provisioned in a network from the **Edges** screen. The **Edges** screen lists all the provisioned edges in a network and also allows you to provision a new edge by clicking the **New Edge** button on the top right-hand corner of the screen. You can also select an edge from here and perform various actions such as change local credential, delete edge, assign profile, assign software image, assign edge license, update alerts and so on using the **Actions** drop-down menu.

Edge	Certificates	Profile	Operator Profile	Device	Biz Pol	Firewall	Operator Ale	Software Ver.	Build No
b1-edge1		Quick Start ...	3.3.2					4.0.0	R400-20
b2-edge1		Quick Start ...	3.3.2					4.0.0	R400-20
b3-edge1		Quick Start ...	3.3.2					4.0.0	R400-20
b4-edge1		Quick Start ...	3.3.2					4.0.0	R400-20
b5-edge1		Quick Start ...	3.3.2					4.0.0	R400-20

Note If you are logged in using a user ID that has Customer Support privileges, you will only be able to view SD-WAN Orchestrator objects. You will not be able to create new objects or configure/update existing ones.

The following table provides details for each field displayed on the **Edges** screen.

Most of the column headers have a sorting feature that lists items in the column in alphabetical order, numerical order, or by type. (The Device, Biz Policy, Firewall, Alerts, and Operator Alerts columns do not have this feature). Click the column headers that have this feature to sort the list.

Option	Description
Edge	Displays the name of the Edge. Click the Edge column header to sort the Edge list in alphabetical order. The Edge name is also a link; click the link to open the Chapter 17 View or Modify Edge Information screen. Select the checkbox next to the name of the Edge to select the Edge.
Certificates	Displays an Edge's current and expired certificates. Click the View link next to the number of certificates for more information.
Profile	Lists the Profile assigned to the Edge. The Profile name is also a link; clicking the link opens the Profile Overview page. NOTE: If an Edge Staging Profile is displayed due to Zero Touch Provisioning, this profile is used by a newly assigned Edge, but has not been configured with a production Profile. Enterprise Admins must manually assign a Profile to these Edges. See section titled, <i>Assign a Profile (Change a Profile)</i> for instructions on how to manually assign a profile to an Edge.
Operator Profile	This column is visible to only Operators. The Operator Profile is the template assigned to the customer the moment the customer is created by the Operators. It includes the software image, application maps, Gateway selection, and the management settings of the Edge. Operator-level Admins can change the Operator Profile for specific Edges. Enterprise Admins have read-only access. The Operator Profile name is also a link; clicking the link opens the <i>Operator Profiles</i> screen.
HA	Selecting the HA checkbox enables the Active Standby HA option.
Device	Displays a blue  icon if Edge specific configurations have been configured. Displays a gray  icon to indicate that all settings (if any) have been inherited from the Profile. To navigate to the Device settings screen, click the icon in the Device column, and then click the Device tab.
Biz Policy	Displays a blue  icon if Business Policy rules have been configured. Displays a gray  icon to indicate that all rules (if any) have been inherited from the Profile. To navigate to the Business Policy screen, click the icon in the Biz Policy column and then click the Business Policy tab.
Firewall	Displays a blue  icon if Firewall rules have been configured. Displays a gray  icon to indicate that all rules (if any) have been inherited from the Profile. Displays a red line across the icon  if the Firewall is deactivated. When the Firewall is deactivated, it indicates that it has been turned off in an Edge's profile configuration. To turn the Firewall on, go the profile configuration (Configure > Profiles > Firewall tab). To navigate to the Firewall screen, click the icon in the Firewall column and then click the Firewall tab.
Alerts	If Customer alerts are enabled for the Edge, the Alerts checkbox will be checked in this column. Click the name of the Edge in the Edge column to open the Chapter 17 View or Modify Edge Information to activate or deactivate Customer alerts.
Operator Alerts	If Operator alerts are enabled for the Edge, the Operator Alerts checkbox will be checked in this column. Click the name of the Edge in the Edge column to open the Chapter 17 View or Modify Edge Information to activate or deactivate Operator alerts.
Software Version	Displays the software version of the Edge.
Factory Software Version	When the Edge is shipped from the factory, it is shipped with a default software version.

Option	Description
Build Number	Displays the build number of an activated Edge.
Model	Displays the model type of the Edge.
Serial Number	Displays the serial number of the Edge. Assigning a serial number to an Edge is optional. If a serial number is not assigned to the Edge, this field will be blank.
Created	Displays the date and time the Edge was provisioned.
Activated	Displays the date and time the Edge was activated.
Last Contact	The last date and time the Edge communicated with the SD-WAN Orchestrator.
Column (Cols)	Click the Cols button to select the options you want to display in the Enterprise Edges list (See image above).
Reset View	Resets the Enterprise Edges list to the default view. (This removes filters and resets any options that were selected from the Cols button drop-down menu to the default view).
Refresh	Refreshes the Enterprise Edges list with current data from the server.
CSV	To export the content displayed in the Enterprise Edges list, click the CSV button.
Selected	Indicates how many Edges are selected from the Edge column. Click the Selected button to select all or deselect all of the Edges listed in the Edge column.
Actions	<p>Lists the actions that you can perform on the selected Edge. Based on the user roles and privileges, the supported actions will vary. For an enterprise user, the following actions are supported:</p> <ul style="list-style-type: none"> ■ New Edge - Creates a new edge. ■ Local Credentials - Assigns local configuration credentials for the selected edge. ■ Delete Edge - Deletes the selected edges. ■ Assign Profile - Changes the profile for the selected edges. ■ Stage to Bastion - Stages a not activated Edge to the Bastion Orchestrator. ■ Unstage from Bastion - Removes a staged Edge from the Production Orchestrator. ■ Promote to Production - Promotes a staged Edge to the Production Orchestrator. <p>Note Stage to Bastion, Unstage from Bastion, and Promote to Production options are available only when the Bastion Orchestrator feature is enabled using the <code>session.options.enableBastionOrchestrator</code> system property.</p> <p>For more information, see <i>Bastion Orchestrator Configuration Guide</i> available at https://docs.vmware.com/en/VMware-SD-WAN/index.html.</p> <ul style="list-style-type: none"> ■ Assign Software Image - Changes or updates the software image assigned to edges. For steps, see Assign Software Image. <p>Note This option is available only for Enterprise Super users with Edge Image Management feature-enabled.</p> <ul style="list-style-type: none"> ■ Assign Edge License - Assigns a license type to a selected edge. <p>Note Superuser Administrators and Standard Administrators can assign a license type to an edge.</p> <ul style="list-style-type: none"> ■ Update Alerts - Activates or deactivates Edge alert notifications for Customers.

Option	Description
New Edge	Opens the Provision New Edge dialog to provision a new Edge. For more information, see Provision a New Edge .
Help	Access the online help for this feature by clicking the Question Mark icon.

Assign Software Image


As an enterprise super user, after you have provisioned the edges, you can change or update the software image assigned to the edges using **Assign Software Image** under the **Actions** drop-down menu in the **Edges** screen.

To update a software image for an edge, perform the following steps.

Procedure

- 1 In the Enterprise portal, click **Configure > Edges**.
- 2 In the **Edges** screen, select an edge or multiple edges for which you want to update the software image.
- 3 Click **Actions** and from the drop-down menu, select **Assign Software Image**.

The **Assign Software Image** screen appears.



- 4 From the **Software** drop-down menu, select a software image to update the selected edges and click **Update**.

A warning message alerting the user about service disruption appears.

10.81.114.0 says

Assigning a different Software Image may cause certain service disruptions. Would you like to continue?



- 5 Click **OK** to continue.

Note If no software image is set for an edge, the edge will inherit the software image assigned to the customer.

Reset Edges to Factory Settings

SD-WAN Edges are required to be reset to factory settings for several reasons, some of which are as follows:

- When you repurpose 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 VMware SD-WAN 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 [How to Factory Reset a VMware SD-WAN Edge](#).

Manage Edges with New Orchestrator UI

Edges inherit the configurations from the associated profile. You can choose to override the settings for a specified Edge.

You can manage the Edges using the new Orchestrator UI.

- 1 In the Enterprise portal, click the **Open New Orchestrator UI** option available at the top of the Window.
- 2 Click **Launch New Orchestrator UI** in the pop-up window.

3 The UI opens in a new tab displaying the monitoring and configuring options.

In the new Orchestrator UI, click the **Configure** tab.

Note The **Configure** tab is available only when the Operator has enabled the option. If the tab is not available for you, contact your Operator.

1 Click **Configure > Edges**.

2 The **Edges** page displays the existing Edges with the following details:

Option	Description
Name	The name of the Edge. Click the link to the Edge to modify the Edge configurations. See Configure Edges with new Orchestrator UI .
Certificates	The current and expired certificates of the Edge.
Profile	The Profile assigned to the Edge. Click the link to the Profile to modify the Profile configurations. See Configure Profiles with new Orchestrator UI .
Operator Profile	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	The analytics details of the Edge.
HA	Displays whether High Availability is enabled for the Edge.
Device	Click the View link to modify the configurations of the Edge. See Configure Edges with new Orchestrator UI .
Alerts	Displays whether Customer alerts are enabled or not for the Edge.
Operator Alerts	Displays whether Operator alerts are enabled or not for the Edge.
Software Version	The software version of the Edge.
Factory Software Version	The Factory Software Version of the Edge. When the Edge is shipped from the factory, it is shipped with a default software version.
Build Number	The build number of the Edge, when the Edge is activated.
Model	The model type of the Edge.
Serial Number	The serial number of the Edge. Assigning a serial number to an Edge is optional. If a serial number is not assigned to the Edge, this field will be blank.
Created	The date and time when the Edge was provisioned.

Option	Description
Software Updated	The date and time when the Software was updated on the Edge.
Activated	The date and time when the Edge was activated.
Custom Info	The custom information associated with the Edge.
Last Contact	The last date and time when the Edge communicated with the SD-WAN Orchestrator.
Edge License	The Edge license assigned to the Edge.

3 Select one or more Edges to perform the following activities:

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.
Update Alerts	Allows to turn on or turn off the alerts sent to the Customer. To configure the alerts, see Chapter 24 Configure Alerts . You can view the alerts in the Monitor > Alerts tab.

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.
Update Alerts	Allows to turn on or turn off the alerts sent to the Customer. To configure the alerts, see Chapter 24 Configure Alerts . You can view the alerts in the Monitor > Alerts tab.

Click **More** to configure the following:

Option	Description
Update Operator Alerts	Allows to turn on or turn off the alerts sent to the Operator. To configure the alerts, see Chapter 24 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.

Option	Description
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	<p>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.</p> <hr/> <p>Note Refer to the <i>SD-WAN Gateway Migration FAQs</i>, <i>Important Caveats</i>, and <i>Allow List Limitation</i> sections in the KB article for complete details.</p> <hr/> <p>Note The Rebalance Gateways option is available only for Operator users.</p> <hr/>
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.

Configure Edges with new Orchestrator UI

You can configure the Edges using the new Orchestrator UI.

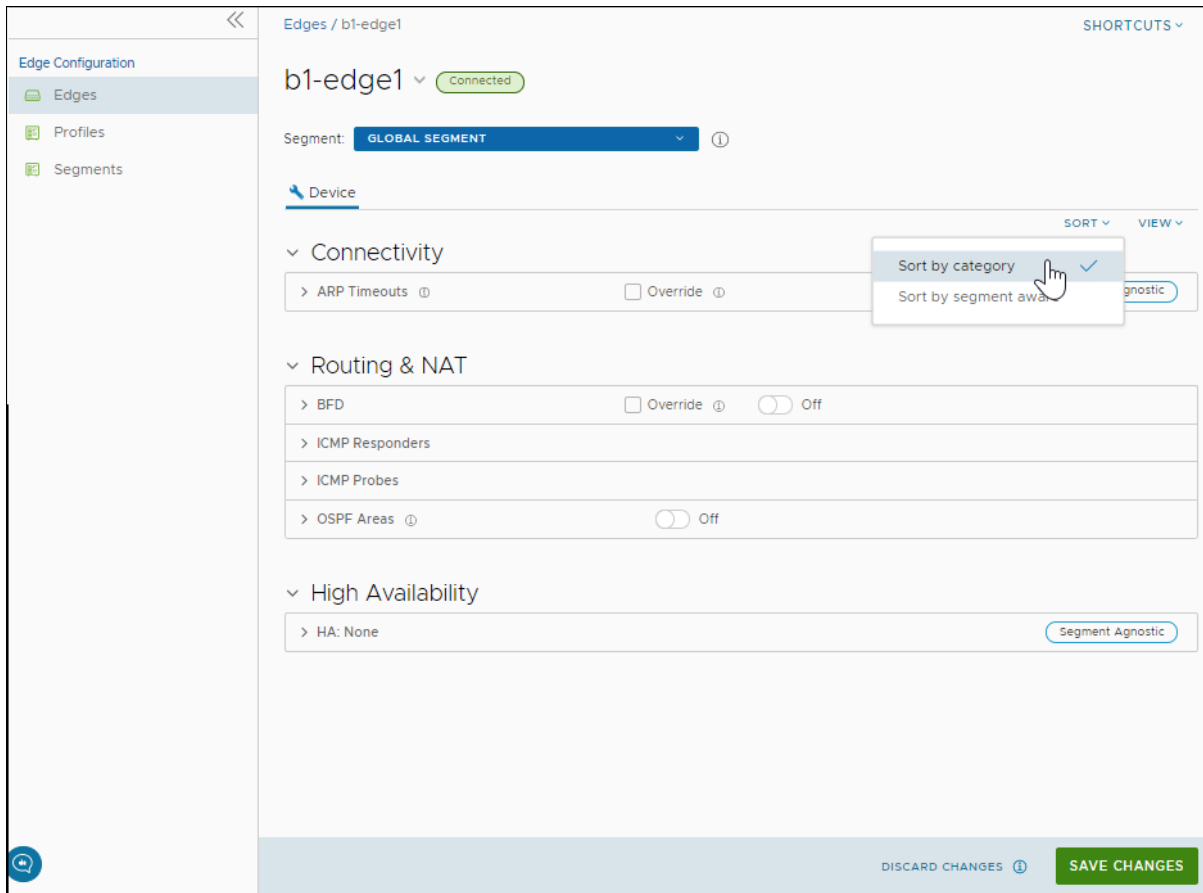
- 1 In the Enterprise portal, click the **Open New Orchestrator UI** option available at the top of the Window.
- 2 Click **Launch New Orchestrator UI** in the pop-up window.
- 3 The UI opens in a new tab displaying the monitoring and configuring options.

In the new Orchestrator UI, click the **Configure** tab.

Note The **Configure** tab is available only when the Operator has enabled the option. If the tab is not available for you, contact your Operator.

To configure a specific Edge:

- 1 Click **Configure > Edges**.
- 2 The **Edges** page displays the existing Edges.
- 3 Click the link to an Edge or click the **View** link in the **Device** column of the Edge.
- 4 The configuration options for the selected Edge are displayed.



- 5 Click **View** to expand or collapse the view of available settings.
- 6 You can also view the configuration settings sorted by category or segmentation. 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.

The following settings are available when you choose to sort by category:

Connectivity

- **ARP Timeouts** – By default, the Edge inherits the ARP settings from the associated Profile. Select the **Override** and **Override default ARP Timeouts** checkboxes to modify the values. For more information, see [Configure Layer 2 Settings for Edges](#).
- **Wi-Fi Radio** – Activate or deactivate Wi-Fi Radio and configure the band of radio frequencies. For more 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.

Routing & NAT

- **BFD** – By default, the Edge inherits the BFD configuration settings from the associated Profile. If required, you can select the **Override** checkbox to modify the settings. For more information, see [Configure BFD](#).

- **ICMP Responders** – Select this option to configure ICMP Responders that respond to ICMP probes from a specified IP address. See [Configure ICMP Probes/Responders](#).
- **ICMP Probes** – Select this option to configure ICMP probes that check for the network continuity by pinging specified IP address at frequent intervals. See [Configure ICMP Probes/Responders](#).
- **OSPF Areas** – 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 more information, see [Enable OSPF](#).

High Availability – You can enable High Availability for the selected Edge. Choose one of the following options:

- **None** – This is the default option where High Availability is not enabled.
- **Active Standby Pair** – Select this option to enable HA on the selected Edge. For more information, see [Activate High Availability](#).
- **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 Edge Clustering](#).
- **VRRP with 3rd party router** – Select this option to configure Virtual Router Redundancy Protocol (VRRP) on the selected Edge to enable next-hop redundancy in the SD-WAN Orchestrator network by peering with third-party CE router. To configure VRRP, see [Configure VRRP Settings](#).

7 After modifying the required settings, click **Save Changes**.

8 Click 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 with new Orchestrator UI](#).
- **Generate Diagnostic Bundle** – Allows to generate Diagnostic Bundle for the selected Edge. See [Diagnostic Bundles with new Orchestrator UI](#).
- **Remote Actions** – Allows to perform the Remote actions for the selected Edge. See [Perform Remote Actions with new Orchestrator UI](#).
- **View Profile** – Navigates to the Profile page, that is associated with the selected Edge.
- **View Gateways** – Displays the Gateways connected to the selected Edge.

View or Modify Edge Information

17

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.

In the Enterprise portal, click **Configure > Edges**. The page displays the existing Edges. Click the link to an Edge. In the **Edge Overview** tab, you can view or modify the following Edge information:

The screenshot shows the 'Edge Overview' configuration page for a device named 'b1-edge1'. The page has a top navigation bar with tabs for 'Edge Overview', 'Device', 'Business Policy', and 'Firewall'. The 'Edge Overview' tab is active. Below the navigation bar, there is a 'Properties' section with various fields and a status summary. The 'Properties' section includes fields for Name (b1-edge1), Description, Custom Info, Enable Pre-Notifications (checked), Enable Alerts (checked), Authentication Mode (Certificate Not Required), and License (Select Edge License). The status summary shows the device is 'Activated' on 'Mon Jul 19, 19:24' with 'Software Version 4.5.0 (build R450-20210719-MH-9bebbbec73)' and 'Local Credentials *****'. There is a 'Modify' button next to the Local Credentials. Below the Properties section, there are four expandable sections: 'Profile', 'Contact & Location', and 'RMA Reactivation', each with a down arrow icon.

Properties	
Name	b1-edge1
Description	
Custom Info	
Enable Pre-Notifications	<input checked="" type="checkbox"/>
Enable Alerts	<input checked="" type="checkbox"/>
Authentication Mode	Certificate Not Required
License	Select Edge License

Status	Activated
Activated	Mon Jul 19, 19:24
Software Version	4.5.0 (build R450-20210719-MH-9bebbbec73)
Local Credentials	***** Modify

Profile

Contact & Location

RMA Reactivation

Identifying a Device Model

To identify a device model, click the down arrow next to the device name. A pop-up window displays, which shows Edge and device model information. A non-WiFi Edge model will contain a "-n" at the end of the model name. See image below.

Edge Info

Activation	Reactivation Pending
Act. Key	
Activated	Mon Aug 09, 16:48:28
Last Contact	Mon Aug 09, 16:48:28
System Up Since	Mon Aug 09, 16:48:28
Service Up Since	
Pre-Notifications	<input checked="" type="checkbox"/>
Authentication	Certificate Acquire

Device Hardware

Model	edge620-n
Serial Number	

Device Software

Current Version	4.5.0 [R450-20210808-
MN-5d80c45f9e]	
Factory Version	
Analytics	None

Edge Properties

The existing details of the selected Edge are displayed. If required, you can modify the information.

Option	Description
Name	Displays the existing name of the Edge.
Description	Displays the existing description of the Edge.
Custom Info	Displays the custom information associated with the Edge.
Enable Pre-Notifications checkbox	By default, this option is enabled, which sends alert notifications for the Edge, to the Operators. The Operators can receive the alerts through Email, SMS, or SNMP traps. To configure the alerts, see Chapter 24 Configure Alerts . You can also view the alerts by clicking Monitor > Alerts .
Enable Alerts checkbox	By default, this option is enabled, which sends alert notifications for the Edge, to the Customers. The Customers can receive the alerts through Email, SMS, or SNMP traps. To configure the alerts, see Chapter 24 Configure Alerts . You can also view the alerts by clicking Monitor > Alerts .

Option	Description
Authentication Mode	<p>Choose the mode of authentication from the following available list:</p> <ul style="list-style-type: none"> ■ Certificate Deactivated: Edge uses a pre-shared key mode of authentication. <hr/> <p>Warning This mode is not recommended for any customer deployments.</p> <hr/> <ul style="list-style-type: none"> ■ 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 SD-WAN 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 SD-WAN Orchestrator and for establishment of VCMP tunnels. <hr/> <p>Note After acquiring the certificate, the option could be updated to Certificate Required, if needed.</p> <hr/> <ul style="list-style-type: none"> ■ 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. <hr/> <p>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.</p> <hr/> <p>Certificate Required mode means that no Edge heartbeats are accepted without a valid certificate.</p> <hr/> <p>Caution Using this mode can cause Edge failures in cases where a customer is unaware of this strict enforcement.</p> <hr/> <p>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 more information, contact your Operator.</p> <hr/> <p>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.</p>
License	Choose an Edge License from the available list. The list displays the licenses assigned to the Enterprise, by the Operator.
View Certificate	This option is displayed when the Edge has a valid certificate. Click the View link to view, export, or revoke the certificate.
Status	<p>Displays the status of the Edge:</p> <ul style="list-style-type: none"> ■ Pending: The Edge has not been activated. ■ Activated: The Edge has been activated. ■ Reactivation Pending: When you click Request Reactivation, the status changes to Reactivation Pending, which indicates that a new or replaced Edge can be activated with the existing configuration. Anyways, this status does not affect the functionality of the Edge.
Activated	Displays the date and time the Edge got activated.
Software Version	Displays the software version and build number of the Edge.

Option	Description
Local Credentials	Displays the credentials for the local UI. The local credentials include a default username, 'admin' and a randomly generated password. Click Modify to update the credentials at the Edge level.
Serial Number	This option is available when the Edge is in Pending state. You can enter the serial number of the Edge, which is optional. If entered, then the number must match the serial number of the Edge when activated.
Activation Key	This option is available when the Edge is in Pending state. The activation key is valid for one month. After one month, the key expires and a warning message is displayed. You can generate a new key by clicking Generate New Activation Key in the warning message.

Assigned Profile and Edge-Specific Overrides

The profile assigned to the Edge and the Edge Specific Overrides & Additions are displayed. 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 added to the selected Edge. A summary of all Edge overrides and additions are displayed in this section.

You can modify the assigned profile by selecting a profile from the drop-down list.

Note When switching to a different profile, the Edge override configurations are not modified.

Note 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 list.

While switching the profiles, check the compatibility between a customer-assigned Operator Profile and an Edge-assigned Enterprise Profile. The following table provides the 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.

Customer Operator Profile Type	Current Edge Enterprise Profile	Selected Edge Enterprise Profile	Result
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.

Edge Contact and Location

The existing contact and location details of the Edge are displayed. You can modify the contact details. To update the location details, click **Set Location**.

In the **Set Edge Location** window, update the location by either searching for the address or entering the address manually.

If the shipping address is different from the Edge location, clear the **Same as above** checkbox next to the shipping address, then enter the shipping contact. To update the Shipping Location, click **Set Location**. In the **Edge Shipping Location** window, update the location by either searching for the address or entering the address manually.

RMA Reactivation

This option is available only for activated Edges. You can initiate an RMA reactivation request to:

- Replace an Edge due to a malfunction
- Upgrade an Edge hardware model

RMA Reactivation

Request Reactivation:

Cancel Reactivation Request

RMA Activation key will expire in a month ⓘ

Send Activation Email...

RMA Edge Attributes

RMA Serial Number

Ex: VC00000490

RMA Model

Edge 5X0 ▼


Optional. If specified, the activated Edge device must match the above specified attributes.

Update


The following table provides the details of the RMA Reactivation options:

Option	Description
Request Reactivation	<p>Generates a new activation key. The status of the Edge changes to Reactivation Pending mode.</p> <p>Note The reactivation key is valid for one month only.</p>
Cancel Reactivation Request	<p>Cancels the RMA reactivation request. When you cancel the request, the status of the Edge changes to Activated mode.</p>
Send Activation Email	<p>Sends an email with activation instructions to the Site Contact. This option does not activate the Edge, but initiates the activation process.</p> <p>When you click this option, a pop-up window appears with the Email details. You can modify the instructions and send the Email.</p> <p>The Email consists of the instructions along with the activation URL. The URL displays the Activation key and the IP address of the SD-WAN Orchestrator.</p>
RMA Serial Number	<p>The serial number of the Edge to be activated. This Edge replaces your current Edge for which you have requested the RMA reactivation.</p> <p>Note If the Serial Number and the Edge model do not match the Edge to be activated, then the activation fails.</p>
RMA Model	<p>The model number of the Edge to be activated. This Edge replaces your current Edge for which you have requested the RMA reactivation.</p> <p>Note</p> <ul style="list-style-type: none"> ■ 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.

The RMA Activation Key is valid for one month. When the key expires, a warning message is displayed. To generate a new key, click **Generate New Activation Key**.



In the **Generate New Activation Key** window, specify the number of days for key to be active, and click **Submit**.



After generating the key, reactivate the Edge with the new key.

After making changes to the Edge details, click **Save Changes**.

Edge Device Configurations

18

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.

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 more information, see [Configure Business Policy for Edges](#).

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.

Edge Device Configurations—A Roadmap

Go to **Configure > Edges > Device** to override configurations 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 common across segments.

The following table provides the list of Edge-level configurations:

Configuration	Type of Configuration	For details, refer to ...
DSL Settings Note Available only for 610 Edge and 610-LTE devices.	Common	Configure DSL Settings
GPON Settings Note Available only for 6X0 Edges.	Common	Configure GPON Settings
Authentication Settings	Segment-aware (Only for Global segment)	Configure Authentication Settings
DNS Settings	Segment-aware (Only for Global segment)	Configure DNS Settings
NetFlow Settings	Segment-aware	Configure Netflow Settings for Edges
LAN-Side NAT Rules	Common	LAN-side NAT Rules at Edge Level
Syslog Settings	Segment-aware	Configure Syslog Settings for Edges
Static Route Settings	Segment-aware	Configure Static Route Settings

Configuration	Type of Configuration	For details, refer to ...
ICMP Probes	Segment-aware	Configure ICMP Probes/Responders
ICMP Responders	Segment-aware	Configure ICMP Probes/Responders
VRRP Settings	Segment-aware	Configure VRRP Settings
Cloud VPN	Segment-aware	Configure Cloud VPN and Tunnel Parameters at the Edge Level
BFD Rules	Common	Configure BFD
OSPF Areas	Segment-aware	Enable OSPF
BGP Settings	Segment-aware	Configure BGP from Edge to Underlay Neighbors
Multicast Settings	Segment-aware	Configure Multicast Settings
Cloud Security Service	Segment-aware	Configure Cloud Security Services for Edges
High Availability	Common	Activate High Availability
VLAN	Common	Configure VLAN for Edges
Loopback Interfaces	Segment-aware	Loopback Interfaces Configuration
Orchestrator Management Traffic	Common	Configure Orchestrator Management Traffic for Edges
Device Settings (Interface and WAN)	Common	Configure Device Settings
Security VNF	Common	Security VNFs
Multi-Source QOS	Common	Activate Multi-Source QOS
L2 Settings	Common	Configure Layer 2 Settings for Edges
SNMP Settings	Common	Configure SNMP Settings for Edges
NTP Servers	Common	Configure NTP Settings for Edges
Visibility Mode	Common	Configure Visibility Mode

Read the following topics next:

- [Configure DSL Settings](#)
- [Configure GPON Settings](#)
- [Configure Netflow Settings for Edges](#)
- [LAN-side NAT Rules at Edge Level](#)
- [Configure Syslog Settings for Edges](#)
- [Configure Static Route Settings](#)
- [Configure ICMP Probes/Responders](#)
- [Configure VRRP Settings](#)
- [Configure Cloud VPN and Tunnel Parameters at the Edge level](#)
- [Configure VLAN for Edges](#)
- [Loopback Interfaces Configuration](#)

- [Configure Orchestrator Management Traffic for Edges](#)
- [Configure Device Settings](#)
- [Configure Wi-Fi Radio Overrides](#)
- [Security VNFs](#)
- [Configure Layer 2 Settings for Edges](#)
- [Configure SNMP Settings for Edges](#)
- [Configure NTP Settings for Edges](#)
- [Configure Edge Activation](#)

Configure DSL Settings

Support is available for xDSL 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 more 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.

DSL Status Run

View the xDSL(ADSL2/VDSL2) modem status connected to SFP interfaces Test Duration: 10.003 seconds

Interfaces							
Name	Mode	Vendor MAC	xDSL Mode	Link Time	Status	Link Rate	Annex
SFP1	DSL	00:0E:AD:00:55:FE	VDSL2	0	Idle	0/0	N/A
SFP2	DSL	00:0E:AD:00:55:AC	VDSL2	49223	Showtime	12045/23407	AnnexA

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.

Configuring SFP

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.

Device Settings: Edge 610

Interface Settings ➕ Add Subinterface ➕ Add Secondary IP ➕ Add WIFI SSID

Actions	Interface		Switch Port Settings		Routed Interface Settings		Segment	Multicast		VNF Insertion
	Override	Interface	Mode	VLANs	Addressing	WAN Overlay		IGMP	PIM	
Edit	✗	GE1	Access	1 - Corporate			Global Segment			
Edit	✗	GE2	Access	1 - Corporate			Global Segment			
Edit	✗	GE3			DHCP	🔌 Auto Detect	all segments			✗
Edit	✗	GE4			DHCP	🔌 Auto Detect	all segments			✗
Edit	✗	GE5			DHCP	🔌 Auto Detect	all segments			✗
Edit	✗	GE6			DHCP	🔌 Auto Detect	all segments			✗
Edit	☑	SFP1			DHCP	🔌 Auto Detect	all segments			✗
Edit	✗	SFP2			DHCP	🔌 Auto Detect	all segments			✗
Edit	☑	WLAN1	Wifi	1 - Corporate			Global Segment			
Edit	✗	WLAN2	Interface disabled							

View the [recommended method](#) to configure interfaces at the profile and edge level.

To Configure SFP:

- 1 Click the **Edit** link in the **Actions** column.

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](#).

Edge 610

Multicast: Multicast is not enabled for the selected segment

RADIUS Authentication: ⓘ
Require User Authentication to access WAN

✗ WAN Overlay must be disabled to configure RADIUS Authentication.

Advertise: ☐

ICMP Echo Response: ☒

NAT Direct Traffic: ☒

Underlay Accounting: ⓘ ☒

Trusted Source: ⓘ ☐

Reverse Path Forwarding: ⓘ Specific ▾

L2 Settings

Autonegotiate: ☒

* MTU: 1500

SFP Settings

SFP Module: DSL ▾

DSL Settings

Mode: VDSL2 ▾

Profile: 17a ▾

Update SFP1 Cancel

- 2 To configure DSL Settings, select the **Override Interface** check box.
- 3 Select the **Interface Enabled** check box.
- 4 In the **SFP Settings** area, there are two options available from the drop-down list, Standard and DSL. Choose **DSL** as the SFP Module.

SFP Settings

SFP Module: DSL ▾

DSL Settings

Mode: ADSL2/2+ ▾

PVC: 0 ▾

VPI: ⓘ 0

VCI: ⓘ 35

PVC VLAN: ⓘ 1

Update SFP1 Cancel

- 5 In the **DSL Settings** area, configure the following:

Option	Description
SFP Module	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.

Option	Description
DSL Mode: VDSL2	<p>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.</p> <p>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.</p>
DSL Mode: ADSL2/2+	<p>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.</p> <p>If you choose ADSL2/2+, configure the following settings:</p> <ul style="list-style-type: none"> ■ 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. <p>Note The 4.3 release introduces four new parameters for PVC VLAN as described below. Click the arrow next to the PVC VLAN text box to display these settings.</p> <ul style="list-style-type: none"> ■ 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.

6 Click **Update SFP1** to save the configuration.

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 SD-WAN Orchestrator

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 **Device Settings** area of the SD-WAN Orchestrator.

Device Settings: Edge 610

Interface Settings [Add Subinterface](#) [Add Secondary IP](#) [Add WIFI SSID](#)

Actions	Interface		Switch Port Settings		Routed Interface Settings		Segment	Multicast		
	Override	Interface	Mode	VLANs	Addressing	WAN Overlay		IGMP	PIM	VNF Insertion
Edit	<input checked="" type="checkbox"/>	GE1	Access	1 - Corporate			Global Segment			
Edit	<input checked="" type="checkbox"/>	GE2	Access	1 - Corporate			Global Segment			
Edit	<input checked="" type="checkbox"/>	GE3			DHCP	Auto Detect	all segments			<input checked="" type="checkbox"/>
Edit	<input checked="" type="checkbox"/>	GE4			DHCP	Auto Detect	all segments			<input checked="" type="checkbox"/>
Edit	<input checked="" type="checkbox"/>	GE5			DHCP	Auto Detect	all segments			<input checked="" type="checkbox"/>
Edit	<input checked="" type="checkbox"/>	GE6			DHCP	Auto Detect	all segments			<input checked="" type="checkbox"/>
Edit	<input checked="" type="checkbox"/>	SFP1			Static CIDR: 192.168.78.185/24 Gateway: 192.168.78.200	Auto Detect	all segments			<input checked="" type="checkbox"/>
Edit	<input checked="" type="checkbox"/>	SFP2			DHCP	Auto Detect	all segments			<input checked="" type="checkbox"/>
Edit	<input checked="" type="checkbox"/>	WLAN1	Interface disabled							
Edit	<input checked="" type="checkbox"/>	WLAN2	Interface disabled							

View the [recommended method](#) to configure interfaces at the profile and edge level.

To Configure GPON ONT SFP from the SD-WAN Orchestrator:

- 1 From the SD-WAN Orchestrator, go to **Configure > Edges**, and then select an 6X0 Edge.
- 2 Click the Device tab, and scroll down to Device Settings for the 6X0 Edge.
- 3 Click the **Edit** link in the **Actions** column for the SFP (SFP1 or SFP 2) Interface as shown in the image above.

The 6X0 Edge device Interface dialog box displays.

Edge 6X0

Interface SFP1

Interface Enabled ☒

Capability Routed

Segments All Segments

Addressing Type DHCP

Static/PPPoE addressing details must be configured individually per edge.

WAN Overlay ☒ Auto-Detect Overlay

OSPF ✕ OSPF not enabled for the selected Segment.

VNF Insertion ✕ VNF insertion is disallowed when an interface is configured for WAN overlays

Multicast Multicast is not enabled for the selected segment

RADIUS Authentication ⓘ ✕ WAN Overlay must be disabled to configure RADIUS Authentication.
Require User Authentication to access WAN

Advertise ☐

ICMP Echo Response ☒

NAT Direct Traffic ☒

Underlay Accounting ⓘ ☒

Trusted Source ⓘ ☐

Reverse Path Forwarding ⓘ Specific

VLAN

L2 Settings

Autonegotiate ☒

* MTU 1500

SFP Settings

SFP Module GPON

GPON Settings

Subscriber Location ID Mode ASCII

Subscriber Location ID 1234567890

Update SFP1 Cancel

- 4 The **Override Interface** checkbox must be checked to configure GPON Settings.
- 5 Check the **Interface Enabled** checkbox.

Note The steps in this procedure are specifically for configuring GPON. For a complete description of all items in the 6X0 Edge device Interface dialog box, see [Configure Interface Settings](#).

- 6 In the **SFP Settings** area, choose GPON as the SFP Module as shown in the image below.

SFP Settings

SFP Module GPON

GPON Settings

Subscriber Location ID Mode ASCII

Subscriber Location ID 118555507

- 7 In the **GPON Settings** area, 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).
- 8 In the **GPON Settings** area, enter the Subscriber Location ID.
- 9 Click the **Update SFP1** button.

To Configure GPON ONT SFP from the Local UI:

- 1 From the Local UI, navigate to **Details Edge Overview**.

The Edge Overview Page displays as shown in the image below.

Edge Overview

Reset Settings Download Diagnostic Bundle Connected Both Unconnected

Routed Interfaces

Wired

GE5

GE6

SFP1

SFP2

Cellular

USB1

USB2

Switched Interfaces

- 2 Select an SFP interface to change its setting.
- 3 In the SFP properties page, select GPON from the **SFP Module** drop-down menu. (See image below).

SFP1 Properties

[<< Return](#)[Save](#)

Status

Link Detected: No

Configuration

(Fields marked with * are required.)

* Addressing:

☒ DHCP
☐ Static
☐ PPPoE

* IP Address:

* Subnet Mask:

* Gateways:

VLAN:

* Autonegotiation:

☒ On

SFP Settings

SFP Module:

✓ Standard

DSL

GPON

©2013-2019 VeloCloud Networks | Version 3.3.2, build 201905201400

The following GPON Settings, the Subscriber Location ID Mode and Subscriber Location ID, display as shown in the image below.

SFP Settings

SFP Module:

GPON

GPON Settings

Subscriber Location ID Mode:

ASCII

Subscriber Location ID:

0

- 4 In the **Subscriber Location ID Mode** drop-down menu, choose either ACII or HEX (image below).

SFP Settings

SFP Module:

GPON

GPON Settings

Subscriber Location ID Mode:

✓ ASCII

Subscriber Location ID:

HEX

- 5 Enter the Subscriber Location ID in the appropriate text box.

Troubleshooting GPON Settings

The GPON diagnostic test is available only for 6X0 devices. See [Performing Remote Diagnostics Tests](#) for more information.

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 **Enable Edge Override** checkbox.

Procedure

- 1 From the SD-WAN Orchestrator, go to **Configure > Edges**.
- 2 Select an Edge you want to override Netflow settings and click the icon under the **Device** column.

The Device Setting page for the selected Edge appears.

Netflow Settings ⓘ

Enable Edge Override ⓘ

Netflow Enabled: ☒

Version: ⓘ v10

Observation ID: 14

Collector	Filter ⓘ	Allow All ⓘ	Source Interface
C-global10.4.1.32		<input checked="" type="checkbox"/>	[none]

Intervals:

- * Flow Stats: 69
- * FlowLink Stats: 62
- * Segment Table: 101
- * Application Table: 103
- * Interface Table: 105
- * Link Table: 95
- * Tunnel Stats: 60

- 3 From the **Configure Segment** drop-down menu, select a profile segment to configure Netflow settings.
- 4 Go to the **Netflow Settings** area and select the **Enable Edge Override** check box.
- 5 Select the **Netflow Enabled** 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 Click **Save Changes**.

Results

After you enable Netflow on the VMware SD-WAN Edge, it periodically sends messages to the configured collector. The contents of these messages are defined using IPFIX templates. For more information on templates, see [IPFIX Templates](#).

LAN-side NAT Rules at Edge Level

LAN-Side NAT 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, VMware introduced a new LAN-side NAT module to NAT VPN routes on the Edge. The primary use cases are as follows:

- Branch overlapping IP due to M&A
- 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.


LAN-side NAT (3.3.2 Release)

Use Case Number One: "Many:1 Source NAT"

In this scenario, a third-party has assigned multiple non-overlapping subnets to a customer's site. The server in the customer's data center recognizes traffic from this third-party by a single IP address at any given site.

The configuration required for Use Case Number One for Version 3.3.2: New rule: LAN-side NAT 192.168.1.0/24 -> 172.16.24.4/32

As shown in the image below, because the NAT rule is a single IP, TCP and UDP traffic will be NAT'ed. Therefore, in this example, 192.168.1.50 becomes 172.16.24.4 with an ephemeral source port for TCP/UDP traffic, ICMP traffic becomes 172.16.24.4 with a custom ICMP ID for reverse lookup, and all other traffic will be dropped.



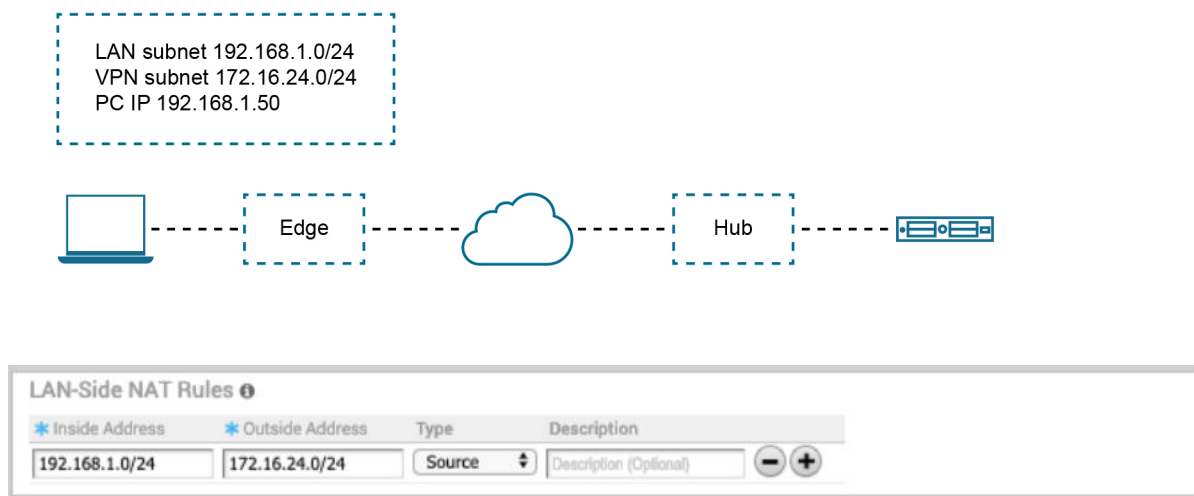
* Inside Address	* Outside Address	Type	Description		
192.168.1.0/24	172.16.24.4/32	Source	Description (Optional)	-	+
192.168.5.0/24	172.16.24.4/32	Source	Description (Optional)	-	+
192.168.7.0/24	172.16.24.4/32	Source	Description (Optional)	-	+

Use Case Number Two: "1:1 Source NAT"

In this scenario, the LAN subnet is 192.168.1.0/24. However, this is an overlapping subnet with other sites. A unique subnet of equal size, 172.16.24.0/24 has been assigned to use for VPN communication at this site. Traffic from the PC must be NAT'ed on the Edge prior to the route lookup, otherwise the source route will match 192.168.1.0/24 which is not advertised from this Edge and traffic will drop.

The configuration required for Use Case Number Two: New rule: LAN-side NAT 192.168.1.0/24 -> 172.16.24.0/24

Because the subnets match in size, all bits matching the subnet mask will be NAT'ed. Therefore, in the image below example, 192.168.1.50 becomes 172.16.24.50.




LAN-side NAT Based on Source or Destination (3.4 Release)

The 3.4 release introduces LAN-side NAT based on Source/Destination support as part of a single rule, in which you can enable NAT only for a subset of traffic based on Source or Destination subnets. See the following use cases for this enhancement below.

Use Case Number One: "Perform SNAT or DNAT with Source or Destination as Match Criteria"

In the illustration example below, the branch should NAT the source IP 10.4.1.1 to 10.200.1.245 only for the traffic destined to 100.1.1.0/24. Similarly, at the DC, the destination IP 100.1.1.9 should be NATed to 10.1.10.9 only if the traffic is received from source 10.200.1.0/24.



See the image below (LAN-side NAT Rules area for the Branch).

Branch:

LAN-Side NAT Rules ⓘ

NAT Source or Destination

Type	* Inside Address	* Outside Address	Source Route	Destination Route	Description
Source	10.4.1.1	10.200.1.245	n/a	100.1.1.0/24	Description (Optional)

See the image below (LAN-side NAT Rules area for the Hub).

Hub:

LAN-Side NAT Rules ⓘ


NAT Source or Destination

Type	* Inside Address	* Outside Address	Source Route	Destination Route	Description
Destination	100.1.1.9	10.1.10.9	10.200.1.0.24	n/a	Description (Optional)

Use Case Number Two: To NAT Both Source and Destination IP on the Packet

Consider the below scenario. In this example, each site in the network is assigned the same subnet so that the Branch LAN is identical at every site. "PC1" and "PC2" have the same IP address and both need to communicate with a server behind the Hub. We need to source NAT the traffic in order to use overlapping IP addresses, e.g. in Edge 1, PCs (192.168.1.0/24) should be NAT'ed to 192.168.10.0/24, in Edge2, PCs (192.168.1.0/24) should be NAT'ed to 192.168.20.0/24.

Also, for security reason, the server behind the Hub with real IP "172.16.0.1" should be presented to PCs as "192.168.100.1," and this IP should not be distributed to SD-WAN between the Hub and Edge, source + destination combination rules on the same Edge are required.



LAN-Side NAT Rules ⓘ

NAT Source or Destination

Type	Inside Address	Outside Address	Source Route	Destination Route	Description
Source	e.g. 10.0.0.0/24	e.g. 192.168.0.0/24	n/a	e.g. 192.168.0.0/24	Description (Optional)

NAT Source and Destination

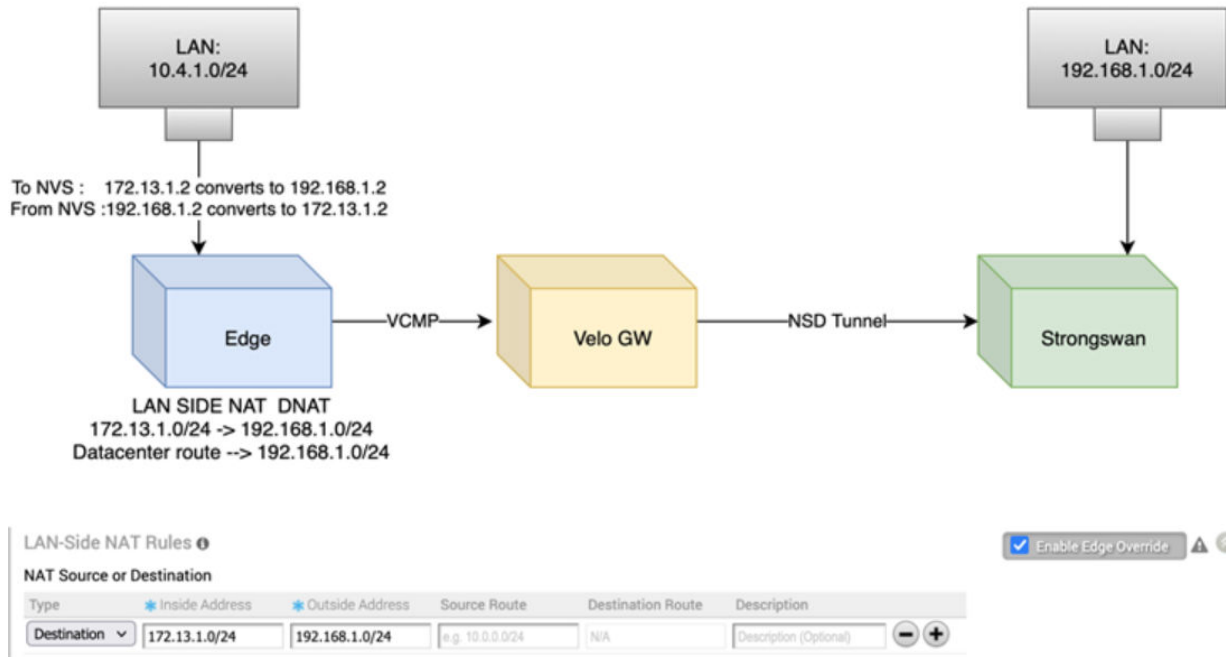
Type	Inside Address	Outside Address	Type	Inside Address	Outside Address	Description
Source	192.168.1.0/24	192.168.10.0/24	Destination	192.168.100.1	172.16.0.1	Description (Optional)

Note LAN-side NAT Rules can be configured at the Profile level or the Edge level. To configure at the Edge level, make sure the **Enable Edge Override** checkbox is checked.

Use Case Number Three: Many-to-Many Destination NAT for a NSD Subnet

As illustrated in the image below, the Edge LAN is 10.4.1.0/24 and the NVS Site subnet is 192.168.1.0/24. The LAN- side DNAT rule was configured to covert 172.13.1.0/24 to 192.168.1.0/24. The Gateway Pushes Data center route of NVS subnet (192.168.1.0/24) to the Edge, which was configured in the VMware Classic Orchestrator. Therefore, when traffic from LAN client (10.4.1.25) is initiated to 172.13.1.2 , 172.13.1.2 will be converted to 192.168.1.2, as per the DNAT rule. From the Edge to Gateway, its VCMP and GW to NSD its via IPSEC tunnel as usual. If the NVS client initiates traffic to 10.4.1.25, the Source IP: 192.168.1.2 will be translated to 172.13.1.2 as per the DNAT rule.

Note For the Reverse traffic, the use case works as SNAT.



Configure Procedure

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.

To apply LAN-Side NAT Rules:

- 1 From the navigational panel, go to **Configure > Edges**.
- 2 In the **Device Settings** tab screen, scroll down to the **LAN-Side NAT Rules** area.
- 3 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).
 - a Enter an address for the **Inside Address** text box.
 - b Enter an address for the **Outside Address** text box.
 - c Enter the Source Route in the appropriate text box.
 - d Enter the Destination Route in the appropriate text box.
 - e Type a description for the rule in the **Description** textbox (optional).
- 4 In the **LAN-side NAT Rules** area, complete the following for NAT Source and Destination: (See the table below for a description of the fields in the steps below).
 - a For the **Source** type, enter the **Inside Address** and the **Outside Address** in the appropriate text boxes.
 - b For the **Destination** type, enter the **Inside Address** and the **Outside Address** in the appropriate text boxes.

- c Type a description for the rule in the **Description** textbox (optional).

LAN-side NAT Rule	Type	Description
Type drop-down menu	Select either Source or Destination	Determine whether this NAT rule should be applied on the source or destination IP address of user traffic.
Inside Address text box	IPv4 address/prefix, Prefix must be 1-32	The "inside" or "before NAT" IP address (if prefix is 32) or subnet (if prefix is less than 32).
Outside Address text box	IPv4 address/prefix, Prefix must be 1-32	The "outside" or "after NAT" IP address (if prefix is 32) or subnet (if prefix is less than 32).
Source Route text box	<ul style="list-style-type: none"> - Optional - IPv4 address/prefix - Prefix must be 1-32 - Default: any 	For destination NAT, specify source IP/subnet as match criteria. Only valid if the type is "Destination."
Destination Route text box	<ul style="list-style-type: none"> - Optional - IPv4 address/prefix - Prefix must be 1-32 - Default: any 	For source NAT, specify destination IP/subnet as match criteria. Only valid if the type is "Source."
Description text box	Text	Custom text box to describe the NAT rule.

For packet sent from **LAN to WAN**, packet **source** addresses match "Inside" is translated to "Outside"

For packet sent from **WAN to LAN**, packet **destination** addresses match "Outside" is translated to "Inside"

For packet sent from **LAN to WAN**, packet **destination** addresses match "Inside" is translated to "Outside"

For packet sent from **WAN to LAN**, packet **source** addresses match "Outside" is translated to "Inside"

For packet sent from **LAN to WAN**, packet **source** addresses match **INSIDE ADDRESS** is translated to "Outside" under "Source", and packet **destination** address match "Inside" is translated to "Outside" under "Destination"

Note Important: If the Inside Prefix is less than the Outside Prefix, support Many:1 NAT in the LAN to WAN direction and 1:1 NAT in the WAN to LAN direction. For example, if the Inside Address = 10.0.5.0/24, Outside Address = 192.168.1.25/32 and type = source, for sessions from LAN to WAN with source IP matching 'Inside Address,' 10.0.5.1 will be translated to 192.168.1.25. For sessions from WAN to LAN with destination IP matching 'Outside Address,' 192.168.1.25 will be translated to 10.0.5.25. Similarly, if the Inside Prefix is greater than Outside Prefix, support Many:1 NAT in the WAN to LAN direction and 1:1 NAT in the LAN to WAN direction. The NAT'ed IP are not automatically advertised, make sure a static route for the NAT'ed IP should be configured and the next hop should be the LAN next hop IP of the source subnet.

LAN-side NAT "Cheat Sheet"

Use Case 1:

- **Traffic direction:** LAN->WAN
- **What needs to be translated:** packet source address
- **Config mapping:**
 - NAT Type = "Source"
 - Orig IP = "Inside Address"
 - NAT IP = "Outside Address"

NAT Type	Inside	Outside	Type	LAN->WAN Behavior
Source	A.0/24	B.0/24	1:1	A.1 translates to B.1, A.2 to B.2, etc.
Source	A.0/24	B.1/32	Many:1	A.1 and A.2 translate to B.1
Source	A.1/32	B.0/24	1:1	A.1 translates to B.1, other B.X are unused

Use Case 2:

- **Traffic direction:** WAN -> LAN
- **What needs to be translated:** packet destination address
- **Config mapping:**
 - NAT Type = "Source"
 - Orig IP = "Outside Address"
 - NAT IP = "Inside Address"

NAT Type	Inside	Outside	Type	WAN->LAN Behavior
Source	A.0/24	B.0/24	1:1	B.1 translates to A.1, B.2 to A.2, etc.
Source	A.0/24	B.1/32	Many:1	B.1 translates to A.1
Source	A.1/32	B.0/24	1:Many	B.1 and B.2 translate to A.1

Use Case 3:

- **Traffic direction:** LAN->WAN
- **What needs to be translated:** packet destination address

- **Config mapping:**

- NAT Type = “Destination”
- Orig IP = “Inside Address”
- NAT IP = “Outside Address”

NAT Type	Inside	Outside	Type	LAN->WAN Behavior
Destination	A.0/24	B.0/24	1:1	A.1 translates to B.1, A.1 to B.2, etc.
Destination	A.0/24	B.1/32	Many:1	A.1 and A.2 translate to B.1
Destination	A.1/32	B.0/24	1:Many	A.1 translates to B.1

Use Case 4:

- **Traffic direction:** WAN->LAN
- **What needs to be translated:** packet source address
- **Config mapping:**
 - NAT Type = “Destination”
 - Orig IP = “Outside Address”
 - NAT IP = “Inside Address”

NAT Type	Inside	Outside	Type	WAN->LAN Behavior
Destination	A.0/24	B.0/24	1:1	B.1 translates to A.1, B.2 to A.2, etc.
Destination	A.0/24	B.1/32	Many:1	B.1 translates to A.1
Destination	A.1/32	B.0/24	1:Many	B.1 and B.2 translate to A.1

Configure Syslog Settings for Edges

In an Enterprise network, SD-WAN Orchestrator supports collection of SD-WAN Orchestrator bound events and firewall logs originating from enterprise SD-WAN 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** checkbox.

To override the Syslog settings at the Edge level, perform the following steps.

Prerequisites

- Ensure that Cloud VPN (branch-to-branch VPN settings) is configured for the SD-WAN Edge (from where the SD-WAN Orchestrator bound events are originating) to establish a path between the SD-WAN Edge and the Syslog collectors. For more information, see [Configure Cloud VPN for Profiles](#).

Procedure

- 1 From the SD-WAN Orchestrator, go to **Configure > Edges**.
The SD-WAN Edge page appears.
- 2 Select an Edge you want to override Syslog settings and click the icon under the **Device** column.
The Device Settings page for the selected Edge appears.
- 3 From the **Configure Segment** drop-down menu, select a profile segment to configure syslog settings. By default, **Global Segment [Regular]** is selected.
- 4 Go to the **Syslog Settings** area and select the **Enable Edge Override** checkbox.
- 5 From the **Source Interface** drop-down list, 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 Click the **+** button to add another Syslog collector or else click **Save Changes**. The syslog settings for the edge will be 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.

Syslog Settings

Facility : local0

Syslog Enabled: ☒

* IP	* Protocol	* Port	* Source Interface	* Roles	* Syslog Level	Tag	All Segments
10.1.1.25	TCP	514	Auto	FIREWALL EVENT	INFO	VMware.SDWAN.FW	<input checked="" type="checkbox"/>
10.1.2.25	TCP	514	Auto	EDGE EVENT	ERROR	VMware.SDWAN.Edge	<input checked="" type="checkbox"/>

Firewall logs are forwarded at INFO level by default

You are at the maximum limit of 2 collectors per segment

Note 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 SD-WAN Orchestrator auto-generated local events to be received at the Syslog collector, you must configure Syslog at the SD-WAN 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](#).

What to do next

On the **Firewall** page of the Edge configuration, enable the **Syslog Forwarding** button if you want to forward firewall logs originating from enterprise SD-WAN 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 more information about Firewall settings at the Edge level, see [Configure Firewall for Edges](#).

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 Enterprise portal, click **Configure > Edges**.
- 2 Click the Device Icon next to an Edge, or click the link to an Edge and then click the **Device** tab.
- 3 In the **Device** tab, scroll down to the **Static Route Settings** section.
- 4 In the **IPv4** tab, you can configure the static routes for IPv4 addresses.

Static Route Settings

IPv4 IPv6

Subnet	Source IP	Next Hop	Interface	VLAN	Cost	Preferred	Advertise	ICMP Probe	Description
10.1.1.0/31	10.0.2.1	10.0.0.2	GE3	<input checked="" type="checkbox"/> 100	0	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	ICMP_1	Description (Optional)

NSD Routes

Subnet	NSD	Gateway	Cost	Preferred	Advertise
1.2.3.70/32	NVS1		0	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
169.254.11.0/24	wsd	GW1	5	<input checked="" type="checkbox"/>	<input type="checkbox"/>
1.2.33.0/24	wsd	GW1	5	<input checked="" type="checkbox"/>	<input type="checkbox"/>

You can click the **IPv6** tab to configure static routes for IPv6 addresses.

Static Route Settings

IPv4 IPv6

Subnet	Source IP	Next Hop	Interface	VLAN	Cost	ICMP Probe	Description
2000:1b96:779a:0000:0000:8ae2:7334:0300	N/A	e.g. 2001:0db8:85a3:0000:0000:8ade:0370:7394	[none]	<input type="checkbox"/>	0	[none]	Description (Optional)

Configure the settings as follows:

Option	Description
Subnet	<p>Enter the IPv4 or IPv6 address of the Static Route Subnet that should be advertised.</p> <p>The IPv6 Subnet supports the following address format:</p> <ul style="list-style-type: none"> IPv6 global unicast address (2001:CAFE:0:2::1) IPv6 unique local address (FD00::1234:BEFF:ACE:E0A4) IPv6 Default (::/0)
Source IP	<p>Enter the corresponding IPv4 or IPv6 address of the selected VLAN. This option is available only when you select the VLAN check box.</p>
Next Hop	<p>Enter the next hop IPv4 or IPv6 address for the static route.</p> <p>The IPv6 next hop supports the following address format:</p> <ul style="list-style-type: none"> 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	<p>Choose the WAN Interface to which the static route would be bounded.</p> <p>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.</p>
VLAN	<p>Select the check box and enter the VLAN ID.</p>
Cost	<p>Enter the cost to apply weightage on the routes. The range is from 0 to 255.</p>

Option	Description
Preferred	<p>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.</p> <p>The static route will be matched only when the corresponding VPN routes are not available.</p> <p>Note This option is not available for IPv6 address type.</p>
Advertise	<p>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.</p> <p>Note This option is not available for IPv6 address type.</p>
ICMP Probe	<p>Choose an ICMP probe from the drop-down menu. The SD-WAN 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.</p> <p>Note This option is not supported for IPv6 address type.</p>
Description	Enter an optional description for the static route.

In addition, you can configure the NSD Static Routes. The NSD Static Routes are configured in the **Network Services** and are 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.

- Click **Save Changes** in the **Device** tab.

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 VMware. The **Device Settings** area provides sections for specifying ICMP Probes and Responders.

ICMP Probes can be specified settings for Name, VLAN Tagging (none, 802.1q, 802.1ad, QinQ (0x8100), or QinQ (0x9100)), C-Tags, S-Tags, Source/Destination/Next Hop IPs, Frequency to send ping requests, and Threshold the value for number of missed pings that will cause route to be marked unreachable.

ICMP Responders can be specified settings for **Name**, **IP Address**, and **Mode** (**Conditional** or **Always**).

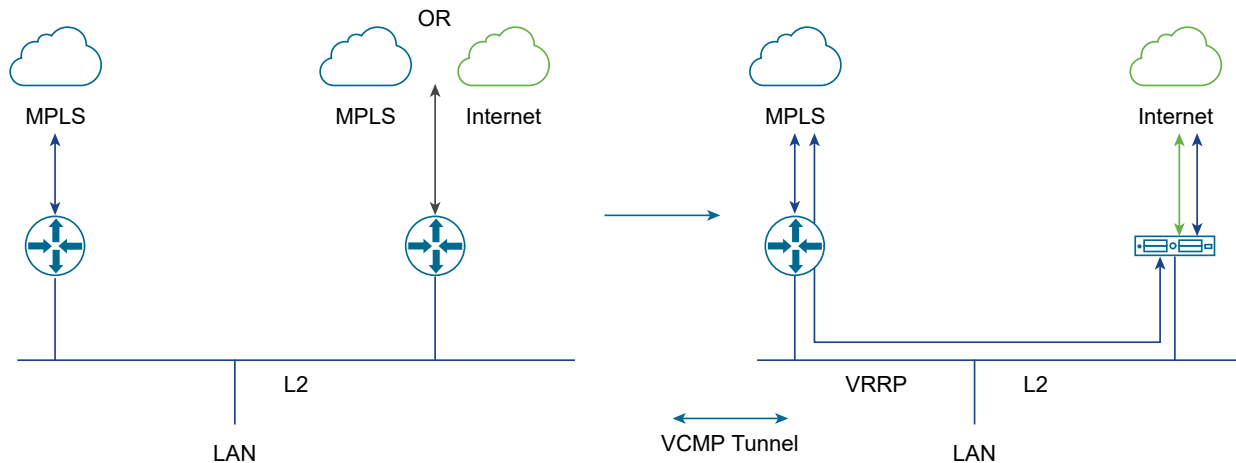
- **Always:** Edge always responds to ICMP Probes.
- **Conditional:** Edge only responds to ICMP Probes when the SD-WAN Overlay is up.

The screenshot shows two configuration panels. The top panel, titled 'ICMP Probes', includes fields for Name, VLAN Tagging (set to 'none'), C-Tag, S-Tag, Source IP, Destination IP, Next Hop IP, Frequency, and Threshold. The bottom panel, titled 'ICMP Responders', includes fields for Name, IP Address, and Mode (set to 'Conditional'). Both panels have a 'Clone' button.

Configure VRRP Settings

You can configure Virtual Router Redundancy Protocol (VRRP) on an Edge to enable next-hop redundancy in the SD-WAN 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.

The following illustration shows a network configured with VRRP:



Prerequisites

Consider the following guidelines before configuring VRRP:

- You can enable VRRP only between the SD-WAN Edge and third party router connected to the same subnet through an L2 switch.
- You can add only one SD-WAN 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.
- SD-WAN Edge must be configured as the primary VRRP device, by setting higher priority, in order to steer the traffic through SD-WAN.

- If the SD-WAN 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 SD-WAN 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 SD-WAN Edge always remains as the primary device.

Procedure

- 1 In the Enterprise portal, click **Configure > Edges**.
- 2 Either click the **Device** Icon corresponding to the Edge, or click the Edge and then click the **Device** tab.
- 3 In the **Device** tab, select the **VRRP with Third-Party Router** checkbox under **High Availability**.

4 In the **VRRP Settings**, configure the following:

Configure Segments

Configure Segment: Global Segment [Regular]

Authentication Settings ☐ Enable Edge Override

DNS Settings ☐ Enable Edge Override

Netflow Settings ☐ Enable Edge Override

LAN-Side NAT Rules ☐ Enable Edge Override

Syslog Settings ☐ Enable Edge Override

Static Route Settings

ICMP Probes

ICMP Responders

VRRP Settings

VRID: 5 Interface: 1 - Corporate Virtual IP: 10.3.0.200 Advertise Interval: 5 Priority: 110 Preempt (Delay): 10 ☒ Clone

Cloud VPN ☒ On

OSPF Areas ☒ On

BGP Settings ☐ Off

Multicast Settings ☐ Off

Cloud Security Service ☐ Off

Gateway Handoff Assignment ☐ Enable Edge Override

High Availability

Type: ☐ None ☐ Active Standby Pair ☐ Cluster ☒ VRRP with Third-Party Router

Segment Name	VRID	Interface	Virtual IP	Advertise Interval	Priority	Preempt (Delay)
Global Segment	5	[VLAN] 1 - Corporate	10.3.0.200	5	110	<input checked="" type="checkbox"/> (10)
Segment1	7	GE5/100	172.19.11.200	2	120	<input checked="" type="checkbox"/> (10)

- VRID** – Enter the VRRP group ID. The range is from 1 to 255.
- 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 checkbox so that SD-WAN Edge can preempt the third-party router which is currently the primary device, after the specified preempt delay.

5 Click **Save Changes**.

Results

In a branch network VLAN, if the Edge goes down, then the clients behind the VLAN are redirected through the backup router.

The SD-WAN Edge that acts as a primary VRRP device becomes the default Gateway for the subnet.

If the SD-WAN Edge loses connectivity with all the SD-WAN Edge/Controllers, then the VRRP priority gets reduced to 10 and the SD-WAN Edge withdraws the routes learned from the SD-WAN 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.

SD-WAN Edge automatically tracks overlay failure to the SD-WAN Edge. When all the overlay paths to the SD-WAN Edge are lost, the VRRP priority of the SD-WAN 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 SD-WAN Edge sends the route withdrawal request to the SD-WAN Edge/Controller and all the remote SD-WAN 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 SD-WAN 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 SD-WAN 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 SD-WAN 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 SD-WAN 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 SD-WAN Edge, when public overlay goes down the Edge becomes Backup. So the conditional backhaul feature cannot be utilized on a VRRP-enabled Edge.

Monitor VRRP Events

You can monitor the events related to changes in VRRP configuration.

In the enterprise portal, click **Monitor > Events**.

To view the events related to VRRP, you can use the filter option. Click the drop-down arrow next to the Search option and choose to filter by the Event column. The following events are available for VRRP:

- VRRP HA updated to primary

- VRRP HA updated out of primary
- VRRP Failed

The following image shows some of the VRRP events.

Event	Segment	Edge	User	Severity	Message
VRRP HA updated to primary	b7-edge1	b7-edge1		Notice	Get into VRRP master state
VRRP HA updated out of primary	b7-edge1	b7-edge1		Notice	Get out of VRRP master state
VRRP HA updated to primary	b7-edge1	b7-edge1		Notice	Get into VRRP master state
VRRP HA updated out of primary	b7-edge1	b7-edge1		Notice	Get out of VRRP master state
VRRP HA updated to primary	b7-edge1	b7-edge1		Notice	Get into VRRP master state
VRRP HA updated to primary	b7-edge1	b7-edge1		Notice	Get into VRRP master state
VRRP HA updated out of primary	b7-edge1	b7-edge1		Notice	Get out of VRRP master state
VRRP HA updated to primary	b7-edge1	b7-edge1		Notice	Get into VRRP master state
VRRP HA updated out of primary	b7-edge1	b7-edge1		Notice	Get out of VRRP master state
VRRP HA updated to primary	b7-edge1	b7-edge1		Notice	Get into VRRP master state
VRRP HA updated to primary	b7-edge1	b7-edge1		Notice	Get into VRRP master state

You can also view the events in the new Orchestrator UI.

Click **Open New Orchestrator UI** and then click **Launch New Orchestrator UI** in the pop-up window. The UI opens in a new tab displaying the monitoring options. Click **Events**. Click the Filter icon in the Search option to filter the VRRP events.

Configure Cloud VPN and Tunnel Parameters at the Edge level

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 the Branch to Non SD-WAN Destination via Edge settings inherited from a Profile and configure Tunnel parameters (WAN link selection and Per tunnel credentials).

- 1 From the SD-WAN Orchestrator, go to **Configure > Edges**.
- 2 Select an Edge you want to override Non SD-WAN Destination settings and click the icon under the **Device** column. The Device Setting page for the selected Edge appears.
- 3 Go to the **Branch to Non SD-WAN Destination via Edge** area and select the **Enable Edge Override** checkbox.

Cloud VPN ⓘ On

Branch to Non SD-WAN Destination via Gateway
Enable: ✕

Branch to Hubs
Enable: ✕

Branch to Branch VPN
Enable: ☒
Isolate Profile ✕

Use Cloud Gateways ☒ Use Hubs for VPN ✕

Dynamic Branch To Branch VPN ✕

To All Edges ☒
To Edges Within Profile ✕

Branch to Non SD-WAN Destination via Edge ☒ Enable Edge Override ⓘ
Enable ☒

Service				Link			
Action	Name	Automation for all public WAN Links	Enable Service	Enable tunnel	Destination Primary Public IP	Destination Secondary Public IP	Action
<div>+</div> <div>-</div>	Azure_New <div> <div></div> <div></div> </div>	N/A	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/> 104.208.31.2...	52.185.65.128	52.185.67.187	Edit Add Del
				<input checked="" type="checkbox"/> 104.208.26.99	52.185.65.128	52.185.67.187	Edit Add Del

- 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.

- Under **Action**, click **Add** to add tunnels. The **Add Tunnel** pop-up window appears.

Add Tunnel

Public Wan Link

104.208.31.249

Local Identification Type

FQDN/Hostname

Local Identification ⓘ

a.com

PSK

.....

Destination Primary Public IP

52.185.65.128

Destination Secondary Public IP

52.185.67.187

Save Changes

Cancel

- 6 Enter the following details for configuring a tunnel to the Non SD-WAN Destination and click **Save Changes**.

Field	Description
Public WAN Link	
Local Identification Type	<p>Select any one of the Local authentication type from the drop-down menu:</p> <ul style="list-style-type: none"> ■ 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 IP address used to communicate with the local gateway.
Local Identification	<p>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 address.</p>
PSK	<p>Enter the Pre-Shared Key (PSK), which is the security key for authentication across the tunnel in the textbox.</p>
Destination Primary Public IP	<p>Enter the Public IP address of the destination Primary VPN Gateway.</p>
Destination Secondary Public IP	<p>Enter the Public IP address of the destination Secondary VPN Gateway.</p>

- 7 Click **Save Changes**.

Configure VLAN for Edges

At an 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, SD-WAN 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.

To configure VLAN settings for an Edge:

- 1 In the Enterprise portal, click **Configure > Edges**.
- 2 Click the Device Icon next to an Edge, or click the link to an Edge and then click the **Device** tab.
- 3 In the **Device** tab, scroll down to the **Configure VLAN** section.

Configure VLAN

+ Add VLAN

+ Add Secondary IP

	Override ⓘ							Multicast		
Action	VLAN DHCP	VLAN	Network	IP Address	Interfaces	DHCP	Segment	IGMP	PIM	VNF Insertion
Edit ⓘ ⚠	<input checked="" type="checkbox"/> ✕	1 - Corporate	10.1.1.0/24	10.1.1.1		Enabled (242)	Global Segment			✕
Edit ⓘ	ⓘ	2 - VLAN-2	10.11.1.0/24	10.11.1.1	GE2 GE8	Enabled (242) ⓘ	Global Segment			✕
Edit Del		Secondary IP	10.12.1.0/24	10.12.1.1						
			10.13.1.0/24	10.13.1.1						
Edit ⓘ	ⓘ	100 - VLAN-100	10.101.1.0/24	10.101.1.1	GE2	Enabled (242) ⓘ	segment1			✕
Edit Del		Secondary IP	10.102.1.0/24	10.102.1.1						
			10.103.1.0/24	10.103.1.1						
Edit ⓘ	ⓘ	101 - VLAN-101	10.111.1.0/24	10.111.1.1	GE2	Enabled (242) ⓘ	segment2			✕
Edit Del		Secondary IP	10.112.1.0/24	10.112.1.1						
			10.113.1.0/24	10.113.1.1						
Edit ⓘ	ⓘ	200 - VLAN-200	10.200.1.0/24	10.200.1.1	GE2 GE8	Enabled (242) ⓘ	segment1			✕
Edit Del		Secondary IP	10.201.1.0/24	10.201.1.1						
			10.202.1.0/24	10.202.1.1						
Edit ⓘ	ⓘ	201 - VLAN-201	10.211.1.0/24	10.211.1.1	GE2 GE8	Enabled (242) ⓘ	segment2			✕
Edit Del		Secondary IP	10.212.1.0/24	10.212.1.1						
			10.213.1.0/24	10.213.1.1						

You can add or edit a VLAN and add multiple secondary IP addresses to the VLAN.

Add VLANs

To add a VLAN, click **Add VLAN**.

VLAN
? ✕

The VLAN is configured for this Edge only and does not inherit any settings from the profile.

* Segment
Global Segment

* VLAN Name

* VLAN Id

Assign Overlapping Subnets
✕ ⓘ

* Edge LAN IP Address

* Cidr Prefix
24

Network

Advertise
☒

ICMP Echo Response
☒

VNF Insertion
✕ VNF insertion requires that the selected segment have a Service VLAN

Multicast
Multicast is not enabled for the selected segment

Fixed IPs

MAC Address	IP	Description		
Ex: aa:bb:cc:dd:ee:ff	Ex: 10.0.2.5	Description (optional)	-	+

LAN Interfaces
N/A

SSID
N/A

DHCP

Type
Enabled Relay Not Enabled

☒ Source from Secondary IP(s)

* Relay Agent IP(s)
 - +

OSPF

Enabled
✕ OSPF not enabled for the selected Segment.

Add VLAN Cancel

Configure the following settings:

Table 18-1.

Option	Description
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.

Table 18-1. (continued)

Option	Description
Assign Overlapping Subnets	<p>The LAN IP Addressing is managed from the assigned Profile of the Edge. When this checkbox 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.</p> <hr/> <p>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.</p> <hr/>
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 checkbox to advertise the VLAN to other branches in the network.
ICMP Echo Response	Select the checkbox to enable the VLAN to respond to ICMP echo messages.
VNF Insertion	Select the checkbox 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	<p>This option is enabled only when you have configured multicast settings for the Edge. You can configure the following multicast settings for the VLAN.</p> <ul style="list-style-type: none"> ■ IGMP ■ PIM <p>Click toggle advanced multicast settings to set the timers:</p> <ul style="list-style-type: none"> ■ 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 the LAN Interfaces for the VLAN.
SSID	Configure the Wi-Fi SSID details for the VLAN.

Table 18-1. (continued)

Option	Description
DHCP Type	<p>Choose one of the following DHCP settings:</p> <p>Enabled – Enables DHCP with the Edge as the DHCP server. Configure the following details:</p> <ul style="list-style-type: none"> ■ 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. <p>Relay – Enables DHCP with the DHCP Relay Agent installed at a remote location. If you choose this option, configure the following:</p> <ul style="list-style-type: none"> ■ Source from Secondary IP(s) – When you select this checkbox, 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. <p>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.</p> <ul style="list-style-type: none"> ■ Relay Agent IP(s) – Specify the IP address of Relay Agent. Click the Plus(+) Icon to add more IP addresses. <p>Not Enabled – Deactivates DHCP.</p>
OSPF	<p>This option is enabled only when you have configured OSPF for the Edge. Select the checkbox and choose an OSPF from the drop-down list.</p>

After configuring the required parameters, click **Add VLAN**.

Edit VLANs

To update the existing VLAN settings inherited from the Profile, click the **Edit** link corresponding to the VLAN.

VLAN

* Segment

Global Segment ▼

Enable Edge Override

* VLAN Name

Corporate ⓘ

* VLAN Id

1 ⓘ

Assign Overlapping Subnets

✖ ⓘ

* Edge LAN IP Address

10.0.1.1

* Cidr Prefix

24

Network

10.0.1.0

Advertise

✓

ICMP Echo Response

✓

VNF Insertion

✖ VNF insertion requires that the selected segment have a Service VLAN

Multicast

Multicast is not enabled for the selected segment

Fixed IPs

MAC Address	IP	Description
02:42:0a:00:01:19	10.0.1.25	Description (optional) - +

LAN Interfaces

GE1 GE2

SSID

There are no Wi-Fi SSIDs configured on this VLAN.

DHCP

Enable Edge Override

Type

Enabled Relay Not Enabled

* DHCP Start ⓘ

10.0.1.13

* Num. Addresses

242

* Lease Time

1 day ▼

Options

Option	Code	Data Type	Value
add an option ▼			

OSPF

Enable Edge Override

Enabled

✖ OSPF not enabled for the selected Segment.

Update VLAN

Cancel

Click the **Enable Edge Override** checkboxes to override the VLAN settings inherited from the Profile.


Note You cannot override the Profile VLAN name and ID.

After modifying the required parameters, click **Update VLAN**.

For Configuring VLANs at the Profile level, see [Configure VLAN for Profiles](#).

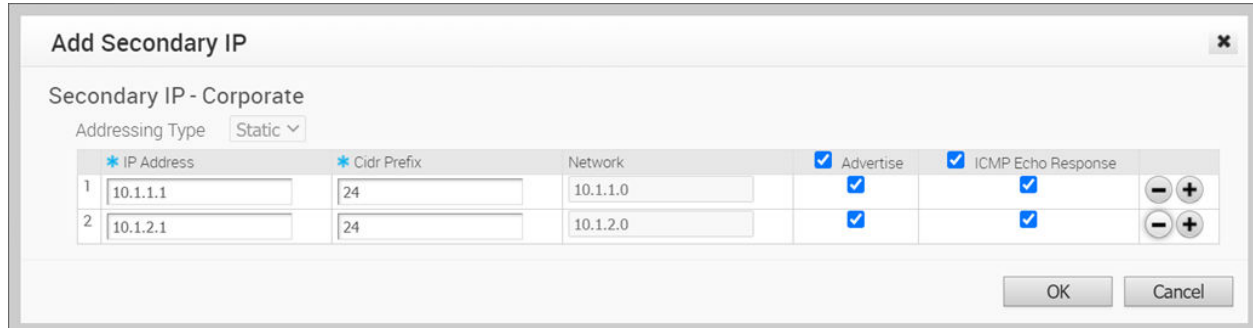
Secondary IP Addresses

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. To add secondary IP addresses to the VLAN, click **Add Secondary IP**.



In the **Add Secondary IP** window, select a VLAN from the drop-down list and click **Next**.

Configure the following settings:



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 checkbox to advertise the secondary IP address network of the VLAN to other branches in the network.
ICMP Echo Response	Select the checkbox to enable the VLAN with the secondary IP address to respond to ICMP echo messages.

Click the Plus (+) Icon to add more IP addresses to the VLAN.

Note You can add up to 16 secondary IP addresses to a VLAN.

Click **OK**.

In the **Device** tab, click **Save Changes** to save the settings.

Loopback Interfaces Configuration

A loopback interface is a logical interface that allows you to assign an IP address, which is used to identify a VMware SD-WAN Edge.

You can configure loopback interfaces only for SD-WAN Edges that are running on version 4.3 and above. The **Configure Loopback Interfaces** area is not available for SD-WAN Edges that are running on version 4.2 or lower. For such Edges, you must configure Management IP address. For details, refer to [Configure the Management IP Address](#).

Loopback Interfaces—Benefits

Following are the benefits of configuring loopback interfaces for an Edge:

- As loopback interfaces are logical interfaces that are always up and reachable, 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.

Loopback Interfaces—Limitations

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 SD-WAN 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.

Configure a Loopback Interface for an Edge

To configure a loopback interface for an Edge:

Prerequisites

For information about the rules and notes that you must consider before you configure a loopback interface, see [Loopback Interfaces—Limitations](#).

Procedure

- 1 Log in to VMware SD-WAN Orchestrator, and then go to **Configure > Edges**.
- 2 Either click the Device icon next to the Edge for which you want to configure the loopback interface or click the Edge link, and then go to the **Device** tab.
- 3 Scroll down to the **Configure Loopback Interfaces** area, and then click **Add Loopback**.
- 4 In the **Add Loopback** modal popup, configure the required loopback settings. For details, refer to [Loopback Interfaces—Field References](#).
- 5 Click **Add Loopback**.

6 Click **Save Changes**.

Results

The loopback interface is listed in the **Configure Loopback Interfaces** area.

At any point in time, you can choose to edit the loopback interface settings, 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 more 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.

What to do next

Once you configure the loopback interface for an Edge, you can select the interface as the source interface for the following services:

Services/Settings	For details, refer to ...
Orchestrator Management Traffic	Configure Orchestrator Management Traffic for Edges
Authentication Settings	Configure Authentication Settings
DNS Settings	Configure DNS Settings
Netflow Settings	Configure Netflow Settings for Edges
Syslog Settings	Configure Syslog Settings for Edges

Services/Settings	For details, refer to ...
BGP Settings	Configure BGP from Edge to Underlay Neighbors
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.

Loopback Interfaces—Field References

The following table lists and describes the fields that must be configured for a loopback interface:

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 checkbox 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 checkbox to advertise the loopback interface to other branches in the network.
OSPF	Select the checkbox and choose an OSPF Area from the drop-down list. The loopback interface IP address is advertised in the selected OSPF area. Note This option is enabled only when you have configured OSPF for the segment that you have selected for the loopback interface.
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** checkboxes for the IPv4 and IPv6 settings, to enable the corresponding addressing type for the Interface. By default, the option is enabled for IPv4 settings.

Related Topics

- [Loopback Interfaces Configuration](#)
- [Configure a Loopback Interface for an Edge](#)

Configure Orchestrator Management Traffic for Edges

You can configure the Orchestrator Management Traffic for the Edge to transmit the traffic to VMware SD-WAN Orchestrator.

To configure the Orchestrator Management Traffic:

Procedure

- 1 Log in to VMware SD-WAN Orchestrator, and then go to **Configure > Edges**.
- 2 Either click the Device icon next to the Edge for which you want to configure the Orchestrator Management Traffic or click the Edge link, and then go to the **Device** tab.
- 3 Scroll down to the **Orchestrator Management Traffic** area, and then from the **Source Interface** drop-down list, 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 VMware SD-WAN Orchestrator.

Results

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.

Configure Device Settings

The Edge **Device Settings** screen provides the ability to do the following tasks:

- Set VLAN Settings
- Override Syslog Settings
- Override Profile Interface Settings
- Add a User Defined WAN Overlay
- Configure NAT for overlapping Network

Configure DHCP Server on Routed Interfaces

You can configure DHCP server on a Routed Interface in an SD-WAN Edge.

To configure DHCP Server settings:

- 1 In the Enterprise portal, click **Configure > Edges**.

- 2 Click the Device Icon next to an Edge, or click the link to the Edge, and then click the **Device** tab.
- 3 Scroll down to the **Device Settings** section and click the DOWN arrow to view the **Interface Settings** for the Edge.
- 4 The **Interface Settings** section displays the existing interfaces available in the Edge.
- 5 Click the **Edit** option for the Routed interface that you want to configure DHCP settings.

Virtual Edge

Interface GE3

Interface Enabled

Capability

Segments

RADIUS Authentication ⓘ

ICMP Echo Response

Underlay Accounting ⓘ

Enable WAN Overlay

VLAN

IP Preference ⓘ

EVDSL Modem Attached

☒

Routed

All Segments

✖ WAN Overlay must be turned off to configure RADIUS Authentication.
Require User Authentication to access WAN

☒

☒

☒

☒ IPv4 ☐ IPv6

☐

☒ Override Interface

IPv4 Settings

Addressing Type

IP Address

CIDR prefix

Gateway

Static

169.254.7.10

29

169.254.7.9

WAN Overlay

OSPF

Multicast

VNF Insertion

Advertise

NAT Direct Traffic

Trusted Source ⓘ

Reverse Path Forwarding ⓘ

Auto-Detect Overlay

 unlock

✖ OSPF not enabled for the selected Segment.

Multicast is not enabled for the selected segment

✖ VNF insertion is disallowed when an interface is configured for WAN overlays

☐

☒

☐

Specific

☒ Active

IPv6 Settings

Addressing Type

IP Address

CIDR prefix

Gateway

DHCP Stateless

N/A

N/A

N/A

WAN Overlay

Trusted Source ⓘ

Reverse Path Forwarding ⓘ

Auto-Detect Overlay

☒

Specific

☒ Active

L2 Settings

Autonegotiate

MTU

Enable LOS Detection

☒

1500

☐

DHCP Server

Type

DHCP Start ⓘ

Num. Addresses

Lease Time

Enabled

 Relay Not Enabled

169.254.7.10

0

1 hour

Options

Option	Code	Data Type	Value
add an option			

Update GE3

Cancel

- In the **IPv4 Settings** section, select the **Addressing Type** as **Static** and enter the IP addresses for the Edge Interface and the Gateway.

Note 31-bit prefixes are supported for IPv4 as per RFC 3021.

7 In the **DHCP Server** section, choose one of the following DHCP settings:

- **Enabled** – Allows the DHCP with the Edge as the DHCP server. Configure the following details:
 - **DHCP Start** – Enter a valid IP address available within the subnet.
 - **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** – Add pre-defined or custom DHCP options from the drop-down menu. The DHCP option is a network service passed to the clients from the DHCP server. For a custom option, enter the code, data type, and value. The table below lists the DHCP options for IPv4 and IPv6:

Table 18-2. DHCP Options for IPv4

Option	Code	Description
Time offset	2	Specifies the offset of the client's subnet in seconds, from Coordinated Universal Time (UTC).
DNS server	6	<p>Lists Domain Name System (RFC 1035) servers available to the client. Servers are listed in order of preference.</p> <hr/> <p>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.</p>
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.

Table 18-2. DHCP Options for IPv4 (continued)

Option	Code	Description
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 18-3. DHCP Options for IPv6

DHCP Option Name	Code	Description
SIP server names	21	Lists the domain names of the SIP outbound proxy servers that the client can use.
SIP server addresses	22	Lists the IPv6 addresses of the SIP outbound proxy servers that the client can use.
DNS Recursive Name Servers	23	Lists IPv6 addresses of DNS recursive name servers to which DNS queries may be sent by the client resolver in order of preference.
Domain list	24	Provides a domain search list for the client, to be used when resolving hostnames through DNS.
NIS servers list	27	Provides an ordered list of NIS servers with IPv6 addresses available to the client.
NIS Domain name	29	Provides the NIS domain name to be used by the client.
SNTP server	31	Provides an ordered list of SNTP servers with IPv6 addresses available to the client.
Information refresh time	32	Specifies the upper bound of the number of seconds from the current time that a client should wait before refreshing information received from the DHCPv6 server, particularly for stateless DHCPv6 scenarios.

Table 18-3. DHCP Options for IPv6 (continued)

DHCP Option Name	Code	Description
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. Click the Plus(+) Icon to add more IP addresses.
- **Not Enabled** – Deactivates DHCP.

For more information on other options in the **Interface Settings** window, see [Configure Interface Settings](#).

Note See also [Tunnel Overhead and MTU](#) for more information.

Enable RADIUS on a Routed Interface

RADIUS can be enabled on any interface that is configured as a routed interface. The SD-WAN Edge supports both username/password (EAP-MD5) and certificate (EAP-TLS) based 802.1x Authentication methods.

Requirements

- 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.

Note RADIUS enabled interfaces do not use DPDK.

Enabling RADIUS on a Routed Interface

- 1 In the Enterprise portal, click **Configure > Edges**.
- 2 Click the Device Icon next to an Edge, or click the link to the Edge, and then click the **Device** tab.
- 3 Scroll down to the **Device Settings** section and click the DOWN arrow to view the **Interface Settings** for the Edge.

- 4 The **Interface Settings** section displays the existing interfaces available in the Edge.
- 5 Click the **Edit** option for the Routed interface that you want to enable RADIUS authentication.

Interface GE3

Interface Enabled ☒

Capability Routed

Segments Global Segment

RADIUS Authentication ☒
 Require User Authentication to access WAN
 Add mac-addresses of devices that are pre-authenticated (allowlist) that should not be forwarded to RADIUS for re-authentication.

Mac Address or OUI	Description
Ex: aa:bb:cc:dd:ee:ff	Description (Optional)

ICMP Echo Response ☒

Underlay Accounting ☒

Enable WAN Overlay ☐

VLAN

- 6 Deactivate the **WAN Overlay** option.
- 7 Select the **RADIUS Authentication** checkbox.
- 8 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 more information on other options in the **Interface Settings** window, see [Configure Interface Settings](#).

Configure Edge LAN Overrides

The LAN settings specified in the Profile can be overridden by selecting the **Override Interface** check box.

See [Chapter 13 Configure a Profile Device](#) for LAN interface configuration parameters.

Edge 500: LAN1

Interface LAN1 ☒ **Override Interface**

Interface Enabled ☒

Capability Switched

Mode Access Port

VLAN's 1 - Corporate

L2 Settings ☒

Autonegotiate ☒

MTU 1500

Update LAN1 Cancel

Configure Edge WAN Overrides

The WAN settings specified in the Profile can be overridden by selecting the **Override Interface** checkbox.

See [Chapter 13 Configure a Profile Device](#) for LAN interface configuration parameters.



Configure Edge WAN Overlay Settings

The WAN 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 VMware SD-WAN Gateways over the Internet, as determined by the SD-WAN 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 SD-WAN 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.

Procedure

- 1 In the SD-WAN Orchestrator portal, click **Configure > Edges**.
- 2 In the **Edges** page, either click the device Icon next to an Edge or click the link to the Edge and click the **Device** tab.
- 3 Scroll down to **WAN Settings**.

WAN Settings + Add User Defined WAN Overlay								
Actions	Type	Name	Address Type	Interfaces	Link Type	Public IP	Pre-Notifications	Alerts
Edit Del	User Defined	GE6_Private	IPv4	GE6	Private Wired		✕	✕
Edit Del	User Defined	GE7_Private	IPv4	GE7	Private Wired		✕	✕
Edit Del	Auto Detect	169.254.9.3	IPv4	GE3	Public Wired	169.254.9.3	☑	☑
Edit Del	Auto Detect	169.254.7.10	IPv6	GE4	Public Wired	fe00:1:1:2::2	☑	☑

- 4 For an existing auto-detected or user-defined WAN Overlay, click **Edit** to modify the settings.
- 5 To create a new Public or Private overlay, click **Add User Defined WAN Overlay**.
- 6 In the **User Defined WAN Overlay** window, 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 18-4. Settings common for Public and Private Overlay

Option	Description
Address Type	<p>Select the IP address type from the drop-down list. You can choose the WAN overlay link to use either IPv4 or IPv6 address.</p> <p>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.</p>
Name	<p>Enter a descriptive name for the public or private WAN overlay link. You can reference this name while choosing a WAN link in a Business Policy.</p> <p>Note WAN overlay name should only consist of ASCII characters. Non-ASCII characters are not supported.</p> <p>See Configure Link Steering Modes.</p>
Pre-Notification 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	<p>The Routed Interfaces enabled with IPv4 WAN Overlay or IPv6 WAN Overlay and set to User Defined Overlay are displayed as checkboxes. The Interfaces displayed are based on the selected Address Type. Select one or more routed interfaces and the current user-defined overlay is attached to the selected interface.</p> <p>Note For the 610-LTE, you can add User Defined WAN overlay on CELL1 or CELL2. The SD-WAN Orchestrator will display both CELL1 and CELL2, irrespective of SIM presence. Therefore, you must be aware of which SIM slot is enabled (Active) and choose that SIM.</p>

Table 18-5. Public Overlay Settings

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 Settings for Public Overlay:

Virtual Edge: new link

User Defined WAN Overlay

Address Type

IPv6

Link Type

Public

Name

GE6_Public

Public IP Address

n. a.

Pre-Notification Alerts

✖

Alerts

✖

Select Interfaces

☐ GE3

☒ GE6

Optional Configuration ⓘ

Source IP Address: ⓘ

Enter IPv4 Or IPv6 Address

Next-Hop IP Address: ⓘ

Enter IPv4 Or IPv6 Address

Custom VLAN

☐

802.1P Setting ⓘ

☐

Advanced

Update Link

Cancel

Table 18-6. Private Overlay Settings

Option	Description
SD-WAN Service Reachable	<p>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:</p> <ul style="list-style-type: none"> ■ A secondary path to the internet for access to internet hosted SD-WAN 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. <p>Note The SD-WAN 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).</p> <p>Note Per-packet or round-robin load balancing will not be performed between the short and long paths.</p> <p>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 VMware SD-WAN service.</p>
Public SD-WAN Addresses	<p>When you select the SD-WAN Service Reachable checkbox, a list of public IPv4 and IPv6 addresses of SD-WAN Gateways and SD-WAN 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.</p> <p>Note Some IP addresses in the list, such as Gateways, may change over time.</p>

The following image shows an example of Settings for Private Overlay:

Table 18-7. Optional Configuration

Option	Description
Source IP Address	<p>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.</p> <p>Source IP address does not have to be pre-configured anywhere but must be routable to and from the selected interface.</p> <p>You can enter IPv4 or IPv6 address to establish WAN overlay with the peer.</p>
Next-Hop IP Address	<p>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.</p> <p>You can enter IPv4 or IPv6 address.</p>
Custom VLAN	<p>Select the checkbox to enable custom VLAN and enter the VLAN ID. The range is 2 to 4094.</p> <p>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.</p>
802.1P Setting	<p>Sets 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.</p> <p>This checkbox is available only when the system property session.options.enable8021PConfiguration must be set to True. By default, this value is False.</p> <p>If this option is not available for you, contact the VMware support of your operations team to enable the setting.</p>

Click **Advanced** to configure the following settings:

Table 18-8. Advanced Settings common for Public and Private Overlay

Option	Description
Bandwidth Measurement	<p>Choose a method to measure the bandwidth from the following options:</p> <ul style="list-style-type: none"> ■ 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 SD-WAN 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: <ul style="list-style-type: none"> 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.
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	<p>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.</p> <hr/> <p>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.</p> <hr/> <p>Note This configuration is not supported with public link CoS.</p> <hr/>

Table 18-8. Advanced Settings common for Public and Private Overlay (continued)

Option	Description
Link Mode	<p>Select the mode of the WAN link from the drop-down. The following options are available:</p> <ul style="list-style-type: none"> ■ 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. <p>Only one interface on an Edge can be put into Backup Mode. When so configured, the interface will be displayed in Monitor > Edges page as Cloud Status: Standby.</p> <hr/> <p>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.</p> <ul style="list-style-type: none"> ■ 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. <p>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.</p> <p>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).</p> <p>For more information, see Configure Hot Standby Link.</p>

Table 18-8. Advanced Settings common for Public and Private Overlay (continued)

Option	Description
	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.
MTU	<p>The SD-WAN 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.</p> <p>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.</p> <p>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.</p> <p>For more information, see Tunnel Overhead and MTU.</p>
Overhead Bytes	<p>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.</p> <p>When you configure the Overhead Bytes, the bytes are additionally accounted for by the QoS scheduler 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.</p>
Path MTU Discovery	Select the checkbox 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 more information, see Tunnel Overhead and MTU .

Table 18-8. Advanced Settings common for Public and Private Overlay (continued)

Option	Description
Configure Class of Service	<p>SD-WAN 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.</p> <p>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.</p> <hr/> <p>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.</p> <hr/> <p>After you select this checkbox, configure the following:</p> <ul style="list-style-type: none"> ■ 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 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. <hr/> <p>Note The Dynamic Bandwidth Adjustment configuration is not supported with public link CoS.</p> <hr/> <p>For more information about how to configure CoS, see Configure Class of Service.</p>
Strict IP precedence	<p>This checkbox is available when you select the Configure Class of Service checkbox.</p> <p>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.</p> <p>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.</p>

Table 18-9. Advanced Settings for Public Overlay

Option	Description
UDP Hole Punching	<p>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.</p> <p>Use Branch to Branch VPN to enable branch to branch tunnels. See Configure a Tunnel Between a Branch and a Branch VPN and Configure Cloud VPN and Tunnel Parameters at the Edge level.</p> <p>Use Remote Diagnostics > List Paths to check that one Edge has built a tunnel to another Edge.</p> <p>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.</p> <p>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 SD-WAN Gateway as destination IP and destination port for creating a VCMP tunnel to this Edge overlay interface.</p> <hr/> <p>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.</p> <hr/> <p>UDP hole punching may introduce additional connectivity issues as remote sites try to use the new UDP dynamic port for VCMP tunnels.</p>
Type	<p>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 Configure Link Steering Modes.</p> <p>Choose Wired or Wireless, to put the overlay into a public wired or wireless transport group.</p>

The following image shows Advanced settings for a Public Overlay:

Advanced Settings

Bandwidth Measurement ⓘ

Do Not Measure (define manually) ⚙

Upstream Bandwidth (Mbps)

100

Downstream Bandwidth (Mbps)

100

Dynamic Bandwidth Adjustment ⓘ

☐

Link Mode ⓘ

Active ⚙ ⚠

MTU

1500

Overhead Bytes

0

Path MTU Discovery

☒

Public Link Configuration

UDP Hole Punching

☐

Type

Wired ⚙

Configure Class of Service

☒

Strict IP Precedence ⓘ :

☐

Class Of Service	DSCP Tags		Bandwidth (%)	Policing	Default Class	
c1	CS0	<input type="range"/>	50	<input type="checkbox"/>	<input checked="" type="radio"/>	⊖ ⊕
c2	AF21	<input type="range"/>	30	<input type="checkbox"/>	<input type="radio"/>	⊖ ⊕
c3	EF	<input type="range"/>	20	<input type="checkbox"/>	<input type="radio"/>	⊖ ⊕

Advanced

Update Link

Cancel

Table 18-10. Advanced Settings for Private Overlay

Option	Description
Private Network Name	<p>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.</p> <p>For example:</p> <p>Edge1 GE1 is attached to <i>private network A</i>. Use <i>private network A</i> for the private overlay attached to GE1.</p> <p>Edge1 GE2 is attached to <i>private network B</i>. Use <i>private network B</i> for the private overlay attached to GE2.</p> <p>Repeat the same attachment and naming for Edge2.</p> <p>When you enable branch to branch or when Edge2 is a hub site:</p> <ul style="list-style-type: none"> ■ 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	<p>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.</p> <hr/> <p>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.</p>

The following image shows Advanced settings for a Private Overlay:

Advanced Settings

Bandwidth Measurement: Measure Bandwidth (Slow Start)

Dynamic Bandwidth Adjustment: ☐

Link Mode: Hot Standby

Minimum Active Links: 1

MTU: 1500

Overhead Bytes: 0

Path MTU Discovery: ☒

Private Network Name: Private_Network_A

Private Link Configuration

Configure Static SLA: ☐

Configure Class of Service: ☒

Strict IP Precedence: ☒

Class Of Service	DSCP Tags	Bandwidth (%)	Policing	Default Class
CoS1	CS5, EF	60	<input checked="" type="checkbox"/>	<input type="radio"/>
CoS2	AF41, CS4	20	<input type="checkbox"/>	<input type="radio"/>
CoS3	AF21, CS2	20	<input type="checkbox"/>	<input checked="" type="radio"/>

Advanced Update Link Cancel

7 Click **Update Link** to save the settings.

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 VMware service.

For hybrid environments that have MPLS-only links or require failover to MPLS links, you can enable the SD-WAN Service Reachable option.

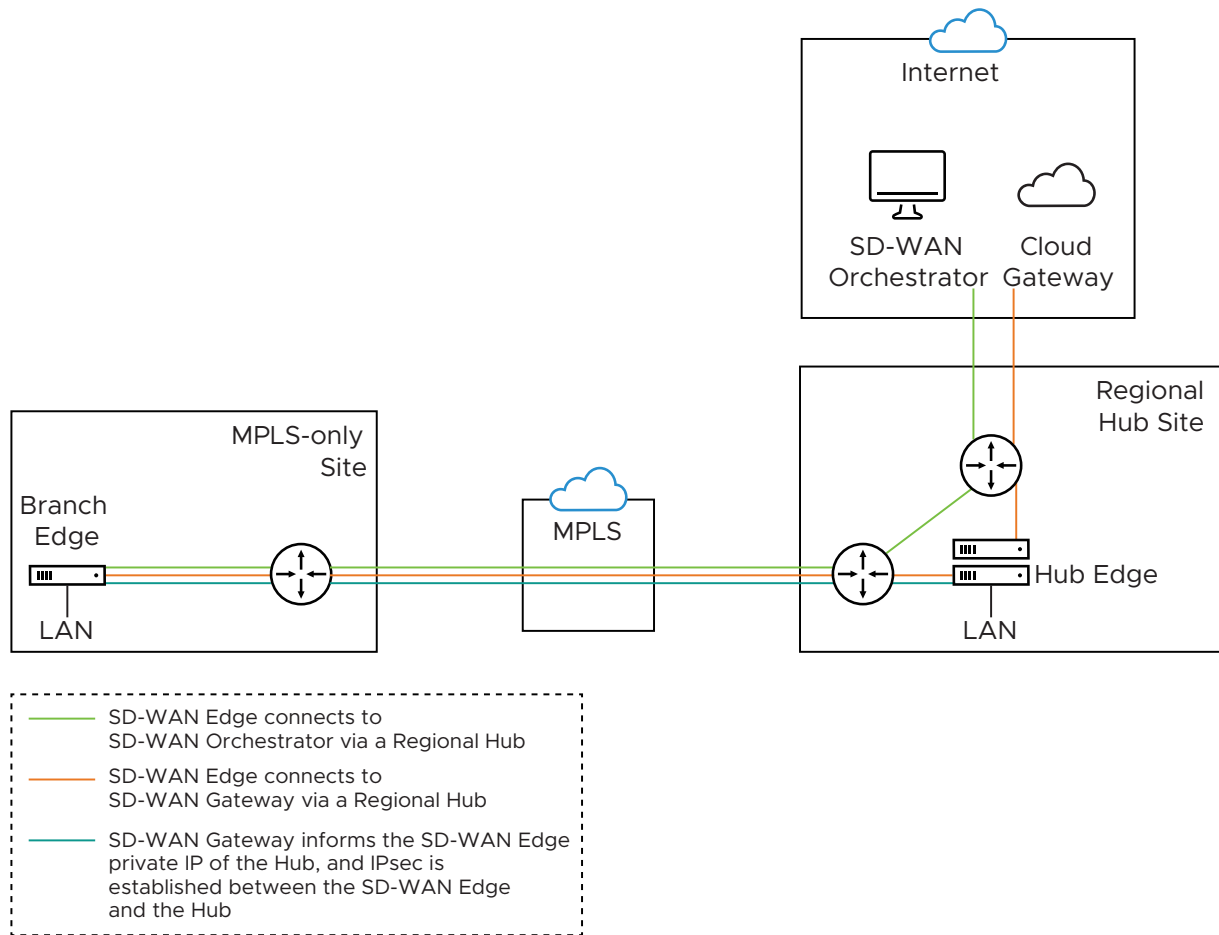
MPLS-only Sites

VMware supports private WAN deployments with a hosted VMware 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 VMware 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 SD-WAN Edge with only MPLS connection.



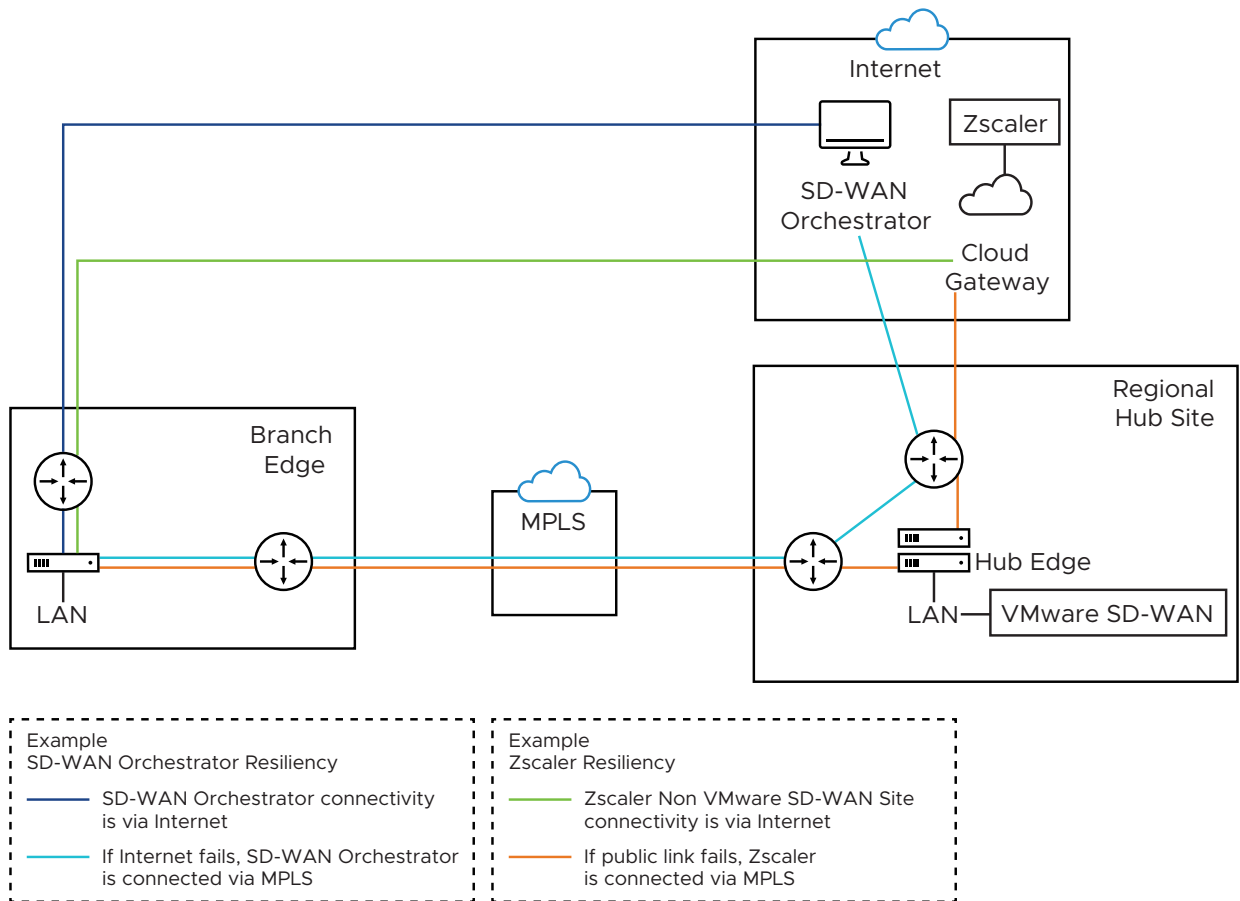
The traffic from the SD-WAN Edge with MPLS-only links is routed to the Orchestrator and Gateway through a Regional Hub, which is able to break out to the public cloud. SD-WAN Service Reachable option allows the Edge to remain online and manageable from the Orchestrator, and allows public internet connectivity 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 Internet Multipath traffic to a private WAN link.

Note Traffic classified as Direct to Internet will not failover to a private WAN link regardless of priority or service class. In other words, traffic where the Network Service is Direct does not failover to a private link even if the Priority is High and the Service Class is Real Time. Only traffic with a Multipath classification fails over to a private link.

The following image illustrates Resiliency of SD-WAN Orchestrator and Non SD-WAN Destination, Zscaler.



- **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 Enterprise portal, click **Configure > Edges**.
- 2 In the Edges page, either click the device Icon next to an Edge or click the link to the Edge and click the **Device** tab.
- 3 Scroll down to **Interface Settings** and **Edit** the Interface connected to the MPLS link.
- 4 In the **Interface** window, select the **User Defined Overlay** checkbox.

Virtual Edge

Interface GE6

Interface Enabled ☒

Capability **Routed**

Segments All Segments

Addressing Type **Static**

IP Address 172.16.1.10

CIDR prefix 29

Gateway 172.16.1.11

WAN Overlay ☒ **User Defined Overlay**

OSPF ✖ OSPF not enabled for the selected Segment.

VNF Insertion ✖ VNF Insertion is disallowed when an interface is configured for WAN overlays

Multicast Multicast is not enabled for the selected segment

RADIUS Authentication ⓘ ✖ WAN Overlay must be disabled to configure RADIUS Authentication.
Require User Authentication to access WAN

Advertise ☐

ICMP Echo Response ☒

NAT Direct Traffic ☒

Underlay Accounting ⓘ ☒

Trusted Source ⓘ ☐

Reverse Path Forwarding ⓘ **Specific**

VLAN

L2 Settings

Autonegotiate ☒

* MTU 1500

DHCP Server

Type **Enabled** **Relay** **Disabled**

Update GE6 **Cancel**

The **SD-WAN Service Reachable** is available only for a **User Defined Overlay** network.

- 5 In the **WAN Settings** section, **Edit** the Interface enabled with **User Defined Overlay**.

- In the **User Defined WAN Overlay** window, select the **SD-WAN Service Reachable** checkbox to deploy sites which only have a private WAN link and/or enable the capability to failover critical Internet traffic to a private WAN link.

When you select the **SD-WAN Service Reachable** checkbox, a list of public IP addresses of SD-WAN Gateways and SD-WAN 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.

- Configure other options as required and click **Update Link** to save the settings.

For more information on other options in the **WAN Overlay** window, see [Configure Edge WAN Overlay Settings](#).

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.

- In the Enterprise portal, click **Configure > Edges**.
- Either click the Device Icon next to an Edge or click the link to the Edge and click the **Device** tab.
- In the **WAN Settings** section, click **Add User Defined WAN Overlay** and then choose the Link Type as required, that is **Public** or **Private**.
- You can also define the CoS for an existing link by clicking **Edit**.
- In the **WAN Overlay** settings, click **Advanced** and select the **Configure Class of Service** checkbox. When you enable this option, the following settings appear and configure them appropriately. You can click the Plus (+) icon to add multiple class of services.
 - **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:** Enter a descriptive name for the class of service. The name can be a combination of alphanumeric and special characters.
- **DSCP Tags:** Click **Set** to assign DSCP tags to the class of service. You can select multiple 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.

6 Click **Update Link** to save the settings.

The following example shows multiple class of services with different set of DSCP tags.

Class of Service	Description	DSCP Tags	Policing
CoS1	Voice	CS5, EF	Activated
CoS2	Video	AF41, CS4	Deactivated
CoS3	File Transfer	AF21, CS2	Deactivated

Private Link Configuration

Configure Static SLA: ☐

Configure Class of Service: ☒

Strict IP Precedence ⓘ: ☒

Class Of Service	DSCP Tags	Bandwidth (%)	Policing	Default Class
CoS 1	CS5, EF	60	<input checked="" type="checkbox"/>	<input type="radio"/>
CoS 2	AF41, CS4	20	<input type="checkbox"/>	<input type="radio"/>
CoS 3	AF21, CS2	20	<input type="checkbox"/>	<input checked="" type="radio"/>

For more information on the WAN Overlay Settings, see [Configure Edge WAN Overlay Settings](#).

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.

Prerequisites

To configure a Hot Standby link on an Edge, ensure that the Edge is upgraded to software image version 4.0.0 or later.

Procedure

- 1 In the SD-WAN Orchestrator portal, click **Configure > Edges**.
- 2 In the **Edges** page, either click the device Icon next to an Edge or click the link to the Edge and click the **Device** tab.
- 3 Scroll down to **WAN Settings**.

WAN Settings + Add User Defined WAN Overlay									
Actions	Type	Name	Address Type	Interface	Link Type	Public IP	Pre-Notifications	Alerts	
Edit Del	User Defined	GE6_Private	IPv4	GE6	Private Wired		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	
Edit Del	User Defined	GE7_Private	IPv4	GE7	Private Wired		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	
Edit Del	Auto Detect	169.254.9.3	IPv4	GE3	Public Wired	169.254.9.3	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	
Edit Del	Auto Detect	169.254.7.10	IPv6	GE4	Public Wired	fa00:1:1:2:2	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	

- 4 For an existing auto-detected or user-defined WAN Overlay, click **Edit** to modify the settings.
- 5 To create a new Public or Private overlay, click **Add User Defined WAN Overlay**.
- 6 In the **User Defined WAN Overlay** window, choose the **Link Type**.

7 Click **Advanced** to configure Hot Standby links.

The screenshot shows the 'Advanced Settings' window for configuring Hot Standby links. The 'Link Mode' is set to 'Hot Standby'. The 'Minimum Active Links' is set to '1'. The 'MTU' is set to '1500'. The 'Overhead Bytes' is set to '0'. The 'Path MTU Discovery' checkbox is checked. The 'Public Link Configuration' section is collapsed. The 'UDP Hole Punching' checkbox is unchecked, and the 'Type' is set to 'Wired'. At the bottom, there are buttons for 'Advanced', 'Update Link', and 'Cancel'.

Select **Hot Standby** from the **Link Mode** drop-down.

Note You cannot enable Hot standby link for a Hub.

Select the **Minimum Active Links** from the drop-down. This option indicates 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 being 1.

8 Configure other options as required and click **Update Link** to save the settings.

Note For more information on other options in the **WAN Overlay** window, see [Configure Edge WAN Overlay Settings](#).

Results

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.

What to do next


You can monitor the Hot Standby links in the monitoring dashboard. See [Monitor Hot Standby Links](#).

Monitor Hot Standby Links

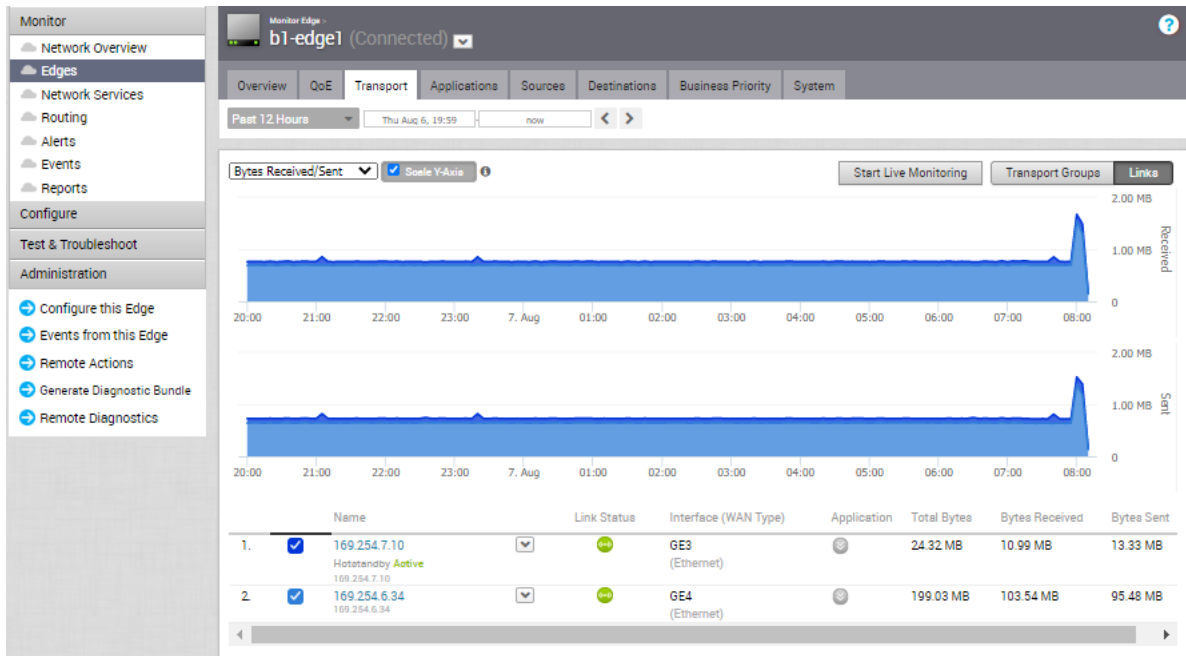
You can monitor the Hot standby links and the corresponding status using the monitoring dashboard.

To monitor the Hot standby links:

- 1 In the Enterprise portal, click **Monitor > Edges**.
- 2 Select the Edge configured with Hot standby link.
- 3 The **Overview** tab displays the links with status.

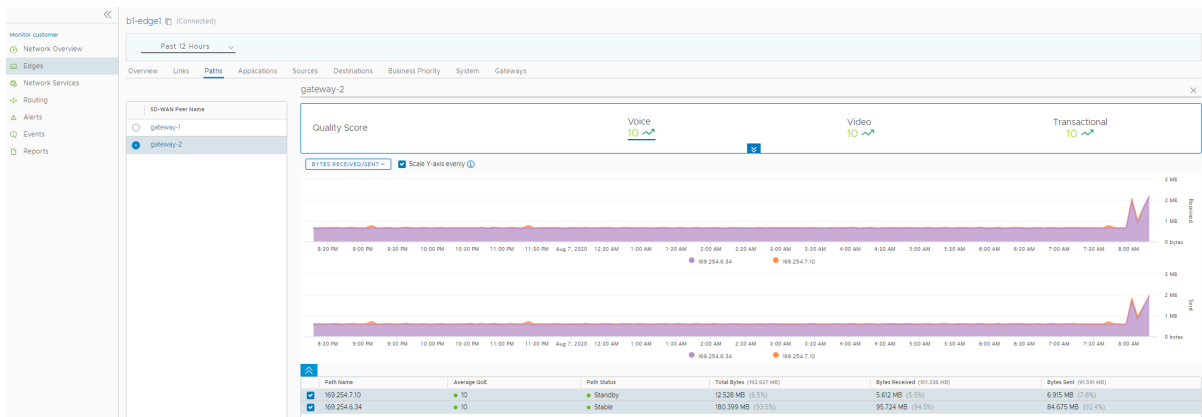


- 4 Click the **Transport** tab to view more information on the links, with graphical representation.



You can also view the status of Hot Standby links in the new Orchestrator UI.

- 1 In the Enterprise portal, click **Open New Orchestrator UI**.
- 2 Click **Launch New Orchestrator UI** in the pop-up window. The UI opens in a new tab displaying the monitoring options.
- 3 Click **Edges** to view the Edges associated with the Enterprise. Click the link to an Edge.
- 4 The **Overview** tab displays the links with status.
- 5 Click the **Links** tab to view more details with graphs.
- 6 Click the **Paths** tab and select an SD-WAN peer to view the status of the paths from the selected Edge.



Configure Wi-Fi Radio Overrides

At the Edge level, you can override the WI-FI Radio settings specified in the Profile by selecting the **Enable Edge Override** checkbox. 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.

To override the WI-FI Radio settings at the Edge level, perform the following steps.

Prerequisites

- 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. Ensure that the correct country of operation for this edge is set in the **Contact & Location** section of the **Edge Overview** configuration page. The address is populated automatically after the Edge is activated; however, you can override the address manually, if needed.

Note The country should be specified using the 2-character ISO 3166-1-alpha-2 notation (for example, US, DE, IN, and so on.)

Procedure

- From the SD-WAN Orchestrator, go to **Configure > Edges**.
- Select an Edge you want to override WI-FI Radio settings and click the icon under the **Device** column.

The **Device Setting** page for the selected Edge appears.

- In the **Configure Segment** drop-down menu, by default, **Global Segment [Regular]** is selected. If needed, you can select a different profile segment from the drop-down menu.
- Go to the **WI-FI Radio Settings** area and select the **Enable Edge Override** checkbox.

The screenshot shows the 'Wi-Fi Radio Settings' configuration page. On the left, there's a sidebar with various settings: 'Radio Enabled' (checked), 'Country' (IN), 'Band' (2.4 GHz, 5 GHz), 'Channel' (Automatic), 'Multi-Source QoS' (36, 40), 'SNMP Settings' (44, 48), 'NTP Servers' (153, 157), and 'Private NTP Servers' (161). The main area displays a table of radio settings with columns for Name, Value, and Action. The 'Enable Edge Override' checkbox is checked in the top right corner.

Name	Value	Action
Multi-Source QoS	36	Enable Edge Override
	40	
SNMP Settings	44	Enable Edge Override
Versions Enabled	48	
	149	
NTP Servers	153	Enable Edge Override
	157	
Private NTP Servers	161	Enabled

- 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.

- 7 If you want to change the location of the Edge, click **Go to Edge Overview to change edge location**. The **Edge Overview** page for the selected Edge appears.
 - a Under **Contact & Location** area, click the **Update Location** link to set the Edge location and click **Save Changes**.
- 8 Click **Save Changes**. The Wi-Fi Radio settings are overridden for the selected Edge.

Note If a country is not set for the edge or the country is invalid, then the radio **Band** is set to **2.4 GHz** and **Channel** is set to **Automatic**.

Security VNFs

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 VMware along with the support matrix:

Table 18-11. Palo Alto Networks Firewall – Support Matrix

VMware SD-WAN Edge Platform	Edge 520v	Edge 620	Edge 640	Edge 680	Edge 840	Edge 2000	Edge 3400	Edge 3800
Recommended VM Series Firewall Models	VM-50 Lite	VM-50 Lite	VM-100	VM-100	VM-100	*	VM-100	VM-100
Number of vCPUs Available for VM-Series Firewall	2	2	2	2	2	*	2	2
Memory Available for VNF	4.5 GB	4.5 GB	6.5 GB	6.5 GB	6.5 GB	*	9 GB	9 GB
Storage Space Available on Edge for VNF	64 GB	64 GB	120 GB	120 GB	120 GB	*	220.2 GB	220.2 GB
Earliest Supported VMware Release	Release 3.2.0 or later	Release 3.4.3 or later	Release 3.4.3 or later	Release 3.4.3 or later	Release 3.2.0 or later	Not supported on any release	Release 4.3.0 or later	Release 4.3.0 or later
Panorama Version	Release 8.0.5	Release 8.0.5	Release 8.0.5	Release 8.0.5	Release 8.0.5	*	Release 8.0.5	Release 8.0.5

Table 18-12. Check Point Firewall – Support Matrix

VMware SD-WAN Edge Platform	Edge 520v	Edge 620	Edge 640	Edge 680	Edge 840	Edge 2000	Edge 3400	Edge 3800
Memory Available for VNF	2 GB	2 GB	4 GB	4 GB	4 GB	*	4 GB	4 GB
Number of vCPUs Available for VNF	2	2	2	2	2	*	2	2
Storage Available on Edge for VNF	64 GB	120 GB	120 GB	120 GB	100 GB	*	220.2 GB	220.2 GB

Table 18-12. Check Point Firewall – Support Matrix (continued)

VMware SD-WAN Edge Platform	Edge 520v	Edge 620	Edge 640	Edge 680	Edge 840	Edge 2000	Edge 3400	Edge 3800
Maximum Throughput of SD-WAN and Checkpoint VNF	100 Mbps	100 Mbps	350 Mbps	500 Mbps	550 Mbps	*	4.9 Gbps	4.9 Gbps
Earliest Supported VMware Release	Release 3.3.2 or later	Release 3.4.3 or later	Release 3.4.3 or later	Release 3.4.3 or later	Release 3.3.2 or later	Not supported on any release	Release 4.3.0 or later	Release 4.3.0 or later
Checkpoint VNF OS Version	Release R80.20.15 R80.20.35 R80.20.60	Release R80.20.15 R80.20.35 R80.20.60	Release R80.20.15 R80.20.35 R80.20.60	Release R80.20.15 R80.20.35 R80.20.60	Release R80.20.15 R80.20.35 R80.20.60	*	Release R80.20.15 R80.20.35 R80.20.60	Release R80.20.15 R80.20.35 R80.20.60
Checkpoint Manager Software Version	Release 80.20 80.20.60	Release 80.20 80.20.60	Release 80.20 80.20.60	Release 80.20 80.20.60	Release 80.20 80.20.60	*	Release 80.20 80.20.60	Release 80.20 80.20.60

Table 18-13. Fortinet Firewall – Support Matrix

VMware SD-WAN Edge Platform	Edge 520v	Edge 620	Edge 640	Edge 680	Edge 840	Edge 2000	Edge 3400	Edge 3800
Recommended VM Series Firewall Models	VM00, VM01, VM01v	VM00, VM01, VM01v	VM00, VM01, VM01v, VM02, VM02v	VM00, VM01, VM01v, VM02, VM02v	VM00, VM01, VM01v, VM02, VM02v	*	VM01, VM02	VM01, VM02
Memory Available for VNF	2 GB	2 GB	4 GB	4 GB	4 GB	*	4 GB	4 GB
Number of vCPUs Available for VNF	2	2	2	2	2	*	2	2
Storage Available on Edge for VNF	64 GB	64 GB	100 GB	100 GB	100 GB	*	220.2 GB	220.2 GB

Table 18-13. Fortinet Firewall – Support Matrix (continued)

VMware SD-WAN Edge Platform	Edge 520v	Edge 620	Edge 640	Edge 680	Edge 840	Edge 2000	Edge 3400	Edge 3800
Maximum Throughput of SD-WAN and FortiGate VNF	100 Mbps	100 Mbps	500 Mbps	500 Mbps	500 Mbps	*	1.5 Gbps	1.5 Gbps
Earliest Supported VMware Release	Release 3.3.1 or later	Release 4.0.0 or later	Release 4.0.0 or later	Release 4.0.0 or later	Release 3.3.1 or later	Not supported on any release	Release 4.3.0 or later	Release 4.3.0 or later
FortiOS Version	Release 6.2.4	Release 6.2.4	Release 6.2.4	Release 6.2.4	Release 6.2.4	*	Release 6.2.4	Release 6.2.4

You can deploy and forward traffic through VNF on an SD-WAN Edge.

Configure VNF Management Service

VMware 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.

Procedure

- 1 In the Enterprise portal, click **Configure > Network Services**.
- 2 In the **Services** page, scroll down to the **VNFs** section and click **New**.

- 3 In the **VNF Service Management Configuration** window, enter a descriptive name for the security VNF service and select a VNF Type from the drop-down list.



The image shows a software window titled "VNF Service Management Configuration". It contains two required fields, each marked with a blue asterisk: "Name" and "VNF Type". The "Name" field is a text input box. The "VNF Type" field is a dropdown menu currently showing "Select". Below these fields is a message: "Select a VNF type to continue." In the bottom right corner of the window is a "Cancel" button. The window has a standard title bar with a question mark icon and a close button.

4 Configure the settings based on the selected VNF Type.

- a For the VNF Type **Palo Alto Networks Firewall**, configure the following:

The screenshot shows the 'VNF Service Management Configuration' window. It has a title bar with a question mark icon and a close button. The window contains the following fields:

- Name:** Palo Alto Networks Management Server West Coast
- VNF Type:** Palo Alto Networks Firewall
- Primary Panorama IP Address:** 172.16.3.52
- Secondary Panorama IP Address:** (empty field)
- Panorama Auth Key:** (masked with dots and an eye icon to toggle visibility)

At the bottom right, there are two buttons: 'Save Changes' (green) and 'Cancel' (gray).

- 1 **Primary Panorama IP Address** – Enter the primary IP address of the Panorama server.
- 2 **Secondary Panorama IP Address** – Enter the secondary IP address of the Panorama server.
- 3 **Panorama Auth Key** – Enter the authentication key configured on the Panorama server. VNF uses the Auth Key to login and communicate with Panorama.
- 4 Click **Save Changes**.

After configuring Palo Alto Networks as VNF Type, define the VNF licenses. These licenses will be applied to one or more VNF configured Edges.

- 1 In the **Services** page, scroll down to the **VNF Licenses** section and click **New**.
- 2 In the **VNF License Configuration** window, configure the following:

The screenshot shows the 'VNF License Configuration' window. It has a title bar with a question mark icon and a close button. The window contains the following fields:

- Name:** VM-50 License
- VNF Type:** Palo Alto Networks Firewall
- License Server API Key:** (masked with dots and an eye icon to toggle visibility)
- Auth Code:** V5073094

Below the 'Auth Code' field, there is a 'Test' button and a green checkmark followed by the word 'Valid'. At the bottom right, there are two buttons: 'Save Changes' (green) and 'Cancel' (gray).

- **Name** – Enter a descriptive 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 SD-WAN Orchestrator uses this key to communicate with the Palo Alto Networks license server.
- **Auth Code** – Enter the authorization code purchased from Palo Alto Networks.

- Click **Test** to validate the configuration.

3 Click **Save Changes**.

You can apply the VNF licenses while configuring **Palo Alto Networks Firewall** as a VNF Type on Edges.

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.

- b For the VNF Type **Check Point Firewall**, configure the following:

The screenshot shows the 'VNF Service Management Configuration' dialog box. The title bar includes a help icon and a close button. The dialog contains the following fields and controls:

- Name**: CheckPoint_VNF
- VNF Type**: Check Point Firewall
- Primary Check Point Mgmt Server IP**: 172.24.2.26
- SIC Key for Mgmt Server Access**: Password field with 8 dots and an eye icon.
- Admin Password**: Password field with 8 dots and an eye icon.
- VNF Image Location**: Check_Point_R80_SMB_fw1
- Image Version**: R77.20.87(8f8f,sha-1) (dropdown menu)
- File Checksum Type**: sha-1
- File Checksum**: 8f8f42784818f473c36b26d2
- Download Type**: Radio buttons for https and s3 (s3 is selected).
- AccessKeyId**: Empty text field.
- SecretAccessKey**: Password field with an eye icon.
- Region**: eu-central-1 (dropdown menu)

At the bottom right, there are two buttons: **Save Changes** (green) and **Cancel** (gray).

- 1 **Primary Check Point Mgmt Server IP** – Enter the Check Point Smart Console IP address that will connect to the Check Point Firewall.
- 2 **SIC Key for Mgmt Server Access** – Enter the password used to register the VNF to the Check Point Smart Console.
- 3 **Admin Password** – Enter the administrator password.
- 4 **VNF Image Location** – Enter the image location from where the SD-WAN Orchestrator will download the VNF image.
- 5 **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.extrImageInfos`.

- 6 **File Checksum Type** – Specifies the method used to validate the VNF image and is automatically populated after you select an image version.
 - 7 **File Checksum** – Specifies 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.extralImageInfos`.
 - 8 **Download Type** – Choose the type of the image. For **https**, enter the username and password. For **s3**, enter the AccessKeyId, SecretAccessKey, and choose the Region.
 - 9 Click **Save Changes**.
- c For the VNF Type **Fortinet Firewall**, configure the following:

VNF Service Management Configuration

- Name: Fortinet
- VNF Type: Fortinet Firewall
- Fortinet Mgmt Server IP: 192.168.33.100
- Fortimanager Serial Number: FMG-VMTM-100554942
- Registration Password: [masked]
- VNF Image Location: zsu-p3s/forti.5.3.0
- Image Version: 6.4.0(6caa,sha-1)
- File Checksum Type: sha-1
- File Checksum: 6caad8af9d60dcde20f9
- Download Type: ☐ https ☒ s3
- AccessKeyId: 5a064p228b43a3ea10
- SecretAccessKey: [masked]
- Region: us-east-2


Buttons: Save Changes, Cancel

- 1 **Fortinet Mgmt Server IP** – Enter the IP address of the FortiManager to connect to the FortiGate.
- 2 **Fortimanager Serial Number** – Enter the serial number of FortiManager.
- 3 **Registration Password** – Enter the password used to register the VNF to the FortiManager.
- 4 **VNF Image Location** – Enter the image location from where the SD-WAN Orchestrator will download the VNF image.
- 5 **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.extralImageInfos`.
- 6 **File Checksum Type** – Specifies the method used to validate the VNF image and is automatically populated after you choose an image version.
- 7 **File Checksum** – Specifies 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.extralImageInfos`.

- 8 **Download Type** – Choose the type of the image. For **https**, enter the username and password. For **s3**, AccessKeyId, SecretAccessKey, and choose the Region.
- 9 Click **Save Changes**.

Results

The **VNFs** section displays the created VNF services. The following image shows an example of VNF Type as Check Point Firewall.



Name	Type	Used By
✓ Check Point Firewall	Check Point Security Firewall	

What to do next

You can configure security VNF for an Edge to direct the traffic through the VNF management services. See

- [Configure Security VNF without HA](#)
- [Configure Security VNF with High Availability](#)

Configure Security VNF without HA

You can deploy and forward traffic through VNF on the SD-WAN Edge, using third-party firewalls.

Only an Operator can enable the Security VNF configuration. If the Security VNF option is not available for you, contact your Operator.

Prerequisites

Ensure that you have the following:

- SD-WAN Orchestrator and activated SD-WAN Edge running software versions that support deploying a specific security VNF. For more information on the supported software versions and Edge platforms, refer to the Support Matrix in [Security VNFs](#).
- Configured VNF Management service. For more information, see [Configure VNF Management Service](#).

Procedure

- 1 In the Enterprise portal, click **Configure > Edges**.
- 2 In the **Edges** page, either click the **Device** Icon next to an Edge or click the link to an Edge and click the **Device** tab.
- 3 In the **Device** tab, scroll down to the **Security VNF** section and click **Edit**.



- 4 In the **Edge VNF Configuration** window, check the **Deploy** checkbox.
- 5 Configure the following in **VM Configuration**:
 - 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.

- e **Security VNF** – Choose a pre-defined VNF management service from the drop-down list. You can also click **New VNF Service** to create a new VNF management service. For more information, see [Configure VNF Management Service](#).

The following image shows an example of **Check Point Firewall** as the Security VNF type.

The image shows the 'Edge VNF Configuration' dialog box. The 'Deploy' checkbox is checked. Under 'VM Configuration', the 'VLAN' is set to '100 - VLAN-100', 'VM-1 IP' is '10.100.1.2', and 'VM-1 Hostname' is 'VM-1'. The 'Deployment State' has two radio buttons: 'Image Downloaded and Powered On' (selected) and 'Image Downloaded and Powered Off'. The 'Security VNF' dropdown is set to 'CPM'. At the bottom right are 'Update' and 'Cancel' buttons.

If you choose **Palo Alto Networks Firewall** as Security VNF, configure the following additional settings:

The image shows the 'Edge VNF Configuration' dialog box with 'Palo Alto Networks Firewall' selected as the Security VNF. The 'Deploy' checkbox is checked. Under 'VM Configuration', the 'VLAN' is set to '1 - Corporate', 'VM-1 IP' is '10.0.1.2', and 'VM-1 Hostname' is 'VM-1'. The 'Deployment State' has two radio buttons: 'Powered On' and 'Powered Off' (selected). The 'Security VNF' dropdown is set to 'Palo Alto Networks Management Server West Coast'. Below this, there is a section for additional settings: 'License' is set to 'VM-50 License', 'Device Group Name' is 'Demo_Group', and 'Config Template Name' is 'Demo_template'. At the bottom right are 'Update' and 'Cancel' buttons.

- **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.

If you choose **Fortinet Firewall**, configure the following additional settings:

The image shows the 'Edge VNF Configuration' dialog box. It has a 'Deploy' checkbox which is checked. Under 'VM Configuration', there are three fields: 'VLAN' set to '1 - Corporate', 'VM-1 IP' set to '10.0.5.50', and 'VM-1 Hostname' set to 'VM640'. Below these is the 'Deployment State' section with two radio buttons: 'Image Downloaded and Powered On' (selected) and 'Image Downloaded and Powered Off'. Under 'Security VNF', there is a dropdown menu set to 'Fortinet624'. Below this is a section for 'VM Cores' (set to '2') and 'Inspection Mode' (with 'proxy' and 'flow' radio buttons, where 'flow' is selected). At the bottom of this section is a text area labeled 'Drop your license file or paste your license file's content' containing a license key: '-----BEGIN FGT VM LICENSE----- QAAAAMJkk+tOIICJbnb7TnoHAQMOXq1AM5CssQxd7hh/d86w/j7FEe1jUXLTw9H2'. At the very bottom of the dialog are 'Update' and 'Cancel' buttons.

- **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.

f Click **Update**.

Results

The configuration details are displayed in the **Security VNF** section.



What to do next

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](#).

Configure Security VNF with High Availability

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.

Prerequisites

Ensure that you have the following:

- SD-WAN Orchestrator and activated SD-WAN Edge running software version 4.0.0 or later. For more information on the supported Edge platforms, refer to the Support Matrix in [Security VNFs](#).

- Configured Check Point Firewall VNF Management service. For more information, see [Configure VNF Management Service](#).

Note VMware supports only Check Point Firewall VNF on Edges with HA.

Procedure


- 1 In the Enterprise portal, click **Configure > Edges**.
- 2 Either click the **Device** Icon next to an Edge, or click the link to an Edge and then click the **Device** tab.
- 3 In the **Device** tab, navigate to the **High Availability** section and choose the **Active Standby Pair**.
- 4 Navigate to the **Security VNF** section and click **Edit**.

The screenshot displays the configuration page for an edge device. It is divided into several sections:

- High Availability:** Contains radio buttons for 'Type' (None, Active Standby Pair, Cluster, VRRP with Third-Party Router) and a text field for 'HA Interface' (GE1).
- Configure VLAN:** Includes a 'Management IP' field with the value 10.0.3.2 and a checkbox for 'Enable Edge Override'.
- Device Settings: Virtual Edge:** A section header with a dropdown arrow.
- Security VNF:** Features an 'Edit' button and a message stating 'No security VNFs have been configured for this Edge.'

- 5 In the **Edge VNF Configuration** page, click **Deploy**.

6 Configure the following in **VM Configuration**:



The image shows a dialog box titled "Edge VNF Configuration". It has a "Deploy:" checkbox which is checked. Below it is the "VM Configuration" section. It contains several fields: "VLAN:" with a dropdown menu showing "100 - VLAN-100", "VM-1 IP:" with a text box containing "10.100.2.31", "VM-2 IP:" with a text box containing "10.100.2.32", "VM-1 Hostname:" with a text box containing "VNF-1", and "VM-2 Hostname:" with a text box containing "VNF-2". There is also a "Deployment State:" section with two radio buttons: "Image Downloaded and Powered On" (which is selected) and "Image Downloaded and Powered Off". At the bottom of the configuration section is a "Security VNF:" dropdown menu showing "CPM". At the bottom right of the dialog are "Update" and "Cancel" buttons.

- a **VLAN** – Choose a VLAN, to be used for the VNF management, from the drop-down list.
- b **VM-1 IP, VM-2 IP** – Enter the IP addresses of the VM1 and VM2. Ensure that the IP addresses are in the subnet range of the chosen VLAN.
- c **VM-1 Hostname, VM-2 Hostname** – Enter the names for the VM hosts.
- 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.
- e **Security VNF** – Choose a pre-defined Check Point Firewall VNF Management service from the drop-down list. You can also click **New VNF Service** to create a new VNF management service. For more information, see [Configure VNF Management Service](#).
- f Click **Update**.

Results

The **Security VNF** section displays the configured details:

Security VNF
Edit

VM Configuration

Deployment State: Powered On
VLAN: 100 - VLAN-100
VM-1 IP: 10.100.2.31
VM-1 Hostname: VNF-1
VM-2 IP: 10.100.2.32
VM-2 Hostname: VNF-2

Security VNF:

CPM
Check Point Firewall

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 more information on HA configuration, see [Configure High Availability \(HA\)](#).

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.

What to do next

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](#).

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:

Procedure

- 1 In the Enterprise portal, click **Configure > Segments**.
- 2 In the **Segments** page, enter the Service VLAN ID for each segment.

Monitor
Configure
Edges
Profiles
Object Groups
Segments
Overlay Flow Control
Network Services
Alerts & Notifications
Customer
Test & Troubleshoot
Administration

Segments Save Changes

Segment Name	Description	Type	Service VLAN	Delegate To Partner	Delegate To Customer	
Global Segment	Default segment for traffic that	Regular	1	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	ⓘ +
segment1		Regular	100	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	ⓘ +
segment2		Regular	101	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	ⓘ +

- 3 Click **Save Changes**.

Results

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.

What to do next

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](#).

Configure VLAN with VNF Insertion

You can insert the security VNF into both the VLAN as well as routed interface.

Prerequisites

Ensure that you have created a security VNF and configured the settings. See [Configure Security VNF without HA](#) 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](#).

Procedure

- 1 In the Enterprise portal, click **Configure > Edges**.
- 2 In the **Edges** page, either click the **Device** icon next to an Edge or click the link to an Edge and click the **Device** tab.
- 3 In the **Device** tab, scroll down to the **Configure VLAN** section.
- 4 Click the **Edit** link of the VLAN to which you want to insert the VNF.

- 5 In the **VLAN** window, select the **VNF Insertion** checkbox to insert the VNF into VLAN. This option redirects traffic from a specific VLAN to the VNF.

VLAN

Segment: segment1 ☒ Enable Edge Override

VLAN Name: VLAN-100

VLAN Id: 100

Assign Overlapping Subnets: ☒

Edge LAN IP Address: 10.100.1.1

Cidr Prefix: 24

Network: 10.100.1.0

Advertise: ☒

ICMP Echo Response: ☒

VNF Insertion: ☒

Multicast: Multicast is not enabled for the selected segment

Fixed IPs:

MAC Address	IP	Description
00:ba:be:73:02:fa	10.100.1.100	Description (optional)

LAN Interfaces: GE2

SSID: There are no Wi-Fi SSIDs configured on this VLAN.

DHCP ☐ Enable Edge Override

Type: Enabled

DHCP Start: 10.100.1.13

Num. Addresses: 242

Lease Time: 1 day

DHCP Options: not set

OSPF ☐ Enable Edge Override

Enabled: ☒ OSPF not enabled for the selected Segment.

Update VLAN Cancel

- 6 Click **Update VLAN**.

Results

The **Configure VLAN** section displays the status of the VNF insertion.

Action	VLAN	DHCP	Network	IP Address	Interfaces	DHCP	Segment	Multicast	VNF Insertion
Edit	1 - Corporate	<input checked="" type="checkbox"/>	10.0.1.0/24	10.0.1.1	GE1, GE2	Enabled (242)	Global Segment		<input checked="" type="checkbox"/>
Edit	100 - VLAN-100	<input checked="" type="checkbox"/>	10.100.1.0/24	10.100.1.1	GE2	Enabled (242)	segment1		<input checked="" type="checkbox"/>
Edit	101 - VLAN-101	<input checked="" type="checkbox"/>	10.101.1.0/24	10.101.1.1	GE2	Enabled (242)	segment2		<input checked="" type="checkbox"/>

You can also insert the VNF into Layer 3 interfaces or sub-interfaces. This insertion redirects traffic from the Layer 3 interfaces or subinterfaces to the VNF.

If you choose to use the routed interface, ensure that the trusted source is checked and WAN overlay is turned off on that interface. For more information, see [Configure Interface Settings](#).

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:

- In the Enterprise portal, click **Monitor > Edges**. The list of Edges along with the details of configured VNFs is displayed.

Edge	Status	HA	Links	VM Status	VNF	Cloud Services S...	Gateways	Profile
1 B1-Edge-2000	●	●	↔ 2					View Quick Start Profile
2 B2-Edge-520V	●	●	↔ 2	View	●			View Quick Start Profile
3 B3-Edge-840	●	●	↔ 2	View	●			View Quick Start Profile
4 b4-edge1	●		↔ 2					View Quick Start Profile

- With mouse pointer, hover-over the Icon in the **VNF** column to view additional details of the VNF type.
- Click the **View** link in the **VM Status** column to open the **VNF Virtual Machine Status** window, where you can view the deployment status for the Edge. To view the deployment details, click the **View** link next to **Deployment Details**.

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 image:

Time	VM Status	CPU %	Memory Used (MB)	Storage Used (GB)	Serial Number
Fri Jun 26, 22:54:48 5 days ago	Powered On	5.8	4096	6	BKXFXC2
Fri Jun 26, 22:54:48 5 days ago	Powered On	5.8	4096	6	BKXFXC2
Fri Jun 26, 22:54:48 5 days ago	Powered On	5.8	4096	6	BKXFXC2
Fri Jun 26, 22:54:48 5 days ago	Powered On	5.8	4096	6	BKXFXC2
Fri Jun 26, 22:54:48 5 days ago	Powered On	5.8	4096	6	BKXFXC2

To monitor the status of VNFs and VMs:

- In the Enterprise portal, click **Monitor > Network Services**. The list of Edges along with the details of configured VNFs is displayed.

Edge VNFs			
	Service	Used By	Edge VM Status
1	CPM Check Point Firewall	1 Edge View	Unknown 1 Edge
2	Fortinet Fortinet Firewall	1 Edge View	Unknown 1 Edge
3	PAN Palo Alto Networks Firewall	0	

You can also view the status of VNFs in the new Orchestrator UI.

- 1 In the Enterprise portal, click **Open New Orchestrator UI**.
- 2 Click **Launch New Orchestrator UI** in the pop-up window. The UI opens in a new tab displaying the monitoring options.
- 3 Click **Edges** to view the status of Edges along with the VNFs and VMs.

Edges										
<div> <div>Monitor customer</div> <div> <div>Network Overview</div> <div>Edges</div> <div>Network Services</div> <div>Routing</div> <div>Alerts</div> <div>Events</div> <div>Reports</div> </div> </div>										
<div> <div>Search</div> <div>Map Distribution</div> </div>										
Name	Status	HA	Links	Stable	Degraded	Down	VM Status	VNF	Gateways	Last Contact
B1-Edge-2000	Connected		2	2			Status Deployment Details	CheckPoint	View	Thu Apr 30, 10:36:48
B2-Edge-520V	Connected		2	2			Status Deployment Details	CheckPoint	View	Thu Apr 30, 10:36:53
B3-Edge-840	Connected		2	2			Status Deployment Details	CheckPoint	View	Thu Apr 30, 10:36:49
b4-edge1	Connected		2	2					View	Thu Apr 30, 10:37:00

- 4 Click **Network Services > Edge VNFs** to view the status of VNFs and VMs.

Network Overview

Edges

Network Services

Routing

Alerts

Events

Reports

Application Analytics?

Branch Analytics ?

Non SD-WAN Destinations via Gateway

Cloud Security Service Sites

Edge Clusters

Edge VNFs

Service	Used By	Edge VM Status
CPM Check Point Security Firewall	1 Edge	Powered On (Insertion Enabled) 1 Edge

COLUMNS

1 items

VNF Edge Deployments

Edge Name	Edge VM Status
h3-edge1-E840	Powered On (Insertion Enabled)

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 enterprise portal, click **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

You can also view the events in the new Orchestrator UI.

Click **Launch New Orchestrator UI** in the pop-up window. The UI opens in a new tab displaying the monitoring options. Click **Events**. Click the Filter Icon in the **Search** option to filter the VNF events.

The following image shows some of the VNF events.

Event	User	Segment	Edge	Severity	Time	Message
VNF VM config changed			b6-edge1-E840	Info	Jul 19, 2021, 3:28:52 PM	VNF power on
VNF VM config changed			b6-edge1-E840	Info	Jul 19, 2021, 3:31:42 PM	VNF power on
VNF VM config changed			b6-edge1-E840	Info	Jul 20, 2021, 1:39:27 PM	VNF power on
VNF VM config changed			b6-edge1-E840	Info	Jul 20, 2021, 1:42:05 PM	VNF power on
VNF VM config changed			b6-edge1-E840	Info	Jul 25, 2021, 2:50:22 AM	VNF power on
VNF VM config changed			b6-edge1-E840	Info	Jul 25, 2021, 2:55:55 AM	VNF power on
VNF VM config changed			b6-edge1-E840	Info	Jul 25, 2021, 2:57:13 AM	VNF power on
VNF VM config changed			b6-edge1-E840	Info	Jul 25, 2021, 3:00:31 AM	VNF power on
VNF VM config changed			b6-edge1-E840	Info	Jul 27, 2021, 12:55:33 PM	VNF power on
VNF VM config changed			b6-edge1-E840	Info	Jul 27, 2021, 12:58:15 PM	VNF power on

Configure VNF Alerts

You can configure to receive alerts and notifications related to the VNF events.

In the Enterprise portal, click **Configure > Alerts & Notifications**. In the **Alert Configuration** page, you can select the Alert Types.

The screenshot shows the 'Alert Configuration' page with a 'Save Changes' button and a help icon. Below is a table with columns for 'Select Alerts', 'Alert Type', and 'Notification Delay'.

Select Alerts	Alert Type	Notification Delay
<input type="checkbox"/>	Edge Down ⓘ	3 minutes
<input type="checkbox"/>	Edge Up ⓘ	1 minutes
<input type="checkbox"/>	Link Down ⓘ	3 minutes
<input type="checkbox"/>	Link Up ⓘ	1 minutes
<input type="checkbox"/>	VPN Tunnel Down ⓘ	3 minutes
<input type="checkbox"/>	Edge HA Failover ⓘ	1 minutes
<input type="checkbox"/>	Edge VNF Virtual Machine Deployment ⓘ	0 minutes
<input type="checkbox"/>	Edge VNF Insertion ⓘ	0 minutes
<input type="checkbox"/>	Edge CSS tunnel up ⓘ	3 minutes
<input type="checkbox"/>	Edge CSS tunnel down ⓘ	3 minutes
<input type="checkbox"/>	Edge VNF Image Download Event ⓘ	0 minutes

To receive alerts for VNF events, select the following Alert Types:

- **Edge VNF Virtual Machine Deployment** – 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.

You can view the alert notifications in the **Monitor > Alerts** page.

To view the alerts related to VNF, you can use the filter option. Click the drop-down arrow next to the **Search** option and choose to filter by the Type.

The following image shows some of the VNF alerts.

Monitor

Network Overview

Edges

Network Services

Routing

Alerts

Events

Firewall Logs

Reports

Configure

Test & Troubleshoot

Administration

Alerts

Past 7 Days

Thu Jun 25, 18:44

now

<

<

<

>

>

>

Search

Cols

Reset View

Refresh

Refresh

Refresh

Refresh

CSV

☐ Include Operator Alerts

Display 81 items

Trigger Time	Notification Time	Category	Type	Description	Status	
Sat Jun 27, 02:55:42	Thu Jul 02, 18:47:12	Customer	VNF_INSERTION_ENABLED	Edge b4-6X0-1	Closed	
Sat Jun 27, 02:55:42	Thu Jul 02, 18:47:12	Customer	VNF_VM_DEPLOYED_AND_POWERED_OFF	6c261793-5e91-429b-83f3-ddb731064e44 Link up ...	Closed	
Sat Jun 27, 02:55:42	Thu Jul 02, 18:47:12	Customer	VNF_VM_POWERED_ON	6c261793-5e91-429b-83f3-ddb731064e44 Link up ...	Closed	
Sat Jun 27, 02:55:32	Thu Jul 02, 18:47:12	Customer	VNF_INSERTION_ENABLED	Edge b4-6X0-1	Closed	
Sat Jun 27, 02:55:32	Thu Jul 02, 18:47:12	Customer	VNF_VM_DEPLOYED_AND_POWERED_OFF	1e662489-066f-445d-8be8-00b682f29a29 Link up ...	Closed	
Sat Jun 27, 02:55:32	Thu Jul 02, 18:47:12	Customer	VNF_VM_POWERED_ON	1e662489-066f-445d-8be8-00b682f29a29 Link up ...	Closed	
Sat Jun 27, 02:47:13	Thu Jul 02, 18:47:12	Customer	VNF_INSERTION_DISABLED	Edge b4-6X0-1	Closed	
Sat Jun 27, 02:47:13	Thu Jul 02, 18:47:12	Customer	VNF_VM_POWERED_OFF	fecedb94-a962-4abc-9478-92f5cd019c10 Link up ...	Closed	
Sat Jun 27, 02:47:13	Thu Jul 02, 18:47:12	Customer	VNF_VM_DELETED	fecedb94-a962-4abc-9478-92f5cd019c10 Link up ...	Closed	
Sat Jun 27, 02:47:02	Thu Jul 02, 18:47:12	Customer	VNF_INSERTION_DISABLED	Edge b4-6X0-1	Closed	
Sat Jun 27, 02:47:02	Thu Jul 02, 18:47:12	Customer	VNF_VM_POWERED_OFF	35cf5583-969f-4c81-be3e-bfc7b71ea516 Link up ...	Closed	
Sat Jun 27, 02:47:02	Thu Jul 02, 18:47:12	Customer	VNF_VM_DELETED	35cf5583-969f-4c81-be3e-bfc7b71ea516 Link up ...	Closed	
Sat Jun 27, 02:14:44	Thu Jul 02, 18:47:12	Customer	VNF_INSERTION_ENABLED	Edge b4-6X0-1	Closed	
Sat Jun 27, 02:14:44	Thu Jul 02, 18:47:12	Customer	VNF_VM_DEPLOYED_AND_POWERED_OFF	35cf5583-969f-4c81-be3e-bfc7b71ea516 Link up ...	Closed	
Sat Jun 27, 02:14:44	Thu Jul 02, 18:47:12	Customer	VNF_VM_POWERED_ON	35cf5583-969f-4c81-be3e-bfc7b71ea516 Link up ...	Closed	
Sat Jun 27, 02:14:35	Thu Jul 02, 18:47:12	Customer	VNF_VM_DELETED	35cf5583-969f-4c81-be3e-bfc7b71ea516 Link up ...	Closed	
Sat Jun 27, 02:14:15	Thu Jul 02, 18:47:12	Customer	VNF_VM_DELETED	35cf5583-969f-4c81-be3e-bfc7b71ea516 Link up ...	Closed	

You can also view the alerts in the new Orchestrator UI.

Click **Launch New Orchestrator UI** in the pop-up window. The UI opens in a new tab displaying the monitoring options. Click **Alerts**. Click the Filter Icon in the **Search** option to filter the VNF alerts.

Configure Layer 2 Settings for Edges


At the Edge level, you can override the Layer 2 settings inherited from a Profile by selecting the **Enable Edge Override** checkbox.

To override the ARP timeouts values at the Edge-level, perform the following steps:

Procedure

- 1 From the SD-WAN Orchestrator, go to **Configure > Edges**.
- 2 Select an Edge you want to override L2 settings and click the icon under the **Device** column. The Device Setting page for the selected Edge appears.
- 3 Go to the **L2 Settings** area and select the **Enable Edge Override** checkbox.

L2 Settings

☒ **Enable Edge Override** 

Override default ARP Timeouts ☒

ARP Stale Timeout: Hours Minutes

ARP Dead Timeout: Hours Minutes

ARP Cleanup Timeout: Hours Minutes

- 4 Select the **Override default ARP Timeouts** checkbox and then override the various ARP timeouts inherited from the Profile as follows:

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 Layer 2 Settings for Profiles](#).

Note To set the default ARP timeout values at the Edge level, unselect the **Override default ARP Timeouts** checkbox.

- 5 Click **Save Changes**.

What to do next

You can override the default ARP timeouts at the Profile-level. For more information, see [Configure Layer 2 Settings for Profiles](#).

Configure SNMP Settings for Edges

SNMP is a commonly used protocol for network monitoring and MIB is a database associated with SNMP to manage entities. SNMP can be enabled by selecting the desired SNMP version as described in the steps below. At the Edge Level, you can override the SNMP settings specified in the Profile by selecting the **Enable Edge Override** check box.

Note SD-WAN 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 SD-WAN Orchestrator, which in turn generates traps based on the alerts configured for the received events.

Before you begin:

- To download the SD-WAN Edge MIB: go to the **Remote Diagnostic** screen (**Test & Troubleshooting > Remote Diagnostics**) and run MIB for SD-WAN Edge. Copy and paste 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 the above-mentioned MIBs are available on the Remote Diagnostics page.

About this task: At the Edge level, you can override the SNMP settings specified in the Profile by selecting the **Enable Edge 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 rather than Edge exception level.

Supported MIBs

- SNMP MIB-2 System
- SNMP MIB-2 Interfaces
- VELOCLOUD-EDGE-MIB

Procedure to Configure SNMP Settings at Edge Level:

- 1 Obtain the VELOCLOUD-EDGE-MIB on the Remote Diagnostic screen of the SD-WAN Orchestrator.
- 2 Install all MIBs required by VELOCLOUD-EDGE-MIB.
- 3 From the SD-WAN Orchestrator, go to **Configure > Edges**.
- 4 Select an Edge you want to configure SNMP settings for, and click the **Device** icon under the Device column.

The **Configuration Edges** screen for the selected Edge appears.

- 5 Scroll down to the **SNMP Settings** area and check the **Enable Edge Override** check box. You can choose between two versions, v2c or v3.



- 6 For a SNMP v2c config follow the steps below:
 - a Check the **v2c** check box.
 - b Type in a Port in the **Port** textbox. The default setting is 161.
 - c In the **Community** textbox, type in a word or sequence of numbers that will act as a 'password' that will allow you access to the SNMP agent.

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.

- d For Allowed IPs:
 - Check the **Any** check box to allow any IP to access the SNMP agent.
 - To restrict access to the SNMP agent, uncheck the **Any** check box and enter the IP address(es) that will be allowed access to the SNMP agent.

SNMP Settings

SNMP Version: v2c

Port: 161

Community:

Allowed IPs: Allowed IP

- 7 For a SNMP v3 config, which provides added security support follow the steps below:
 - a Type in a port in the **Port** text box. 161 is the default setting.
 - b Type in a user name and password in the appropriate text boxes. Ensure that the Privacy Password is same as the Authentication Password configured on the Edge.
 - c Check the **Privacy** check box if you want your packet transfer encrypted.
 - d If you have checked the **Privacy** check box, choose **DES** or **AES** from the **Algorithm** drop-down menu.

Note Algorithm **AES** indicates **AES-128**.

SNMP Settings

SNMP Version: v3

Port: 161

Name: admin

Password:

Privacy: ☒

Algorithm: DES

- 8 Configure Firewall Settings. After you have configured SNMP Settings, go to Firewall settings (**Configure > Profiles > Firewall**) to configure the Firewall settings that will enable your SNMP settings.

Note

- SNMP interface monitoring is supported on DPDK enabled interfaces for 3.3.0 and later releases.
 - For software versions below 5.x, the supported authentication method is **MD5** by default.
-

Configure NTP Settings for Edges

As an enterprise Administrator, at the Edge level, you can override the Network Time Protocol (NTP) settings specified in the Profile by selecting the **Enable Edge Override** checkbox. By default, at the Edge level, the NTP Servers are deactivated.

To override NTP settings at the Edge-level, perform the following steps.

Prerequisites

NTP has the following prerequisites:

- To configure an SD-WAN 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 SD-WAN Edge NTP Server configuration has the following limitations:

- NTP Clients can synchronize to LAN/loopback IP address of the SD-WAN Edge as NTP server but cannot synchronize to WAN IP address.
- NTP synchronization from another segment to LAN interface is not supported.

Procedure

- 1 From the SD-WAN Orchestrator, go to **Configure > Edges**.
- 2 Select an Edge you want to override NTP and click the icon under the **Device** column.
The Device Settings page for the selected Edge appears.
- 3 Go to the **NTP** area and select the **Enable Edge Override** checkbox.

NTP ✓ Enable Edge Override ⚠

Edge as NTP Client

Source Interface Auto

Private NTP Servers Enabled ⓘ ☒

Servers

IP Address or DNS Name	
10.1.1.1	– +

Edge as NTP Server

Enabled ⓘ ☒

Authentication None MD5

Keys

Trusted Key #	Key Value
1	1

- 4 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.

- 5 Override the other NTP settings specified in the Profile associated with the Edge by following the Step 3 and 4 in [Configure NTP Settings for Profiles](#).
- 6 Click **Save Changes**. The NTP settings for the Edge will be overridden.


What to do next

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 more information about how to run remote diagnostic tests on an Edge, see [Remote Diagnostics](#).

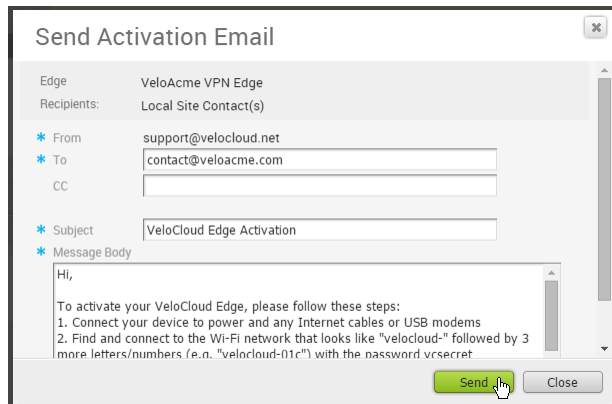
Configure Edge Activation

This section describes how to initiate Edge activation.

Once an Edge configuration has been saved, it is assigned an activation key. Edge activation begins by clicking the **Send Activation Email** link on the **Edge Overview** Tab.



A **Send Activation Email** dialog box appears with a suggested email to be sent to a Site Contact. Simple instructions are provided for the Site Contact to connect and activate Edge hardware. Specify additional instructions in the email for connecting specific site WAN and LAN networks to the Edge.



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 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 [Remote Diagnostics](#).

Edge Software Image Management

19

Read the following topics next:

- [Edge Software Image Management Overview](#)
- [Enable Edge Software Image Management](#)
- [Edge Image Assignment and Access](#)
- [Upgrade SD-WAN Edges](#)

Edge Software Image Management Overview

The Edge Software Image Management feature provides Enterprise Super Users the ability to upgrade SD-WAN Edge firmware without relying on VMware Support or the Partner.

Traditionally, whenever a new Edge image is published by VMware SD-WAN, the Enterprise Administrators will have to request the VMware support or the Partner to upgrade the software on their enterprise Edges. The VMware 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 Super Users the ability to upgrade SD-WAN Edge firmware without relying on VMware 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.

Enable Edge Software Image Management

The Edge software image management feature is deactivated by default for customers. Only an Operator (or VMware Support) can activate this feature for a Direct Enterprise and the Partner. In turn, the Partners can enable this feature for their Partner Enterprise customers. The feature can be enabled during or after the customer creation. The Enterprises with Edge software image

management deactivated must engage with VMware Support or Partner for Edge firmware upgrades.

Enable Edge Software 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 Super User to manage the available list of software images assigned for an Enterprise by enabling the **Manage Software Image** checkbox under **Customer Configuration** in the **New Customer** screen. For more information, see the *Create New Customer* section in the *VMware SD-WAN Operator Guide*.

Enable Edge Software 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 Super user to manage the available list of software images for the customer by enabling the **Manage Software Image** checkbox under **Customer Configuration** in the **New Customer** screen. 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 more information, see the *Create New Customer* section in the *VMware SD-WAN Partner Guide*.

Enable Edge Software Image Management for Existing Customer

As an Operator user or a Partner Administrator, you can enable the Edge software image management feature for an existing customer from **Configure > Customer > Edge Image management** area. When the feature is enabled, the default software image is the only assigned software image for the customer. You can assign additional software images post enabling the feature.

For more information, see the *Configure Customers* and *Manage Edge Software Images* sections in the *VMware SD-WAN Operator Guide*.

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 VMware upgrades a hosted Orchestrator to a newer version of VMware 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 VMware 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 VMware SD-WAN Edge software versions and recommended releases, refer <https://kb.vmware.com/s/article/80741>.

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 VMware 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 more information, see the *Software Images* and *Manage Operator Profiles* section in the *VMware SD-WAN Operator Guide*.

Upgrade SD-WAN Edges

With the Edge software image management feature enabled, Enterprises can upgrade a specific Edge or a set of Edges, or all Edges.

Upgrade All Edges

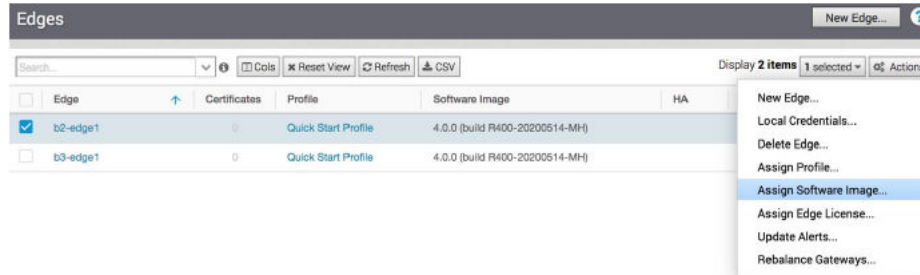
To upgrade all the Edges of an Enterprise, under **Administration > System Settings**, change the default software image used by the Enterprise.

Software Image

Software Image	Is Default?
4.3.0 (build R430-20210421-MH-d6fe9703cb) Description: Used by: edge(x)	<input checked="" type="radio"/>
3.3.0 (build R330-20190723-GA) Description: Used by: edge(x)	<input type="radio"/>

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 > Actions > Assign Software Image**.



For more information, see [Manage Edges](#) and [Assign Software Image](#).

SD-WAN Gateway Migration

20

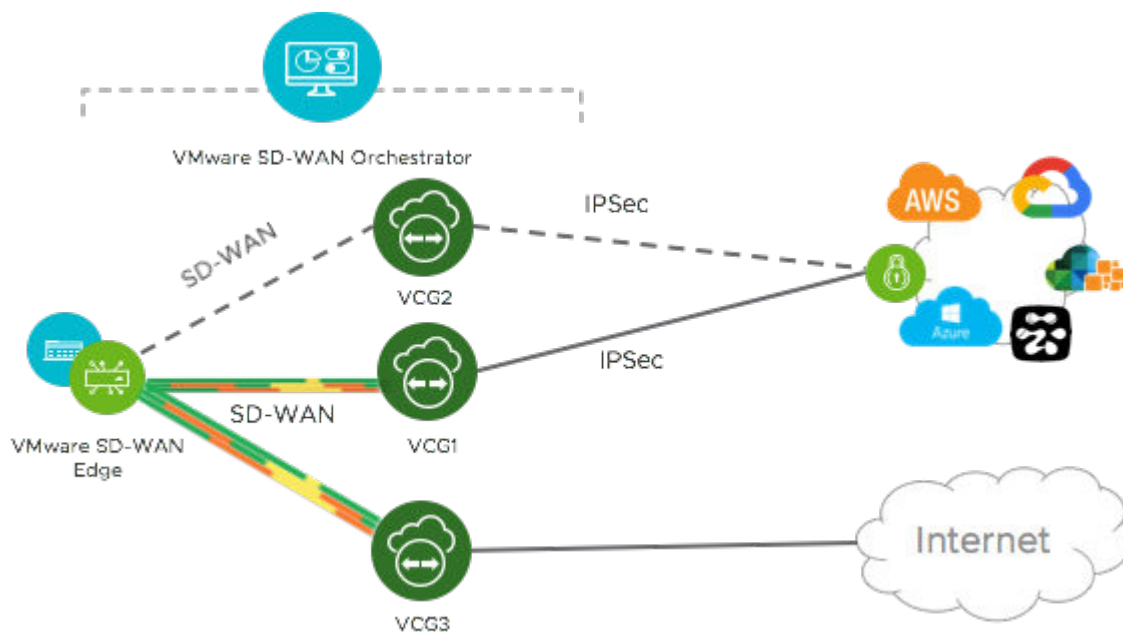
VMware SD-WAN Orchestrator provides a self-service migration functionality that allows you to migrate from your existing Gateway to a new Gateway without your Operator's support.

Gateway migration may be required in the following scenarios:

- Achieve operational efficiency.
- Decommission old Gateways.

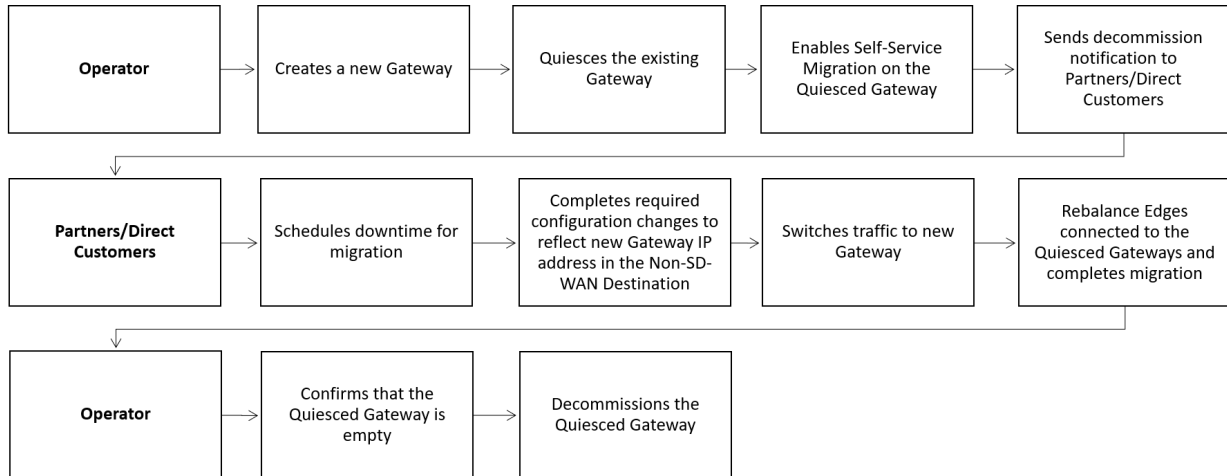
Gateways are configured with specific roles. For example, a Gateway with data plane role is used to forward data plane traffic from source to destination. Similarly, a Gateway with Control Plane role is called a Super Gateway and is assigned to an Enterprise. Edges within the Enterprise are connected 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 more information about the Gateway roles, see the "Configure Gateways" section in the VMware SD-WAN Operator Guide available at [VMware SD-WAN Documentation](#).

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



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:



Read the following topics next:

- [SD-WAN Gateway Migration - Limitations](#)
- [Migrate Quiesced Gateways](#)
- [What to do When Switch Gateway Action Fails](#)

SD-WAN Gateway Migration - Limitations

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

- Self-service migration is not supported on Partner Gateways.
- There will be a minimum service disruption based on the time taken to switch 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 Gateway is quiesced, the redundant Gateway cannot be the replacement Gateway for the quiesced Gateway.

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. You can migrate your quiesced Gateways using the new Orchestrator UI only.

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:

Prerequisites

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

Procedure

- 1 In the Enterprise portal, click the **Open New Orchestrator UI** option available at the top of the Window.
- 2 Click **Launch New Orchestrator UI** in the pop-up window. The UI opens in a new tab.
- 3 Go to **Settings > Gateway Migration**. The list of quiesced Gateways appears.
- 4 Click **Start** for the quiesced Gateway from which you want to migrate to the new Gateway.
- 5 Make the required configuration to all the NSDs that are configured through the quiesced Gateway.
 - a Click 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 new Gateway 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 click 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.

- 6 Select each NSD and click **Switch Gateway** to switch the traffic from the quiesced Gateway to the 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 configuration changes to the NSD.

Note This confirmation is not applicable for NSDs configured automatically.

- b Click **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". If the Switch Gateway action fails, see [What to do When Switch Gateway Action Fails](#).

- c Click **Next**.

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

- 7 Rebalance either all Edges or the required Edges that are connected to the quiesced Gateway so that the Edges get reassigned to the new Gateway.

- 8 Click **Finish** to complete the Gateway migration.

Results

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

What to do When Switch Gateway Action Fails

During the Gateway migration, when the Switch Gateway action for an NSD fails, perform the following steps to troubleshoot the issue:

Procedure

- 1 In the Enterprise portal, launch the new Orchestrator UI, and then 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 click **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 click **Undo Switch Gateway**.

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

- 3 Click **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.

What to do next

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

An object group consists of a range of IP addresses or Port numbers. 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 ports, by including the object groups in the rule definitions.

You can create Address groups to save the range of valid IP addresses and Port groups for the range of port numbers. 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.

Read the following topics next:

- [Configure Address Groups](#)
- [Configure Port Groups](#)
- [Configure Business Policies with Object Groups](#)
- [Configure Firewall Rules with Object Groups](#)

Configure Address Groups

Address Groups can store a range of IP addresses with different options and/or domain names.

Procedure

- 1 In the Enterprise portal, click **Configure > Object Groups**.
- 2 In the **Address Groups** tab, click **Actions > New**.

- 3 In the **Configure Address Group** window, enter a Name and Description for the Address Group.
- 4 Under **IP Address Ranges**, enter the range of IP Addresses by selecting the Prefix or Mask options: **CIDR prefix**, **Subnet mask**, or **Wildcard mask**, as required.

Configure Address Group

* Name:

Description:

IP Address Ranges

* IP Address	Prefix/Mask	Prefix/Mask Value	
<input type="text" value="10.10.1.1"/>	None ▼	<input type="text"/>	– +
<input type="text" value="109.20.1.0"/>	CIDR prefix ▼	<input type="text" value="24"/>	– +
<input type="text" value="10.0.2.0"/>	Subnet mask ▼	<input type="text" value="255.255.255.0"/>	– +
<input type="text" value="11.1.1.20"/>	Wildcard mask ▼	<input type="text" value="0.0.0.255"/>	– +

Domains

* Domain Name

<input type="text" value="salesforce.com"/>	– +
<input type="text" value="vmware.com"/>	– +

- 5 Under **Domains**, 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.
- 6 Click **Create**.

What to do next

- You can define a business policy and a firewall rule with the Address Group. For more information, see:
 - [Configure Firewall Rules with Object Groups](#)
 - [Configure Business Policies with Object Groups](#)

- You can modify the IP addresses and domain names in an Address Group by clicking **Actions** > **Update** in the **Address Groups** tab.
- To delete an Address Group, click **Actions** > **Delete**. Before deleting an Address Group, ensure to exclude the Address Group from the business policies or firewall rules.

Configure Port Groups

Port Groups can store a range of TCP and UDP port numbers.

Procedure

- 1 In the Enterprise portal, click **Configure** > **Object Groups**.
- 2 In the **Port Groups** tab, click **Actions** > **New**.
- 3 In the **Configure Port Group** window, enter a Name and Description for the Port Group.
- 4 Select the protocol as TCP or UDP and enter the corresponding port numbers as required.

- 5 Click **Create**.

What to do next

You can define a business policy or a firewall rule with the Port Group, to include the range of port numbers.

You can add or modify the port numbers in a Port Group by clicking **Actions** > **Update** in the Port Groups tab.

If you want to delete a Port Group, ensure to exclude the Port Group from the business policies or firewall rules.

Configure Business Policies with Object Groups

While configuring business policies, you can select the existing object groups to match the source or destination. This includes the range of IP addresses or port numbers available in the object groups in the business policy definition.

For more information on business policies, see [Create Business Policy Rules](#).

Procedure

- 1 In the Enterprise portal, click **Configure > Profiles**.
- 2 Select a profile from the list and click the **Business Policy** tab.
- 3 Click **New Rule** or **Actions > New Rule**.
- 4 Enter a name for the business rule.
- 5 In the **Match** area, click **Object Group** for the source.
- 6 Select the relevant Address Group and Port Group from the drop-down list.

If the selected address group contains any domain names, then they would be ignored when matching for the source.

- 7 If required, you can select the Address and Port Groups for the destination as well.

- 8 Choose the other actions as required and click **OK**.

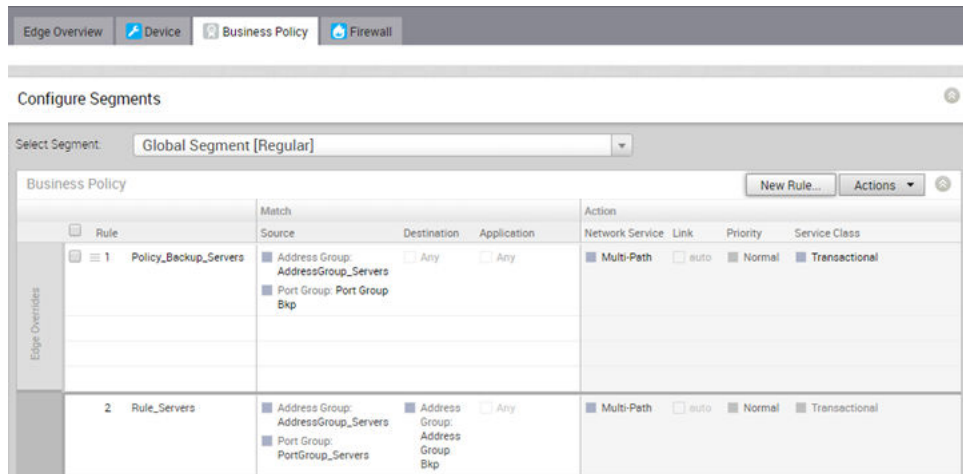
Results

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 click the **Business Policy** tab.
- 2 Click **New Rule** or **Actions > New Rule**.

- 3 Define the rule with relevant object groups and other actions.

The Business Policy tab of the Edge displays the policies from the associated profile along with the policies specific to the Edge.



Note By default, the business policies are assigned to the global segment. If required, you can choose a segment from the **Select Segment** drop-down and create business policies specific to the selected segment.

What to do next

You can modify the object groups with additional IP addresses and port numbers. The changes are automatically included in the business policies that use the object groups.

Configure Firewall Rules with Object Groups

While configuring Firewall rules, you can select the existing object groups to match the source or destination. This includes the range of IP addresses or port numbers available in the object groups in the rules.

For more information on Firewall Rules, see [Configure Firewall Rules](#).

Procedure

- 1 In the Enterprise portal, click **Configure > Profiles**.
- 2 Select a profile from the list and click the **Firewall** tab.
- 3 Click **New Rule** or **Actions > New Rule**.
- 4 Enter a name for the Firewall rule.
- 5 In the **Match** area, click **Object Group** for the source.
- 6 Select the relevant Address Group and Port Group from the drop-down list.

If the selected address group contains any domain names, then they would be ignored when matching for the source.

- 7 If required, you can select the Address and Port Groups for the destination as well.

Configure Rule

Rule Name:

Match

Source:

Address Group:

Port Group:

Destination:

Application:

Action

Firewall:

- 8 Choose required Action and click **OK**.

Results

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.

- 1 Navigate to **Configure > Edges**, select an Edge, and click the **Firewall** tab.
- 2 Click **New Rule** or **Actions > New Rule**.
- 3 Define the rule with relevant object groups and other actions.

The Firewall tab of the Edge displays the Firewall rules from the associated profile along with the rules specific to the Edge.

Edge Overview ☒ Device ☒ Business Policy ☒ **Firewall**

Firewall Status: ☒ On

Configure Segments

Configure Segment:

Firewall Rules		Match	Destination	Application	Action
Rule	Source	Destination	Application	Action	
1 Firewall_Servers	Address Group: AddressGroup_Servers Port Group: PortGroup_Servers	<input type="checkbox"/> Any	<input type="checkbox"/> Any	<input type="button" value="Allow"/>	
2 Firewall_Backup	Address Group: Address Group Bkp Port Group: Port Group Bkp	<input type="checkbox"/> Any	<input type="checkbox"/> Any	<input type="button" value="Allow"/>	

Note By default, the firewall rules are assigned to the global segment. If required, you can choose a segment from the **Select Segment** drop-down and create firewall rules specific to the selected segment.

What to do next

You can modify the object groups with additional IP addresses and port numbers. The changes are automatically included in the Firewall rules that use the object groups.


Site Configurations

22

Topologies for data centers that include a SD-WAN Hub and VMware branch configurations (Gold, Silver, and Bronze branches) are configured using both MPLS and Internet connections. Legacy branch configurations (those without a SD-WAN 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 the Gold, Silver, and Bronze 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.



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.

Read the following topics next:

- [Data Center Configurations](#)
- [Configure Branch and Hub](#)

Data Center Configurations

An SD-WAN Edge in a data center can act as a hub to direct traffic to/from branches. The SD-WAN 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.

The following describes the various designs with different options of how SD-WAN Edge can be inserted into the topology.

Option	Description
Hub 1	Data Center or regional hub site with SD-WAN Edge deployed in two-arm topology.
Hub 2	Data Center or regional hub site with SD-WAN Edge deployed in one-arm topology (same interface carries multiple WAN links).
Legacy 1 and 2	Classic MPLS sites.
Silver 1	SD-WAN Edge is deployed off-path. SD-WAN Edge creates overlay across both MPLS and Internet paths. Traffic is first diverted to the SD-WAN Edge.
Silver 2	SD-WAN Edge is deployed in-path as the default gateway. It is always the default gateway. This topology is simpler but makes SD-WAN Edge a single point of failure and may require HA.
Bronze 1	Dual-Internet site (one of the links is behind a NAT router).

Configure Branch and Hub

This section provides an overview of configuring SD-WAN Edge in a two-arm configuration.

Overview

To configure the SD-WAN Edge in a two-arm configuration:


- 1 Configure and activate Hub 1
- 2 Configure and activate the Silver 1 site
- 3 Enable branch-to-hub tunnel (Silver 1 to Hub 1)
- 4 Configure and activate Bronze 1 site
- 5 Configure and activate Hub 2
- 6 Configure and activate Silver 2 site

The following sections describe the steps in more detail.

Configure and Activate Hub 1

This step helps you understand the typical workflow of how to bring up SD-WAN Edge at the hub location. SD-WAN Edge is deployed with two interfaces (one interface for each WAN link).


You will use the Virtual Edge as a hub. Below is an example of the wiring and IP address information.



Configure and Activate Hub 1 SD-WAN Edge to Reach the Internet

Because this is the data center/hub site, it is unlikely that the SD-WAN Edge can get its WAN IP using DHCP. Thus, you will need to first enable the SD-WAN Edge to connect to the Internet through the data center firewall so that SD-WAN Edge can be activated.

1. Configure a PC with a static IP **192.168.2.100/24** and gateway **192.168.2.1** which is the default LAN setting for accessing a SD-WAN Edge. Connect the PC to the SD-WAN Edge LAN interface.
2. From the PC, browse to <http://192.168.2.1> (the local Web interface of the SD-WAN Edge). Click the link **review the configuration**.



3. Configure the GE2 static WAN IP and default gateway of the SD-WAN Edge so that it can reach the Internet.

Click **Save** and provide login/password of **admin/admin**.

Typically at the data center/hub site, the static IP address will be assigned to you and the enterprise IT admin will configure the firewall to translate the SD-WAN Edge WAN IP to a Public IP and also filter the appropriate traffic (outbound: TCP/443, inbound: UDP/2426, UDP/500, UDP/4500).

- At this point, the Internet status should show Connected.

After configuration of the SD-WAN Edge static WAN IP address and associated firewall configuration is complete, the SD-WAN Edge Internet status shows "Connected".



Enable the Virtual SD-WAN Edge in Default Profile

- Login to the SD-WAN Orchestrator.
- The default VPN profile allows the activation of the SD-WAN Edge 500.

Activate Hub 1 SD-WAN Edge

- Go to **Configure > Edges** and add a new SD-WAN Edge. Specify the correct model and the profile (we use the Quick Start VPN Profile).
- Go to the hub SD-WAN 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.
- Copy the activation URL and paste that to the browser on the PC connected to the SD-WAN Edge or just click on the activation URL from the PC browser.
- Click on **Activate** button.

- Now the **DC1-VCE** data center hub should be up. Go to **Monitor > Edges**. Click the **Edge Overview** tab. The public WAN link capacity is detected along with the correct public IP **71.6.4.9** and ISP.



- Go to **Configure > Edges** and select **DC1-VCE**. Go to the **Device** tab and scroll down to the **Interface Settings**.

You will see that the registration process notifies the SD-WAN Orchestrator of the static WAN IP address and gateway that was configured through the local UI. The configuration on the VMware will be updated accordingly.

- Scroll down to the **WAN Settings** section. The Link Type should be automatically identified as **Public Wired**.

Configure the Private WAN Link on Hub 1 SD-WAN Edge

- Configure the private MPLS Edge WAN interface directly from the SD-WAN 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.
- Under **WAN Settings**, click the **Add User Defined WAN Overlay** button (see the following screen capture).
- 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. Click 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.

- 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.
- Validate that the WAN link is configured and save the changes.

You are done with configuring the SD-WAN Edge on the hub. You will not see the User Defined MPLS overlay that you just added until you enable a branch SD-WAN Edge.

For more information, see:


- [Configure Interface Settings](#)
- [Configure Edge WAN Overlay Settings](#)

Configure Static Route to LAN Network Behind L3 Switch

Add a static route to the **172.30.0.0/24** subnet through the L3 switch. You need to specify the interface GE3 to use for routing to the next hop. Make sure you enable the Advertise checkbox so other SD-WAN Edge can learn about this subnet behind L3 switch. For more information, see [Configure Static Route Settings](#).

Configure and Activate Silver 1 Site


This step helps you understand the typical workflow of how to insert the SD-WAN Edge at a Silver site. The SD-WAN 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.



Activate the Silver 1 Site Branch SD-WAN Edge

In this example, we assume that the SD-WAN Edge gets its public IP address using DHCP, so there is no configuration required. SD-WAN Edge ships with default configuration to use DHCP on all routed interfaces.

- 1 Create a new Edge **SILVER1-DCE** and select the appropriate Model and configuration profile.



- 2 Activate this SD-WAN Edge by connecting a PC to its LAN or Wi-Fi.

- 3 The SD-WAN Edge should now be active in the SD-WAN Orchestrator with one public link. We can now configure the private WAN link.

Configure the Private WAN Link on the Silver 1 Site SD-WAN Edge

At this point, we need to build the IP connectivity from the SD-WAN Edge towards the L3 switch.

- 1 Go to **Configure > Edges**, select the **SILVER1-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.
- 2 Under the **WAN Settings** section, click **Add User Defined WAN Overlay**.
- 3 Define the WAN overlay for the MPLS path. Select the **Link Type** as **Private**. Specify the next-hop IP (10.12.1.2) of the WAN link in the IP Address field. Choose the GE3 as the Interface. Click the **Advanced** button. **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.
- 4 Set the Bandwidth Measurement to **Measure Bandwidth**. This will cause the branch SD-WAN Edge to run a bandwidth test with the hub SD-WAN Edge just like what happens when it connects to the SD-WAN Gateway.
- 5 Validate that the WAN link is configured and save the changes.

Configure Static Route to LAN Network Behind L3 Switch

Add a static route to **192.168.128.0/24** through the L3 switch. You need to specify the Interface GE3. Make sure you enable the Advertise checkbox so other SD-WAN Edge learn about this subnet behind L3 switch.

Enable Branch to Hub Tunnel (Silver 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 SD-WAN Gateway over the Internet path and not the tunnel to the hub. We will need to enable 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.

Enable Cloud VPN and Edge to SD-WAN Hub tunnel

- 1 Step 1: Go to the **Configure > Profiles**, select **Quick Start VPN Profile** and go to the **Device** tab. Enable the Cloud VPN and do the following.
 - Under **Branch to Hubs**, check the **Enable** checkbox.
 - Under **Branch to Branch VPN**, check the **Enable** checkbox.

- Under **Branch to Branch VPN**, uncheck the Use Cloud Gateways checkbox. Doing this will deactivate the data plane through the SD-WAN 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.

Click the button **Select Hubs**. Next, move the **DC1-VCE** to the right. This will designate the **DC1-VCE** to be a SD-WAN Hub. Click the **DC1-VCE** in the Hubs, and click both **Enable Backhaul Hubs** and **Enable B2B 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 SD-WAN Hubs and used for Branch to Branch VPN hubs.


- 2 At this point, the direct tunnel between the branch and the hub SD-WAN Edge should come up. The debug command will now also show the direct tunnel between the branch and the hub. The below example is from the **SILVER1-VCE**. Note that the additional tunnels to **71.6.4.9** and **172.31.2.1**. These are the direct tunnels to the hub SD-WAN Edge (GE2 over public Internet and GE3 over private link).

Configure and Activate Bronze 1 Site

This step helps create a Bronze site--a dual Internet site with one DIA and one broadband. Below is an example of the wiring and IP address information. The **BRONZE1-VCE** SD-WAN Edge LAN and activate the SD-WAN Edge. There is no configuration required on the WAN because it uses DHCP for both WAN interfaces.

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.



Configure the Hub 2 SD-WAN Edge to Reach the Internet

- 1 Connect a PC to the SD-WAN Edge and use the browser to point to <http://192.168.2.1>.
- 2 Configure the hub SD-WAN Edge to reach the Internet by configuring the first WAN interface, GE2.

Configuration

Changes may require the link to briefly go offline.
(Fields marked with * are required.)

* Addressing:	<input type="radio"/> DHCP <input checked="" type="radio"/> Static <input type="radio"/> PPPoE
* IP Address:	172.29.0.2
* Subnet Mask:	255.255.255.0
* Gateways:	172.29.0.4
* Autonegotiation:	<input checked="" type="radio"/> On <input type="radio"/> Off

Add the Hub 2 SD-WAN Edge to the SD-WAN Orchestrator and Activate

In this step, you will create the second hub SD-WAN Edge, called **DC2.VCE**.

- 1 On the SD-WAN Orchestrator, go to **Configure > Edges**, select **New Edge** to add a new SD-WAN Edge.
- 2 Go to **Configure > Edges**, select the SD-WAN 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 SD-WAN 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.


- 3 At this point, you need to create the overlay. Under **WAN Settings**, click **Add User Defined WAN Overlay**.
- 4 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**.
- 5 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.
- 6 Activate the SD-WAN 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.

Add the Hub 2 SD-WAN Edge to the Hub List in the Quick Start VPN Profile

- 1 Go to **Configure > Profiles** and select the profile **Quick Start VPN**.
- 2 Go to the **Device** tab and add this new SD-WAN Edge to a list of hubs.

Configure and Activate Silver 2 Site

This step helps you create a Silver site--a hybrid site, which has the SD-WAN Edge behind CE router as well as SD-WAN Edge being the default router for the LAN. Below is an example of the wiring and IP address information for each hardware.



Connect a PC to the SD-WAN Edge LAN or Wi-Fi and use the browser to point to <http://192.168.2.1>.

Configure Dynamic Routing with OSPF or BGP

23

This section describes how to configure dynamic routing with OSPF or BGP.

SD-WAN 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 SD-WAN Edge. The Overlay Flow Control (OFC) enables enterprise-wide route visibility and control for ease of programming and for full and partial overlay.

VMware 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 the section titled [OSPF/BGP Redistribution](#).

Note In the 3.2 release, both BGP and OSPF can be enabled in a SD-WAN Edge at a time.

Read the following topics next:

- [Enable OSPF](#)
- [Configure BGP](#)
- [OSPF/BGP Redistribution](#)
- [BFD Settings](#)
- [Overlay Flow Control](#)

Enable OSPF


Open Shortest Path First (OSPF) can be enabled only on a LAN interface as a 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.

To enable OSPF, perform the steps on this procedure:


- 1 Configure OSPF for VPN profiles.
 - a Go to **Configure > Profile**.

- b Click the **Device** icon corresponding to the VPN profile for which you want to configure OSPF.

The **Configure Segments** screen appears.




- c In the **OSPF Areas** section, turn **ON** the **OSPF Areas** toggle button.
- d Configure the redistribution settings for OSPF areas.
- 1 From the **Default Route** drop-down menu, choose an OSPF route type (E1 or E2) to be used for default route.
 - 2 From the **Advertise** drop-down menu, 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.
 - 3 If applicable, check the **Overlay Prefixes** checkbox.
 - 4 Optionally, to enable injection of BGP routes into OSPF, select the **BGP** checkbox. BGP routes can be redistributed into OSPF, so if this is applicable, enter or choose the configuration options as follows:
 - a In the **Set Metric** textbox, 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.
 - b 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 the **ID** text box, enter an **OSPF Area ID**.
 - 6 In the **Name** textbox, enter a descriptive name for your area.

- 7 By default, the **Normal** type is selected. Only **Normal** type is supported at this time.
 - 8 Add additional areas, if necessary, by clicking .
- 2 Configure routed interface settings for the OSPF-enabled Edge device.

Note SD-WAN Orchestrator supports OSPF **Point to Point** network mode at the Edge and Profile level.

- a In the **Configure Segments** screen, scroll down to the **Device Settings** area of the Edge device for which you want to configure interface and OSPF settings.
- b Click the expand icon corresponding to the Edge.
- c In the **Interface Settings** area, click the **Edit** link of your interface. The Interface Setting screen for the Edge device appears.



Virtual Edge

Interface GE6

Interface Enabled ☒

Capability **Routed**

Segments **All Segments**

Addressing Type **Static**

IP Address **172.16.1.10**

CIDR prefix **29**

Gateway **172.16.1.11**

WAN Overlay ☒ **User Defined Overlay**

OSPF ☒

OSPF Area **100 - OSPF1**

[toggle advance ospf settings](#)

Custom Settings

Hello Timer **10** seconds

Dead Timer **40** seconds

Enable BFD ☒

Enable MD5 Authentication ☒

Key ID **1**

Password

Interface Path Cost **1**

MTU **1380**

Mode **Broadcast**

Passive ☐

VNF Inception ☒ VNF inception is disallowed when an interface is configured for WAN overlays

Multicast ☐ Multicast is not enabled for the selected segment

RADIUS Authentication ☒ Require User Authentication to access WAN

Advertise ☒

ICMP Echo Response ☒

NAT Direct Traffic ☒

Underlay Accounting ☒

Trusted Source ☐

Reverse Path Forwarding **Specific**

VLAN

L2 Settings

Autonegotiate ☒

MTU **1500**

DHCP Server

Type **Enabled** **Relay** **Disabled**

Update GE6 **Cancel**

- d Select the **OSPF** checkbox.
- e From the **OSPF Area** drop-down menu, select an OSPF area.
- f Click the **toggle advance ospf settings** link to configure advanced OSPF settings.
 - 1 Create filters for **Inbound Route Learning** and **Route Advertisement**. For more information, see [Route Filters](#).
 - 2 Click the **Customs Settings** tab and configure the following OSPF settings.
 - a In the **Hello Timer** text box, enter the OSPF Hello time interval in seconds. The allowable range is 1 through 255.
 - b In the **Dead Timer** text box, enter the OSPF Dead time interval in seconds. The allowable range is 1 through 65535.
 - c Select **Enable BFD** to enable subscription to existing BFD session for OSPF.
 - d Select the **Enable MD5 Authentication** checkbox to enable MD5 authentication.
 - e In the **Interface Path Cost** text box, enter the OSPF cost for the interface path.
 - f In the **MTU** text box, enter the Maximum Transmission Unit (MTU) value of the interface.
 - g From the **Mode** drop-down menu, select **Broadcast** or **Point to Point** as the OSPF network type mode. The default OSPF mode is **Broadcast**.
 - h Select the **Passive** checkbox to enable OSPF Passive mode.
 - i Click the **Update** button.
- 3 Click **Save Changes**.

The **Confirm Changes** dialog box appears requesting you to confirm the OSPF areas you want to enable. It also displays how many Edges are affected.

Note If you have Edges that are not associated with the OSPF configuration at the Profile level, then you must configure at the Edge level from **Configure > Edges > Device > Interface Settings area**.

Route Filters

There are two different types of routing: inbound and outbound.

- Inbound routing includes preferences that can be learned or ignored from OSPF and installed into the Overlay Flow Control.
- Outbound Routing indicates what prefixes can be redistributed into the OSPF.

Edge 500: INTERNET2

Interface: INTERNET2

Interface Enabled: ☒

Capability: Routed

Addressing Type: DHCP ▾

Static/PPPoE addressing details must be configured individually per edge.

WAN Overlay: ☒ User Defined Overlay ▾

OSPF: ☒

OSPF Area: 1 - BRANCHES ▾

[toggle advance ospf settings](#)

Custom Settings Inbound Route Learning **Route Advertisement**

Default Action: Advertise ▾

Route Filters:

Route	Action
172.17.1.0/25	Ignore ▾

NAT Direct Traffic: ☐

VLAN:

L2 Settings

Autonegotiate: ☒

* MTU: 1500

Update INTERNET2 Cancel

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.

VMware supports 4-Byte ASN BGP as follows:

- As the ASN of SD-WAN 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


You can configure the BGP per segment for a Profile or an Edge. This section provides steps on how to configure BGP with Underlay Neighbors.

VMware supports 4-Byte ASN BGP. See [Configure BGP](#), for more information.

To configure BGP:

Procedure

- 1 In the Enterprise portal, click **Configure > Profiles**.
- 2 Click the **Device** icon for a profile, or select a profile and click the **Device** tab.
- 3 In the **Device** tab, scroll down to the **BGP Settings** section, click the slider to **ON** position, and then click **Edit**.



4 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.

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 keepalive timer in seconds, which is the duration between the keepalive 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 keepalive 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	<p>Enter the community string to be treated as uplink routes.</p> <p>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.</p> <p>Enter the value in number format ranging from 1 to 4294967295 or in AA:NN format.</p>

- b Click **Add Filter** 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.

In the **Create BGP Filter** window, set the rules for the filter.

Create BGP Filter

Filter Name:

Rules

Match			Action	
Type	Value	Exact Match	Type	Set
Community	100.101	<input checked="" type="checkbox"/>	Permit	Community 12345.11 Community Additive <input checked="" type="checkbox"/>

OK Cancel

Option	Description
Filter Name	Enter a descriptive name for the BGP filter.
Match Type and Value	<p>Choose the type of the routes to be matched with the filter:</p> <ul style="list-style-type: none"> ■ 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.
Set	<p>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:</p> <ul style="list-style-type: none"> ■ 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 checkbox to enable the additive option, which appends the community value to existing communities.

Option	Description
	<ul style="list-style-type: none"> ■ Metric: The matching traffic is routed to the path with the specified metric value. ■ AS-Path-Prepend: Allows prepending multiple entries of Autonomous System (AS) to a BGP route.

Click the Plus(+) Icon to add more matching rules for the filter.

Click **OK**.

Repeat the procedure to create more BGP filters.

The configured filters are displayed in the **BGP Editor** window.

Caution 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.

- c Configure the following settings for IPv4 addressing type.

The screenshot shows the BGP Editor window with the following configuration details:

- Local ASN:** 100
- BGP Settings:**
 - Router ID: (empty)
 - Hold Timers: e.g., 180
 - Keep Alive: e.g., 60
 - Uplink Community: 00:00
- Filter List:**

	Filter Name	Rule Match	Value	Exact Match	Rule Action	Set
Edit	1. Inbound_Corp	Prefix for IPv4	10.1.1.1/24	<input checked="" type="checkbox"/>	Permit	Local Preference 100000
Edit	2. Outbound_Corp	Community	100:101	<input checked="" type="checkbox"/>	Permit	Community 12345:11 Community Additive: Enabled
- IPv4 Neighbors:**

	Neighbor IP	ASN	Inbound Filter	Outbound Filter	Additional Options	
1.	10.0.0.5	200	Inbound_Corp	Outbound_Corp	Max-hop: 2 Uplink: <input checked="" type="checkbox"/> Allow AS: <input checked="" type="checkbox"/> Default Route: <input checked="" type="checkbox"/> Enable BFD: <input checked="" type="checkbox"/> Keep Alive: 60 Hold Timer: 180 Connect Time: 120 MD5 Auth: <input checked="" type="checkbox"/>	Clone
- Route Redistribution:**
 - Overlay Prefix: ☒
 - Turn off AS-PATH Carry Over: ☒
 - Connected Routes: ☒
 - OSPF: ☒
 - Set Metric: 20
 - Default Route: ☒
 - Advertise: Conditional
- Networks:** e.g., 10.10.10/21
- Route Propagation:**
 - Overlay Prefixes Over Uplink: ☒

Option	Description
Neighbor IP	Enter the IPv4 address of the BGP neighbor
ASN	Enter the ASN of the neighbor
Inbound Filter	Select an Inbound filter from the drop-down list
Outbound Filter	Select an Outbound filter from the drop-down list
Additional Options – Click the view all link to configure the following additional settings:	
Local IP	<p>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. If you do not enter any value, the IP address of the physical Interface is used as the source IP address.</p> <p>Important For eBGP, this field is available only when Max-hop count is greater than 1. For iBGP, it is always available as iBGP is inherently multi-hop.</p>
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-

Option	Description
	WAN Hub), by propagating routes learnt 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.
Max-hop	<p>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.</p> <p>Important 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.</p>
Allow AS	Select the checkbox to allow the BGP routes to be received and processed even if the Edge detects its own ASN in the AS-Path.
Default Route	The Default Route adds a network statement in the BGP configuration to advertise the default route to the neighbor.
Enable BFD	Enables subscription to existing BFD session for the BGP neighbor.
Keep Alive	Enter the keepalive timer in seconds, which is the duration between the keepalive 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 keepalive 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 checkbox 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	<p>Enter a password for MD5 authentication. The password should not contain the character \$ followed by numbers. For example \$1, \$123, password\$123 are invalid inputs.</p> <p>Note The MD5 authentication fails if the password has the character \$ followed by numbers.</p>

Click 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 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 more information, see [Multi-hop BGP Routes](#).

- d Click **Advanced** to configure the following advanced settings, which are globally applied to all the BGP neighbors with IPv4 addresses.

Option	Description
Overlay Prefix	Select the checkbox to redistribute the prefixes learned from the overlay.
Turn off AS-Path carry over	By default, this should be left unchecked. Select the checkbox 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. Warning When the AS-PATH Carry Over is deactivated, tune your network to avoid routing loops.
Connected Routes	Select the checkbox to redistribute all the connected Interface subnets.
OSPF	Select the checkbox 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 checkbox 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 checkbox 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. 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:

Default Route Selection		
Global	Per BGP Neighbor	Advertising Options
Yes	Yes	The per BGP neighbor configuration overrides the global configuration and hence default route is always advertised to the BGP peer.
Yes	No	BGP redistributes the default route to its neighbor only when the Edge learns an explicit default route through the overlay or underlay network.

Default Route Selection		
Global	Per BGP Neighbor	Advertising Options
No	Yes	Default route is always advertised to the BGP peer.
No	No	The default route is not advertised to the BGP peer.

- e Click 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)
- f Configure the other settings as required.

BGP Editor

Local ASN: 100

BGP Settings

Filter List

Filter Name	Type	Value	Exact Match	Rule Action
1. Inbound_Corp	Prefix for IPv4	10.1.1.1/24	68	Permit Local Preference: 100000
2. Outbound_Corp	Community	100:101	✓	Permit Community: 12345:11 Community Additive: Enabled

IPv4 IPv6

Neighbors

Neighbor IP	ASN	Inbound Filter	Outbound Filter	Additional Options
1. 2000:1b98:779a:0000:0000:8ae2:7334:0370	100	[none]	Outbound_Corp	no options view all Uplink: <input type="checkbox"/> Allow AS: <input checked="" type="checkbox"/> Default Route: <input checked="" type="checkbox"/> Enable BFD: <input checked="" type="checkbox"/> Keep Alive: [e.g. 10] Hold Timer: [e.g. 10] Connect: [e.g. 10] MD5 Auth: <input type="checkbox"/> MD5 Password: <input type="text"/>

Route Redistribution

Connected Routes: ☒
Default Route: ☒
Advertise: Conditional

Networks

[e.g. 2001:0d58:85a3:0000:0000:8a7e:0370:7334/64]

Advanced OK Cancel

Note The Local IP address configuration is not available for IPv6 address type.

- g Click **Advanced** to configure the following advanced settings, which are globally applied to all the BGP neighbors with IPv6 addresses.

Option	Description
Connected Routes	Select the checkbox to redistribute all the connected Interface subnets.
Default Route	Select the checkbox to redistribute the default route only when Edge learns the BGP routes through overlay or underlay. When 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. Click the Plus (+) Icon to add more network addresses.

- h Click **OK**.

Results

The **BGP Settings** section displays the BGP configuration settings.

Click **Save Changes** in the **Device** screen to save the configuration.

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. If required, you can override the configuration for a specific Edge as follows:

- 1 In the Enterprise portal, click **Configure > Edges**.
- 2 Click the Device Icon next to an Edge, or click the link to an Edge and then click the **Device** tab.
- 3 In the **Device** tab, scroll down to the **BGP Settings** section.
- 4 Select the **Enable Edge Override** checkbox, and then turn on the BGP Settings.
- 5 Click **Edit** to modify the BGP configuration settings for the selected Edge.

- 6 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.

Note

- 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.

- 7 Click **Save Changes** in the **Device** screen to save the modified configuration.

You can also configure BGP for Non SD-WAN Destination Neighbors in an Edge. For more information, see [Configure BGP Over IPsec from Edge to Non SD-WAN Neighbors](#).

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.

About this task:

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, VMware 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.

VMware SD-WAN supports 4-Byte ASN BGP. See [Configure BGP](#), for more 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.




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.



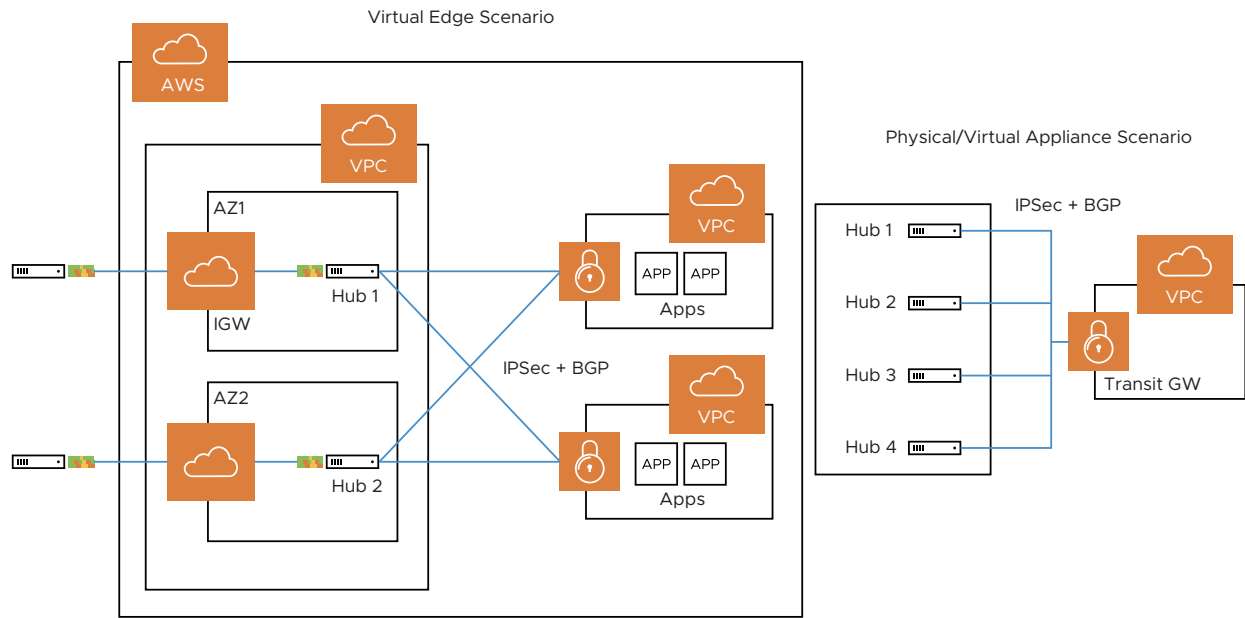
Use Case 3: Edge Connecting to Both AWS and Azure VPN Gateways (Hybrid Cloud)

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.



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.



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 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.

Procedure

To enable BGP with Non SD-WAN neighbors:

- 1 In the Enterprise portal, click **Configure > Edge** and select an SD-WAN Edge.
- 2 Click the **Device** tab.
- 3 In the **Device** tab, scroll down to the **BGP Settings** section and select the **Enable Edge Override** checkbox.

- 4 Click the slider to the **ON** position, and then click the **Edit** button.

BGP Settings On Edit Enable Edge Override

BGP Settings
 Local ASN: 100
 Router ID: not set
 Keep Alive: not set
 Hold Timers: not set
 Uplink Community: not set

Filter List		Rule Match			Rule Action	
Filter Name	Type	Value	Exact Match	Type	Set	
1. Inbound_Corp	Prefix for IPv4	10.1.1.1/24	<input checked="" type="checkbox"/>	Permit	Local Preference 100000	
2. Outbound_Corp	Community	100.101	<input checked="" type="checkbox"/>	Permit	Community 12345.11 Community Additive: Enabled	

IPv4 **IPv6**

Neighbors

Neighbor IP	ASN	Inbound Filter	Outbound Filter	Additional Options
1. 10.0.0.5	200	Inbound_Corp	Outbound_Corp	Max-hop: 2 Uplink: <input checked="" type="checkbox"/> Allow AS: <input checked="" type="checkbox"/> Default Route: <input checked="" type="checkbox"/> Enable BFD: <input checked="" type="checkbox"/> Keep Alive: 60 Hold Timer: 180 Connect Time: 120 MD5 Auth: <input checked="" type="checkbox"/>

NSD Neighbors: not set

Advanced Settings

Route Redistribution
 Overlay Prefix: ☒
 Turn off AS-PATH Carry Over: ☒
 Connected Routes: ☒
 OSPF: ☒
 Set Metric: 20
 Default Route: ☒
 Advertise: Conditional

Networks: not set

Route Propagation
 Overlay Prefixes Over Uplink: ☒

- 5 In the **BGP Editor** window, configure the following settings:
- Enter the local Autonomous System Number (ASN) and then configure the following in the **BGP Settings** section.

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 keepalive timer in seconds, which is the duration between the keepalive messages that are sent to the peer. The range is from 0 to 65535 seconds. The default value is 60 seconds.

Option	Description
Hold Timer	Enter the hold timer in seconds. When the keepalive 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	<p>Enter the community string to be treated as uplink routes.</p> <p>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.</p> <p>Enter the value in number format ranging from 1 to 4294967295 or in AA:NN format.</p>

- b Click the **Add Filter** button to create one or more filters. Filters are applied to the neighbor to deny or change the attributes of the route. The same filter can be used for multiple neighbors, including both Underlay Neighbors and NSD Neighbors.
- c In the **Create BGP Filter** area, set the rules for the filter.

Create BGP Filter

Filter Name:

Rules:

Match			Action	
Type	Value	Exact Match	Type	Set
Community	100.101	<input checked="" type="checkbox"/>	Permit	Community 12345.11

Community Additive ☒

OK Cancel

Option	Description
Filter Name	Enter a descriptive name for the BGP filter.
Match Type and Value	<p>Choose the type of the routes to be matched with the filter:</p> <ul style="list-style-type: none"> ■ 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.

Option	Description
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.
Set	<p>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:</p> <ul style="list-style-type: none"> ■ 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 checkbox 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 prepending multiple entries of Autonomous System (AS) to a BGP route.

To add more matching rules to the filter, click the Plus (+) icon.

Click **OK** to create the filter.

The configured filters are displayed in the **BGP Editor** window.

- 6 Configure Underlay Neighbors for IPv4 and IPv6 addresses, as required. For more information, see [Configure BGP from Edge to Underlay Neighbors](#).
- 7 Configure NSD Neighbors as follows:

Note The 4.3 and later releases support Non SD-WAN (NSD) neighbors. All global settings will be shared by both the Underlay and NSD neighbors, and the filter list can also be used for both types of neighbors. Ensure that you have configured the [Prerequisites](#) before configuring NSD Neighbors.

BGP Editor

Local ASN: 100

BGP Settings

Filter List

	Filter Name	Rule Match	Value	Exact Match	Rule Action	Set
1.	Inbound_Corp	Prefix for IPv4	10.1.1.1/24	<input checked="" type="checkbox"/>	Permit	Local Preference 100000
2.	Outbound_Corp	Community	100:101	<input checked="" type="checkbox"/>	Permit	Community 12345:11 Community Additive: Enabled

IPv4 IPv6

Neighbors

	Neighbor IP	ASN	Inbound Filter	Outbound Filter	Additional Options
1.	10.0.0.5	200	Inbound_Corp	Outbound_Corp	Max-hop: 2 Uplink: <input checked="" type="checkbox"/> Allow AS: <input checked="" type="checkbox"/> Default Route: <input checked="" type="checkbox"/> Enable BFD: <input checked="" type="checkbox"/> Keep Alive: 60 Hold Timer: 180 Connect Time: 120 MD5 Auth: <input checked="" type="checkbox"/> MD5 Password: *****

NSD Neighbors

	NSD Name	Link Name	Tunnel Type	Neighbor IP	ASN	Inbound Filter	Outbound Filter	Additional Options
1.	[none]	[none]	[none]	[e.g. 10.0.13.37]	[e.g. 100]	[none]	[none]	no options view all

Advanced OK Cancel

- a In the **NSD Neighbors** section, configure the following settings:

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 SD-WAN 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 filter from the drop-down list.
Outbound Filter	Select an Outbound filter from the drop-down list.
Additional Options – Click the view all link to configure the following additional settings:	

Option	Description
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 learnt 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 checkbox 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. 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 deactivated by default and this field is not configurable.
Allow AS	Select the checkbox to allow the BGP routes to be received and processed even if the Edge detects its own ASN in the AS-Path.
Default Route	The Default Route adds a network statement in the BGP configuration to advertise the default route to the neighbor.
Enable BFD	Enables subscription to 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 keepalive timer in seconds, which is the duration between the keepalive 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 keepalive 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.

Option	Description
Connect	Enter the time interval to try a new TCP connection with the peer if it detects the TCP session is not passive. The default value is 120 seconds.
MD5 Auth	Select the checkbox to enable BGP MD5 authentication. 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	<p>Enter a password for MD5 authentication. The password should not contain the character \$ followed by numbers. For example \$1, \$123, password\$123 are invalid inputs.</p> <p>Note The MD5 authentication fails if the password has the character \$ followed by numbers.</p>

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 more information, see [Multi-hop BGP Routes](#).

- b Click **Advanced** to configure the following settings.

Note Advanced Settings are shared across both the underlay BGP neighbors and NSD BGP neighbors.

Option	Description
Overlay Prefix	Select the checkbox to redistribute the prefixes learned from the overlay.
Turn off AS-Path carry over	<p>By default, this should be left unchecked. Select the checkbox 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.</p> <p>Warning When the AS-PATH Carry Over is turned off, tune your network to avoid routing loops.</p>
Connected Routes	Select the checkbox to redistribute all the connected Interface subnets.
OSPF	Select the checkbox to enable OSPF redistribute into BGP.

Option	Description
Set Metric	When you enable OSPF, enter the BGP metric for the redistributed OSPF routes. The default value is 20.
Default Route	Select the checkbox 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 checkbox 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.

Default Route Selection		
Global	Per BGP Neighbor	Advertising Options
Yes	Yes	The per BGP neighbor configuration overrides the global configuration and hence default route is always advertised to the BGP peer.
Yes	No	BGP redistributes the default route to its neighbor only when the Edge learns an explicit default route through the overlay or underlay network.
No	Yes	Default route is always advertised to the BGP peer.
No	No	The default route is not advertised to the BGP peer.

- Click **OK** to save the configured filters and NSD Neighbors.

The **BGP Settings** section displays the configured settings.

- Click **Save Changes** in the **Device** screen to save the configuration.

Configure BGP over IPsec from Gateways

You can configure BGP Settings for SD-WAN Gateways over IPSec tunnels.

Only eBGP is supported with BGP over IPsec.

To configure the BGP settings for a Gateway:

Prerequisites

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.

Ensure that you have configured the following:

- Create a Non SD-WAN Destination via Gateway for one of the following sites:
 - Configure a Non VMware SD-WAN Site 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 via Gateway
 - Configure a Non SD-WAN Destination of Type Generic IKEv2 Router via Gateway
 - Configure a Microsoft Azure Non SD-WAN Destination
- 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 more information on **Distributed Cost Calculation**, refer to the **Configure Distributed Cost Calculation** section in the *VMware SD-WAN Operator Guide* available at: <https://docs.vmware.com/en/VMware-SD-WAN/index.html>.

Procedure

- 1 In the Enterprise portal, click **Configure > Network Services**.
- 2 In the **Non SD-WAN Destinations via Gateway** area, click the **Edit** link in the **BGP** column that corresponds to the Non SD-WAN Destination.

Non SD-WAN Destinations via Gateway								
					New...	Delete...	Actions ▾	
Name	Servers	Tunnels	Pre-Notifications ⓘ	Alerts ⓘ	Used By	Segment	BGP	BFD
<input type="checkbox"/> NSD_AWS	Type: AWS VPN Gateway Primary: 54.183.9.191 Secondary: 54.183.9.192	Not enabled	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	1 Profile	Global Segment	Edit	Edit
<input type="checkbox"/> CPM	Type: Check Point Primary: 1.2.3.4 Secondary: none	Not enabled	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	0			
<input type="checkbox"/> Gen_IKEv1_NSD	Type: Generic IKEv1 Router (Route Based VPN) Primary: 8.36.116.14 Secondary: none	Not enabled	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	1 Profile	Global Segment	Edit	Edit
<input type="checkbox"/> Zscaler_NSD	Type: Zscaler Primary: 54.183.9.191 Secondary: none	Not enabled	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	0			

3 In the **BGP Editor** window, click the slider to **ON** to configure the BGP settings.

- a Click **Add Filter** 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.

In the **Create BGP Filter** window, set the rules for the filter.

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: <ul style="list-style-type: none"> ■ Prefix: Choose to match with a prefix and enter the 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.
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: <ul style="list-style-type: none"> ■ 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.

Option	Description
	<ul style="list-style-type: none"> ■ Metric: The matching traffic is routed to the path with the specified metric value. ■ AS-Path-Prepend: Allows prepending multiple entries of Autonomous System (AS) to a BGP route.

Click the Plus(+) Icon to add more matching rules for the filter.

Click **OK**.

Repeat the procedure to create more BGP filters.

The configured filters are displayed in the **BGP Editor** window.

- b In the **BGP Editor** window, configure the BGP settings for the Primary and Secondary Gateways.

Note The Secondary Gateway option is available only if you have configured a secondary Gateway for the corresponding Non SD-WAN Destination.

The screenshot shows the **BGP Editor** window. At the top, it says "BGP Enabled: On". Below this is a "Filter List" table with two filters: "Inbound_Corp" and "Outbound_Corp".

Filter Name	Rule Match	Type	Value	Exact Match	Rule Action	Type	Set
1. Inbound_Corp	Prefix		10.1.1.1/24	<input checked="" type="checkbox"/>	Permit	Local Preference	100000
2. Outbound_Corp	Community		100.101	<input checked="" type="checkbox"/>	Permit	Community	12345.11

Below the filter list, the "Primary Cloud Gateway: gateway-1" section is shown. It includes fields for "Local ASN" (100) and "Router ID" (0.0.0.0). Under "Neighbors", there are two entries:

Tunnel Type	Neighbor IP	ASN	Inbound Filter	Outbound Filter	Additional Options
1. Primary	10.0.10.35	101	Inbound_Corp	Outbound_Corp	Max-hop: 1 Local IP: 10.1.1.0 Max-hop: 3 Allow AS: <input checked="" type="checkbox"/> Default Route: <input checked="" type="checkbox"/> Enable BFD: <input checked="" type="checkbox"/> Keep Alive: 10 Hold Timer: 10 Connect: 10
2. Secondary	10.0.20.45	102	Inbound_Corp	Outbound_Corp	Max-hop: 1 Local IP: 10.1.2.1 Max-hop: 1 Allow AS: <input checked="" type="checkbox"/> Default Route: <input checked="" type="checkbox"/> Enable BFD: <input checked="" type="checkbox"/> Keep Alive: 10 Hold Timer: 10 Connect: 10

At the bottom right, there are "Save" and "Cancel" buttons.

Option	Description
Local ASN	Enter the local Autonomous System Number (ASN)
Router ID	Enter the BGP Router ID

Option	Description
Neighbor IP	Enter the IP address of the BGP neighbor
ASN	Enter the ASN of the neighbor
Inbound Filter	Select an Inbound filter from the drop-down list
Outbound Filter	Select an Outbound filter from the drop-down list
Additional Options – Click 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. 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.
Allow AS	Select the checkbox 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 keepalive timer in seconds, which is the duration between the keepalive 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 keepalive 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.

- c Click **OK** to save the changes.

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](#)

Monitor BGP Events

You can view the events related to the BGP sessions.

In the Enterprise portal, click **Monitor > Events**.

To view the events related to BGP, 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 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

You can also view the events in the new Orchestrator UI.

Click **Launch New Orchestrator UI** in the pop-up window. The UI opens in a new tab displaying the monitoring options.

Click **Events**. Click the Filter Icon in the **Search** option to filter the BGP events.

The following image shows some of the BGP events.

Monitor Configure Diagnostics							
Events							
<div> <div>Past 12 Hours</div> <div>Q Search</div> <div> <div>Filter Icon</div> <div>CSV</div> <div>Message contains B</div> <div>CLEAR ALL</div> </div> </div>							
Event	User	Segment	Edge	Severity	Time	Message	
BGPv6 session established to edge neighbor		Global Segment	b1-edge1	Info	Jul 16, 2021, 12:30:40 PM	BGPv6 session up for edge [b1-edge1] to neighbor IP: [fd00:1:1:4::1]	
BGP session established to edge neighbor		Global Segment	b1-edge1	Info	Jul 16, 2021, 12:29:40 PM	BGP session up for edge [b1-edge1] to neighbor IP: [172.16.1.3]	
BGP session established to edge neighbor		Global Segment	b1-edge1	Info	Jul 16, 2021, 12:29:40 PM	BGP session up for edge [b1-edge1] to neighbor IP: [172.16.1.11]	
BGPv6 session established to edge neighbor		Global Segment	b1-edge1	Info	Jul 16, 2021, 12:29:40 PM	BGPv6 session up for edge [b1-edge1] to neighbor IP: [fd00:1:1:3::1]	
BGP session established to edge neighbor		Global Segment	b1-edge1	Info	Jul 16, 2021, 12:29:40 PM	BGP session up for edge [b1-edge1] to neighbor IP: [172.17.1.3]	
BGP session established to edge neighbor		Global Segment	b1-edge1	Info	Jul 16, 2021, 12:29:40 PM	BGP session up for edge [b1-edge1] to neighbor IP: [172.17.1.11]	
BGPv6 session established to edge neighbor		Global Segment	b1-edge1	Info	Jul 16, 2021, 12:29:40 PM	BGPv6 session up for edge [b1-edge1] to neighbor IP: [fd00:1:2:3::1]	
BGPv6 session established to edge neighbor		Global Segment	b1-edge1	Info	Jul 16, 2021, 12:29:40 PM	BGPv6 session up for edge [b1-edge1] to neighbor IP: [fd00:1:2:4::1]	
Edge BGP neighbor unavailable		Global Segment	b1-edge1	Info	Jul 16, 2021, 12:29:39 PM	BGP session down for edge [b1-edge1] to neighbor IP: [172.16.1.3]	
Edge BGP neighbor unavailable		Global Segment	b1-edge1	Info	Jul 16, 2021, 12:29:39 PM	BGP session down for edge [b1-edge1] to neighbor IP: [172.16.1.11]	
Edge BGPv6 neighbor unavailable		Global Segment	b1-edge1	Info	Jul 16, 2021, 12:29:39 PM	BGPv6 session down for edge [b1-edge1] to neighbor IP: [fd00:1:1:3::1]	
Edge BGPv6 neighbor unavailable		Global Segment	b1-edge1	Info	Jul 16, 2021, 12:29:39 PM	BGPv6 session down for edge [b1-edge1] to neighbor IP: [fd00:1:1:4::1]	

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 Enterprise portal, click **Test & Troubleshoot > 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 that 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 more information on Remote Diagnostics, see [Performing Remote Diagnostics Tests](#).

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.

The redistribution behavior is standardized 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.

As an example, consider the redistribution of the prefix 192.168.1.0/24 into BGP. Let us 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 configure 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.

Note The OSPF External Type-1 (OE1) and OSPF External Type-2 (OE2) redistribution metrics are calculated as follows:

- The OE1 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 42 if the route is learned via a Gateway or a Hub Edge.
 - The OE2 redistribution metric is calculated by taking the original route metric and adding the non-preferred metric constant, where the non-preferred metric constant is 8388607. This is why you would observe a very high metric value for an OE2 route type on Edge peers.
-

See [Enable OSPF](#) and [Configure BGP from Edge to Underlay Neighbors](#) for more information.

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. VMware 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.

VMware supports BFD for the following routing protocols:

- BGP on Edges and Partner Gateways
- OSPF on Edges

Configure BFD

VMware SD-WAN allows to configure BFD sessions to detect route failures between two connected entities.

To configure a BFD session:

Procedure

- 1 In the Enterprise portal, click **Configure > Profiles**.
- 2 Click the **Device** Icon for a profile, or select a profile and click the **Device** tab.
- 3 In the **Device** tab, scroll down to the **BFD Rules** section and click the slider to **ON** position.
- 4 Configure the following settings:

	Peer Address	Local Address	Multihop	Timers
1	172.21.1.1	172.21.1.20	<input checked="" type="checkbox"/>	Detect Multiplier: 3 Receive Interval (ms): 300 Transmit Interval (ms): 300
2	172.21.4.1	172.21.4.20	<input type="checkbox"/>	Detect Multiplier: 3 Receive Interval (ms): 300 Transmit Interval (ms): 300

- 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 click 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.

Note 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.
- 5 Click the Plus (+) Icon to add details of more peers.

6 Click **Save Changes**.

Results

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 as follows:

- 1 In the Enterprise portal, click **Configure > Edges**.
- 2 Click the Device Icon next to an Edge, or click the link to an Edge and then click the **Device** tab.
- 3 In the **Device** tab, scroll down to the **BFD Rules** section.
- 4 Select the **Enable Edge Override** check box to modify the BFD configuration settings for the selected Edge.

BFD Rules On Enable Edge Override

	Peer Address	Local Address	Multihop	Timers	
1	172.21.1.1	172.21.1.20	<input checked="" type="checkbox"/>	Detect Multiplier ⓘ 3 Receive Interval (ms) ⓘ 300 Transmit Interval (ms) ⓘ 300	<input type="button" value="-"/> <input type="button" value="+"/>
2	172.21.4.1	172.21.4.20	<input type="checkbox"/>	Detect Multiplier ⓘ 3 Receive Interval (ms) ⓘ 300 Transmit Interval (ms) ⓘ 300	<input type="button" value="-"/> <input type="button" value="+"/>

What to do next

VMware SD-WAN supports configuring BFD for BGP and OSPF.

- To enable BFD for BGP, see [Configure BFD for BGP](#).
- To enable BFD for OSPF, see [Configure BFD for OSPF](#).
- To view the BFD sessions, see [Monitor BFD Sessions](#).
- To view the BFD events, see [Monitor BFD Events](#).
- For troubleshooting and debugging BFD, see [Troubleshooting BFD](#).

Configure BFD for BGP

You can configure BFD for BGP on SD-WAN Edges.

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](#).

The following procedure describes how to enable BFD for an already configured BGP session on an Edge. To configure BGP settings, see [Configure BGP from Edge to Underlay Neighbors](#).

To enable BFD for BGP on partner Gateways, you must be an Operator super user. For more information, see the **Configure Partner Handoff** section in the *VMware SD-WAN Operator Guide*.

Procedure

- 1 In the Enterprise portal, click **Configure > Profiles**.
- 2 Click the **Device** Icon for a profile, or select a profile and click the **Device** tab.
- 3 In the **Device** tab, scroll down to the **BGP Settings** section and click **Edit**.
- 4 In the **BGP Editor** window, click **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 more information about NSD Neighbors, see section titled, [Configure BGP Over IPsec from Edge to Non SD-WAN Neighbors](#).

Note A single-hop BFD session is not supported for BGP over IPsec from the SD-WAN Edge.

- 5 Configure the other settings as required and click **OK**.

Results

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 as follows:

- 1 In the Enterprise portal, click **Configure > Edges**.
- 2 Click the Device Icon next to an Edge, or click the link to an Edge and then click the **Device** tab.
- 3 In the **Device** tab, scroll down to the **BGP Settings** section.
- 4 Select the **Enable Edge Override** check box and you can modify the BGP settings for the selected Edge.

When a BGP neighbor receives an update that BFD session is down, the corresponding BGP session immediately goes down and the routes learnt through the BGP peer are flushed without waiting for the expiry of keepalive timer.

Configure BFD for OSPF

You can configure BFD for OSPF on Edges.

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](#).

The following procedure describes how to enable BFD for an already configured OSPF session on an Edge Interface. To configure OSPF settings, see [Enable OSPF](#).

To configure the Interface settings, see [Configure Interface Settings](#).

Procedure

- 1 In the Enterprise portal, click **Configure > Profiles**.
- 2 Click the **Device** Icon for a Profile, or select a Profile and click the **Device** tab.
- 3 In the **Device** tab, scroll down to the **Device Settings** section of an Edge.
- 4 In the **Interface Settings** section, click the **Edit** option for an Interface.
- 5 In the **Interface** window, select the **OSPF** checkbox and choose the **OSPF Area** from the drop-down list.
- 6 Click **toggle advance ospf settings** and in the **Custom Settings** tab, select the **Enable BFD** checkbox.
- 7 Configure the other settings as required and click **Update**.

Results

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 as follows:

- 1 In the Enterprise portal, click **Configure > Edges**.
- 2 Click the Device Icon next to an Edge, or click the link to an Edge and then click the **Device** tab.
- 3 In the **Device** tab, scroll down to the **Device Settings** section and click the **Edit** option for an Interface.
- 4 In the **Interface** window, select the **Override Interface** checkbox and you can modify the Interface settings for the selected Edge.

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.

Configure BFD for Gateways

You can configure BFD Settings for SD-WAN Gateways over IPSec tunnels.

To configure BFD settings for a Gateway:

Prerequisites

Ensure that you have configured the following:

- Create a Non SD-WAN Destination via Gateway for one of the following sites:
 - [Configure a Non VMware SD-WAN Site 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 via Gateway](#)
- [Configure a Non SD-WAN Destination of Type Generic IKEv2 Router via Gateway](#)
- [Configure a Microsoft Azure Non SD-WAN Destination](#)
- Associate the Non SD-WAN Destination to a Profile See [Configure a Tunnel Between a Branch and a Non SD-WAN Destinations via Gateway](#).

Procedure

- 1 In the Enterprise portal, click **Configure > Network Services**.
- 2 In the **Non SD-WAN Destinations via Gateway** area, click the **Edit** link in the **BFD** column that corresponds to the Non SD-WAN Destination.

Non SD-WAN Destinations via Gateway								
					New...	Delete...	Actions ▾	
Name	Servers	Tunnels	Pre-Notifications ⓘ	Alerts ⓘ	Used By	Segment	BGP	BFD
<input type="checkbox"/> NSD_AWS	Type: AWS VPN Gateway Primary: 54.183.9.191 Secondary: 54.183.9.192	Not enabled	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	1 Profile	Global Segment	Edit	Edit
<input type="checkbox"/> CPM	Type: Check Point Primary: 1.2.3.4 Secondary: none	Not enabled	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	0			
<input type="checkbox"/> Gen_IKEv1_NSD	Type: Generic IKEv1 Router (Route Based VPN) Primary: 8.36.116.14 Secondary: none	Not enabled	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	1 Profile	Global Segment	Edit	Edit
<input type="checkbox"/> Zscaler_NSD	Type: Zscaler Primary: 54.183.9.191 Secondary: none	Not enabled	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	0			

- 3 In the **BFD Editor** window, click the slider to **ON** to configure the BFD settings for the Primary and Secondary Gateways.

Note The Secondary Gateway option is available only if you have configured a secondary Gateway for the corresponding Non SD-WAN Destination.

BFD Editor

BFD Enabled: On

Primary Cloud Gateway: gateway-1

	Tunnel Type	Peer Address	Local Address	Multihop	Timers
1	Primary	<input type="text" value="10.0.0.12"/>	<input type="text" value="10.0.100.12"/>	<input type="checkbox"/>	Detect Multiplier <input type="text" value="3"/> Receive Interval (ms) <input type="text" value="300"/> Transmit Interval (ms) <input type="text" value="300"/>
2	Secondary	<input type="text" value="10.0.1.12"/>	<input type="text" value="10.0.101.12"/>	<input type="checkbox"/>	Detect Multiplier <input type="text" value="3"/> Receive Interval (ms) <input type="text" value="300"/> Transmit Interval (ms) <input type="text" value="300"/>

- a **Peer Address** – Enter the IP address of the remote peer to initiate a BFD session.
- b **Local Address** – Enter a locally configured IP address for the peer listener. This address is used to send the packets.
- c **Multihop** – This option is not supported for the Gateways.
- d **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.
- e **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.
- f **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.

Monitor BFD Sessions

You can monitor the BFD sessions on Edges and Gateways.

To monitor the BFD sessions:

- 1 In the Enterprise portal, click **Monitor > Routing > BFD**.
- 2 The Page displays the BFD sessions on Edge and Gateway. You can click the drop-down arrow next to the **Search** option and choose to filter the details by different categories.

The screenshot shows the VMware SD-WAN Administration Guide interface. On the left is a navigation menu with sections: Monitor (Network Overview, Edges, Network Services, Routing, Alerts, Events, Firewall Logs, Reports), Configure, Test & Troubleshoot, and Administration. The 'Routing' section is selected. The main content area is titled 'Routing' and has tabs for 'Multicast' and 'BFD'. The 'BFD' tab is active, showing 'Edge BFD Sessions' and 'Gateway BFD Sessions'.

Edge BFD Sessions

Search: [] Refresh [] Display 24 items

Edge	Segment	Peer Address	Local Address	State	Remote Timers	Local Timers	Events	Session Time
b1-hub1	Global Segment	172.21.4.1	172.21.4.2	UP	rx: 300ms / tx: 300ms	rx: 300ms / tx: 300ms	120 View	4 hour(s), 37 minute(s), 57 sec
b1-hub1	Global Segment	172.21.1.1	172.21.1.2	UP	rx: 300ms / tx: 300ms	rx: 300ms / tx: 300ms	120 View	6 hour(s), 4 minute(s), 45 sec
b1-hub2	Global Segment	172.21.1.1	172.21.1.10	UP	rx: 300ms / tx: 300ms	rx: 300ms / tx: 300ms	120 View	5 hour(s), 51 minute(s), 27 sec
b1-hub2	Global Segment	172.21.4.1	172.21.4.10	UP	rx: 300ms / tx: 300ms	rx: 300ms / tx: 300ms	120 View	4 hour(s), 37 minute(s), 40 sec
b1-hub3	Global Segment	172.21.1.1	172.21.1.20	UP	rx: 300ms / tx: 300ms	rx: 300ms / tx: 300ms	120 View	2 hour(s), 17 minute(s), 26 sec
b1-hub3	Global Segment	172.21.4.1	172.21.4.20	UP	rx: 300ms / tx: 300ms	rx: 300ms / tx: 300ms	120 View	2 hour(s), 17 minute(s), 26 sec
b2-hub1	Global Segment	172.22.4.1	172.22.4.2	UP	rx: 300ms / tx: 300ms	rx: 300ms / tx: 300ms	120 View	1 hour(s), 34 minute(s), 25 sec

Gateway BFD Sessions

Search: [] Refresh [] Display 0 items

Gateway	Segment	Peer Address	Local Address	State	Remote Timers	Local Timers
---------	---------	--------------	---------------	-------	---------------	--------------

The page displays 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

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

You can also view the BFD sessions in the new Orchestrator UI.

- 1 In the Enterprise portal, click **Open New Orchestrator UI**.
- 2 Click **Launch New Orchestrator UI** in the pop-up window. The UI opens in a new tab displaying the monitoring options.
- 3 Click **Routing > BFD**. You can click the Filter Icon next to the **Search** option and choose to filter the details by different categories.

Network Overview

Edges

Network Services

Routing

Alerts

Events

Reports

Application Analytics

Branch Analytics

Multicast Groups

PIM Neighbors

BGP Edge Neighbor State

BFD

BGP Gateway Neighbor State

Edge BFD Sessions

Q Search

Edge	Segment	Peer Address	Local Address	State	Remote Timers	Local Timers	Events	Session Time
b1-hub3	Global Segment	11.99.1	172.211.20	Down	rx: 300ms / tx: 300ms	rx: 300ms / tx: 300ms	110 View	
b1-hub2	Global Segment	11.99.1	172.211.10	Down	rx: 300ms / tx: 300ms	rx: 300ms / tx: 300ms	104 View	
b1-hub1	Global Segment	11.99.1	172.211.2	Down	rx: 300ms / tx: 300ms	rx: 300ms / tx: 300ms	120 View	12 minute(s), 29 second(s)
b4-hub-edge2000	Global Segment	1.4.1.1	1.4.1.100	Down	rx: 1000ms / tx: 1000ms	rx: 300ms / tx: 300ms	87 View	22 second(s)
b4-hub-edge2000	segment1	1.4.1.1	1.4.1.100	Down	rx: 1000ms / tx: 1000ms	rx: 300ms / tx: 300ms	33 View	15 minute(s), 10 second(s)
b4-hub-edge2000	segment2	1.4.12.1	1.4.1.102	Down	rx: 1000ms / tx: 1000ms	rx: 300ms / tx: 300ms	120 View	21 hour(s), 56 minute(s), 55 second(s)
b9-edge1_E540	Global Segment	1.9.1.1	1.9.1.100	Down	rx: 1000ms / tx: 1000ms	rx: 300ms / tx: 300ms	120 View	1 day(s), 14 hour(s), 44 minute(s), 33 second(s)
b1-hub2	Global Segment	172.211.1	172.211.10	Down	rx: 1000ms / tx: 1000ms	rx: 300ms / tx: 300ms	120 View	

Columns

Refresh

31 items

Gateway BFD Sessions

Q Search

Gateway	Segment	Peer Address	Local Address	State	Remote Timers	Local Timers	Events	Session Time
---------	---------	--------------	---------------	-------	---------------	--------------	--------	--------------

No BFD events available for selected enterprise

Monitor BFD Events

You can view the events related to the BFD sessions.

In the Enterprise portal, click **Monitor > Events**.

To view the events related to BFD, 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 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

You can also view the events in the new Orchestrator UI.

Click **Launch New Orchestrator UI** in the pop-up window. The UI opens in a new tab displaying the monitoring options.

Click **Events**. Click the Filter Icon next to the **Search** option to filter the BFD events.

The following image shows some of the BFD events.

Network Overview

Edges

Network Services

Routing

Alerts

Events

Firewall Logs

Reports

Past 12 Hours

Q Search

CSV

Message contains

CLEAR ALL

Event	User	Segment	Edge	Severity	Time	Message
Edge BFDv6 neighbor unavailable		Global Segment	bl-hub1	Info	Jul 14, 2021, 4:41:29 PM	BFDv6 session down for edge [bl-hub1] to peer: [f600:172:211::1]
BFDv6 session established to edge neighbor		Global Segment	bl-hub1	Info	Jul 14, 2021, 3:19:26 PM	BFDv6 session up for edge [bl-hub1] to peer: [f600:172:211::1]
Edge BFDv6 neighbor unavailable		Global Segment	bl-hub1	Info	Jul 14, 2021, 3:18:52 PM	BFDv6 session down for edge [bl-hub1] to peer: [f600:172:211::1]
BFDv6 session established to edge neighbor		Global Segment	bl-hub1	Info	Jul 14, 2021, 9:45:59 AM	BFDv6 session up for edge [bl-hub1] to peer: [f600:172:211::1]
Edge BFDv6 neighbor unavailable		Global Segment	bl-hub1	Info	Jul 14, 2021, 9:44:59 AM	BFDv6 session down for edge [bl-hub1] to peer: [f600:172:211::1]
BFDv6 session established to edge neighbor		Global Segment	bl-hub1	Info	Jul 14, 2021, 9:12:55 AM	BFDv6 session up for edge [bl-hub1] to peer: [f600:172:211::1]
Edge BFDv6 neighbor unavailable		Global Segment	bl-hub1	Info	Jul 14, 2021, 3:00:39 AM	BFDv6 session down for edge [bl-hub1] to peer: [f600:172:211::1]
BFDv6 session established to edge neighbor		Global Segment	bl-hub1	Info	Jul 14, 2021, 1:38:31 AM	BFDv6 session up for edge [bl-hub1] to peer: [f600:172:211::1]
Edge BFDv6 neighbor unavailable		Global Segment	bl-hub1	Info	Jul 14, 2021, 1:36:35 AM	BFDv6 session down for edge [bl-hub1] to peer: [f600:172:211::1]
BFDv6 session established to edge neighbor		Global Segment	bl-hub1	Info	Jul 14, 2021, 1:06:26 AM	BFDv6 session up for edge [bl-hub1] to peer: [f600:172:211::1]
Edge BFDv6 neighbor unavailable		Global Segment	bl-hub1	Info	Jul 14, 2021, 1:05:57 AM	BFDv6 session down for edge [bl-hub1] to peer: [f600:172:211::1]
BFDv6 session established to edge neighbor		Global Segment	bl-hub1	Info	Jul 14, 2021, 12:52:51 AM	BFDv6 session up for edge [bl-hub1] to peer: [f600:172:211::1]
Edge BFDv6 neighbor unavailable		Global Segment	bl-hub1	Info	Jul 13, 2021, 10:39:11 PM	BFDv6 session down for edge [bl-hub1] to peer: [f600:172:211::1]
BFDv6 session established to edge neighbor		Global Segment	bl-hub1	Info	Jul 13, 2021, 10:27:45 PM	BFDv6 session up for edge [bl-hub1] to peer: [f600:172:211::1]

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 Enterprise portal, click **Test & Troubleshoot > 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 that 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. Click **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. Click **Run** to view the details of counters of the BFD peers.
 - **Troubleshoot BFD - Show BFD Setting** – Click **Run** to view the details of BFDv4 settings and status of neighbors.
 - **Troubleshoot BFD6 - Show BFD6 Setting** – Click **Run** to view the details of BFDv6 settings and status of neighbors.

For more information on Remote Diagnostics, see [Performing Remote Diagnostics Tests](#).

Overlay Flow Control

The **Overlay Flow Control** page displays a summarized view of all the routes in your network.

For the 4.3 release, a new NSD bucket has been introduced 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.

In the Enterprise portal, click **Configure > Overlay Flow Control**.

The **Overlay Flow Control** page displays the following details:

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.

- **Edit** – Click 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 more information on **Distributed Cost Calculation**, refer to the **Configure Distributed Cost Calculation** section in the *VMware SD-WAN Operator Guide* available at: <https://docs.vmware.com/en/VMware-SD-WAN/index.html>.

Note To enable the **Distributed Cost Calculation** feature, check with your supporting partner. If you are directly supported by VMware, [contact the support team](#).

Click **Refresh Routes** which makes the Edges and Gateways to 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.

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 more information, see [Configure Subnets](#).

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.

Procedure

- 1 In the Enterprise portal, click **Configure > Overlay Flow Control**.
- 2 In the **Overlay Flow Control** page, click **Edit** and the **Edit Global Configs** window is displayed.

The screenshot shows the 'Edit Global Configs' window with the following sections:

- Preferred VPN Exits:**
 - Eligible:** A list box (currently empty).
 - Default Priority:** A list with checkboxes and priority up/down arrows:
 - ☐ Edge
 - ☐ Router
 - ☐ NSD
 - ☐ Partner ...
 - ☐ Hub
- Global Advertise Flags:**

Edge	Hubs	Partner Gateways	NSD via Edge
Assigned	Assigned	Assigned	Assigned
Static Routes	Static Routes	Static Routes	Static Routes
Connected Routes	Connected Routes	BGP	BGP
OSPF	OSPF	Advertise External & Internal	Advertise Internal
Advertise External	Advertise External		Advertise External
Advertise InterArea	Advertise InterArea		Advertise Uplink Routes
Advertise IntraArea	Advertise IntraArea		
BGP	BGP		NSD via Gateway
Advertise Internal	Advertise Internal		Assigned
Advertise External	Advertise External		Static Routes
Advertise Uplink Routes	Advertise Uplink Routes		BGP
			Advertise Internal
			Advertise External

At the bottom right, there are **Update** and **Close** buttons.

3 You can update the following settings:

- a In the **Preferred VPN Exits** area, click the **UP** and **DOWN** arrows to modify the priorities.
- b In the **Global Advertise Flags** area, select the relevant checkboxes to modify the advertise actions for the routes.
- c Click **Update** to save the changes.

Results

The updated settings are displayed in the **Overlay Flow Control** page.

Configure Subnets

In the **Overlay Flow Control** window, you can update the priorities of the destinations for the learned routes in the subnets.

Procedure

- 1 In the Enterprise portal, click **Configure > Overlay Flow Control**.
- 2 The bottom panel of the **Overlay Flow Control** window displays the subnets with the following details:

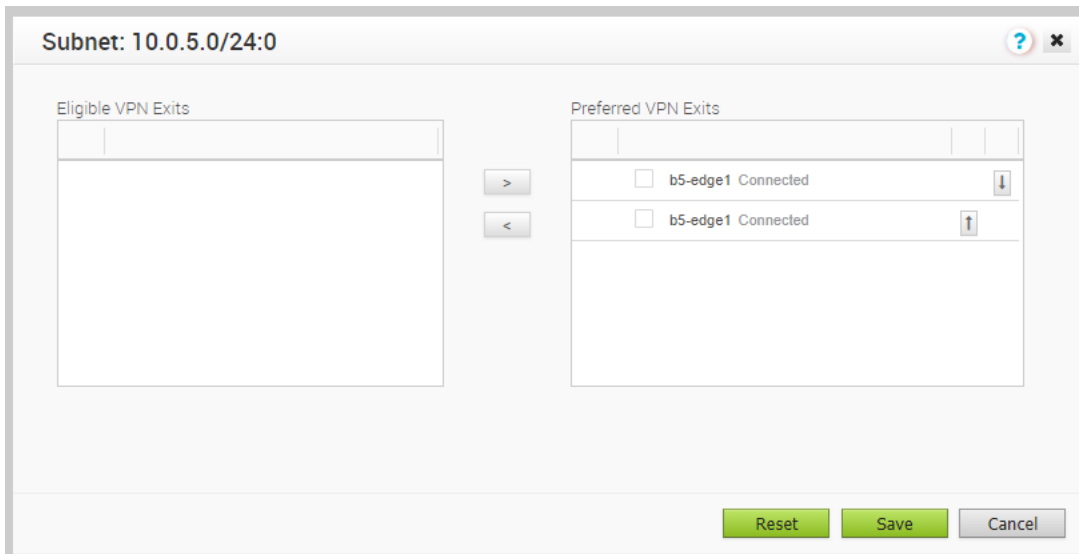
<div> <div>Search</div> <div>Cols</div> <div>Reset View</div> <div>Refresh</div> <div>CSV</div> </div> <div>Display 7 items. 0 selected Actions</div>						
<input type="checkbox"/>	Modify	Segment	Subnet	Preferred VPN Exits	Route Type	Last Update
<input type="checkbox"/>	Edit	Global Segment	10.0.1.0/24	b1-edge1	Connected	
<input type="checkbox"/>	Edit	Global Segment	10.0.2.0/24	b2-edge1	Connected	
<input type="checkbox"/>	Edit	Global Segment	10.0.3.0/24	b3-edge1	Connected	
<input type="checkbox"/>	Edit	Global Segment	10.0.4.0/24	b4-edge1	Connected	
<input type="checkbox"/>	Edit	Global Segment	10.0.5.0/24	b5-edge1	Connected	
<input type="checkbox"/>	Edit	Global Segment	172.16.1.0/29	b1-edge1	Connected	
<input type="checkbox"/>	Edit	Global Segment	172.16.5.0/29	none	Connected (b5-edge1)	

Option	Description
Modify	Displays the option to edit the subnet. The option displays a superscript number, which indicates the number of Edges and Gateways that learned the corresponding route.
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.

Select one or more subnets and click the **Actions** to perform the following activities:

- **Edit Subnet** – Modify the preferred destinations and prioritize them.
- **Pin Learned Route Preference** – Pins the preferences of the selected learned route.
- **Reset Learned Route Preference** – Resets 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 and is used for cleaning up stale routes. The option is available only when **Configure Distributed Cost Calculation** is turned off.

- 3 Click the **Edit** 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.



- b In the **Preferred VPN Exits** panel, click the **UP** and **DOWN** arrows to change the priorities and click **Save**.
 - c You can reset the cost calculation for the subnets when there are pinned routes available. Click **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 The **Reset** option is available only when Distributed Cost Calculation is enabled.

For more information on Distributed Cost Calculation, refer to the **Configure Distributed Cost Calculation** section in the *VMware SD-WAN Operator Guide* available at: <https://docs.vmware.com/en/VMware-SD-WAN/index.html>.

Overlay Flow Control

The **Overlay Flow Control** page displays a summarized view of all the routes in your network.

For the 4.3 release, a new NSD bucket has been introduced 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.

In the **SD-WAN** Service of the Enterprise portal, click **Configure > Overlay Flow Control**.

To configure the Overlay Flow Control settings, perform the following steps:

- 1 In the **SD-WAN** Service of the Enterprise portal, click **Configure > Overlay Flow Control**.

Overlay Flow Control

IPv4 IPv6

Refresh Routes

Force edges to recalculate preferences of learned IPv4 routes and send them to vco. This change may cause service disruption to this enterprise temporarily.

Type "YES" to confirm

YES

REFRESH ROUTES

VRF Global Routing Preferences

> Preferred VPN Exits ⓘ

> Global Advertise Flags ⓘ

Routes List

> Routes List ⓘ

Q Search ⓘ **CSV**

EDIT SUBNET PIN LEARNED ROUTE PREFERENCE RESET LEARNED ROUTE PREFERENCE

<input type="checkbox"/>	IPv4 Subnet	Preferred VPN Exits ⓘ	Route Type ⓘ	Segment	Last Update ⓘ	Created On
<input type="checkbox"/>	10.0.1.0/24	b1-edge1	Connected	Global Segment		
<input type="checkbox"/>	172.16.1.0/29	none	Connected (b1-edge1)	Global Segment		
<input type="checkbox"/>	1.1.0.1/32	b1-edge1	Connected	Global Segment		
<input type="checkbox"/>	1.1.0.2/32	b1-edge1	Connected	Global Segment		
<input type="checkbox"/>	10.0.2.0/24	b2-edge1	Connected	Global Segment		
<input type="checkbox"/>	1.2.0.1/32	b2-edge1	Connected	Global Segment		
<input type="checkbox"/>	10.0.3.0/24	b3-edge1	Connected	Global Segment		
<input type="checkbox"/>	1.3.0.1/32	b3-edge1	Connected	Global Segment		
<input type="checkbox"/>	10.0.4.0/24	b4-edge1	Connected	Global Segment		
<input type="checkbox"/>	1.4.0.1/32	b4-edge1	Connected	Global Segment		
<input type="checkbox"/>	10.0.5.0/24	b5-edge1	Connected	Global Segment		
<input type="checkbox"/>	172.16.5.0/29	none	Connected (b5-edge1)	Global Segment		
<input type="checkbox"/>	1.5.0.1/32	b5-edge1	Connected	Global Segment		

The Overlay Flow Control page displays the following details:

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 clicking Edge Subnet in the Overlay Flow Control page.

2 In the **Overlay Flow Control** page, you can configure the following settings:

- **Edit** – Click 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 more information on **Distributed Cost Calculation**, refer to the **Configure Distributed Cost Calculation** section in the *VMware SD-WAN Operator Guide* available at: <https://docs.vmware.com/en/VMware-SD-WAN/index.html>.

Note To enable the **Distributed Cost Calculation** feature, check with your supporting partner. If you are directly supported by VMware, [contact the support team](#).

- Type **YES** and then click **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

VRF Global Routing Preferences

Preferred VPN Exits (0)

Default Priority

[EDIT](#)

Order	Header
1.	NSD
2.	Edge
3.	Partner Gateway
4.	Router
5.	Hub

Global Advertise Flags (0)

Edge Assigned	Hubs Assigned	Partner Gateways Assigned	NSD via Edge Assigned
<input checked="" type="checkbox"/> Connected Routes	<input checked="" type="checkbox"/> Connected Routes	<input checked="" type="checkbox"/> Static Routes	<input checked="" type="checkbox"/> Static Routes
<input checked="" type="checkbox"/> Static Routes	<input checked="" type="checkbox"/> Static Routes	<input checked="" type="checkbox"/> BGP	<input checked="" type="checkbox"/> BGP
<input checked="" type="checkbox"/> BGP	<input checked="" type="checkbox"/> BGP	<input checked="" type="checkbox"/> Advertise External & Internal	<input checked="" type="checkbox"/> Advertise External
<input checked="" type="checkbox"/> Advertise External	<input checked="" type="checkbox"/> Advertise External		<input type="checkbox"/> Advertise Internal
<input checked="" type="checkbox"/> Advertise Internal	<input checked="" type="checkbox"/> Advertise Internal		<input type="checkbox"/> Advertise Uplink Routes
<input type="checkbox"/> Advertise Uplink Routes	<input checked="" type="checkbox"/> Advertise Uplink Routes		NSD via Gateway Assigned
<input type="checkbox"/> OSPF	<input type="checkbox"/> OSPF		<input checked="" type="checkbox"/> Static Routes
<input type="checkbox"/> Advertise External	<input type="checkbox"/> Advertise External		<input checked="" type="checkbox"/> BGP
<input checked="" type="checkbox"/> Advertise InterArea	<input checked="" type="checkbox"/> Advertise InterArea		<input checked="" type="checkbox"/> Advertise External
<input checked="" type="checkbox"/> Advertise IntraArea	<input checked="" type="checkbox"/> Advertise IntraArea		<input type="checkbox"/> Advertise Internal

routed.

- Click **Preferred VPN Exits** to prioritize the VPN Exits.
- Click **Edit** and use the **UP** and **DOWN** arrows to modify the priorities.

Edit Preferred VPN

Eligible

<input type="checkbox"/> Eligible VPN Exits
<input type="checkbox"/> NSD

1 item

Preferred

<input type="checkbox"/> Preferred VPN Exits
<input type="checkbox"/> Edge
<input type="checkbox"/> Partner Gateway
<input type="checkbox"/> Router
<input type="checkbox"/> Hub

4 items

[CANCEL](#) [UPDATE](#)

- 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 click the IPv4 or IPv6 tab to view the corresponding subnets. The following image shows IPv6 subnets. For more information, see [Configure Subnets](#).

The screenshot shows the 'Overlay Flow Control' configuration window with the 'IPv6' tab selected. The 'Routes List' section is expanded, displaying a table of learned routes. The table has columns for 'IPv6 Subnet', 'Preferred VPN Exits', 'Route Type', 'Segment', 'Last Update', and 'Created On'. There are 12 rows of data, each representing a different IPv6 subnet and its associated route information.

IPv6 Subnet	Preferred VPN Exits	Route Type	Segment	Last Update	Created On
fd00:0001:0001:0000:0000:0000:0000:0000/64	b1-edge1	Connected	Global Segment		
fd00:ffff:ffff:0001:0000:0000:0000:0001/128	none	Connected (b1-edge1)	Global Segment		
fd00:ffff:ffff:0002:0000:0000:0000:0001/128	none	Connected (b1-edge1)	Global Segment		
fd00:0002:0001:0000:0000:0000:0000:0000/64	b2-edge1	Connected	Global Segment		
fd00:ffff:ffff:0003:0000:0000:0000:0001/128	none	Connected (b2-edge1)	Global Segment		
fd00:0003:0001:0000:0000:0000:0000:0000/64	b3-edge1	Connected	Global Segment		
fd00:ffff:ffff:0004:0000:0000:0000:0001/128	none	Connected (b3-edge1)	Global Segment		
fd00:0004:0001:0000:0000:0000:0000:0000/64	b4-edge1	Connected	Global Segment		
fd00:ffff:ffff:0005:0000:0000:0000:0001/128	none	Connected (b4-edge1)	Global Segment		
fd00:ffff:ffff:0006:0000:0000:0000:0001/128	none	Connected (b5-edge1)	Global Segment		

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 more information, see [Configure Subnets](#).

Configure Alerts

24

SD-WAN Orchestrator allows 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 Enterprise portal, click **Configure > Alerts & Notifications** to configure the alerts.

Select the events for which the alerts need to be sent, and enter the notification delay time in minutes under **Select Alerts**.

You can use the `EDIT_ALERT_CONFIGURATION` event to record the changes to the enterprise alert configurations.

5-site Recently Viewed Superuser Operator Help super@velocloud.net

Monitor

Configure

- Edges
- Profiles
- Object Groups
- Segments
- Overlay Flow Control
- Network Services
- Alerts & Notifications**
- Customer
- Test & Troubleshoot
- Administration

Alert Configuration Save Changes ?

Select Alerts

Select Alerts	Alert Type	Notification Delay
<input checked="" type="checkbox"/>	Edge Down ⓘ	3 minutes
<input checked="" type="checkbox"/>	Edge Up ⓘ	1 minutes
<input checked="" type="checkbox"/>	Link Down ⓘ	3 minutes
<input checked="" type="checkbox"/>	Link Up ⓘ	1 minutes
<input type="checkbox"/>	VPN Tunnel Down ⓘ	3 minutes
<input type="checkbox"/>	Edge HA Failover ⓘ	1 minutes
<input type="checkbox"/>	Edge VNF Virtual Machine Deployment ⓘ	0 minutes
<input type="checkbox"/>	Edge VNF Insertion ⓘ	0 minutes
<input type="checkbox"/>	Edge CSS tunnel up ⓘ	3 minutes
<input type="checkbox"/>	Edge CSS tunnel down ⓘ	3 minutes
<input type="checkbox"/>	Edge VNF Image Download Event ⓘ	0 minutes

Customers

Admin	User Role	Email ⓘ	SMS ⓘ
5_site_operator@velocloud.net	Superuser	<input checked="" type="checkbox"/> 5_site_operator@velocloud.net	<input type="checkbox"/> (not set) Test

Email Addresses

Add a comma separated list of emails

Phone Numbers

Name	Phone
<input type="text"/>	<input type="text"/>

SNMP Traps

Version	Hostname / IP Address	Port	Version Specific Attributes
<input checked="" type="checkbox"/> v2c	<input type="text"/> 10.20.1.1	<input type="text"/> 162	Community: <input type="text"/> public Test

Webhooks

URL	Code ⓘ	Secret	JSON Payload Template ⓘ
<input checked="" type="checkbox"/> <input type="text"/> https://www.velocloud.net	<input type="text"/> 200	<input type="text"/>	<pre>{ "alertTime": "{{alertTime}}", "alertType": "{{alertType}}", "customer": "{{customer}}", ... }</pre> Test

Under **Customers**, the contact details of existing admin users are displayed. You can select the checkboxes for Email and SMS to send alerts to the corresponding users.

The alerts are sent to both the operators team managing the entire SD-WAN Orchestrator and to the customers.

Alerts that go to operators are called Pre-Notification Alerts as they are sent immediately. Customer or Enterprise alerts are subject to delays as configured by the Enterprise Admin.

For example, a Link Down alert may go to both an operator configured destination and to customer configured destinations. Assume that a link is down for a minute and the customer configures the **Link Down** Alert delay for 2 minutes. If the Pre-Notification Alerts are enabled for this link, the Orchestrator will send an Operator alert for Link Down, but the customer would not get an alert as it fell under the configured delay.

SNMP Traps

Simple Network Management Protocol (SNMP) Traps are notifications sent to an SNMP Agent to indicate that an event has occurred. SD-WAN Orchestrator sends SNMP Traps corresponding to the existing alerts like Edge Down and Edge Up. You can select the SNMP version and enter the corresponding details under **SNMP Traps**.

Note Currently, only SHA-1 and AES-128 algorithms are supported for SNMP v3 Trap.

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.

SD-WAN 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 SD-WAN 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.

Configure Webhooks

In the **Alert Configuration** window, you can enter the following details under **Webhooks**.

Option	Description
URL	Enter a valid HTTPS URL. This serves as the target application for the webhooks.
Code	<p>Enter an expected HTTP response status code for each webhook recipient. By default, the SD-WAN Orchestrator expects webhook recipients to respond to HTTP POST requests with a status code as HTTP 200.</p> <p>When SD-WAN 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 <code>ALERT_DELIVERY_FAILED</code> customer event. This event helps to identify when a webhook recipient server may fail to function as expected.</p>

Option	Description
Secret	<p>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 <code>X-Webhook-Signature</code> HTTP header, along with a version parameter, which identifies the signature algorithm and a timestamp.</p> <pre data-bbox="826 447 1265 499">X-Webhook-Signature: v=<signature-version>&t=<timestamp>&s=<hmac></pre> <p>The recipient interprets the components as follows:</p> <ul style="list-style-type: none"> ■ v: Version of the algorithm used to produce the signature. The only supported value is 1. ■ t: Millisecond-precision epoch timestamp corresponding to the time at which the request is issued. ■ s: HMAC computed by SD-WAN Orchestrator. The HMAC is computed as follows: <code>HMAC-SHA256(request-body + '.' + timestamp, secret)</code>. <p>The message used to compute the HMAC is formed by concatenating the request body, a single period, and the value of the timestamp parameter that appears in the signature header. The specific HMAC algorithm used to produce the code is HMAC-SHA256.</p> <p>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 SD-WAN Orchestrator.</p>
JSON Payload Template	<p>This is a required field.</p> <p>SD-WAN Orchestrator delivers alert notifications to each webhook recipient, through a JSON payload contained within the body of an outgoing HTTP POST request. SD-WAN 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.</p> <p>Webhook payload templates support the following placeholder variables:</p> <ul style="list-style-type: none"> ■ alertTime: Time at which the alert got triggered. ■ alertType: The type of the alert, like <code>EDGE_DOWN</code>, <code>LINK_UP</code>, <code>VNF_VM_DEPLOYED</code>. ■ customer: Name of the customer to whom the notification is sent. ■ customerLogicalId: The logical ID of the customer to whom the notification is sent. ■ deviceLogicalId: The logical ID of the Edge to which the alert is applied.

Option	Description
	<ul style="list-style-type: none"> ■ entityAffected: Name of the entity, like Edge or link or VNF, to which the alert is applied. ■ lastContact: The time at which the affected Edge most recently communicated with the SD-WAN Orchestrator. This is applicable only for the Edge alerts. ■ message: A brief message describing the event that triggered the alert. ■ VCO: Hostname or public IP of the SD-WAN Orchestrator from which the notification is sent.

The following example shows a sample JSON payload template:

```
{
  "alertTime": "{{alertTime}}",
  "alertType": "{{alertType}}",
  "customer": "{{customer}}",
  "customerLogicalId": "{{customerLogicalId}}",
  "deviceLogicalId": "{{deviceLogicalId}}",
  "entityAffected": "{{entityAffected}}",
  "lastContact": "{{lastContact}}",
  "message": "{{message}}",
  "vco": "{{vco}}"
}
```

You can click the plus (+) icon to add more target URLs and the corresponding details.

Click **Test** to check the Webhook alerts.

Whenever an alert is triggered, an alert message along with relevant information is sent to the target URL.

Testing and Troubleshooting

25

The SD-WAN Orchestrator Test & Troubleshoot functionality provides tools to test the status of the VMware services, perform remote Edge actions, and gather debugging information for an Edge.

In the Enterprise portal, click **Test & Troubleshoot** to access and perform the testing and troubleshooting options.

Read the following topics next:

- [Remote Diagnostics](#)
- [Run Remote Diagnostics with new Orchestrator UI](#)
- [Performing Remote Diagnostics Tests](#)
- [Remote Actions](#)
- [Perform Remote Actions with new Orchestrator UI](#)
- [Diagnostic Bundles](#)
- [Diagnostic Bundles with new Orchestrator UI](#)

Remote Diagnostics

VMware SD-WAN supports bi-directional communication with the VMware SD-WAN 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 VMware SD-WAN Orchestrator.

The WebSocket communication involves the following two WebSocket connections for passing WebSocket messages from a Web browser to a VMware SD-WAN 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:

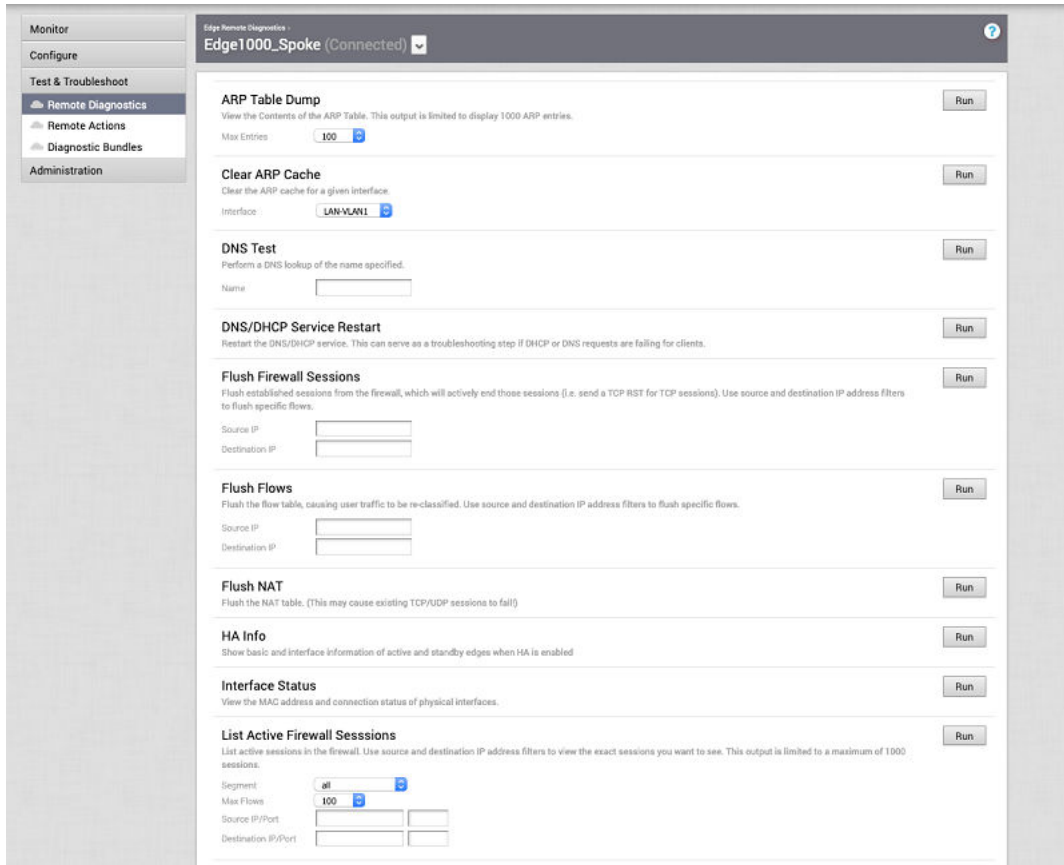
- `network.portal.websocket.address` - Allows to set an alternate address/DNS hostname to access the UI from a browser if the browser address is not the same as the value of `network.public.address` system property. By default, the `network.portal.websocket.address` system property is not set.
- `session.options.websocket.portal.idle.timeout` - 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.

To run Remote Diagnostics tests on an Edge, perform the following steps.

Procedure

- 1 In the Enterprise portal, click **Test & Troubleshoot** and click **Remote Diagnostics**. The **Remote Diagnostics** page displays all the active Edges.
- 2 Search for an Edge that you want to troubleshoot by using the **Filter** option, and click **Apply**.
- 3 Select an Edge to troubleshoot.

The Edge enters live mode and displays all the possible Remote Diagnostics tests than you can run on the Edge.



- 4 Choose an appropriate Remote Diagnostics test to run on the Edge and click **Run**. The diagnostic information is fetched from the Edge and displayed in the **Edge Remote Diagnostics** screen.

For more information about all the supported Remote Diagnostics tests, see [Performing Remote Diagnostics Tests](#).


Run Remote Diagnostics with new Orchestrator UI

VMware SD-WAN Orchestrator enables you to run various Remote Diagnostic tests on a selected Edge.

For more information on Remote Diagnostics, see [Remote Diagnostics](#).

To run Remote Diagnostics on an Edge:

- 1 In the Enterprise portal, click the **Open New Orchestrator UI** option available at the top of the Window.
- 2 Click **Launch New Orchestrator UI** in the pop-up window.
- 3 The UI opens in a new tab displaying the monitoring and configuring options.
- 4 In the new UI, click the **Diagnostics** tab.
- 5 The **Remote Diagnostics** page displays the existing Edges.



Name	Status	Model	Software version	Build Number	Created	Software updated
s1-edge1	Connected	Virtual Edge	4.3.0	R430-2020425-MH-e3962b7cf9	Apr 26, 2021, 12:58:47 AM	
s2-edge1	Connected	Virtual Edge	4.3.0	R430-2020425-MH-e3962b7cf9	Apr 26, 2021, 12:58:52 AM	
s3-edge1	Connected	Virtual Edge	4.3.0	R430-2020425-MH-e3962b7cf9	Apr 26, 2021, 12:58:55 AM	
s4-edge1	Connected	Virtual Edge	4.3.0	R430-2020425-MH-e3962b7cf9	Apr 26, 2021, 12:58:59 AM	
s5-edge1	Connected	Virtual Edge	4.3.0	R430-2020425-MH-e3962b7cf9	Apr 26, 2021, 12:59:02 AM	

- 6 Click the link to an Edge.
- 7 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.
- 8 Choose an appropriate Remote Diagnostics test to run on the Edge and click **Run**. The diagnostic information is fetched from the Edge and displayed in the screen.

For more information about all the supported Remote Diagnostics tests, see [Performing Remote Diagnostics Tests](#).

You can also run the Remote Diagnostics test using the **Shortcuts** option available in the **Configure > Edges** or **Monitor > Edges** pages.

For more information, see:

- [Configure Edges with new Orchestrator UI](#)
- [Monitor Edges](#)

Performing Remote Diagnostics Tests

Describes all the possible remote diagnostics tests that you can run on an Edge to obtain diagnostic information. The diagnostic information contains Edge-specific logs for analysis.

VMware SD-WAN Orchestrator allows you to run various remote diagnostics test on a selected Edge from the **Test & Troubleshoot > Remote Diagnostics** menu.

The following are the supported remote diagnostics tests:

- [ARP Table Dump](#)
- [Clear ARP Cache](#)
- [DNS Test](#)
- [DNS/DHCP Service Restart](#)
- [List item.](#)
- [Flush Firewall Sessions](#)
- [Flush Flows](#)
- [Flush NAT](#)
- [Gateway](#)
- [GPON Status](#)

- HA Info
- IPv6 Clear ND Cache
- IPv6 ND Table Dump
- IPv6 RA Table Dump
- IPv6 Route Table Dump
- Interface Status
- LTE Modem Information
- List item.
- List Active Firewall Sessions
- List Active Flows
- List Clients
- List Paths
- MIBs for Edge
- NAT Table Dump
- NTP Dump
- Ping IPv6 Test
- Ping Test
- Reset USB Modem
- Route Table Dump
- Source Interface Dump
- System Information
- Traceroute
- Troubleshoot BFD - Show BFD/BFDv6 Peer Status
- Troubleshoot BFD - Show BFD/BFDv6 Peer counters
- Troubleshoot BFD - Show BFD Setting
- Troubleshoot BFDv6 - Show BFDv6 Setting
- Multi-hop BGP Routes
- Troubleshoot BGP - List BGP Redistributed Routes
- Troubleshoot BGP - List BGP Routes
- Troubleshoot BGP - List Routes per Prefix
- Troubleshoot BGP - Show BGP Neighbor Advertised Routes
- Troubleshoot BGP - Show BGP Neighbor Learned Routes

- [Troubleshoot BGP - Show BGP Neighbor Received Routes](#)
- [Troubleshoot BGP - Show BGP Neighbor details](#)
- [Troubleshoot BGP - Show BGP Routes per Prefix](#)
- [Troubleshoot BGP - Show BGP Summary](#)
- [Troubleshoot BGP - Show BGP Table](#)
- [Troubleshoot BGPv6 - Show BGPv6 Neighbor Advertised Routes](#)
- [Troubleshoot BGPv6 - Show BGPv6 Neighbor Learned Routes](#)
- [Troubleshoot BGPv6 - Show BGPv6 Neighbor Received Routes](#)
- [Troubleshoot BGPv6 - Show BGPv6 Neighbor details](#)
- [Troubleshoot BGPv6 - Show BGPv6 Routes per Prefix](#)
- [Troubleshoot BGPv6 - Show BGPv6 Summary](#)
- [Troubleshoot BGPv6 - Show BGPv6 Table](#)
- [Troubleshoot OSPF - List OSPF Redistributed Routes](#)
- [Troubleshoot OSPF - List OSPF Routes](#)
- [Troubleshoot OSPF - Show OSPF Database](#)
- [Troubleshoot OSPF - Show OSPF Database for E1 Self-Originate Routes](#)
- [Troubleshoot OSPF - Show OSPF Neighbors](#)
- [Troubleshoot OSPF - Show OSPF Route Table](#)
- [Troubleshoot OSPF - Show OSPF Setting](#)
- [USB Port Status](#)
- [VPN Test](#)
- [WAN Link Bandwidth Test](#)

ARP Table Dump

Run this test to view the contents of the ARP table. The output is limited to display 1000 ARP entries.

ARP Table Dump

View the Contents of the ARP Table. This output is limited to display 1000 ARP entries.

Max Entries

100 ▼

Run

Test Duration: 1.002 seconds

Stale Timeout: 2min Dead Timeout: 25min Cleanup Timeout: 240min			
LAN-VLAN1			
10.0.1.25	00:ba:be:71:0d:7b	ALIVE	6s
LAN-VLAN100			
10.100.1.100	00:ba:be:71:0d:7b	ALIVE	6s
LAN-VLAN101			
10.101.1.100	00:ba:be:71:0d:7b	ALIVE	5s
GE3			
169.254.7.9	00:ba:be:16:40:2c	ALIVE	1s
169.254.7.12	00:ba:be:29:43:07	REFRESH	212s
GE4			
169.254.6.33	00:ba:be:39:a6:86	ALIVE	1s
GE5			
172.17.1.3	00:ba:be:0a:aa:e9	ALIVE	1s
172.18.1.3	00:ba:be:0a:aa:e9	ALIVE	1s
172.16.1.3	00:ba:be:0a:aa:e9	ALIVE	1s

Clear ARP Cache

Run this test to clear the ARP cache entries for the specified interface.

Clear ARP Cache

Clear the ARP cache for a given interface.

Interface

GE1 ▼

Run

Test Duration: 0.982 seconds

The ARP cache has been cleared for the selected interface.

DNS Test

Run this test to perform a DNS lookup of the specified domain name.

DNS Test

Perform a DNS lookup of the name specified.

Name

google.com

Run

Test Duration: 1.002 seconds

google.com
172.217.14.206

DNS/DHCP Service Restart

Run this test to restart the DNS/DHCP service. This can serve as a troubleshooting step if DHCP or DNS requests are failing for clients.

DNS/DHCP Service Restart

Restart the DNS/DHCP service. This can serve as a troubleshooting step if DHCP or DNS requests are failing for clients.

Test Duration: 1.001 seconds

DNS/DHCP service has been restarted.

DSL Status

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. Run this test to show the DSL status, which includes information such as Mode (Standard or DSL), Profile, xDSL Mode, etc. as shown in the image below.

DSL Status

View the xDSL(ADSL2/VDL2) modem status connected to SFP interfaces

Test Duration: 10.003 seconds

Interfaces

Name	Mode	Vendor MAC	xDSL Mode	Link Time	Status	Link Rate	Annex
SFP1	DSL	00:0E:AD:00:55:FE	VDSL2	0	Idle	0/0	N/A
SFP2	DSL	00:0E:AD:00:55:AC	VDSL2	49223	Showtime	12045/23407	AnnexA

Flush Firewall Sessions

Run this test to reset established sessions from the firewall. Running this test on an Edge not only flushes the firewall sessions, but actively send a TCP RST for the TCP-based sessions.

Flush Firewall Sessions

Flush established sessions from the firewall, which will actively end those sessions (i.e. send a TCP RST for TCP sessions). Use source and destination IP address filters to flush specific flows.

Source IP

Destination IP

Test Duration: 2.002 seconds

12 active firewall sessions have been flushed from the system.

Flush Flows

Run this test to flush the flow table, causing user traffic to be re-classified. Use source and destination IP address filters to flush specific flows.

Flush Flows

Flush the flow table, causing user traffic to be re-classified. Use source and destination IP address filters to flush specific flows.

Source IP

Destination IP

Run

Test Duration: 1.001 seconds

26 flows have been flushed from the system.

Flush NAT

Run this test to flush the NAT table.

Flush NAT

Flush the NAT table. (This may cause existing TCP/UDP sessions to fail!)

Run

Test Duration: 1.001 seconds

All NAT entries have been flushed from the system.

Gateway

Run this test by choosing whether cloud traffic should or should not use the Gateway Service.

Note This does not affect the routing of VPN traffic.

Gateway

Choose whether cloud traffic should or should not use the Gateway Service. Note: This does not affect the routing of VPN traffic.

Cloud Traffic Routing

Always use Gateway Service ▼

Run

Test Duration: 1.001 seconds

Cloud traffic will all be sent to the VeloCloud Gateway Service. This is intended for debugging and will not persist across restart/reboot!

GPON Status

Run this test on any selected 6x0 Edge device to view the GPON SFP status, including Vendor MAC, Host Link Status, Link Rate, TX and RX power, and Optical Status.

GPON Status

View the GPON sfp status

Run

Test Duration: 4.004 seconds

GPON Interfaces

Name	Vendor MAC	Host Link Status	Link Rate	TX power	RX power	Optical Status
SFP1	18:5a:58:16:31:63	up	1000Mb/s	17128	233	No LOS ,No TX Fault

HA Info

Run this test to view basic and interface information of active and standby Edges when HA is enabled.

HA Info

Show basic and interface information of active and standby edges when HA is enabled

Run

Test Duration: 8.024 seconds

Active and Standby Edge Info

Edge Type	Edge Serial Number	No. of LANs	No. of WANs
ACTIVE	JFS8PK2	1	1
STANDBY	4FS8PK2	1	1

Active and Standby Interfaces

Logical Name	Interface IP	Nexthop IP	Interface State (Active)	Interface State (Standby)
SFP1	0.0.0.0	0.0.0.0	LOCAL_DOWN	LOCAL_DOWN
GE5	192.168.0.197	192.168.0.1	LOCAL_UP	LOCAL_DOWN
SFP2	0.0.0.0	0.0.0.0	LOCAL_DOWN	LOCAL_DOWN
GE3	0.0.0.0	0.0.0.0	LOCAL_DOWN	LOCAL_DOWN
GE4	0.0.0.0	0.0.0.0	LOCAL_DOWN	LOCAL_DOWN
GE6	0.0.0.0	0.0.0.0	LOCAL_DOWN	LOCAL_DOWN
CELL1	100.235.130.246	100.235.130.245	USE_PEER	USED_BY_PEER

IPv6 Clear ND Cache

Run this test to clear the cache from the ND for the selected interface.

IPv6 Clear ND Cache

Clear the IPv6 ND cache for a given interface.

Interface

GE3 ▼

Run

Test Duration: 1.001 seconds

The ND cache has been cleared for the selected interface.

IPv6 ND Table Dump

Run this test to view the IPv6 address details of Neighbor Discovery (ND) table.

IPv6 ND Table Dump

View the Contents of the IPv6 ND Table. This output is limited to display 1000 ND entries.

Max Entries

100 ▼

Run

Test Duration: 1.001 seconds

IPv6 ND Cache

GE5

```
fd00:1:2:3::1 00:50:56:81:9d:03 STALE
fd00:1:3:3::1 00:50:56:81:9d:03 STALE
fd00:1:1:3::1 00:50:56:81:9d:03 STALE
```

GE6

```
fd00:1:3:4::1 00:50:56:81:25:64 STALE
fd00:1:1:4::1 00:50:56:81:25:64 STALE
fd00:1:2:4::1 00:50:56:81:25:64 STALE
```

IPv6 RA Table Dump

Run this test to view the details of the IPv6 RA table.

IPv6 RA Table Dump

View the Contents of the IPv6 RA Table

Run

Test Duration: 1.001 seconds

No RA results found

IPv6 Route Table Dump

Run this test to view the contents of the IPv6 Route Table.

IPv6 Route Table Dump
View the contents of the IPv6 Route Table.

Segment:

Test Duration: 2.007 seconds

Segmented Route Table

Address	Segment	Netmask	Type
fe80::7f02:9e35:b20d:fa8	Global Segment	:::::ffff:ffff:ffff:ffff	N/A
fe80::c411:cbb7:c3e8:d0ea	Global Segment	:::::ffff:ffff:ffff:ffff	N/A
fe80::b928:b9f1:83ca:fa1	Global Segment	:::::ffff:ffff:ffff:ffff	N/A
fe80::9b88:40f5:e408:742c	Global Segment	:::::ffff:ffff:ffff:ffff	N/A
fe80::405:9424:fd19:df15	Global Segment	:::::ffff:ffff:ffff:ffff	N/A
fe80::2d9e:c39f:3036:7fb	Global Segment	:::::ffff:ffff:ffff:ffff	N/A
fd00:172:21:7::1	Global Segment	:::::ffff:ffff:ffff:ffff	N/A
fd00:172:21:4::2	Global Segment	:::::ffff:ffff:ffff:ffff	N/A
fd00:172:21:1::2	Global Segment	:::::ffff:ffff:ffff:ffff	N/A
2d00:1:1:15::51	Global Segment	:::::ffff:ffff:ffff:ffff	N/A
2d00:1:1:14::51	Global Segment	:::::ffff:ffff:ffff:ffff	N/A
fe80::	Global Segment	:::::ffff:ffff:ffff:ffff	Connected
fe80::	Global Segment	:::::ffff:ffff:ffff:ffff	Connected
fe80::	Global Segment	:::::ffff:ffff:ffff:ffff	Connected
fe80::	Global Segment	:::::ffff:ffff:ffff:ffff	Connected
fd00:192:168:100::	Global Segment	:::::ffff:ffff:ffff:ffff	BGP
fd00:192:168:46::	Global Segment	:::::ffff:ffff:ffff:ffff	BGP
fd00:172:24:1::	Global Segment	:::::ffff:ffff:ffff:ffff	BGP
fd00:172:21:22::	Global Segment	:::::ffff:ffff:ffff:ffff	BGP
fd00:172:21:21::	Global Segment	:::::ffff:ffff:ffff:ffff	BGP
fd00:172:21:12::	Global Segment	:::::ffff:ffff:ffff:ffff	BGP
fd00:172:21:11::	Global Segment	:::::ffff:ffff:ffff:ffff	BGP
fd00:172:21:7::	Global Segment	:::::ffff:ffff:ffff:ffff	Connected
fd00:172:21:4::	Global Segment	:::::ffff:ffff:ffff:ffff	Connected
fd00:172:21:2::	Global Segment	:::::ffff:ffff:ffff:ffff	BGP
fd00:172:21:1::	Global Segment	:::::ffff:ffff:ffff:ffff	Connected

Interface Status

Run this test to view the MAC address and connection status of physical interfaces.

Interface Status

View the MAC address and connection status of physical interfaces.

Run

Test Duration: 2.002 seconds

Routed Interfaces

Name	MAC Address	Link Detected	IP Address	Netmask	Speed	Autonegotiation	RX errors	T
GE3	F0:8E:DB:6F:8E:82	true	169.254.7.10	255.255.255.248	10000 Mbps, full duplex	off	0	0
GE4	F0:8E:DB:6F:8E:83	true	169.254.6.34	255.255.255.248	10000 Mbps, full duplex	off	0	0
GE5	F0:8E:DB:6F:8E:84	true	172.16.1.2	255.255.255.248	10000 Mbps, full duplex	off	0	0
GE6	F0:8E:DB:6F:8E:85	true	172.16.1.10	255.255.255.248	10000 Mbps, full duplex	off	0	0
GE7		false	N/A	N/A	N/A	N/A	-1	-1
GE8		false	N/A	N/A	N/A	N/A	-1	-1

Modem Interfaces

Name	Link Detected	IP Address	Netmask	Signal Quality	Operator Name	RX errors	TX errors	Collisi
------	---------------	------------	---------	----------------	---------------	-----------	-----------	---------

Switch Ports

Name	MAC Address	Link Detected	Speed	RX errors	TX errors	Collisions
GE1	00:BA:BE:13:E0:02	true	10000 Mbps, full duplex	0	0	0
GE2	F0:8E:DB:6F:8E:01	true	10000 Mbps, full duplex	0	0	0

LTE Modem Information

Run this test on a selected Edge that has an integrated LTE module, such as 510-LTE or 610-LTE, to collect diagnostic details such as Modem information, Connection information, Location information, Signal information, and Status information for the internal LTE modem.

LTE Modem Information

This will fetch diagnostic information for the internal LTE modem.

Run

Interface

CELL1 ▾

Test Duration: 6.006 seconds

LTE CELL1

Modem Information

```
{
  "Manufacturer": "Sierra Wireless, Incorporated",
  "Model": "EM7511",
  "Modem identifier": "353587100789907",
  "Firmware Revision": "swI9X50C_01.07.02.00 6c91bc jenkins 2018/06/13 23 08 16",
  "Hardware Revision": "10001",
  "Supported capabilities": "gsm-umts, lte",
  "Current capabilities": "gsm-umts, lte",
  "own number": "NA",
  "state": "connected",
  "Failed reason": "--",
  "Power state": "on",
  "current modes": "allowed 2g, 3g, 4g; preferred 4g",
  "imei": "353587100789907",
  "operator code": "310260",
  "operator name": "T-Mobile",
  "registration state": "home",
  "signal quality(%)": "52"
}
```

Connection Information

```
{
  "Bearer": "Available",
  "Connected": "yes",
  "Suspended": "no",
  "Interface": "wwan0",
  "APN": "",
  "IP type": "--",
  "user": "--",
  "password": "NA",
  "IP method": "static",
  "IP address": "100.232.152.201",
  "Gateway": "100.232.152.202",
  "DNS": "10.177.0.34",
  "MTU": "1430",
  "Stats Duration": "24359",
  "Rx bytes": "106396",
  "Tx bytes": "59484"
}
```

Location Information

```
{
  "operator code": "310",
  "operator name": "260",
  "Location area code": "FFFF",
  "tracking area code": "3A69",
  "cell id": "02CB0705"
}
```

Signal Information

```
{
  "Serving": {
    "EARFCN": "5035",
    "MCC": "310",
    "MNC": "260",
    "TAC": "14953",
    "CID": "02CB0705",
    "Bd": "12",
    "D": "2",
    "U": "2",
    "SNR": "4",
    "PCI": "334",
    "RSRQ": "-11.8",
    "RSRP": "-107.4",
    "RSSI": "-81.6",
    "RXLV": "16"
  },
  "IntraFreq": {
    "PCI": "334",
    "RSRQ": "-11.8",
    "RSRP": "-107.4",
    "RSSI": "-81.6",
    "RXLV": "16"
  }
}
```

Status Information

```
response: 'lgSTATUS:
Current Time: 24389
Reset Counter: 1
System mode: LTE
LTE band: B12
LTE Rx chan: 5035
LTE SSC1 state: NOT ASSIGNED
LTE SSC2 state: NOT ASSIGNED
LTE SSC3 state: NOT ASSIGNED
LTE SSC4 state: NOT ASSIGNED
EMM state: Registered
RRC state: RRC Idle
IMS reg state: No Srv

Temperature: 51
Mode: ONLINE
PS state: Attached
LTE bw: 5 MHz
LTE Tx chan: 23035

Normal Service

PCC RxM RSSI: -80
PCC RxM RSRP: -106
PCC RxM RSSI: -81
PCC RxM RSRP: -109
Tx Power: --
TAC: 3a69 (14953)
RSRQ (dB): -12.1
Cell ID: 02cb0705 (46860037)
SINR (dB): 4.2'
```

Debug Information

```
{
  "STATUS": "ERROR",
  "REASON": "Debug data not available"
}
```

LTE SIM Switchover

For 610-LTE devices only, run this test to switch active SIMs. Both SIMs must be inserted to run this test. The test will take approximately four to five minutes.

LTE Switch SIM Slot

Switch Active SIM. Note: Both SIM should be inserted and works only in 610 LTE

Run

SIM1: CELL1 status(active) slot(Inserted) Isp(Verizon) simIsp(Verizon)

SIM2: CELL2 status(inactive) slot(NotInserted) Isp(Unknown) simIsp(Unknown)

After the test is successful, you can check the status of the current active interface in the SD-WAN Orchestrator under the **Monitor -> Edges -> Overview** tab.

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 limit the number of sessions returned by using filters: source and destination IP address, source and destination port, and Segment.

List Active Firewall Sessions

List active sessions in the firewall. Use source and destination IP address filters to view the exact sessions you want to see. This output is limited to a maximum of 1000 sessions.

Run

Segment:

Max Flows:

Source IP/Port:

Destination IP/Port:

Test Duration: 5.002 seconds

Segment	Src IP	Dst IP	Protocol	Src Port	Dst Port	Application	Firewall Policy	TCP State	Bytes Sent	Bytes
Global Segment	10.5.1.25	10.2.1.25	ICMP	N/A	N/A	icmp	AllowAny	N/A	672	672
Global Segment	10.5.1.25	10.5.1.25	TCP	36720	22	ssh	AllowAny	ESTABLISHED	3441	4153

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.

The Remote Diagnostics output displays the following information: Segment name, Source IP, Source Port, Destination IP, Destination Port, Protocol, Application, Firewall Policy, current TCP state of any flows, Bytes Received/Sent, and Duration. There are 11 distinct TCP states as defined in RFC 793:

- LISTEN - represents waiting for a connection request from any remote TCP and port. (This state is not shown in a Remote Diagnostic output).
- SYN-SENT - represents waiting for a matching connection request after having sent a connection request.
- SYN-RECEIVED - represents waiting for a confirming connection request acknowledgment after having both received and sent a connection request.
- ESTABLISHED - represents an open connection, data received can be delivered to the user. The normal state for the data transfer phase of the connection.

- FIN-WAIT-1 - represents waiting for a connection termination request from the remote TCP, or an acknowledgment of the connection termination request previously sent.
- FIN-WAIT-2 - represents waiting for a connection termination request from the remote TCP.
- CLOSE-WAIT - represents waiting for a connection termination request from the local user.
- CLOSING - represents waiting for a connection termination request acknowledgment from the remote TCP.
- LAST-ACK - represents waiting for an acknowledgment of the connection termination request previously sent to the remote TCP (which includes an acknowledgment of its connection termination request).
- TIME-WAIT - represents waiting for enough time to pass to be sure the remote TCP received the acknowledgment of its connection termination request.
- CLOSED - represents no connection state at all.

List Active Flows

Run this test to list active flows in the system. Use source and destination IP address filters to view the exact flows you want to see. This output is limited to a maximum of 1000 flows.

List Active Flows

List active flows in the system. Use source and destination IP address filters to view the exact flows you want to see. This output is limited to a maximum of 1000 flows.

Run

Segment

Max Flows

Source IP/Port

Destination IP/Port

Test Duration: 1.002 seconds

Src IP	Dst IP	Segment	Protocol	Src Port	Dst Port	DSCP	Application	Link Policy	Route	B
10.0.1.25	10.0.1.1	Global Segment	TCP	59520	179	0	bgp	N/A	Routed	N.
10.0.1.25	108.59.2.24	Global Segment	UDP	123	123	0	ntp	Loadbalance	Direct to Cloud	N.
10.100.1.100	10.100.1.1	segment1	TCP	46392	179	0	bgp	N/A	Routed	N.
10.0.1.25	47.190.36.235	Global Segment	UDP	123	123	0	ntp	Loadbalance	Direct to Cloud	N.
10.0.1.25	10.0.1.1	Global Segment	TCP	60182	179	0	bgp	N/A	Routed	N.
10.0.1.25	184.105.182.15	Global Segment	UDP	123	123	0	ntp	Loadbalance	Direct to Cloud	N.
10.0.1.25	103.38.120.36	Global Segment	UDP	123	123	0	ntp	Loadbalance	Direct to Cloud	N.
10.0.1.25	3.217.79.242	Global Segment	UDP	123	123	0	ntp	Loadbalance	Direct to Cloud	N.
10.101.1.100	10.101.1.1	segment2	TCP	32838	179	0	bgp	N/A	Routed	N.
10.0.1.25	23.152.160.126	Global Segment	UDP	123	123	0	ntp	Loadbalance	Direct to Cloud	N.
10.0.1.25	162.159.200.123	Global Segment	UDP	123	123	0	ntp	Loadbalance	Direct to Cloud	N.
10.0.1.25	204.11.201.10	Global Segment	UDP	123	123	0	ntp	Loadbalance	Direct to Cloud	N.
10.0.1.25	46.4.88.180	Global Segment	UDP	123	123	0	ntp	Loadbalance	Direct to Cloud	N.
10.0.1.25	69.10.161.7	Global Segment	UDP	123	123	0	ntp	Loadbalance	Direct to Cloud	N.
10.0.1.25	85.214.38.116	Global Segment	UDP	123	123	0	ntp	Loadbalance	Direct to Cloud	N.
10.101.1.100	10.101.1.1	segment2	TCP	60408	179	0	bgp	N/A	Routed	N.
10.0.1.25	198.255.68.106	Global Segment	UDP	123	123	0	ntp	Loadbalance	Direct to Cloud	N.
10.0.1.25	84.2.44.19	Global Segment	UDP	123	123	0	ntp	Loadbalance	Direct to Cloud	N.
10.0.1.25	73.189.219.4	Global Segment	UDP	123	123	0	ntp	Loadbalance	Direct to Cloud	N.
10.0.1.25	64.79.100.197	Global Segment	UDP	123	123	0	ntp	Loadbalance	Direct to Cloud	N.
10.100.1.100	10.100.1.1	segment1	TCP	45726	179	0	bgp	N/A	Routed	N.
10.0.1.25	129.134.29.123	Global Segment	UDP	123	123	0	ntp	Loadbalance	Direct to Cloud	N.

List Clients

Run this test to view the complete list of clients.

List Clients

View the full list of clients.

Run

Test Duration: 0.977 seconds

Address	MAC Address	Hostname	Lease Expiry (UTC)	Wireless Connection
10.101.1.100	00:ba:be:52:ff:b3	vc-client1	2020-05-13T07:57:00	

List Paths

Run this test to view the list of active paths between local WAN links and each peer.

List Paths

View the list of active paths between local WAN links and each peer.

Run

Peer

Gateway ▼

Test Duration: 0.982 seconds

WAN Link	Local IP	Remote IP	State	VPN	Bandwidth (tx/rx)	Latency (tx/rx)	Jitter (tx/rx)	Loss (tx/rx)	Bytes (tx/rx)	Uptime
169.254.7.10	169.254.7.10	169.254.10.2	WAITING_FOR_LINK_BW	UP	0.00 Kbps 0.00 Kbps	0 ms 0 ms	0.0 ms 0.0 ms	0.0% 0.0%	11.68 MB 12.29 MB	12h
169.254.6.34	169.254.6.34	169.254.10.2	WAITING_FOR_LINK_BW	UP	99.18 Mbps 187.77 Mbps	0 ms 0 ms	0.0 ms 0.0 ms	0.0% 0.0%	5.71 MB 5.64 MB	12h

MIBs for Edge

Run this test to dump Edge MIBs.

MIBs for Edge

Dump Edge MIBs.

VELOCLOUD-MIB: the root MIB of all VeloCloud specified MIBs and required for installing VELOCLOUD-EDGE-MIB.

VELOCLOUD-MIB-EDGE: the MIB specified for Edge device.

Run

MIB

VELOCLOUD-MIB ▼

Test Duration: 1.001 seconds

```

-----
-- VeloCloud MIB Definitions --
-- Contains:
--   .velocloud(45346)
--   .orchestrator(1)
--   .edge(2)
--   .gateway(3)
-----

VELOCLOUD-MIB DEFINITIONS ::= BEGIN

IMPORTS
    MODULE-IDENTITY, enterprises FROM SNMPv2-SMI
;

velocloud MODULE-IDENTITY
    LAST-UPDATED "201908020000Z"
    ORGANIZATION "VMware Corporation"
    CONTACT-INFO
        "postal: VMware Corporation
        World Headquarters
        3401 Hillview Avenue
        Palo Alto, CA 943043
        USA

        web: www.velocloud.com
        email: contact@velocloud.com"
    DESCRIPTION "Top-level infrastructure of the VeloCloud enterprise MIB tree"
    REVISION "201908020000Z"
    DESCRIPTION "Implementation of VeloCloud Edge MIB Objects"
    REVISION "201701180000Z"
    DESCRIPTION "Implementation of VCO MIB Objects"
    REVISION "201701130000Z"
    DESCRIPTION "Initial definition of VeloCloud MIB Objects"
    ::= { enterprises 45346 }

modules
    OBJECT IDENTIFIER ::= { velocloud 1 }

END

```

NAT Table Dump

Run this test to view the contents of the NAT Table. Use the destination IP address filter to view the exact entries you want to see. This output is limited to a maximum of 1000 entries.

NAT Table Dump

View the contents of the NAT Table. Use the destination IP address filter to view the exact entries you want to see. This output is limited to a maximum of 1000 entries.

Destination IP
 Max Entries

Run

Test Duration: 1.002 seconds

Src IP	Dst IP	Protocol	Src Port	Dst Port	NAT Src IP	NAT Src Port
10.0.1.1	10.81.113.73	TCP	52847	443	169.254.6.34	20128
10.0.1.1	10.81.113.73	TCP	35131	443	169.254.6.34	20180
10.0.1.1	10.81.113.73	TCP	36223	443	169.254.6.34	20137
10.0.1.1	10.81.113.73	TCP	34237	443	169.254.6.34	20042
10.0.1.1	10.81.113.73	TCP	32849	443	169.254.6.34	20098
10.0.1.1	10.81.113.73	TCP	60325	443	169.254.6.34	20065
10.0.1.1	10.81.113.73	TCP	59807	443	169.254.6.34	20222
10.0.1.1	10.81.113.73	TCP	44951	443	169.254.6.34	20246
10.0.1.1	10.81.113.73	TCP	51359	443	169.254.6.34	20095
10.0.1.1	10.81.113.73	TCP	33831	443	169.254.6.34	20087
10.0.1.1	10.81.113.73	TCP	50905	443	169.254.6.34	20192
10.0.1.1	10.81.113.73	TCP	43031	443	169.254.6.34	20110
10.0.1.1	10.81.113.73	TCP	42383	443	169.254.6.34	20191
10.0.1.1	10.81.113.73	TCP	36413	443	169.254.6.34	20077
10.0.1.1	10.81.113.73	TCP	49821	443	169.254.6.34	20155
10.0.1.1	10.81.113.73	TCP	40481	443	169.254.6.34	20245
10.0.1.1	10.81.113.73	TCP	40295	443	169.254.6.34	20032
10.0.1.1	10.81.113.73	TCP	40849	443	169.254.6.34	20064
10.0.1.1	10.81.113.73	TCP	33217	443	169.254.6.34	20148
10.0.1.1	10.81.113.73	TCP	59567	443	169.254.6.34	20091
10.0.1.1	10.81.113.73	TCP	44711	443	169.254.6.34	20217

NTP Dump

Run this test to view the current date and time on Edge and NTP information.

NTP Dump

Current date/time on Edge and NTP information

Run

Test Duration: 1.004 seconds

Edge	
Date/Time	Thu Jul 16 14:04:59 UTC 2020
NTP	
System Peer	104.194.8.227:123
System Peer Mode	client
Leap Indicator	00
Stratum	3
Precision	-23
Root Delay	27.603
Root Dispersion	55.854
Reference ID	104.194.8.227
Reference Time	e2badb7c.14b3dfef Thu, Jul 16 2020 13:58:20.080
System Jitter	3.492954
Clock Jitter	0.302
Clock Wander	0.036
Broadcast Delay	-50.000
Auth Delay	0.000

Ping IPv6 Test

Run a ping test to the specified IPv6 destination.

Ping IPv6 Test

Run a ping IPv6 test to the destination specified.

Destination

Ping From

Test Duration: 8.003 seconds

fd00:1:1:2::1: Reachable
 Src Addr: fd00:1:1:2::2 Min RTT: 0ms, Max RTT: 1ms, Avg RTT: 0.28571428571429ms
 Success Rate: 100% (Packets transmitted: 7, Packets received: 7)

Ping Test

Run a ping test to the specified IPv4 destination.

Ping Test

Run a ping test to the destination specified.

Segment

Destination

Ping From

Test Duration: 8.005 seconds

10.0.1.25: Reachable
 Min RTT: 0ms, Max RTT: 1ms, Avg RTT: 0.28571428571429ms
 Success Rate: 100% (Packets transmitted: 7, Packets received: 7)

Reset USB Modem

Run this test on a selected Edge interface to reset an unworking USB modem connected to the given interface. Note that not all USB modems support this type of remote reset.

Reset USB Modem

This will attempt to reset an unworking USB modem connected to the given interface. Note that not all USB modems support this type of remote reset.

Interface

Test Duration: 55.115 seconds

The restart command has been issued to the selected interface.

Route Table Dump

Run this test to view the contents of the Route Table.

Route Table Dump

View the contents of the Route Table.

Segment

Test Duration: 2.003 seconds

Segmented Route Table

Address	Segment	Netmask	Type	Cost	Reachable	Next Hop	Lost Reason
172.16.1.10	Global Segment	255.255.255.255	N/A	0	TRUE	GE6	LR_NO_ELECTION
172.16.1.2	Global Segment	255.255.255.255	N/A	0	TRUE	GE5	LR_NO_ELECTION
169.254.129.2	Global Segment	255.255.255.255	N/A	0	TRUE		LR_NO_ELECTION
169.254.129.1	Global Segment	255.255.255.255	Cloud	0	TRUE		LR_NO_ELECTION
169.254.7.10	Global Segment	255.255.255.255	N/A	0	TRUE	GE3	LR_NO_ELECTION
169.254.6.34	Global Segment	255.255.255.255	N/A	0	TRUE	GE4	LR_NO_ELECTION
1.1.0.2	Global Segment	255.255.255.255	Connected	0	TRUE	LO2	LR_NO_ELECTION
1.1.0.1	Global Segment	255.255.255.255	Connected	0	TRUE	LO1	LR_NO_ELECTION
172.16.1.8	Global Segment	255.255.255.248	Connected	0	TRUE	GE6	LR_NO_ELECTION
172.16.1.0	Global Segment	255.255.255.248	Connected	0	TRUE	GE5	LR_NO_ELECTION
169.254.7.8	Global Segment	255.255.255.248	Connected	0	TRUE	GE3	LR_NO_ELECTION
169.254.6.32	Global Segment	255.255.255.248	Connected	0	TRUE	GE4	LR_NO_ELECTION
10.0.0.0	Global Segment	255.255.255.0	Connected	0	TRUE	br-network1	LR_NO_ELECTION
0.0.0.0	Global Segment	0.0.0.0	Cloud	0	TRUE	Cloud Gateway	LR_NO_ELECTION
0.0.0.0	Global Segment	0.0.0.0	Cloud	5	TRUE	GE3	LR_VIA_VELO_ROUTE
0.0.0.0	Global Segment	0.0.0.0	Cloud	6	TRUE	GE4	LR_METRIC
0.0.0.0	Global Segment	0.0.0.0	Cloud	7	TRUE	GE5	LR_METRIC
0.0.0.0	Global Segment	0.0.0.0	Cloud	8	TRUE	GE6	LR_METRIC
172.17.1.10	segment1	255.255.255.255	N/A	0	TRUE	GE6	LR_NO_ELECTION
172.17.1.2	segment1	255.255.255.255	N/A	0	TRUE	GE5	LR_NO_ELECTION
172.16.1.10	segment1	255.255.255.255	N/A	0	TRUE	GE6	LR_NO_ELECTION
169.254.7.10	segment1	255.255.255.255	N/A	0	TRUE	GE3	LR_NO_ELECTION
169.254.6.34	segment1	255.255.255.255	N/A	0	TRUE	GE4	LR_NO_ELECTION
1.1.100.2	segment1	255.255.255.255	Connected	0	TRUE	LO2100	LR_NO_ELECTION
1.1.100.1	segment1	255.255.255.255	Connected	0	TRUE	LO1100	LR_NO_ELECTION
172.17.1.8	segment1	255.255.255.248	Connected	0	TRUE	GE6	LR_NO_ELECTION
172.17.1.0	segment1	255.255.255.248	Connected	0	TRUE	GE5	LR_NO_ELECTION
172.16.1.8	segment1	255.255.255.248	Connected	0	TRUE	GE6	LR_NO_ELECTION
169.254.7.8	segment1	255.255.255.248	Connected	0	TRUE	GE3	LR_NO_ELECTION
169.254.6.32	segment1	255.255.255.248	Connected	0	TRUE	GE4	LR_NO_ELECTION
10.100.0.0	segment1	255.255.255.0	Connected	0	TRUE	br-network100	LR_NO_ELECTION
0.0.0.0	segment1	0.0.0.0	Cloud	0	TRUE	Cloud Gateway	LR_NO_ELECTION
0.0.0.0	segment1	0.0.0.0	Cloud	5	TRUE	GE3	LR_VIA_VELO_ROUTE
0.0.0.0	segment1	0.0.0.0	Cloud	6	TRUE	GE4	LR_METRIC
0.0.0.0	segment1	0.0.0.0	Cloud	8	TRUE	GE6	LR_METRIC
172.18.1.10	segment2	255.255.255.255	N/A	0	TRUE	GE6	LR_NO_ELECTION

The **Lost Reason** column specifies the code for the reason why a route loses the routing preference calculation logic to the next preferred route, on both Edges and Gateways.

Note An unresolved route, learnt over multi-hop BGP, might point to an intermediate interface. For more information, see [Multi-hop BGP Routes](#).

The following table lists the reason codes for an Edge and the corresponding descriptions:

Reason Code	Description
PR_UNREACHABLE	In case of overlay routes, the remote peer, which is either Gateway or Edge, is not reachable.
IF_DOWN	Egress Interface is down.
INVALID_IFIDX	Egress Interface if-index for this route is invalid.
SLA_STATE_DOWN	State given by IP SLA tracking is down.
HA_STANDBY	When the local Edge is a Standby, all routes synced from the active are marked as reachable for operational convenience.
LOCAL_MGMT	Management routes are always reachable.
LOOPBACK	Loopback IP address is always reachable.
SELF_ROUTE	Self IP routes are always reachable.
RECUR_UNRES	Recursive routes are marked as reachable so that recursive resolution can be done for operational convenience.
VPN_VIA_NAT	vpnViaNat routes are always reachable.
SLA_STATE_UP	State given by IP SLA tracking is up.
IF_RESOLVED	Egress interface is up and resolved.
PR_REACHABLE	In case of overlay routes, the remote peer, which is either Gateway or Edge, is reachable.
LR_NO_ELECTION	Best route.
LR_NP_SWAN_VS_VELO	Predecessor is selected because it is a non-preferred static WAN route (route configured with preferred flag set to false) when compared to the current route which is a via velocloud route.
LR_NP_SWAN_VS_DEFRT	Predecessor is selected because it is a non-preferred static WAN route when compared to the current route which is default route.
LR_NP_ROUTE_TYPE	Predecessor is selected because its route type is better when compared to the current route. Also, one of the routes being compared is a non-preferred route in this case.
LR_BGP_LOCAL_PREF	Both routes are learnt using BGP. The predecessor is selected because it has a higher local preference than the current route.
LR_BGP_ASPATH_LEN	Both routes are learnt using BGP. Predecessor is selected because it has a lower AS path value than the current route.
LR_BGP_METRIC	Both routes are learnt using BGP. Predecessor is selected because it has a lower metric value than the current route.
LR_EXT_OSPF_INTER	Predecessor is selected because it is a route learnt from OSPF with an inter or intra area metric when compared to

Source Interface Dump

Run this test to view the list of source interfaces used by various services for a segment.

Source Interface Dump

View the Source Interfaces selected by services

Segment

all

Run

Test Duration: 6.003 seconds

Source Interface Table

Segment	Service	Source Interface	Source Selection	Source IP	Destination IP	Destination Port	Destination route
Global Segment	SSH	VLAN-1	Automatic	10.0.1.1	10.0.1.25	51674	10.0.1.0/24 via any
Global Segment	SSH	VLAN-1	Automatic	10.0.1.1	10.0.1.25	51311	10.0.1.0/24 via any
Global Segment	DNS	VLAN-1	Automatic	10.0.1.1	10.112.16.144	53	0.0.0.0/0 via gateway-
Global Segment	DNS	VLAN-1	Automatic	10.0.1.1	10.148.20.5	53	0.0.0.0/0 via gateway-
Global Segment	SSH	GE3	Automatic	169.254.7.10	169.254.10.2	16304	0.0.0.0/0 via gateway-
Global Segment	NTP	VLAN-1	Automatic	10.0.1.1	0.0.0.0	123	0.0.0.0/0 via gateway-
Global Segment	SSH	VLAN-1	Automatic	10.0.1.1	10.0.1.25	51809	10.0.1.0/24 via any
Global Segment	SSH	VLAN-1	Automatic	10.0.1.1	10.0.1.25	51537	10.0.1.0/24 via any
Global Segment	SSH	VLAN-1	Automatic	10.0.1.1	10.0.1.25	63778	10.0.1.0/24 via any
Global Segment	SSH	VLAN-1	Automatic	10.0.1.1	10.0.1.25	51369	10.0.1.0/24 via any
Global Segment	SSH	VLAN-1	Automatic	10.0.1.1	10.0.1.25	51391	10.0.1.0/24 via any
Global Segment	SSH	VLAN-1	Automatic	10.0.1.1	10.0.1.25	60887	10.0.1.0/24 via any
Global Segment	SSH	GE3	Automatic	169.254.7.10	169.254.10.2	16294	0.0.0.0/0 via gateway-
Global Segment	SSH	VLAN-1	Automatic	10.0.1.1	10.0.1.25	63546	10.0.1.0/24 via any
Global Segment	SD-WAN Mgmt	VLAN-1	Automatic	10.0.1.1	169.254.8.2	443	0.0.0.0/0 via gateway-
Global Segment	SSH	VLAN-1	Automatic	10.0.1.1	10.0.1.25	63074	10.0.1.0/24 via any
Global Segment	SSH	VLAN-1	Automatic	10.0.1.1	10.0.1.25	64008	10.0.1.0/24 via any
Global Segment	SSH	VLAN-1	Automatic	10.0.1.1	10.0.1.25	63764	10.0.1.0/24 via any
Global Segment	SSH	GE3	Automatic	169.254.7.10	169.254.10.2	16504	0.0.0.0/0 via gateway-

System Information

Run this test to view system information such as system load, recent WAN stability statistics, monitoring services. WAN stability statistics include the number of times individual VPN tunnels and WAN links lost connectivity for at least 700 milliseconds. The tunnel disconnects do not include the count of direct IPsec connections.

System Information

Run

View system information such as system load, recent WAN stability statistics, Monitoring Services. WAN stability statistics include the number of times individual VPN tunnels and WAN links lost connectivity for at least 700 milliseconds.

Test Duration: 1.002 seconds

System Load	
CPU	1% (Last 30 seconds)
CPU	1% (Last 5 minutes)
Current Memory	48%
Current Flow Count	97
Handoff Queue Drops	0
169.254.7.10 Stability Statistics	
Public IP Address	169.254.7.10
Tunnel Disconnects	0 (Last Hour)
Link Disconnects	0 (Last Hour)
Tunnel Disconnects	10 (Last 19 hours)
Link Disconnects	0 (Last 19 hours)
169.254.6.34 Stability Statistics	
Public IP Address	169.254.6.34
Tunnel Disconnects	0 (Last Hour)
Link Disconnects	0 (Last Hour)
Tunnel Disconnects	10 (Last 19 hours)
Link Disconnects	0 (Last 19 hours)
GE6_Private Stability Statistics	
Public IP Address	172.16.1.10
Tunnel Disconnects	0 (Last Hour)
Link Disconnects	0 (Last Hour)
Tunnel Disconnects	3 (Last 19 hours)
Link Disconnects	1 (Last 19 hours)

Traceroute

Run a traceroute via the Gateway or directly out any of the WAN interfaces to the destination specified.

Traceroute

Run

Run a traceroute via the Gateway or directly out any of the WAN interfaces to the destination specified.

Destination

Traceroute Using

Test Duration: 5.987 seconds

```

traceroute to 10.101.1.100 (10.101.1.100), 30 hops max, 60 byte packets
 1 169.254.7.9 (169.254.7.9) 0.090 ms 0.054 ms 0.043 ms
 2 169.254.6.9 (169.254.6.9) 0.075 ms 0.053 ms 0.050 ms
 3 192.168.0.100 (192.168.0.100) 0.068 ms 0.046 ms 0.066 ms
 4 169.254.249.21 (169.254.249.21) 0.423 ms 0.351 ms 169.254.249.9 (169.254.249.9) 0.266 ms
 5 10.75.12.18 (10.75.12.18) 6.241 ms 10.75.12.14 (10.75.12.14) 7.276 ms 10.75.12.18 (10.75.12.18) 7.222 ms
 6 10.75.12.13 (10.75.12.13) 8.462 ms 6.598 ms 10.75.12.17 (10.75.12.17) 7.562 ms
 7 * * *
 8 * * *
 9 * * *
10 * * *
11 * * *
12 * * *
13 * * *
14 * * *
15 * * *
16 * * *
17 * * *
18 * * *
19 * * *
20 * * *
21 * * *
22 * * *
23 * * *
24 * * *
25 * * *
26 * * *
27 * * *
28 * * *
29 * * *
30 * * *

```

Troubleshoot BFD - Show BFD/BFDv6 Peer Status

Run this test to show all the status of BFD peers with IPv4 or IPv6 address.

Troubleshoot BFD - Show BFD / BFDv6 Peer Status

Use peer and local IPv4 or IPv6 address filters to show the status of specific BFD peers

Run

Segment
Peer IP
Local IP

Test Duration: 1.001 seconds

```

BFD Peer:
  peer fd00:172:21:1::1 multihop local-address fd00:172:21:1::2 vrf [vc:0:1]
  ID: 659075785
  Remote ID: 0
  Status: down
  Downtime: 7 minute(s), 30 second(s)
  Diagnostics: ok
  Remote diagnostics: ok
  Local timers:
    Receive interval: 300ms
    Transmission interval: 300ms
    Echo transmission interval: 50ms
  Remote timers:
    Receive interval: 1000ms
    Transmission interval: 1000ms
    Echo transmission interval: 0ms

```

Troubleshoot BFD - Show BFD/BFDv6 Peer counters

Run this test to view all the counters of BFD peers with IPv4 or IPv6 address.

Troubleshoot BFD - Show BFD / BFDv6 Peer counters

Use peer and local IPv4 or IPv6 address filters to show the counters of specific BFD peers

Run

Segment
Peer IP
Local IP

Test Duration: 1.001 seconds

```

peer fd00:172:21:1::1 multihop local-address fd00:172:21:1::2 vrf [vc:0:1]
Control packet input: 0 packets
Control packet output: 0 packets
Echo packet input: 0 packets
Echo packet output: 0 packets
Session up events: 0
Session down events: 0
Zebra notifications: 2

```

Troubleshoot BFD - Show BFD Setting

Run this test to view BFD setting and neighbor status.

Troubleshoot BFD - Show BFD Setting

Show BFD setting and neighbor status

Run

Test Duration: 1.001 seconds

Seg Name	Peer Address	Local Address	State	Multihop	Detect Multiplier	Receive Interval	Transmit Interval
Global Segment	1.1.99.1	172.21.1.2	UP	true	3	300	300
Global Segment	172.21.4.1	172.21.4.2	UP	false	3	300	300

Troubleshoot BFDv6 - Show BFDv6 Setting

Run this test to view BFDv6 setting and neighbor status.

Troubleshoot BFDv6 - Show BFDv6 Setting							
Show BFDv6 setting and neighbor status							
Test Duration: 2.002 seconds							
Seg ID	Peer Address	Local Address	State	Multihop	Detect Multiplier	Receive Interval	Transmit Interval
Global Segment	2000:1b98:779a::8ae2:7334:370	2000:1b98:779a::8ae2:7334:310	DOWN	true	3	300	300

Multi-hop BGP Routes

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.

You can view the unresolved routes pointing to intermediate interface in the following Remote Diagnostics tests:

- [Troubleshoot BGP - List BGP Redistributed Routes](#)
- [Troubleshoot BGP - List BGP Routes](#)
- [Troubleshoot BGP - List Routes per Prefix](#)
- [Route Table Dump](#)

Troubleshoot BGP - List BGP Redistributed Routes

Run this test to view routes redistributed to BGP neighbors.

Troubleshoot BGP - List BGP Redistributed Routes						
See routes redistributed to BGP neighbors						
Segment <input type="text" value="all"/>						
Test Duration: 1.018 seconds						
Address	Netmask	Metric Type	Next Hop IP	Interface	Seg Name	Communities
115.115.19.143	255.255.255.255	OE2	172.16.1.3	GE5	Global Segment	N/A
115.115.19.134	255.255.255.255	OE2	172.16.1.3	GE5	Global Segment	N/A
115.115.18.234	255.255.255.255	OE2	172.16.1.3	GE5	Global Segment	N/A
115.115.18.216	255.255.255.255	OE2	172.16.1.3	GE5	Global Segment	N/A
115.115.17.43	255.255.255.255	OE2	172.16.1.3	GE5	Global Segment	N/A
115.115.17.20	255.255.255.255	OE2	172.16.1.3	GE5	Global Segment	N/A
115.115.16.174	255.255.255.255	OE2	172.16.1.3	GE5	Global Segment	N/A
115.115.19.124	255.255.255.255	OE2	172.16.1.3	GE5	Global Segment	N/A
115.115.18.58	255.255.255.255	OE2	172.16.1.3	GE5	Global Segment	N/A
115.115.18.57	255.255.255.255	OE2	172.16.1.3	GE5	Global Segment	N/A
115.115.17.181	255.255.255.255	OE2	172.16.1.3	GE5	Global Segment	N/A
115.115.16.151	255.255.255.255	OE2	172.16.1.3	GE5	Global Segment	N/A
115.115.16.71	255.255.255.255	OE2	172.16.1.3	GE5	Global Segment	N/A
115.115.16.37	255.255.255.255	OE2	172.16.1.3	GE5	Global Segment	N/A
115.115.16.20	255.255.255.255	OE2	172.16.1.3	GE5	Global Segment	N/A
115.115.15.234	255.255.255.255	OE2	172.16.1.3	GE5	Global Segment	N/A

Note An unresolved route, learnt over multi-hop BGP, might point to an intermediate interface. For more information, see [Multi-hop BGP Routes](#).

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 BGP Routes

Run

Use IPv4 or IPv6 prefix to filter the specific BGP routes from neighbors, leave prefix empty to see all

Segment

all

Prefix

Test Duration: 1.002 seconds

Address	Netmask	Metric Type	Next Hop IP	Advertise
1.1.0.1	255.255.255.255	E	172.16.1.3	true
1.2.0.1	255.255.255.255	E	172.16.1.11	true
1.3.0.1	255.255.255.255	E	172.16.1.11	true
1.5.0.1	255.255.255.255	E	172.16.1.11	true
10.0.5.0	255.255.255.0	E	172.16.1.11	true
172.16.1.0	255.255.255.248	E	172.16.1.3	true
172.16.1.8	255.255.255.248	E	172.16.1.11	true
172.16.1.16	255.255.255.248	E	172.16.1.3	true
172.16.1.32	255.255.255.248	E	172.16.1.11	true
172.16.2.0	255.255.255.248	E	172.16.1.11	true
172.16.2.16	255.255.255.248	E	172.16.1.11	true
172.16.2.24	255.255.255.248	E	172.16.1.11	true
172.16.3.0	255.255.255.248	E	172.16.1.11	true
172.16.3.8	255.255.255.248	E	172.16.1.11	true
172.16.5.8	255.255.255.248	E	172.16.1.11	true
172.16.5.32	255.255.255.248	E	172.16.1.11	true
172.16.101.0	255.255.255.248	E	172.16.1.11	true
172.16.102.0	255.255.255.248	E	172.16.1.11	true
172.16.201.0	255.255.255.248	E	172.16.1.11	true
fd00:1:1:3::	ffff:ffff:ffff:ffff::	E	fe80::250:56ff:fea5:e553	false
fd00:1:1:4::	ffff:ffff:ffff:ffff::	E	fe80::250:56ff:fea5:c566	false
fd00:1:1:a003::	ffff:ffff:ffff:ffff::	E	fe80::250:56ff:fea5:e553	false
fd00:1:1:b004::	ffff:ffff:ffff:ffff::	E	fe80::250:56ff:fea5:c566	false
fd00:2:1:3::	ffff:ffff:ffff:ffff::	E	fe80::250:56ff:fea5:c566	false
fd00:2:1:a003::	ffff:ffff:ffff:ffff::	E	fe80::250:56ff:fea5:c566	false
fd00:2:1:b003::	ffff:ffff:ffff:ffff::	E	fe80::250:56ff:fea5:c566	false
fd00:3:1:2::	ffff:ffff:ffff:ffff::	E	fe80::250:56ff:fea5:c566	false

Note An unresolved route, learnt over multi-hop BGP, might point to an intermediate interface, as shown in the above image. For more information, see [Multi-hop 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 - List Routes per Prefix

Use IPv4 or IPv6 prefix to show all the Overlay and Underlay routes for the specific prefix and the details

Segment: Prefix:

Test Duration: 2.002 seconds

Address	Netmask	Metric Type	Next Hop IP	AsPath
172.0.0.0	255.0.0.0	E	172.21.4.1	1000 100 4001 4
172.99.121.0	255.255.255.252	E	172.21.11.1	1000 1001 1500
172.99.121.0	255.255.255.252	E	any	2000 2001 1500
172.99.121.0	255.255.255.252	E	any	2000 2001 1500
172.99.121.0	255.255.255.252	E	any	2000 2001 1500
172.99.121.0	255.255.255.252	E	any	2000 2001 1500
172.99.121.0	255.255.255.252	E	any	30011 12001 500
0.0.0.0	0.0.0.0	E	13.1.1.1	N/A
0.0.0.0	0.0.0.0	E	172.21.4.1	N/A
0.0.0.0	0.0.0.0	E	11.1.1.2	N/A
0.0.0.0	0.0.0.0	E	11.1.2.2	N/A
0.0.0.0	0.0.0.0	E	13.1.1.1	N/A
0.0.0.0	0.0.0.0	E	172.21.4.1	N/A
0.0.0.0	0.0.0.0	E	11.1.1.2	N/A
0.0.0.0	0.0.0.0	E	11.1.2.2	N/A

Note An unresolved route, learnt over multi-hop BGP, might point to an intermediate interface. For more information, see [Multi-hop BGP Routes](#).

Troubleshoot BGP - Show BGP Neighbor Advertised Routes

Run this test to view the BGP routes advertised to a neighbor.

Troubleshoot BGP - Show BGP Neighbor Advertised Routes

Show the BGP routes advertised to a neighbor

Segment: Neighbor IP:

Test Duration: 1.002 seconds

```

BGP table version is 21, local router ID is 10.0.1.2, vrf id 1
Default local pref 100, local AS 1
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, r RIB-failure, S Stale, R Removed
Nexthop codes: @NNN nexthop's vrf id, < announce-nh-self
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network        Next Hop        Metric LocPrf Weight Path
*> 10.0.1.0/24     0.0.0.0          0         32768 ?
*> 10.0.2.0/24     0.0.0.0          42        32768 ?
*> 10.0.3.0/24     0.0.0.0          42        32768 ?
*> 10.0.4.0/24     0.0.0.0          42        32768 ?
*> 10.0.5.0/24     0.0.0.0          42        32768 ?
*> 172.16.1.8/29   172.16.1.11      1 100 i
*> 172.16.1.32/29  172.16.1.11      1 100 i
*> 172.16.2.0/29   172.16.1.11      1 100 21 i
*> 172.16.2.16/29  172.16.1.11      1 100 21 i
*> 172.16.2.24/29  172.16.1.11      1 100 i
*> 172.16.3.0/29   172.16.1.11      1 100 i
*> 172.16.3.8/29   172.16.1.11      1 100 i
*> 172.16.5.8/29   172.16.1.11      1 100 i
*> 172.16.5.32/29  172.16.1.11      1 100 i
*> 172.16.101.0/29 172.16.1.11      1 100 i
*> 172.16.102.0/29 172.16.1.11      1 100 i
*> 172.16.201.0/29 172.16.1.11      1 100 111 i

Total number of prefixes 17

```

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 Learned Routes

Show all the accepted BGP routes learned from a neighbor after filters

Run

Neighbor IP

172.16.1.11

Test Duration: 1.001 seconds

```
BGP table version is 21, local router ID is 10.0.1.2, vrf id 1
Default local pref 100, local AS 1
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, r RIB-failure, S Stale, R Removed
Nexthop codes: @NNN nexthop's vrf id, < announce-nh-self
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network        Next Hop           Metric LocPrf Weight Path
*> 172.16.1.8/29   172.16.1.11           0             1 100 i
*> 172.16.1.32/29  172.16.1.11           0             1 100 i
*> 172.16.2.0/29   172.16.1.11           0             1 100 21 i
*> 172.16.2.16/29  172.16.1.11           0             1 100 21 i
*> 172.16.2.24/29  172.16.1.11           0             1 100 i
*> 172.16.3.0/29   172.16.1.11           0             1 100 i
*> 172.16.3.8/29   172.16.1.11           0             1 100 i
*> 172.16.5.8/29   172.16.1.11           0             1 100 i
*> 172.16.5.32/29  172.16.1.11           0             1 100 i
*> 172.16.101.0/29 172.16.1.11           0             1 100 i
*> 172.16.102.0/29 172.16.1.11           0             1 100 i
*> 172.16.201.0/29 172.16.1.11           0             1 100 111 i
```

Displayed 12 routes and 17 total paths

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 Received Routes

Show all the BGP routes learned from a neighbor before filters

Run

Segment

Global Segment ▼

Neighbor IP

172.16.1.11

Test Duration: 1.002 seconds

```
BGP table version is 0, local router ID is 10.0.1.2, vrf id 1
Default local pref 100, local AS 1
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, r RIB-failure, S Stale, R Removed
Nexthop codes: @NNN nexthop's vrf id, < announce-nh-self
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network        Next Hop           Metric LocPrf Weight Path
*> 10.0.1.0/24      172.16.1.11           0             1 100 1 ?
*> 10.0.2.0/24      172.16.1.11           0             1 100 1 ?
*> 10.0.3.0/24      172.16.1.11           0             1 100 1 ?
*> 10.0.4.0/24      172.16.1.11           0             1 100 1 ?
*> 10.0.5.0/24      172.16.1.11           0             1 100 1 ?
*> 172.16.1.8/29    172.16.1.11           0             1 100 i
*> 172.16.1.32/29   172.16.1.11           0             1 100 i
*> 172.16.2.0/29    172.16.1.11           0             1 100 21 i
*> 172.16.2.16/29   172.16.1.11           0             1 100 21 i
*> 172.16.2.24/29   172.16.1.11           0             1 100 i
*> 172.16.3.0/29    172.16.1.11           0             1 100 i
*> 172.16.3.8/29    172.16.1.11           0             1 100 i
*> 172.16.5.8/29    172.16.1.11           0             1 100 i
*> 172.16.5.32/29   172.16.1.11           0             1 100 i
*> 172.16.101.0/29  172.16.1.11           0             1 100 i
*> 172.16.102.0/29  172.16.1.11           0             1 100 i
*> 172.16.201.0/29  172.16.1.11           0             1 100 111 i
```

Total number of prefixes 17

Troubleshoot BGP - Show BGP Neighbor details

Run this test to view the details of BGP neighbor.

Troubleshoot BGP - Show BGP Neighbor details

Run

Show the details of BGP neighbor

Segment

Global Segment ▼

Neighbor IP

172.16.1.11

Test Duration: 1.002 seconds

```

BGP neighbor is 172.16.1.11, remote AS 100, local AS 1, external link
Hostname: vc-b1-cel
BGP version 4, remote router ID 1.1.1.3, local router ID 10.0.1.2
BGP state = Established, up for 06:45:57
Last read 00:00:01, Last write 00:00:01
Hold time is 3, keepalive interval is 1 seconds
Neighbor capabilities:
  4 Byte AS: advertised and received
  AddPath:
    IPv4 Unicast: RX advertised IPv4 Unicast and received
    Route refresh: advertised and received(old & new)
    Address Family IPv4 Unicast: advertised and received
    Hostname Capability: advertised (name: vc-edge,domain name: n/a) received (name: vc-b1-cel,domain name: n/a)
    Graceful Restart Capability: advertised and received
    Remote Restart timer is 120 seconds
    Address families by peer:
      none
  Graceful restart information:
    End-of-RIB send: IPv4 Unicast
    End-of-RIB received: IPv4 Unicast
    Local GR Mode : Helper*
    Remote GR Mode : Helper
    R bit : False
  Timers :
    Configured Restart Time(sec) : 120
    Received Restart Time(sec) : 120
  IPv4 Unicast :
    F bit : False
    End-of-RIB Received : Yes
    End-of-RIB Send : Yes
    EoRSentAfterUpdate : No
    Timers:
      Configured Stale Path Time(sec) : 360
  Message statistics:
    Inq depth is 0
    Outq depth is 0
    Sent      Rcvd
    Opens:    1      1
    Notifications: 0      0
    Updates:  10      9
    Keepalives: 24354 24354
    Route Refresh: 0      0
    Capability:  0      0
    Total:      24365 24364
  Minimum time between advertisement runs is 0 seconds

For address family: IPv4 Unicast
Update group 1, subgroup 1
Packet Queue length 0
Inbound soft reconfiguration allowed
Community attribute sent to this neighbor(all)
12 accepted prefixes

Connections established 1; dropped 0
Last reset never
Local host: 172.16.1.10, Local port: 60782
Foreign host: 172.16.1.11, Foreign port: 179
Nexthop: 172.16.1.10
Nexthop global: ::
Nexthop local: ::
BGP connection: shared network
BGP Connect Retry Timer in Seconds: 120
Read thread: on Write thread: on

```

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 Routes per Prefix

Show all the BGP routes for the prefix and their attributes

Run

Prefix

172.16.3.0

Test Duration: 1.002 seconds

Segment0:

```
BGP routing table entry for 172.16.3.0/29
Paths: (1 available, best #1, table [vc:0:1])
  Advertised to non peer-group peers:
    172.16.1.11
    100
    172.16.1.11 from 172.16.1.11 (1.1.1.3)
      Origin IGP, Default local pref 100, weight 1, valid, external, best
      Last update: Mon Jun  1 08:06:07 2020
```

Segment1:

% Network not in table

Troubleshoot BGP - Show BGP Summary

Run this test to view the existing BGP neighbor and received routes.

Troubleshoot BGP - Show BGP Summary

Show the existing BGP neighbor and received routes

Run

Test Duration: 1.002 seconds

Instance [vc:0:1]:

IPv4 Unicast Summary:

```
BGP view name [vc:0:1]
BGP router identifier 10.0.1.2, local AS number 1 vrf-id 1
BGP table version 21
RIB entries 33, using 5544 bytes of memory
Peers 1, using 22 KiB of memory
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
172.16.1.11	4	100	24657	24658	0	0	0	06:50:50	12

Total number of neighbors 1

Instance [vc:1:2]:

IPv4 Unicast Summary:

```
BGP view name [vc:1:2]
BGP router identifier 10.100.1.1, local AS number 1 vrf-id 2
BGP table version 17
RIB entries 25, using 4200 bytes of memory
Peers 1, using 22 KiB of memory
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
172.17.1.11	4	100	24656	24656	0	0	0	06:50:49	12

Total number of neighbors 1

Troubleshoot BGP - Show BGP Table

Run this test to view the BGP table.

Troubleshoot BGP - Show BGP Table

Show the BGP table

Run

Segment

Global Segment ▼

Test Duration: 1.001 seconds

```

BGP table version is 21, local router ID is 10.0.1.2, vrf id 1
Default local pref 100, local AS 1
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, r RIB-failure, S Stale, R Removed
Nexthop codes: @NNN nexthop's vrf id, < announce-nh-self
Origin codes: i - IGP, e - EGP, ? - incomplete

```

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 10.0.1.0/24	0.0.0.0	0		32768	?
*> 10.0.2.0/24	0.0.0.0	42		32768	?
*> 10.0.3.0/24	0.0.0.0	42		32768	?
*> 10.0.4.0/24	0.0.0.0	42		32768	?
*> 10.0.5.0/24	0.0.0.0	42		32768	?
*> 172.16.1.8/29	172.16.1.11	0		1	100 i
*> 172.16.1.32/29	172.16.1.11			1	100 i
*> 172.16.2.0/29	172.16.1.11			1	100 21 i
*> 172.16.2.16/29	172.16.1.11			1	100 21 i
*> 172.16.2.24/29	172.16.1.11			1	100 i
*> 172.16.3.0/29	172.16.1.11			1	100 i
*> 172.16.3.8/29	172.16.1.11			1	100 i
*> 172.16.5.8/29	172.16.1.11			1	100 i
*> 172.16.5.32/29	172.16.1.11			1	100 i
*> 172.16.101.0/29	172.16.1.11			1	100 i
*> 172.16.102.0/29	172.16.1.11			1	100 i
*> 172.16.201.0/29	172.16.1.11			1	100 111 i

Displayed 17 routes and 17 total paths

Troubleshoot BGPv6 - Show BGPv6 Neighbor Advertised Routes

Run this test to view the BGPv6 routes advertised to a neighbor.

Troubleshoot BGPv6 - Show BGPv6 Neighbor Advertised Routes

Show the BGPv6 routes advertised to a neighbor

Run

Segment

Global Segment ▼

Neighbor IP

fd00:172:21::1

Test Duration: 1.001 seconds

```

BGP table version is 224, local router ID is 16.1.0.2, vrf id 1
Default local pref 100, local AS 1000
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, r RIB-failure, S Stale, R Removed
Nexthop codes: @NNN nexthop's vrf id, < announce-nh-self
Origin codes: i - IGP, e - EGP, ? - incomplete

```

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 2110::/64	fd00:172:21::2			1	16101 i
*> 2110:010:1::/64	fd00:172:21::2			1	16101 i
*> 2110:010:2::/64	fd00:172:21::2			1	16101 i
*> 2110:010:3::/64	fd00:172:21::2			1	16101 i
*> 2110:010:4::/64	fd00:172:21::2			1	16101 i
*> 2110:010:5::/64	fd00:172:21::2			1	16101 i
*> 2110:010:6::/64	fd00:172:21::2			1	16101 i
*> 2110:010:7::/64	fd00:172:21::2			1	16101 i

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 Learned Routes

Show all the accepted BGPv6 routes learned from a neighbor after filters

Segment:

Neighbor IP:

Test Duration: 2.002 seconds

```

BGP table version is 224, local router ID is 16.1.0.2, vrf id 1
Default local pref 100, local AS 1000
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, r RIB-failure, S Stale, R Removed
Next hop codes: @NNN nexthop's vrf id, < announce-nh-self
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network        Next Hop           Metric LocPrf Weight Path
*  fd00:172:21:1::/64
    fe80::250:56ff:fe93:25c5      0           1 1001 ?
*> fd00:172:21:21::/64
    fe80::250:56ff:fe93:25c5      0           1 1001 ?
*> fd00:172:21:11::/64
    fe80::250:56ff:fe93:25c5      0           1 1001 ?
*> fd00:172:21:12::/64
    fe80::250:56ff:fe93:25c5      0           1 1001 ?
*> fd00:172:21:21::/64
    fe80::250:56ff:fe93:25c5      0           1 1001 ?
*> fd00:172:21:22::/64
    fe80::250:56ff:fe93:25c5      0           1 1001 ?

Displayed 6 routes and 220 total paths

```

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 Received Routes

Show all the BGPv6 routes received from a neighbor before filters

Segment:

Neighbor IP:

Test Duration: 1.002 seconds

```

BGP table version is 0, local router ID is 16.1.0.2, vrf id 1
Default local pref 100, local AS 1000
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, r RIB-failure, S Stale, R Removed
Next hop codes: @NNN nexthop's vrf id, < announce-nh-self
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network        Next Hop           Metric LocPrf Weight Path
*> 2120::/64
    fd00:172:21:1::1              1 1001 1000 16102 i
*> 2120:0:0:1::/64
    fd00:172:21:1::1              1 1001 1000 16102 i
*> 2120:0:0:2::/64
    fd00:172:21:1::1              1 1001 1000 16102 i
*> 2120:0:0:3::/64
    fd00:172:21:1::1              1 1001 1000 16102 i
*> 2120:0:0:4::/64
    fd00:172:21:1::1              1 1001 1000 16102 i
*> 2120:0:0:5::/64
    fd00:172:21:1::1              1 1001 1000 16102 i
*> 2120:0:0:6::/64
    fd00:172:21:1::1              1 1001 1000 16102 i
*> 2120:0:0:7::/64
    fd00:172:21:1::1              1 1001 1000 16102 i
*> 2120:0:0:8::/64
    fd00:172:21:1::1              1 1001 1000 16102 i
*> 2120:0:0:9::/64
    fd00:172:21:1::1              1 1001 1000 16102 i
*> 2120:0:0:a::/64
    fd00:172:21:1::1              1 1001 1000 16102 i
*> 2120:0:0:b::/64
    fd00:172:21:1::1              1 1001 1000 16102 i
*> 2120:0:0:c::/64
    fd00:172:21:1::1              1 1001 1000 16102 i
*> 2120:0:0:d::/64
    fd00:172:21:1::1              1 1001 1000 16102 i

```

Troubleshoot BGPv6 - Show BGPv6 Neighbor details

Run this test to view the details of BGPv6 neighbor.

Troubleshoot BGPv6 - Show BGPv6 Neighbor details
Show the details of BGPv6 neighbor

Segment:

Neighbor IP:

Test Duration: 1.001 seconds

```

BGP neighbor is fd00:172:21:1::1, remote AS 1001, local AS 1000, external link
BGP version 4, remote router ID 172.21.1.1, local router ID 16.1.0.2
BGP state = Established, up for 03:11:56
Last read 00:00:03, Last write 00:00:01
Hold time is 1s, keepalive interval is 5 seconds
Configured hold time is 1s, keepalive interval is 5 seconds
Neighbor capabilities:
  * Byte AS: advertised and received
  AddPath:
    IPv6 Unicast: RX advertised IPv6 Unicast
    Route refresh: advertised and received(old & new)
  Address Family IPv6 Unicast: advertised and received
  Hostname Capability: advertised (name: vc-edge, domain name: n/a) not received
  Graceful Restart Capability: advertised
  Graceful restart information:
    Local GR Mode : Helper
    Remote GR Mode : Disable
    R bit : False
  Timers :
    Configured Restart Time(sec) : 120
    Received Restart Time(sec) : 0
  Message statistics:
    Inq depth is 0
    Outq depth is 0
    Sent      Rcvd
    Opens:    1      1
    Notifications: 0      0
    Updates:   10     16
    Keepalives: 2544   2482
    Route Refresh: 0      0
    Capability: 0      0
    Total:     2555   2499
  Minimum time between advertisement runs is 0 seconds

For address family: IPv6 Unicast
  Update group 5, subgroup 3
  Packet queue length 0
  Inbound soft reconfiguration allowed
  Community attribute sent to this neighbor(all)
  6 accepted prefixes

Connections established 1; dropped 0
Last reset never
Local host: fd00:172:21:1::2, Local port: 56918
Foreign host: fd00:172:21:1::1, Foreign port: 179
Next hop: 172.21.1.2
Next hop global: fd00:172:21:1::2
Next hop local: fe80:f7c219e35b2cd:fa8
BGP connection: shared network
BGP Connect Retry Timer in Seconds: 120
Estimated round trip time: 1 ms
Read thread: on Write thread: on

```

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 Routes per Prefix
Show all the BGPv6 routes for the prefix and their attributes

Prefix:

Test Duration: 1.001 seconds

```

Segment0:
BGP routing table entry for fd00:172:21:1::/64
Paths: (1 available, best #1, table [vc:0:1])
Advertised to non-peer-group peers:
fd00:172:21:1::1 fd00:172:21:4::1 fd00:172:21:7::2
1001
fd00:172:21:1::1 from fd00:172:21:1::1 (172.21.1.1)
(fe80:1250:5eff:fe93:25c5) (used)
Origin incomplete, metric 0, Default local pref 100, weight 1, valid, external, best
Last update: Fri Jul 9 20:07:21 2021

Segment1:
% Network not in table

```

Troubleshoot BGPv6 - Show BGPv6 Summary

Run this test to view the existing BGPv6 neighbor and received routes.

Troubleshoot BGPv6 - Show BGPv6 Summary

Show the existing BGPv6 neighbor and received routes

Run

Test Duration: 1.002 seconds

```

Instance [vc:0:1]:
IPv6 Unicast Summary:
BGP view name [vc:0:1]
BGP router identifier 16.1.0.2, local AS number 1000 vrf-id 1
BGP table version 224
RIB entries 837, using 137 KiB of memory
Peers 3, using 65 KiB of memory

Neighbor    V    AS  MsgRcvd  MsgSent  TblVer  InQ  OutQ  Up/Down  State/PfxRcd
Fd00::172:21:1::1  4    1001  2493    2549      0      0      0  03:31:27      6
Fd00::172:21:4::1  4      100  12397   12698      0      0      0  03:31:27      9
Fd00::172:21:7::1  4    16101  2560    2546      0      0      0  03:31:25    200

Total number of neighbors 3

Instance [vc:1:2]:
% No BGP neighbors found

```

Troubleshoot BGPv6 - Show BGPv6 Table

Run this test to view the details of BGPv6 table.

Troubleshoot BGPv6 - Show BGPv6 Table

Show the BGPv6 table

Segment

Run

Test Duration: 1.002 seconds

```

Instance [vc:0:1]:
BGP table version is 224, local router ID is 16.1.0.2, vrf id 1
Default local pref 100, local AS 1000
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, r RIB-failure, S Stale, R Removed
NextHop codes: @NNN nextHop's vrf id, < announce-nh-self
Origin codes: i - IGP, e - EGP, ? - incomplete

Network        Next Hop        Metric LocPrf Weight Path
>> 2110::/64    fe80::11         1 161001 1
>> 2110:0:0:11::/64 fe80::11       1 161001 1
>> 2110:0:0:12::/64 fe80::11       1 161001 1
>> 2110:0:0:13::/64 fe80::11       1 161001 1
>> 2110:0:0:14::/64 fe80::11       1 161001 1
>> 2110:0:0:15::/64 fe80::11       1 161001 1
>> 2110:0:0:16::/64 fe80::11       1 161001 1
>> 2110:0:0:17::/64 fe80::11       1 161001 1
>> 2110:0:0:18::/64 fe80::11       1 161001 1
>> 2110:0:0:19::/64 fe80::11       1 161001 1
>> 2110:0:0:1a::/64 fe80::11       1 161001 1
>> 2110:0:0:1b::/64 fe80::11       1 161001 1
>> 2110:0:0:1c::/64 fe80::11       1 161001 1
>> 2110:0:0:1d::/64 fe80::11       1 161001 1
>> 2110:0:0:1e::/64 fe80::11       1 161001 1
>> 2110:0:0:1f::/64 fe80::11       1 161001 1
>> 2110:0:0:10::/64 fe80::11       1 161001 1
>> 2110:0:0:11::/64 fe80::11       1 161001 1
>> 2110:0:0:12::/64 fe80::11       1 161001 1
>> 2110:0:0:13::/64 fe80::11       1 161001 1
>> 2110:0:0:14::/64 fe80::11       1 161001 1
>> 2110:0:0:15::/64 fe80::11       1 161001 1
>> 2110:0:0:16::/64 fe80::11       1 161001 1
>> 2110:0:0:17::/64 fe80::11       1 161001 1
>> 2110:0:0:18::/64 fe80::11       1 161001 1
>> 2110:0:0:19::/64 fe80::11       1 161001 1
>> 2110:0:0:1a::/64 fe80::11       1 161001 1
>> 2110:0:0:1b::/64 fe80::11       1 161001 1
>> 2110:0:0:1c::/64 fe80::11       1 161001 1

```

Troubleshoot OSPF - List OSPF Redistributed Routes

Run this test to view all the routes redistributed to OSPF neighbor.

Troubleshoot OSPF - List OSPF Redistributed Routes

Show all the routes redistributed to OSPF neighbor

Run

Test Duration: 1.017 seconds

Address	Netmask	Metric Type	Next Hop IP	Cost	Interface
115.115.19.143	255.255.255.255	OE2	172.16.1.3	1	GE5
115.115.19.134	255.255.255.255	OE2	172.16.1.3	1	GE5
115.115.18.234	255.255.255.255	OE2	172.16.1.3	1	GE5
115.115.18.216	255.255.255.255	OE2	172.16.1.3	1	GE5
115.115.17.43	255.255.255.255	OE2	172.16.1.3	1	GE5
115.115.17.20	255.255.255.255	OE2	172.16.1.3	1	GE5
115.115.16.174	255.255.255.255	OE2	172.16.1.3	1	GE5
115.115.19.124	255.255.255.255	OE2	172.16.1.3	1	GE5
115.115.18.58	255.255.255.255	OE2	172.16.1.3	1	GE5
115.115.18.57	255.255.255.255	OE2	172.16.1.3	1	GE5
115.115.17.181	255.255.255.255	OE2	172.16.1.3	1	GE5
115.115.16.151	255.255.255.255	OE2	172.16.1.3	1	GE5

Troubleshoot OSPF - List OSPF Routes

Run this test to view the OSPF routes from neighbors for the specified Prefix. Displays all the OSPF routes from the neighbors if the Prefix is not specified.

Troubleshoot OSPF - List OSPF Routes

Show the specific OSPF routes from neighbors, leave prefix empty to see all

Prefix

Test Duration: 2.025 seconds

Address	Netmask	Metric Type	Nbr ID	OSPF Cost	Overlay Preference	Advertise	Interface
115.115.15.143	255.255.255.255	OE2	1.1.1.2	1	64	false	GE5
115.115.15.144	255.255.255.255	OE2	1.1.1.2	1	64	false	GE5
115.115.15.145	255.255.255.255	OE2	1.1.1.2	1	64	false	GE5
115.115.15.146	255.255.255.255	OE2	1.1.1.2	1	64	false	GE5
115.115.15.147	255.255.255.255	OE2	1.1.1.2	1	64	false	GE5
115.115.15.148	255.255.255.255	OE2	1.1.1.2	1	64	false	GE5
115.115.15.149	255.255.255.255	OE2	1.1.1.2	1	64	false	GE5
115.115.15.150	255.255.255.255	OE2	1.1.1.2	1	64	false	GE5
115.115.15.151	255.255.255.255	OE2	1.1.1.2	1	64	false	GE5
115.115.15.152	255.255.255.255	OE2	1.1.1.2	1	64	false	GE5
115.115.15.153	255.255.255.255	OE2	1.1.1.2	1	64	false	GE5
115.115.15.154	255.255.255.255	OE2	1.1.1.2	1	64	false	GE5
115.115.15.155	255.255.255.255	OE2	1.1.1.2	1	64	false	GE5

Troubleshoot OSPF - Show OSPF Database

Run this test to view the OSPF link state database summary.

Troubleshoot OSPF - Show OSPF Database

Show the OSPF link state database summary

Test Duration: 1.003 seconds

```

OSPF Router with ID (10.0.1.2)

  Router Link States (Area 0.0.0.1)

  Link ID        ADV Router    Age  Seq#          CkSum  Link count
  1.1.1.2        1.1.1.2      779  0x800000014  0x26a2  2
  10.0.1.2       10.0.1.2    1015 0x80000000e  0x6049  1

  Net Link States (Area 0.0.0.1)

  Link ID        ADV Router    Age  Seq#          CkSum
  172.16.1.3     1.1.1.2     1039 0x80000000c  0x126c

  AS External Link States

  Link ID        ADV Router    Age  Seq#          CkSum  Route
  0.0.0.0        10.0.1.2    1055 0x80000000d  0x5d5c  E2 0.0.0.0/0 [0x0]
  10.0.1.0       10.0.1.2    305  0x80000000f  0x48ed  E1 10.0.1.0/24 [0x0]
  10.0.2.0       10.0.1.2    1105 0x80000000e  0xe41e  E1 10.0.2.0/24 [0x0]
  10.0.3.0       10.0.1.2    1015 0x80000000e  0xd928  E1 10.0.3.0/24 [0x0]
  10.0.4.0       10.0.1.2    1025 0x80000000e  0xc32e  E1 10.0.4.0/24 [0x0]
  10.0.5.0       10.0.1.2    1025 0x80000000e  0xc33c  E1 10.0.5.0/24 [0x0]
  115.115.15.143 1.1.1.2      749  0x80000000c  0xe93f  E2 115.115.15.143/32 [0x0]
  115.115.15.144 1.1.1.2      909  0x80000000c  0xdf48  E2 115.115.15.144/32 [0x0]
  115.115.15.145 1.1.1.2      849  0x80000000c  0xd551  E2 115.115.15.145/32 [0x0]
  115.115.15.146 1.1.1.2      889  0x80000000c  0xcb5a  E2 115.115.15.146/32 [0x0]
  115.115.15.147 1.1.1.2      779  0x80000000c  0xc163  E2 115.115.15.147/32 [0x0]
  115.115.15.148 1.1.1.2      859  0x80000000c  0xb76c  E2 115.115.15.148/32 [0x0]
  115.115.15.149 1.1.1.2      869  0x80000000c  0xad75  E2 115.115.15.149/32 [0x0]
  115.115.15.150 1.1.1.2      799  0x80000000c  0xa37e  E2 115.115.15.150/32 [0x0]
  115.115.15.151 1.1.1.2      829  0x80000000c  0x9987  E2 115.115.15.151/32 [0x0]
  115.115.15.152 1.1.1.2      839  0x80000000c  0x8f90  E2 115.115.15.152/32 [0x0]
  115.115.15.153 1.1.1.2      869  0x80000000c  0x8599  E2 115.115.15.153/32 [0x0]
  115.115.15.154 1.1.1.2      789  0x80000000c  0x7ba2  E2 115.115.15.154/32 [0x0]
  115.115.15.155 1.1.1.2      779  0x80000000c  0x71ab  E2 115.115.15.155/32 [0x0]

```

Troubleshoot OSPF - Show OSPF Database for E1 Self-Originate Routes

Run this test to view the E1 LSA's self-originated routes that are advertised to OSPF router by the Edge.

Troubleshoot OSPF - Show OSPF Database for E1 Self-Originate Routes

Show the E1 LSA's self-originated by the VCE that are advertised to OSPF

Test Duration: 1.002 seconds

Run

```

OSPF Router with ID (10.0.1.2)
  AS External Link States

  LS age: 1197
  Options: 0x2 : *|---|---|E|
  LS Flags: 0xb
  LS Type: AS-external-LSA
  Link State ID: 0.0.0.0 (External Network Number)
  Advertising Router: 10.0.1.2
  LS Seq Number: 8000000d
  Checksum: 0x3d5c
  Length: 36

  Network Mask: /0
  Metric Type: 2 (Larger than any link state path)
  TOS: 0
  Metric: 0
  Forward Address: 0.0.0.0
  External Route Tag: 0

  LS age: 447
  Options: 0x2 : *|---|---|E|
  LS Flags: 0xb
  LS Type: AS-external-LSA
  Link State ID: 10.0.1.0 (External Network Number)
  Advertising Router: 10.0.1.2
  LS Seq Number: 8000000f
  Checksum: 0x48e4
  Length: 36

  Network Mask: /24
  Metric Type: 1
  TOS: 0
  Metric: 0
  Forward Address: 0.0.0.0
  External Route Tag: 0

  LS age: 1247
  Options: 0x2 : *|---|---|E|
  LS Flags: 0xb
  LS Type: AS-external-LSA
  Link State ID: 10.0.2.0 (External Network Number)
  Advertising Router: 10.0.1.2
  LS Seq Number: 8000000e
  Checksum: 0xe41e
  Length: 36

  Network Mask: /24
  Metric Type: 1
  TOS: 0
  Metric: 42
  Forward Address: 0.0.0.0
  External Route Tag: 0

  LS age: 1157
  Options: 0x2 : *|---|---|E|
  LS Flags: 0xb
  LS Type: AS-external-LSA
  Link State ID: 10.0.3.0 (External Network Number)
  Advertising Router: 10.0.1.2
  LS Seq Number: 8000000e
  Checksum: 0xd928
  Length: 36

```

Troubleshoot OSPF - Show OSPF Neighbors

Run this test to view all the OSPF neighbors and associated information.

Troubleshoot OSPF - Show OSPF Neighbors

Show all the OSPF neighbors and associated info

Run

Test Duration: 1.001 seconds

Neighbor ID	Pri	State	Dead Time	Address	Interface	RXmtL	RqstL	DBsml
1.1.1.2	1	Full/DR	36.885s	172.16.1.3	GE5:172.16.1.2	0	0	0

Troubleshoot OSPF - Show OSPF Route Table

Run this test to view the existing OSPF route table.

Troubleshoot OSPF - Show OSPF Route Table

Show the existing OSPF route table

Run

Test Duration: 1.005 seconds

```

===== OSPF network routing table =====
N 172.16.1.0/29      [1] area: 0.0.0.1
                        directly attached to GE5
N 172.16.1.16/29     [11] area: 0.0.0.1
                        via 172.16.1.3, GE5

===== OSPF router routing table =====
R 1.1.1.2            [1] area: 0.0.0.1, ASBR
                        via 172.16.1.3, GE5

===== OSPF external routing table =====
N E2 115.115.15.143/32 [1/20] tag: 0
                        via 172.16.1.3, GE5
N E2 115.115.15.144/32 [1/20] tag: 0
                        via 172.16.1.3, GE5
N E2 115.115.15.145/32 [1/20] tag: 0
                        via 172.16.1.3, GE5
N E2 115.115.15.146/32 [1/20] tag: 0
                        via 172.16.1.3, GE5
N E2 115.115.15.147/32 [1/20] tag: 0
                        via 172.16.1.3, GE5
N E2 115.115.15.148/32 [1/20] tag: 0
                        via 172.16.1.3, GE5
N E2 115.115.15.149/32 [1/20] tag: 0
                        via 172.16.1.3, GE5
N E2 115.115.15.150/32 [1/20] tag: 0
                        via 172.16.1.3, GE5
N E2 115.115.15.151/32 [1/20] tag: 0
                        via 172.16.1.3, GE5

```

Troubleshoot OSPF - Show OSPF Setting

Run this test to view the OSPF setting and neighbor status.

Troubleshoot OSPF - Show OSPF Setting

Show OSPF setting and neighbor status

Run

Test Duration: 1.002 seconds

Area	Network Info	Authentication	Cost	Hello Timer	Dead Timer	Interface	MD5
1	172.16.1.0/29	0	1	10	40	GE5	0

USB Port Status

Run this test to view the status of USB ports on an Edge.

USB ports status

View USB ports status on Edge

Run

Test Duration: 2.005 seconds

BIOS	
Version	3.43.0.9-8
CMOS	
Disabled	True
Offset	82
Value	0x00
GRUB	
Disabled	True
Value	disable_usb=1

VPN Test

Select a segment from the drop-down menu and click **Run** to test VPN connectivity to each peer.

VPN Test

Use ping to test VPN connectivity to each peer.

Run

Segment

Global Segment ▼

Test Duration: 3.002 seconds

Edge Name	Result	Latency(millisecs)
b5-edge1	Pass	3
b2-edge1	Pass	3
b3-edge1	Pass	3
b4-edge1	Pass	3

When the VPN test is run, the Edge selects the Source and Destination IP and initiates the tunnel request. The selected Source and Destination IP should meet the following criteria:

- It should be a connected route IP
- It should be reachable and the routes should be advertised

When the Edge cannot select a valid IP as the Source IP to initiate the tunnel request, the VPN Test will fail with the following error.

```
Branch-to-Branch vpn is disabled. Please enable it before running the test
```

WAN Link Bandwidth Test

Run the bandwidth test on a specified WAN link. This test has the benefit of being non-disruptive in multi-link environments. Only the link under test is blocked for user traffic. This means that you can re-run the test on a specific link and the other link(s) will continue to serve user traffic.

WAN Link Bandwidth Test

Force a re-test the bandwidth of a WAN link.

Run

WAN Link

GE6_Private ▼

Test Duration: 1.001 seconds

Bandwidth test has been queued. When the test completes, the new measurements will be shown on [Edge Overview](#).

As the bandwidth test is run when the tunnel reconnects after a period of instability, there have been occasions in the field where the link has recovered enough for tunnel connectivity, but not enough to accurately measure the bandwidth of the WAN link. To address these scenarios, if the bandwidth test fails or measures a significantly reduced value, the last known “good” measurement will be used and a re-test of the link will be scheduled for 30 minutes after the tunnel is established to ensure a proper measurement.

Note For WAN link over 900Mbps, it is recommended that the user define the bandwidth of the WAN link.

Remote Actions

You can perform actions like Restarting services, Rebooting, or deactivating an Edge remotely, from the Enterprise portal.

You can perform the remote actions only on Edge that are in **Connected** state.

- 1 In the Enterprise portal, click **Test & Troubleshoot > Remote Actions**.
- 2 The **Remote Edge Actions** page displays all the connected Edges. Search for an Edge if necessary using the **Filter**, and click **Apply**.
- 3 Click the link to a connected Edge.

In the **Edge Remote Actions** window, click the relevant action. The action is performed on the selected Edge.



- 4 You can perform the following actions:

Action	Description
Identify	Randomly flash lights on the selected Edge to identify the device.
Restart Service	Restarts the VMware services on the selected Edge.
Reboot	Reboots the selected Edge.
Shutdown	Power off the selected Edge.
Deactivate	Resets the device configuration to its factory default state.

Note The actions may take up to a minute to run on the device.

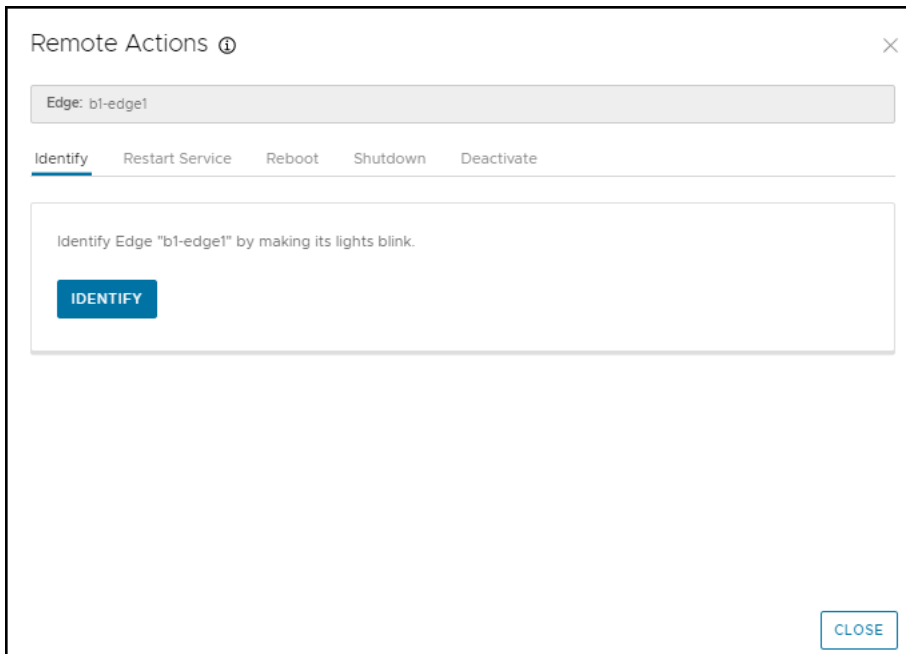
Perform Remote Actions with new Orchestrator UI

You can perform actions like Restarting services, Rebooting, or deactivating an Edge remotely, from the Enterprise portal.

You can perform the remote actions only on Edge that are in **Connected** state.

To perform the remote actions:

- 1 In the Enterprise portal, click the **Open New Orchestrator UI** option available at the top of the Window.
- 2 Click **Launch New Orchestrator UI** in the pop-up window.
- 3 The UI opens in a new tab displaying the monitoring and configuring options.
- 4 In the new UI, you perform the remote actions for an Edge using the **Shortcuts** option available in the **Configure > Edges** or **Monitor > Edges** pages. See [Configure Edges with new Orchestrator UI](#) and [Monitor Edges](#).
- 5 Click the **Shortcuts > Remote Actions** to open the **Remote Actions** dialog.



- 6 You can perform the following actions:

Action	Description
Identify	Randomly flash lights on the selected Edge to identify the device.
Restart Service	Restarts the VMware SD-WAN services on the selected Edge.
Reboot	Reboots the selected Edge.
Shutdown	Power off the selected 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.

Note The actions may take up to a minute to run on the device.

Diagnostic Bundles

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.

In the Enterprise portal, click **Test & Troubleshooting > Diagnostic Bundles**.



The **Diagnostic Bundles** window allows to request for the following bundles:

- **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. See [Request Packet Capture Bundle](#).
- **Diagnostic Bundle** – The Diagnostic bundle is a collection of all the configuration and logs from a specific Edge. Only Operators can request Diagnostic bundles. See [Request Diagnostic Bundle](#).

The generated bundles are displayed in the **Diagnostic Bundles** window. To download the bundle files, see [Download Diagnostic Bundle](#).

The **Diagnostic Bundles** 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.

In the Enterprise portal, click **Test & Troubleshooting > Packet Capture**.



Click **Request PCAP Bundle** to generate Packet Capture bundle, which is a collection of the packet data of the network. See [Request Packet Capture Bundle](#).

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.

If you are an Operator user, in the Enterprise portal, click **Test & Troubleshooting > Diagnostic Bundles**.

If you are a Partner user or an Enterprise user, click **Test & Troubleshooting > Packet Capture**.

- 1 Click **Request PCAP Bundle**.
- 2 In the **Request PCAP Bundle** window that appears, configure the following:

The screenshot shows a window titled "Request PCAP Bundle". Inside the window, there are four configuration fields:

- Target:** A dropdown menu with "b1-edge1" selected.
- Interface:** A dropdown menu with "GE3" selected.
- Duration:** A dropdown menu with "10 seconds" selected.
- Reason for Generation:** A text input field containing "PCAP from b1-edge1 and GE3".

 At the bottom right of the window are two buttons: a green "Submit" button and a grey "Close" button.

- **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.

- 3 Click **Submit**.

The window displays the details of the bundle being generated, along with the status. To download the generated bundle, see [Download Diagnostic Bundle](#).


Request Diagnostic Bundle

A Diagnostic bundle is a collection of configuration files, logs, and related events from a specific Edge.

In the Enterprise portal, click **Test & Troubleshooting > Diagnostic Bundles**.

Note The **Diagnostic Bundles** option is available only for an Operator user.

- 1 Click **Request Diagnostic Bundle**.
- 2 In the **Request Diagnostic Bundle** window, configure the following:



The dialog box titled "Request Diagnostic Bundle" contains the following fields and controls:

- Target:** A dropdown menu with "b5-edge1" selected.
- Reason for Generation:** A text area containing "Diagnostic data from b5-edge1".
- Core Limit:** A dropdown menu with "No Limit" selected. The dropdown list is open, showing options: "No Limit", "3", "2", "1", and "No Cores".
- Buttons:** "Advanced", "Submit", and "Close".

- **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.
- If required, click the **Advanced** button and choose a value from the **Core Limit** drop-down list. The Core Limit is used to reduce the size of the uploaded bundle when the Internet connectivity is experiencing issues.

3 Click **Submit**.

The **Diagnostic Bundles** window displays the details of the bundle being generated, along with the status. To download the generated bundle, see [Download Diagnostic Bundle](#).

Download Diagnostic Bundle

If you are an Operator user, in the Enterprise portal, click **Test & Troubleshooting > Diagnostic Bundles**.

If you are a Partner user or an Enterprise user, click **Test & Troubleshooting > Packet Capture**.

The generated bundles are displayed in the window.

Monitor

Configure

Test & Troubleshoot

Remote Diagnostics

Remote Actions

Diagnostic Bundles

Administration

Diagnostic Bundles

Request PCAP Bundle...

Request Diagnostic Bundle...

?

Search

ⓘ

Cols

✕ Reset View

↻ Refresh

📄 CSV

Display 2 items0 selected⚙ Actions

<input type="checkbox"/>		Request Status	Type	Edge	Reason for Generation	User	Generated	⬇	Cleanup Date
<input type="checkbox"/>	👤	Complete	Diagnostics	b5-edge1	Diagnostic data from b5-edge1	super@velocloud.net	Thu Jun 04, 15:33:03		Mon Aug 03
<input type="checkbox"/>	👤	Complete	PCAP	b1-edge1	PCAP from b1-edge1 and GE3	super@velocloud.net	Thu Jun 04, 15:02:35		Mon Aug 03

To download a generated bundle, click the **Complete** link or select the bundle and click **Actions > Download Diagnostic Bundle**. The bundle is downloaded as a ZIP file.

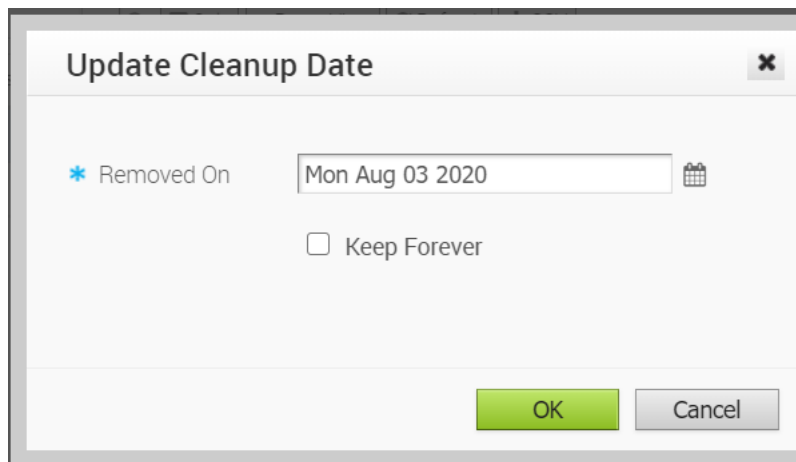
You can send the downloaded bundle to a VMware Support representative for debugging the data.

Delete Diagnostic Bundle

If you are an Operator user, in the Enterprise portal, click **Test & Troubleshooting > Diagnostic Bundles**.

If you are a Partner user or an Enterprise user, click **Test & Troubleshooting > Packet Capture**.

The completed bundles get deleted automatically on the date displayed in the **Cleanup Date** column. You can click the link to the Cleanup Date to modify the Date.



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

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

To delete a bundle manually, select the bundle and click **Actions > Delete Diagnostic Bundle**.

Diagnostic Bundles with new Orchestrator UI


Diagnostic bundles allow 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 using the new Orchestrator UI:

- 1 In the Enterprise portal, click the **Open New Orchestrator UI** option available at the top of the Window.
- 2 Click **Launch New Orchestrator UI** in the pop-up window.
- 3 The UI opens in a new tab displaying the monitoring and configuring options.

- 4 In the new UI, click the **Diagnostics** tab.
- 5 Click **Diagnostic Bundles** to request the following bundles:
 - **PCAP Bundle** – The Packet Capture bundle is a collection of the packet data of the network. See [Request Packet Capture Bundle](#).
 - **Diagnostic Bundle** – The Diagnostic bundle is a collection of all the configuration and logs from a specific Edge. See [Request Diagnostic Bundle](#).

The generated bundles are displayed in the **Diagnostic Bundles** window.



The screenshot shows the 'Diagnostic Bundles' window. On the left is a sidebar with 'Diagnostics', 'Remote Diagnostics', and 'Diagnostic Bundles'. The main area has a search bar and buttons: '+ REQUEST PCAP BUNDLE', '+ REQUEST DIAGNOSTIC BUNDLE', 'DOWNLOAD BUNDLE', 'DELETE', and '... MORE'. Below these is a table with the following data:

<input type="checkbox"/>	Request Status	Type	Edge	Reason for Generation	User	Generated Date	Cleanup Date
<input type="checkbox"/>	Pending	Diagnostics	b2-edge1	For Troubleshooting	super@velocloud.net	Aug 18, 2021, 11:57:22 PM	Oct 17, 2021
<input type="checkbox"/>	In Progress	PCAP	b1-edge1	For troubleshooting	super@velocloud.net	Aug 18, 2021, 11:56:02 PM	Oct 17, 2021

To download the bundle files, see [Download Diagnostic Bundle](#).

To delete the bundle files, see [Delete Diagnostic Bundle](#).

To download the details of generated bundles, click **More > Download CSV**. The details are downloaded in a CSV file.

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 Enterprise portal, click the **Open New Orchestrator UI** option available at the top of the Window.
- 2 Click **Launch New Orchestrator UI** in the pop-up window.
- 3 The UI opens in a new tab displaying the monitoring and configuring options.
- 4 In the new UI, click the **Diagnostics** tab.
- 5 Click **Diagnostic Bundles > Request Diagnostic Bundle**.
- 6 In the **Request Diagnostic Bundle** window, configure the following:

Table 25-1.

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. To download the generated bundle, see [Download Diagnostic Bundle](#).

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 Enterprise portal, click the **Open New Orchestrator UI** option available at the top of the Window.
- 2 Click **Launch New Orchestrator UI** in the pop-up window.
- 3 The UI opens in a new tab displaying the monitoring and configuring options.
- 4 In the new UI, click the **Diagnostics** tab.
- 5 Click **Diagnostic Bundles > Request PCAP Bundle**.
- 6 In the **Request PCAP Bundle** window that appears, configure the following:



The image shows a 'Request PCAP Bundle' dialog box with a close button (X) in the top right corner. It contains four fields: 'Target' with a dropdown menu showing 'b1-edge1', 'Interface' with a dropdown menu showing 'GE5', 'Duration' with a dropdown menu showing '5 seconds', and 'Reason for Generation' with a text area containing 'For troubleshooting'. At the bottom right are 'CLOSE' and 'SUBMIT' buttons.

Table 25-2.

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. To download the generated bundle, see [Download Diagnostic Bundle](#).

Download Diagnostic Bundle

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

- 1 In the Enterprise portal, click the **Open New Orchestrator UI** option available at the top of the Window.
- 2 Click **Launch New Orchestrator UI** in the pop-up window.
- 3 The UI opens in a new tab displaying the monitoring and configuring options.
- 4 In the new UI, click the **Diagnostics** tab and click **Diagnostic Bundles**.

The generated bundles are displayed in the window.

To download a generated bundle, click the **Complete** link or select the bundle and click **Download Bundle**. The bundle is downloaded as a ZIP file.

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

Delete Diagnostic Bundle

To delete a Diagnostic or PCAP bundle:

- 1 In the Enterprise portal, click the **Open New Orchestrator UI** option available at the top of the Window.
- 2 Click **Launch New Orchestrator UI** in the pop-up window.
- 3 The UI opens in a new tab displaying the monitoring and configuring options.
- 4 In the new UI, click the **Diagnostics** tab and click **Diagnostic Bundles**.

The generated bundles are displayed in the window.

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

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

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

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

The **Administration** option in the Enterprise portal allows you to configure the System settings, Authentication information, create Admin users, and manage Edge licenses.

In the Enterprise portal, click **Administration** to configure the following:

- **System Settings**– Configure the user information and enterprise authentication. See [System Settings](#).
- **Administrators**– Create or modify admin users with different role privileges. See [Manage Admin Users](#).
- **Edge Licensing**– View and generate a report of Edge licenses. See [Edge Licensing](#).

Read the following topics next:

- [System Settings](#)
- [Manage Admin Users](#)
- [Roles](#)
- [Edge Licensing](#)

System Settings

The **System Settings** option allows you to configure administrator settings along with the authentication details.

In the Enterprise portal, click **Administration > System Settings** to configure the following:

- **General Information**– Configure the user details, enable Edge configuration updates, configure privacy settings, and enter the contact information. See [Configure Enterprise Information](#).
- **Authentication**– Configure authentication mode and view the API tokens. See [Configure Enterprise Authentication](#).

Configure Enterprise Information

You can configure the user information, software images, Edge updates, privacy settings, and contact details for the enterprise users using the **General Information** tab under **Administration > System Settings**.

In the Enterprise portal, click **Manage Customers** and select an enterprise customer. Then go to **Administration > System Settings**. The **System Settings** page appears. You can configure the following in the **General Information** tab.

The screenshot shows the VMware SD-WAN System Settings page with the General Information tab selected. The left sidebar contains navigation links: Monitor, Configure, Test & Troubleshoot, Administration, System Settings (selected), Administrators, Role Customization, Edge Licensing, and Zero Touch Provisioning. The main content area is titled 'System Settings' and includes a 'Save Changes' button. The General Information section contains fields for Name (mars), Account Number (MAR-UT7C6P4), Domain (mars), and Description. Below these are several checkboxes for authentication and security settings, including 'Enable Two Factor Authentication', 'Require Two Factor Authentication', 'Enable Self Service Password Reset' (checked), 'Require Two Factor Authentication for Password Reset', 'Enable Pre-Notifications' (checked), 'Enable Alerts' (checked), and 'Default Edge Authentication' (Certificate Acquire). The Edge Configuration section has a checkbox for 'Updates' (checked) and a section for 'Enabled on Orchestrator Upgrade'. The Privacy Settings section includes checkboxes for 'Support Access' (checked), 'Grant Access to VeloCloud Support', 'Grant View Sensitive Data to VeloCloud Support', 'Grant User Management Access to VeloCloud Support', and 'Enforce PCI'. The Contact Information section contains fields for Contact Name, Contact Email, Phone, Mobile, Direct Address, City, State, ZIP/Postcode, and Country. The Zero Touch Provisioning Sign Up section includes a checkbox for 'DD' and a 'Submit' button.

System Settings Save Changes ?

General Information **Authentication**

General Information

- Name: mars
- Account Number: MAR-UT7C6P4
- Domain: mars
- Description:
- Enable Two Factor Authentication: ☐
- Require Two Factor Authentication: ☐
- Enable Self Service Password Reset: ☒
- Require Two Factor Authentication for Password Reset: ☐
- Enable Pre-Notifications: ☒
- Enable Alerts: ☒
- Default Edge Authentication: Certificate Acquire

Edge Configuration

Updates: ☒ **Enabled**
 When Enabled is on (normal), Edge configuration updates are communicated to an Edge on its next heartbeat. Configuration changes may cause Edge software restart. When the Orchestrator is upgraded, Enabled is set to the value of Enabled on Orchestrator Upgrade.

☐ **Enabled on Orchestrator Upgrade**
 Orchestrator upgrade may update Edge configurations, which when communicated to the Edge may cause software restart. Setting Enabled on Orchestrator Upgrade off allows the operator to choose when after Orchestrator upgrade to set Enabled to resume having Edge configuration updates communicated to Edges.

Privacy Settings

Support Access: ☒ **Grant Access to VeloCloud Support**
 When activated, VeloCloud Support will be granted access to view, configure and troubleshoot this Customer's Edges. It will also allow VeloCloud Support to build and customize roles for the customer. As a security consideration, VeloCloud Support will not be granted access to view user identifiable information.

☒ **Grant View Sensitive Data to VeloCloud Support**
 When enabled, VeloCloud Support will be able to view configuration passwords in plaintext.

☒ **Grant User Management Access to VeloCloud Support**
 When enabled, VeloCloud Support will be able to assist in user management including creating users, resetting password, etc. VeloCloud Support will be granted access to view all user identifiable information.

Enforce PCI: ☒ **Enforce PCI Compliance**
 When enabled, all users (including VeloCloud Support) will be unable to access sensitive Customer data including PCAPs, etc.

Contact Information

Contact Name:

Contact Email:

Phone:

Mobile: (201) 555-0123

Direct Address:

City:

State:

ZIP/Postcode:

Country:

Zero Touch Provisioning Sign Up

In order to use push fulfillment you should include the SID entered above in your VMware transactions for SDWAN. This will allow VMware to track your inventory correctly. If any issues are discovered with inventory please contact support with the order number, copy of order ack and the serial numbers shipped to you but are not showing up in your inventory. Please note only inventory shipped after you sign up for Zero Touch Provisioning will show on your account.

☒ **SID**

General Information

Option	Description
Name	The existing username is displayed. If required, you can modify the name.
Account Number	<p>The existing account number is displayed. If required, you can modify the number.</p> <hr/> <p>Note This option is available only for Operator and Partner users.</p> <hr/>
Domain	The existing domain name is displayed and you can modify the domain, if required.
Description	Enter a description for the customer.
Enable Two Factor Authentication	<p>Select the checkbox to enable two-factor authentication with SMS for Operators, MSP, and Enterprises. You can enable authentication at the Customer/MSP level or at the Operator level.</p> <p>Ensure that you have provided valid mobile numbers for all admin users before enabling two-factor authentication. You can enter the mobile numbers by selecting the users in the Administration > Administrators screen. See Also Manage Admin Users.</p>
Require Two Factor Authentication	Select the checkbox to mandate the user login with two-factor authentication. After enabling the two-factor authentication, when you try to login with your user credentials, you also need to enter the six-digit pin that you receive as SMS in your mobile.
Enable Self Service Password Reset	<p>By default, this option is selected, which enables you to reset your password in the login page of the Orchestrator.</p> <p>When you try to reset your password in the login page, you are prompted to enter a username. Ensure that you enter a valid email address as the username. Once you submit the username, you receive an email with a link to reset the password. Click the link to setup a new password.</p>
Require Two Factor Authentication for Password Reset	<p>Select this option to enable two-factor authentication to reset your password. You can select this checkbox only when the Enable Two Factor Authentication option is already selected.</p> <p>If this option is enabled, when you try to reset your password in the Login page of the Orchestrator, you are redirected to an Authentication page. The Authentication page prompts you to enter the one-time code that you receive as SMS in your mobile. After validating the code, you are redirected to the Password page to setup a new password.</p>
Enable Pre-Notifications	Select the checkbox to enable pre-notification alerts.

Option	Description
Enable Alerts	Select the checkbox to enable the alerts. You can configure the alert types using the Chapter 24 Configure Alerts option.
Default Edge Authentication	<p>Choose the default option to authenticate the Edges associated to the customer, from the drop-down list.</p> <ul style="list-style-type: none"> ■ Certificate Not Required : Edge uses a pre-shared key mode of authentication. ■ Certificate Acquire: This option is selected by default, and instructs the Edge to acquire a certificate from the certificate authority of the SD-WAN 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 SD-WAN Orchestrator and for establishment of VCMP tunnels. <hr/> <p>Note After acquiring the certificate, the option can be updated to Certificate Required.</p> <hr/> <ul style="list-style-type: none"> ■ Certificate Required: Edge uses the PKI certificate. Operators can change the certificate renewal time window for Edges using system properties. For more information, contact your Operator.

Edge Configuration

Choose the following options to communicate the updates to the Edge configurations to an Edge:

- **Enabled**– Select this option to communicate the configuration updates to an Edge during the next heartbeat. The changes in the configuration may restart the software in the corresponding Edge. By default, this option is selected.
- **Enabled on Orchestrator Upgrade** – Select this option to communicate the updates in the configurations to the Edges when the Orchestrator is upgraded. This may restart the software in the corresponding Edges.

Privacy Settings

- **Support Access** – Choose the following options to grant access to the Support team.
 - **Grant Access to VeloCloud Support** – Select this option to grant access to the VMware Support to view, configure, and troubleshoot the Edges connected to the customer. This option also allows VMware Support to build and customize roles for the customer. For security reasons, the Support cannot access or view the user identifiable information.
 - **Grant View Sensitive Data to VeloCloud Support** – Allows VMware Support to view the configuration passwords in plain text.

- **Grant User Management Access to VeloCloud Support** – Select this option to enable the VMware Support to assist in user management. The user management includes options to create users, reset password, and configure other settings. In this case, the Support has access to user identifiable information.
- **Enforce PCI** – Select this option to prevent operations that are disallowed for PCI compliance reasons. Currently the only operation this option prevents is the ability to request PCAP Diagnostic Bundles from the Edge.

Contact Information

The existing contact details are displayed in this section. If required, you can modify the details.

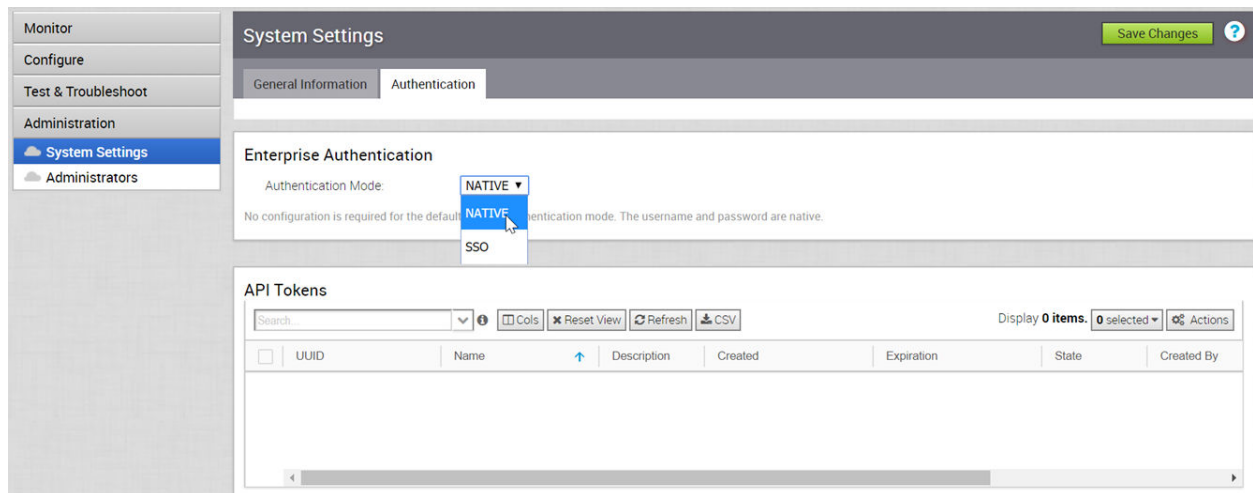
Zero Touch Provisioning Sign Up

You can sign-up for Zero Touch Provisioning by entering the Subscription Identifier in the **SID** field. For more information, see [Activate SD-WAN Edges Using Zero Touch Provisioning](#).

Configure Enterprise Authentication

In the **Authentication** tab, you can setup the authentication mode for the enterprises and view the existing API tokens.

In the Enterprise portal, click **Administration > System Settings > Authentication** to configure the following:



Enterprise Authentication

Choose one of the following from the **Authentication Mode**:

- **NATIVE** – This is the default authentication mode and you can login to the Enterprise with the native username and password. This mode does not require any configuration.

- **SSO** – Single Sign On (SSO) is a session and user authentication service that allows the users to log into the Enterprise with one set of login credentials to access multiple applications. For more information, see [Overview of Single Sign On](#) and [Configure Single Sign On for Enterprise User](#).

API Tokens

You can access the Orchestrator APIs using token-based authentication, irrespective of the authentication mode. You can view the existing API tokens in this section.

The Operator Super User or the User associated with an API token can revoke the token. Select the token and click **Actions > Revoke**. To create and download the API tokens, see [API Tokens](#).

Overview of Single Sign On

The SD-WAN Orchestrator supports a new type of user authentication called Single Sign On (SSO) for all Orchestrator user types: Operator, Partner, and Enterprise.

Single Sign On (SSO) is a session and user authentication service that allows SD-WAN Orchestrator users to log in to the SD-WAN Orchestrator with one set of login credentials to access multiple applications. Integrating the SSO service with SD-WAN Orchestrator improves the security of user authentication for SD-WAN Orchestrator users and enables SD-WAN Orchestrator to authenticate users from other OpenID Connect (OIDC)-based Identity Providers (IDPs). The following IDPs are currently supported:

- Okta
- OneLogin
- PingIdentity
- AzureAD
- VMwareCSP

Configure Single Sign On for Enterprise User

To setup Single Sign On (SSO) authentication for Enterprise user, perform the steps in this procedure.

Prerequisites

- Ensure that you have the Enterprise super user permission.
- Before setting up the SSO authentication, ensure you have set up roles, users, and OpenID connect (OIDC) application for SD-WAN Orchestrator in your preferred identity provider's website. For more information, see [Configure an IDP for Single Sign On](#).

Procedure

- 1 Log in to a SD-WAN Orchestrator application as Enterprise super user, with your login credentials.

2 Click **Administration > System Settings**

The **System Settings** screen appears.

System Settings Save Changes ?

General Information

Name: mars
 Account Number: MAR-UT7C6P4
 Domain: mars
 Description:

Enable Two Factor Authentication: ☐
 Require Two Factor Authentication: ☐
 Enable Self Service Password Reset: ☒
 Require Two Factor Authentication for Password Reset: ☐
 Enable Pre-Notifications: ☒
 Enable Alerts: ☒
 Default Edge Authentication: Certificate Acquire

Edge Configuration

Updates: ☒ **Enabled**
When Enabled is on (normal), Edge configuration updates are communicated to an Edge on its next heartbeat. Configuration changes may cause Edge software restart. When the Orchestrator is upgraded, Enabled is set to the value of Enabled on Orchestrator Upgrade.
☐ **Enabled on Orchestrator Upgrade**
Orchestrator upgrade may update Edge configurations, which when communicated to the Edge may cause software restart. Setting Enabled on Orchestrator Upgrade off allows the operator to choose when after Orchestrator upgrade to set Enabled to resume having Edge configuration updates communicated to Edges.

Privacy Settings

Support Access: ☒ **Grant Access to VeloCloud Support**
When activated, VeloCloud Support will be granted access to view, configure and troubleshoot this Customer's Edges. It will also allow VeloCloud Support to build and customize roles for the customer. As a security consideration, VeloCloud Support will not be granted access to view user identifiable information.
☒ **Grant View Sensitive Data to VeloCloud Support**
When enabled, VeloCloud Support will be able to view configuration passwords in plaintext.
☒ **Grant User Management Access to VeloCloud Support**
When enabled, VeloCloud Support will be able to assist in user management including creating users, resetting password, etc. VeloCloud Support will be granted access to view all user identifiable information.
 Enforce PCI: ☒ **Enforce PCI Compliance**
When enabled, all users (including VeloCloud Support) will be unable to access sensitive Customer data including PCAPs, etc.

Contact Information

Contact Name:
 Contact Email:
 Phone:
 Mobile: (201) 555-0123
 Street Address:
 City:
 State:
 ZIP/Postcode:
 Country:

Zero Touch Provisioning Sign Up

In order to use push fulfillment you should include the DID entered above in your VMware transactions for SDWAN. This will allow VMware to track your inventory correctly. If any issues are discovered with inventory please contact support with the order number, copy of order ack and the serial numbers shipped to you but are not showing up in your inventory. Please note only inventory shipped after you sign up for Zero Touch Provisioning will show on your account.

DID:
 Submit

3 Click the **General Information** tab and in the **Domain** text box, enter the domain name for your enterprise, if it is not already set.

Note To enable SSO authentication for the SD-WAN Orchestrator, you must set up the domain name for your enterprise.

- Click the **Authentication** tab and from the **Authentication Mode** drop-down menu, select **Single Sign-On**.

- From the **Identity Provider template** drop-down menu, select your preferred Identity Provider (IDP) that you have configured for Single Sign On.

Note If you select VMwareCSP as your preferred IDP, ensure to provide your Organization ID in the following format: `/csp/gateway/am/api/orgs/<full organization ID>`.


When you sign in to [VMware CSP console](#), you can view the organization ID you are logged into by clicking on your username. A shortened version of the ID is displayed under the organization name. Click the ID to display the full organization ID.

You can also manually configure your own IDPs by selecting **Others** from the **Identity Provider template** drop-down menu.

- In the **OIDC well-known config URL** text box, enter the OpenID Connect (OIDC) configuration URL for your IDP. For example, the URL format for Okta will be: `https://{oauth-provider-url}/.well-known/openid-configuration`.
- The SD-WAN Orchestrator application auto-populates endpoint details such as Issuer, Authorization Endpoint, Token Endpoint, and User Information Endpoint for your IDP.
- In the **Client Id** text box, enter the client identifier provided by your IDP.
- In the **Client Secret** text box, enter the client secret code provided by your IDP, that is used by the client to exchange an authorization code for a token.

10 To determine user's role in SD-WAN Orchestrator, select one of the options:

- **Use Default Role** – Allows user to configure a static role as default by using the **Default Role** text box that appears on selecting this option. The supported roles are: Enterprise Superuser, Enterprise Standard Admin, Enterprise Support, and Enterprise Read Only.



Note In an SSO configuration setup, if **Use Default Role** option is selected and a default user role is defined, then all the SSO user will be assigned the specified default role. Instead of assigning a user with the default role, a Standard Administrator Super User or Standard Administrator can pre-register a specific user as a Non-Native user and define a specific user role by clicking the **Administration > Administrators** tab in the Enterprise portal. For steps to configure a new Administrator User, see [Create New Admin User](#).

- **Use Identity Provider Roles** – Uses the roles set up in the IDP.
- 11 On selecting the **Use Identity Provider Roles** option, in the **Role Attribute** text box, enter the name of the attribute set in the IDP to return roles.
 - 12 In the **Role Map** area, map the IDP-provided roles to each of the Enterprise user roles, separated by using commas.

Roles in VMware CSP will follow this format: *external/<service definition uuid>/<service role name mentioned during service template creation>*.
 - 13 Update the allowed redirect URLs in OIDC provider website with SD-WAN Orchestrator URL (<https://<Orchestrator URL>/login/ssologin/openidCallback>).
 - 14 Click **Save Changes** to save the SSO configuration.
 - 15 Click **Test Configuration** to validate the entered OpenID Connect (OIDC) configuration.

The user is navigated to the IDP website and allowed to enter the credentials. On IDP verification and successful redirect to SD-WAN Orchestrator test call back, a successful validation message will be displayed.

Results

The SSO authentication setup is complete.

What to do next

[Chapter 5 Log in to VMware SD-WAN Orchestrator Using SSO for Enterprise User.](#)

Configure an IDP for Single Sign On

To enable Single Sign On (SSO) for SD-WAN Orchestrator, you must configure an Identity Partner (IDP) with details of SD-WAN Orchestrator. Currently, the following IDPs are supported: Okta, OneLogin, PingIdentity, AzureAD, and VMware CSP.

For step-by-step instructions to configure an OpenID Connect (OIDC) application for SD-WAN Orchestrator in various IDPs, see:

- [Configure Okta for Single Sign On](#)
- [Configure OneLogin for Single Sign On](#)
- [Configure PingIdentity for Single Sign On](#)
- [Configure Azure Active Directory for Single Sign On](#)
- [Configure VMware CSP for Single Sign On](#)

Configure Okta for Single Sign On

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

Prerequisites

Ensure you have an Okta account to sign in.

Procedure

- 1 Log in to your [Okta](#) account as an Admin user.

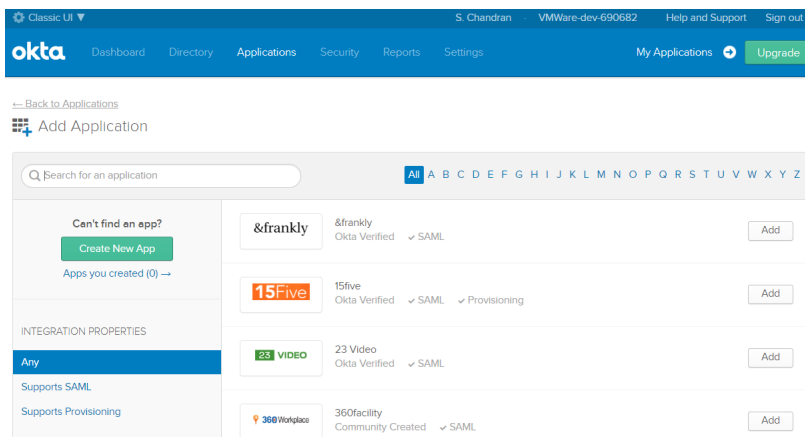
The **Okta** home screen appears.

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

- 2 To create a new application:

- a In the upper navigation bar, click **Applications > Add Application**.

The **Add Application** screen appears.




- b Click **Create New App**.

The **Create a New Application Integration** dialog box appears.


- c From the **Platform** drop-drop menu, select **Web**.
- d Select **OpenID Connect** as the Sign on method and click **Create**.

The **Create OpenID Connect Integration** screen appears.


 Create OpenID Connect Integration


GENERAL SETTINGS

Application name

Application logo (Optional) 

CONFIGURE OPENID CONNECT

Login redirect URIs 

Logout redirect URIs 

- e Under the **General Settings** area, in the **Application name** text box, enter the name for your application.
- f Under the **CONFIGURE OPENID CONNECT** area, in the **Login redirect URIs** text box, enter the redirect URL that your SD-WAN Orchestrator application uses as the callback endpoint.

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

- g Click **Save**. The newly created application page appears.

- h On the **General** tab, click **Edit** and select **Refresh Token** for Allowed grant types, and click **Save**.

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

The screenshot displays the configuration interface for an application in VMware SD-WAN Orchestrator. At the top, there are three tabs: **General**, **Sign On**, and **Assignments**. The **General** tab is selected and highlighted with a green underline.

Below the tabs, there are two main configuration panels:

- General Settings**: This panel has an **Edit** button in the top right corner. It is divided into two sections:
 - APPLICATION**: Contains fields for 'Application label' (VMWare SD-WAN VCO), 'Application type' (Web), and 'Allowed grant types'. Under 'Allowed grant types', there are three options: 'Client acting on behalf of itself' (with a checkbox for 'Client Credentials'), 'Client acting on behalf of a user' (with checkboxes for 'Authorization Code' and 'Refresh Token'), and 'Implicit (Hybrid)' (with a checkbox).
 - LOGIN**: Contains fields for 'Login redirect URIs' (https://vco13-usv1.velocloud.net/login/ssologin/openidCallback), 'Logout redirect URIs', 'Login initiated by' (App Only), and 'Initiate login URI' (https://vco13-usv1.velocloud.net/).
- Client Credentials**: This panel also has an **Edit** button in the top right corner. It contains two fields:
 - Client ID**: A text box containing the value '00apekyj5x5c7h5H60h7' and a copy icon. Below it is a note: 'Public identifier for the client that is required for all OAuth flows.'
 - Client secret**: A text box filled with asterisks, with a toggle icon and a copy icon.

- i Click the **Sign On** tab and under the **OpenID Connect ID Token** area, click **Edit**.
- j From the **Groups claim type** drop-down menu, select **Expression**. By default, Groups claim type is set to **Filter**.

- k In the **Groups claim expression** textbox, enter the claim name that will be used in the token, and an Okta input expression statement that evaluates the token.
- l Click **Save**.

The application is setup in IDP. You can assign user groups and users to your SD-WAN Orchestrator application.

General
Sign On
Assignments

Settings

SIGN ON METHODS

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

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

☒ OpenID Connect

Token Credentials

Edit

Signing credential rotation ⓘ Automatic

OpenID Connect ID Token

Edit

Issuer	https://bokf-sandbox.oktapreview.com
Audience	0oapekyj5x5c7h5H60h7
Claims	Claims for this token include all user attributes on the app profile.
Groups claim type	Expression
Groups claim expression ⓘ	groups Groups.startsWith("active_directory", "VCO_", 100) Using Groups Claim


3 To assign groups and users to your SD-WAN Orchestrator application:

- a Go to **Application > Applications** and click on your SD-WAN Orchestrator application link.
- b On the **Assignments** tab, from the **Assign** drop-down menu, select **Assign to Groups** or **Assign to People**.

The **Assign <Application Name> to Groups** or **Assign <Application Name> to People** dialog box appears.

- c Click **Assign** next to available user groups or users you want to assign the SD-WAN Orchestrator application and click **Done**.

The users or user groups assigned to the SD-WAN Orchestrator application will be displayed.



Results

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

What to do next

Configure Single Sign On in SD-WAN Orchestrator.

Create a New User Group in Okta

To create a new user group, perform the steps on this procedure.

Procedure

- 1 Click **Directory > Groups**.
- 2 Click **Add Group**.

The **Add Group** dialog box appears.

- 3 Enter the group name and description for the group and click **Save**.

Create a New User in Okta

To add a new user, perform the steps on this procedure.

Procedure

- 1 Click **Directory > People**.

- 2 Click **Add Person**.

The **Add Person** dialog box appears.

- 3 Enter all the mandatory details such as first name, last name, and email ID of the user.
- 4 If you want to set the password, select **Set by user** from the **Password** drop-down menu and enable **Send user activation email now**.
- 5 Click **Save**.

An activation link email will be sent your email ID. Click the link in the email to activate your Okta user account.

Configure OneLogin for Single Sign On

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

Prerequisites

Ensure you have an OneLogin account to sign in.

Procedure

- 1 Log in to your [OneLogin](#) account as an Admin user.

The **OneLogin** home screen appears.

2 To create a new application:

- a In the upper navigation bar, click **Apps > Add Apps**.
- b In the **Find Applications** text box, search for “OpenId Connect” or “oidc” and then select the **OpenId Connect (OIDC)** app.

The **Add OpenId Connect (OIDC)** screen appears.

The screenshot displays the OneLogin 'Add OpenId Connect (OIDC)' configuration interface. The top navigation bar includes 'onelogin', 'Users', 'Applications', 'Devices', 'Authentication', 'Activity', 'Security', 'Settings', and 'Developers'. The 'Applications' tab is active. The main content area is titled 'Add OpenId Connect (OIDC)' and includes a 'Cancel' button and a 'Save' button. The configuration section is divided into two main parts: 'Portal' and 'Description'. The 'Portal' section contains a 'Display Name' text box with the value 'OpenId Connect (OIDC)', a 'Visible in portal' toggle switch that is turned on, and two icon upload options: 'Rectangular Icon' and 'Square Icon'. Below these options are instructions for uploading icons. The 'Description' section contains a text area with a 200-character limit.

- c In the **Display Name** text box, enter the name for your application and click **Save**.

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

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

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

- h From the **Authentication Method** drop-down menu, select **POST** as the Token Endpoint and click **Save**.

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

- i On the **Access** tab, choose the roles that will be allowed to login and click **Save**.

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

Results

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

What to do next

Configure Single Sign On in SD-WAN Orchestrator.

Create a New Role in OneLogin

To create a new role, perform the steps on this procedure.

Procedure

- 1 Click **Users > Roles**.

- 2 Click **New Role**.

- 3 Enter a name for the role.

When you first set up a role, the **Applications** tab displays all the apps in your company catalog.

- 4 Click an application to select it and click **Save** to add the selected apps to the role.

Create a New User in OneLogin

To create a new user, perform the steps on this procedure.

Procedure

- 1 Click **Users > Users > New User**.

The **New User** screen appears

- 2 Enter all the mandatory details such as first name, last name, and email ID of the user and click **Save User**.

Configure PingIdentity for Single Sign On

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

Prerequisites

Ensure you have a PingOne account to sign in.

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


Procedure

- 1 Log in to your [PingOne](#) account as an Admin user.

The **PingOne** home screen appears.


2 To create a new application:

- a In the upper navigation bar, click **Applications**.



- b On the **My Applications** tab, select **OIDC** and then click **Add Application**.

The **Add OIDC Application** pop-up window appears.



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

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

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

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

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

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

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

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

- j Under **Group Access**, select all user groups that should have access to your SD-WAN Orchestrator application and click **Done**.

The application will be added to your account and will be available in the **My Application** screen.

Results

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

What to do next

Configure Single Sign On in SD-WAN Orchestrator.

Create a New User Group in PingIdentity

To create a new user group, perform the steps on this procedure.

Procedure

- 1 Click **Users > User Directory**.
- 2 On the **Groups** tab, click **Add Group**
The **New Group** screen appears.
- 3 In the **Name** text box, enter a name for the group and click **Save**.

Create a New User in PingIdentity

To add a new user, perform the steps on this procedure.

Procedure

- 1 Click **Users > User Directory**.
- 2 On the **Users** tab, click the **Add Users** drop-down menu and select **Create New User**.
The **User** screen appears.
- 3 Enter all the mandatory details such as username, password, and email ID of the user.
- 4 Under **Group Memberships**, click **Add**.
The **Add Group Membership** pop-up window appears.
- 5 Search and add the user to a group and click **Save**.

Configure Azure Active Directory for Single Sign On

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

Prerequisites

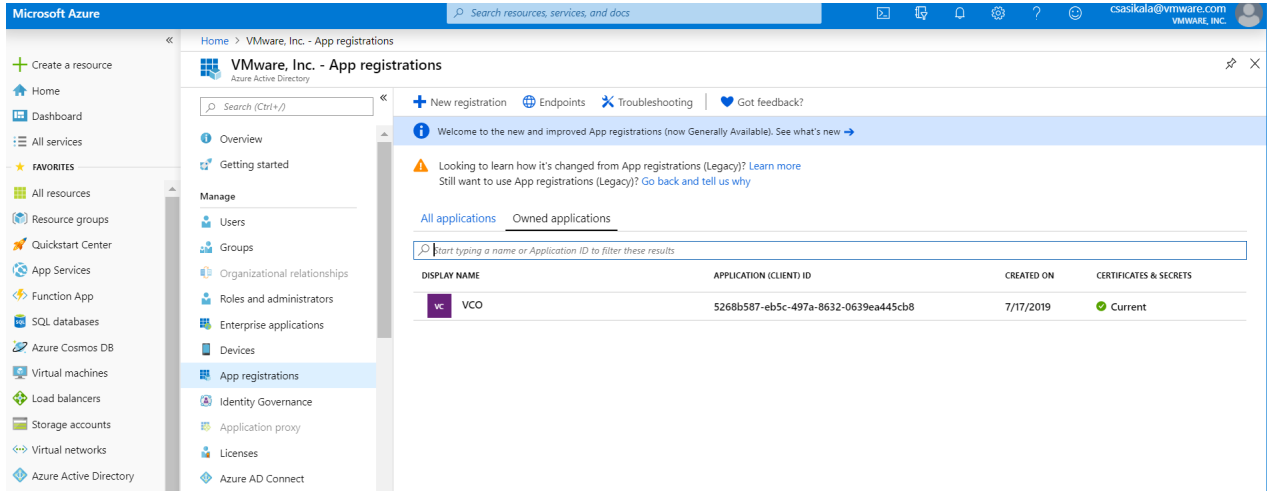
Ensure you have an AzureAD account to sign in.

Procedure

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

2 To create a new application:

- a Search and select the **Azure Active Directory** service.



- b Go to **App registration > New registration**.

The **Register an application** screen appears.

Register an application

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

Supported account types
Who can use this application or access this API?
☒ Accounts in this organizational directory only (Velocloud Networks, inc@velo)
☐ Accounts in any organizational directory
☐ Accounts in any organizational directory and personal Microsoft accounts (e.g. Skype, Xbox, Outlook.com)
[Help me choose...](#)

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

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

Register

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

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

- e Click **Register**.

Your SD-WAN Orchestrator application will be registered and displayed in the **All applications** and **Owned applications** tabs. Make sure to note down the Client ID/ Application ID to be used during the SSO configuration in SD-WAN Orchestrator.

- f Click **Endpoints** and copy the well-known OIDC configuration URL to be used during the SSO configuration in SD-WAN Orchestrator.
- g To create a client secret for your SD-WAN Orchestrator application, on the **Owned applications** tab, click on your SD-WAN Orchestrator application.
- h Go to **Certificates & secrets > New client secret**.

The **Add a client secret** screen appears.

- i Provide details such as description and expiry value for the secret and click **Add**.

The client secret will be created for the application. Note down the new client secret value to be used during the SSO configuration in SD-WAN Orchestrator.


- j To configure permissions for your SD-WAN Orchestrator application, click on your SD-WAN Orchestrator application and go to **API permissions > Add a permission**.

The **Request API permissions** screen appears.

- k Click **Microsoft Graph** and select **Application permissions** as the type of permission for your application.
- l Under **Select permissions**, from the **Directory** drop-down menu, select **Directory.Read.All** and from the **User** drop-down menu, select **User.Read.All**.
- m Click **Add permissions**.

- n To add and save roles in the manifest, click on your SD-WAN Orchestrator application and from the application **Overview** screen, click **Manifest**.

A web-based manifest editor opens, allowing you to edit the manifest within the portal. Optionally, you can select **Download** to edit the manifest locally, and then use **Upload** to reapply it to your application.



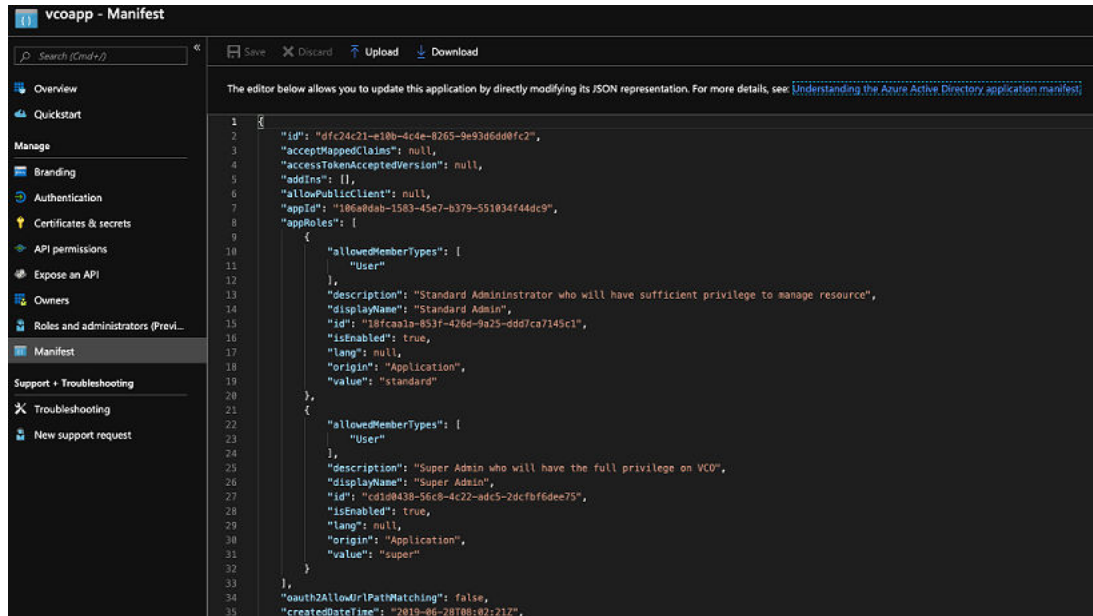
- o In the manifest, search for the `appRoles` array and add one or more role objects as shown in the following example and click **Save**.

Sample role objects

```
{
  "allowedMemberTypes": [
    "User"
  ],
  "description": "Standard Administrator who will have sufficient privilege
to manage resource",
  "displayName": "Standard Admin",
  "id": "18fcaa1a-853f-426d-9a25-ddd7ca7145c1",
  "isEnabled": true,
  "lang": null,
  "origin": "Application",
  "value": "standard"
},
{
  "allowedMemberTypes": [
    "User"
  ],
  "description": "Super Admin who will have the full privilege on SD-WAN
Orchestrator",
  "displayName": "Super Admin",
  "id": "cd1d0438-56c8-4c22-adc5-2dcfbf6dee75",
  "isEnabled": true,
  "lang": null,
  "origin": "Application",
  "value": "superuser"
}
```


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

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



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

Results

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

What to do next

Configure Single Sign On in SD-WAN Orchestrator.

Create a New Guest User in AzureAD

To create a new guest user, perform the steps on this procedure.

Procedure

- 1 Go to **Azure Active Directory > Users > All users**.
- 2 Click **New guest user**.

The **New Guest User** pop-up window appears.

- 3 In the **Email address** text box, enter the email address of the guest user and click **Invite**.

The guest user immediately receives a customizable invitation that lets them to sign into their Access Panel.

- 4 Guest users in the directory can be assigned to apps or groups.

Configure VMware CSP for Single Sign On

To configure VMware Cloud Services Platform (CSP) for Single Sign On (SSO), perform the steps on this procedure.

Prerequisites

Sign in to [VMware CSP console](#) (staging or production environment) with your VMware account ID. If you are new to VMware Cloud and do not have a VMware account, you can create one as you sign up. For more information, see How do I Sign up for VMware CSP section in [Using VMware Cloud](#) documentation.

Procedure

- 1 Contact the VMware Support Provider for receiving a Service invitation URL link to register your SD-WAN Orchestrator application to VMware CSP. For information on how to contact the Support Provider, see <https://kb.vmware.com/s/article/53907> and https://www.vmware.com/support/contacts/us_support.html.

The VMware Support Provider will create and share:

- a Service invitation URL that needs to be redeemed to your Customer organization
- a Service definition uuid and Service role name to be used for Role mapping in Orchestrator

- 2 Redeem the Service invitation URL to your existing Customer Organization or create a new Customer Organization by following the steps in the UI screen.

You need to be an Organization Owner to redeem the Service invitation URL to your existing Customer Organization.

- 3 After redeeming the Service invitation, when you sign in to [VMware CSP console](#), you can view your application tile under **My Services** area in the **VMware Cloud Services** page.

The Organization you are logged into is displayed under your username on the menu bar. Make a note of the Organization ID by clicking on your username, to be used during Orchestrator configuration. A shortened version of the ID is displayed under the Organization name. Click the ID to display the full Organization ID.

- 4 Log in to [VMware CSP console](#) and create an OAuth application. For steps, see [Use OAuth 2.0 for Web Apps](#). Make sure to set Redirect URI to the URL displayed in **Configure Authentication** screen in Orchestrator.

Once OAuth application is created in VMware CSP console, make a note of IDP integration details such as Client ID and Client Secret. These details will be needed for SSO configuration in Orchestrator.

- 5 Log in to your SD-WAN Orchestrator application as Super Admin user and configure SSO using the IDP integration details as follows.
 - a Click **Administration > System Settings**
The **System Settings** screen appears.
 - b Click the **General Information** tab and in the **Domain** text box, enter the domain name for your enterprise, if it is not already set.

Note To enable SSO authentication for the SD-WAN Orchestrator, you must set up the domain name for your enterprise.

- c Click the **Authentication** tab and from the **Authentication Mode** drop-down menu, select **SSO**.
 - d From the **Identity Provider template** drop-down menu, select **VMwareCSP**.
 - e In the **Organization Id** text box, enter the Organization ID (that you have noted down in Step 3) in the following format: `/csp/gateway/am/api/orgs/<full organization ID>`.
 - f In the **OIDC well-known config URL** text box, enter the OpenID Connect (OIDC) configuration URL (<https://console.cloud.vmware.com/csp/gateway/am/api/.well-known/openid-configuration>) for your IDP.
- The SD-WAN Orchestrator application auto-populates endpoint details such as Issuer, Authorization Endpoint, Token Endpoint, and User Information Endpoint for your IDP.
- g In the **Client Id** text box, enter the client ID that you have noted down from the OAuth application creation step.
 - h In the **Client Secret** text box, enter the client secret code that you have noted down from the OAuth application creation step.
 - i To determine user's role in SD-WAN Orchestrator, select either **Use Default Role** or **Use Identity Provider Roles**.
 - j On selecting the **Use Identity Provider Roles** option, in the **Role Attribute** text box, enter the name of the attribute set in the VMware CSP to return roles.
 - k In the **Role Map** area, map the VMwareCSP-provided roles to each of the SD-WAN Orchestrator roles, separated by using commas.

Roles in VMware CSP will follow this format: `external/<service definition uuid>/<service role name mentioned during service template creation>`. Use the same Service definition uuid and Service role name that you have received from your Support Provider.

- 6 Click **Save Changes** to save the SSO configuration.
- 7 Click **Test Configuration** to validate the entered OpenID Connect (OIDC) configuration.

Configure Authentication Save Changes ?

Operator Authentication

Authentication Mode: SSO

Identity Provider template: VMwareCSP

Organization Id: /csp/gateway/am/api/orgs/d94fb648-cbb3-4863-t

OIDC well-known config URL: https://console-stg.cloud.vmware.com/csp/gateway/am/api/.well-known/op

Issuer: https://gaz-preview.csp-vidm-prod.com

Authorization Endpoint: https://console-stg.cloud.vmware.com/csp/gateway/discovery?orgLink=%2

Token Endpoint: https://console-stg.cloud.vmware.com/csp/gateway/am/api/auth/authorize

User Information Endpoint: https://console-stg.cloud.vmware.com/csp/gateway/am/api/userinfo

Client Id: e1UmTD4TPps0h8vak0UMiOf0HCvMw0MDta

Client Secret:

Scopes: openid

☐ Use Default Role ☒ Use Identity Provider Roles

Role Attribute: perms

Role Map

Operator Superuser: external/1e73b58c-475f-4065-95d8-5f

Operator Standard Admin: external/1e73b58c-475f-4065-95d8-5f

Operator Support: support

Operator Business: business

Remember to set <https://13.52.173.235/login/ssologin/openidCallback> as an allowed redirect URL with your IDP application/client

The user is navigated to the VMware CSP website and allowed to enter the credentials. On IDP verification and successful redirect to SD-WAN Orchestrator test call back, a successful validation message will be displayed.

Results

You have completed integrating SD-WAN Orchestrator application in VMware CSP for SSO and can access the SD-WAN Orchestrator application logging in to the VMware CSP console.

What to do next

- Within the organization, manage users by adding new users and assigning appropriate role for the users. For more information, see the *Identity & Access Management* section in [Using VMware Cloud](#) documentation.

Manage Admin Users

The **Administrators** page displays the existing admin users. Standard Administrator Superusers and Standard Administrators can create new admin users with different role privileges and configure API tokens for each admin user.

In the Enterprise portal, click **Administration > Administrators**.

Username	Name	Last Login	Status	Unlocked	Role	Authentication
admin@test.com	admin	Last Login	Enabled	Unlocked	Superuser	Native

Click **Actions** to perform the following activities:

- **New Admin:** Creates new admin users. See [Create New Admin User](#).
- **Modify Admin:** Modifies the properties of the selected admin user. You can also click the link to the username to modify the properties. See [Configure Admin Users](#).
- **Password Reset:** Sends an Email to the selected user with a link to reset the password.
- **Delete Admin:** Deletes the selected users.

Create New Admin User

Standard Administrator Superusers and Standard Administrators can create new admin users.

In the Enterprise portal, click **Administration > Administrators**.

Procedure

- 1 You can create new admin users by clicking either **New Admin**, or **Actions > New Admin**.

- 2 In the **New Admin** window, enter the following details:

- a Enter the user details like username, password, Name, Email, and Phone numbers.
- 1 The username should be in the format of email address, like user@example.com.
 - 2 The password must meet the following requirements:
 - The number of characters must be in the range of 8 to 32.
 - Must have at least one lower case character.
 - Must have at least one number.

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.

- b If you have chosen the authentication mode as Native in [Configure Enterprise Authentication](#), then the type of the user is selected as Native. If you have chosen a different authentication mode, you can choose the type of the user. If you choose the user to be Non-Native, the password option is not available, as it is inherited from the authentication mode.
- c Select the user role from the **Account Role** drop-down list. Once you select a role, the Network and Security functions of the selected role, along with the description, are displayed.

- 3 Click **Create**.

Results

The user details are displayed in the **Administrators** page.

Configure Admin Users

You can configure additional properties and create API tokens for an Enterprise Admin user.

In the Enterprise portal, go to **Administration > Administrators**. To configure an Admin user, click the link to a username or select the user and click **Actions > Modify Admin**.

The existing properties of the selected Admin user are displayed and if required, you can add or modify the following:

The screenshot displays the VMware SD-WAN Administration portal interface. On the left is a navigation sidebar with options: Monitor, Configure, Test & Troubleshoot, Administration, System Settings, Administrators (selected), Role Customization, Edge Licensing, and Zero Touch Provisioning. The main content area is titled 'Administrators - 5_site_operator@velocloud.net' with a 'Save Changes' button. Below the title, there are sections for 'Status' (Enabled/Not Enabled), 'Type' (Native/Non-Native), 'Properties' (Username, Current Password, New Password, Confirm Password, First Name, Last Name, Contact Email, Phone, Mobile Phone), 'User Role' (Superuser), and 'API Tokens' (a table with columns for UUID, Name, Description, Created, Expiration, State, and Created By).

Status

By default, the status is in **Enabled** state. If you choose **Not Enabled**, the user is logged out of all the active sessions.

Type

If you have chosen the authentication mode as **Native** in the [Configure Enterprise Authentication](#), then the type of the user is selected as **Native**. If you have chosen a different authentication mode, you can choose the type of the user. If you choose the user to be **Non-Native**, then you cannot reset the password or modify the user role.

Properties

The existing details such as name, email-id, telephone number, and mobile number of the user are displayed. If needed, you can modify the user details, set a new password, or reset the existing password.

- To set a new password, you must enter the current password correctly in the **Current Password** textbox and the password to be changed in **New Password** and **Confirm Password** textboxes.
- To reset the existing password, click **Password Reset**. An email is sent to the user with a link to reset the password.

Role

The existing type of the user role is displayed. If required, you can choose a different role for the user. The role privileges change accordingly.

API Tokens

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

For Enterprise Read Only Users and MSP Business Specialist users, token-based authentication is not enabled.

Configure API Tokens

Any user can create tokens based on the privileges they have been assigned to their user roles, except the Enterprise Read-Only users and MSP Business Specialist users.

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

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

Manage API Tokens

- In the **API Tokens** section, click **Actions > New API Token**, to create a new token.
- In the **New API Token** window, enter a **Name** and **Description** for the token, and choose the **Lifetime** from the drop-down menu.

The screenshot shows a 'New API Token' window. It has three main input areas: a text field for 'Name' containing '5_Site_API', a text area for 'Description' containing 'To access API tokens', and a dropdown menu for 'Lifetime (in months)' set to '12'. At the bottom right, there are two buttons: 'Create' (highlighted in green) and 'Cancel'.

- Click **Create** and the new token is displayed in the **API Tokens** grid.
- Initially, the status of the token is displayed as **Pending**. To download the token, select the token, and click **Actions > Download API Token**. The status changes to **Enabled**, which means that the API token can be used for API access.
- To deactivate a token, select the token and click **Actions > Revoke API Token**. The status of the token is displayed as **Revoked**.
- When the Lifetime of the token is over, the status changes to **Expired** state.

Only the user who is associated with a token can download it and after downloading, the ID of the token alone is displayed. You can download a token only once.

After downloading the token, the user can send it as part of the Authorization Header of the request to access the Orchestrator API.

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

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

After modifying the settings and API Tokens, click **Save Changes**.

Roles

The Orchestrator consists of different types of roles with different set of privileges.

The roles can be categorized as follows:

- **Functional Roles** – Defined as a set of privileges relevant to a functionality. These privileges are used to carry a certain business process. For more information, see [Functional Roles](#).
- **Composite Roles** – The functional roles from different categories can be grouped to form a composite role. For more information, see [Composite Roles](#).

Functional Roles

Functional Roles are defined as a set of privileges relevant to a functionality.

A functional role can be tagged to one or more of the following services: SD-WAN, Secure Access, Cloud Web Security. These are the group of privileges required by a user to carry a certain business process. For example, a Customer support role in SD-WAN is a functional role required by an SD-WAN user to carry out various support activities. Every service defines such roles based on business functionality that they want to support. These roles are categorized as Network and Security functional roles.

By default, the Orchestrator consists of different functional roles that consist of role privileges based on the requirements. If required, you can customize the role privileges of the functional roles. For more information, see [Role Customization](#).

Composite Roles

Composite roles are a group of functional roles combined from different functional categories.

By default, the following composite roles are available:

Composite Role	Network Functional Role	Security Functional Role
Enterprise Standard Admin	Network Enterprise Admin	Security Enterprise Admin
Enterprise Superuser	Orchestrator Enterprise Superuser	Security Enterprise Superuser
Enterprise Support	Network Enterprise Support	
Enterprise Read Only	Network Enterprise Read Only	
Enterprise Security Admin	Network Enterprise Read Only	Security Enterprise Admin
Enterprise Security Read Only	Network Enterprise Read Only	Security Enterprise Read Only
Enterprise Network Admin	Network Enterprise Admin	Security Enterprise Read Only

You can assign the above roles to a user, while creating a new Enterprise user. See [Create New Admin User](#).

You can also map the composite role while configuring Single Sign on. See [Configure Single Sign On for Enterprise User](#).

To view the existing composite roles along with the description , see [Manage Composite Roles](#).

To create a custom composite role , see [Create New Composite Roles](#).

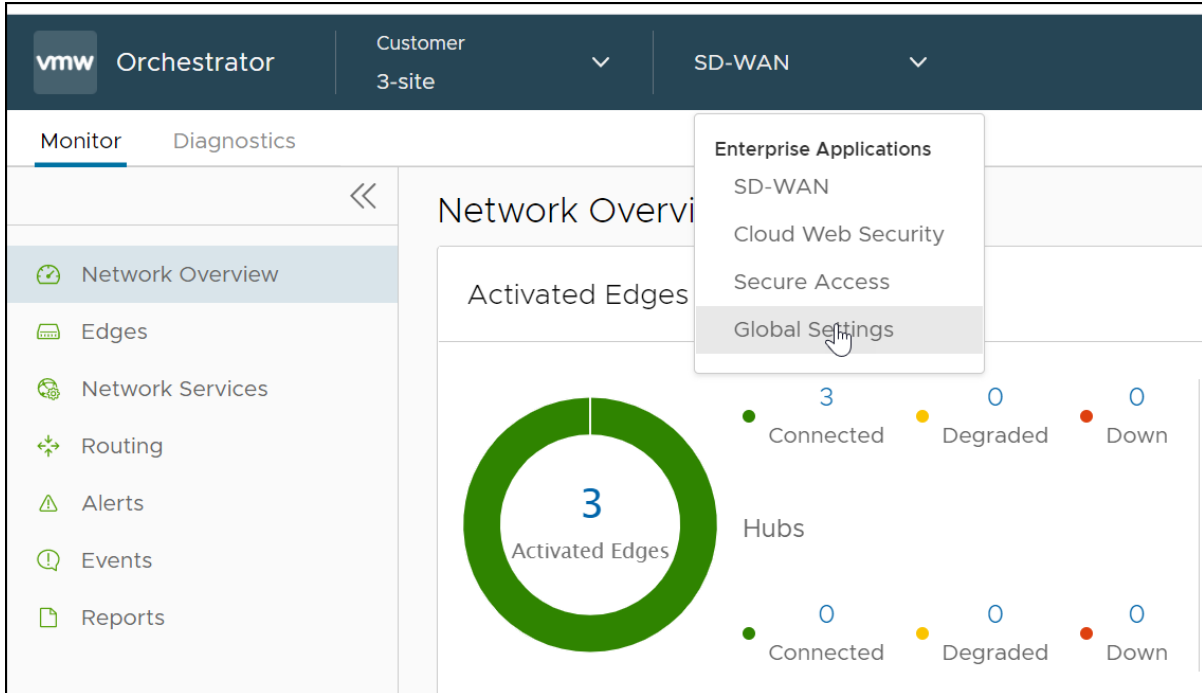
You can also customize the role privileges of the functional roles. For more information, see [Role Customization](#).

Manage Composite Roles

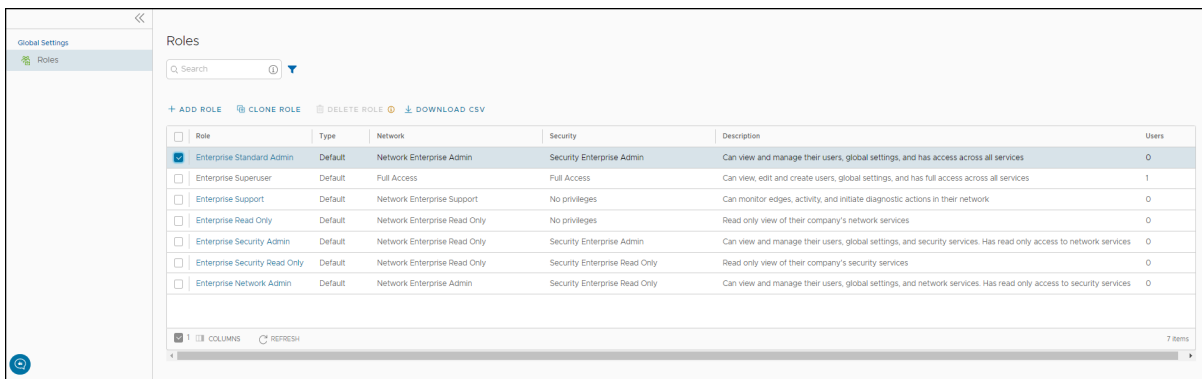
The Orchestrator consists of different functional roles. You can combine the functional roles from these groups to create a composite role.

You can access the composite roles as follows:

- In the Orchestrator portal, click the **Open New Orchestrator UI** option available at the top of the Window.
- Click **Launch New Orchestrator UI** in the pop-up window. The UI opens in a new tab.
- The existing list of customers is displayed in the **Customers** tab. Click the link to a Customer to navigate to the Enterprise portal.
- In the Enterprise portal, click **Enterprise Applications > Global Settings**.



- The **Roles** window opens showing the list of existing roles for the selected Enterprise.



Note You can add, edit, or view the Composite roles only for an Enterprise user.

- You can perform the following activities in the **Roles** window:
 - **Add Role** - Creates a new custom role. See [Create New Composite Roles](#).

- **Clone Role** – Creates a new custom role, by cloning the existing settings from the selected role. You cannot clone the settings of a Super user.
- **Delete Role** – Deletes the selected role. You can delete only custom composite roles. If the role is associated with any user, ensure that you have removed all the users associated with the selected role, before deleting the role.
- **Download CSV** – Downloads the details of the user roles in to a file in CSV format.
- In addition, you can click the link to an existing role to view more details. For custom roles, the **Edit Role** option is available. You cannot edit the default roles. In addition, you cannot edit or view the settings of a Super user.

Create New Composite Roles

To create a new composite role:

Procedure

- 1 Navigate to the **Global Settings** as described in [Manage Composite Roles](#).
- 2 In the **Roles** window, click **Add Role**.
- 3 In the **New Role** window that opens, enter the details for the new custom role as follows:

The screenshot displays the 'New Role' configuration window for a custom role named 'ACME_Admin'. The interface is divided into several sections:

- Global Settings / Roles**: A sidebar on the left with a 'Roles' link.
- Roles / ACME_Admin**: The main header area showing the role name and its type as 'Custom'.
- Role Details**:
 - Role Name**: A text field containing 'ACME_Admin'.
 - Role Description**: A text area containing 'Can manage both the Network and Security features'.
 - Template**: A dropdown menu with a 'SELECT OPTION' button.
- Composite Role Creation**:
 - A descriptive text: 'A composite role is a combination of different functional roles.' and 'The user access privileges to features and services are determined by the user access privileges to features and services.'
 - Network**: A section with a 'Functional Role' list:
 - ☒ **No privileges**: Cannot access any network related features.
 - ☐ **Network Enterprise Admin**: Can access and modify users, global settings, has full access to view and manage edges on network.
 - ☐ **Network Enterprise Support**: Can monitor edges, activity, and initiate diagnostic actions on the customer's network.
 - ☐ **Network Enterprise Read Only**: Has read-only access to the network.
 - Security**: A section with a 'Functional Role' list:
 - ☒ **No privileges**: Cannot access any security related features.
 - ☐ **Security Enterprise Read Only**: Has read-only access to security policies and rules.

Option	Description
Role Name	Enter a name for the new role
Description	Enter a description for the role
Template	Optionally, select an existing role as template from the drop-down list. The functional roles of the selected template are assigned to the new role.
Network	Choose a Network functional role.
Security	Choose a Security function role.

Click **Save**.

Results

The new custom role appears in the **Roles** window. Click the link to the custom role to view the settings. You can click **Edit Role** to modify the settings.

Role Customization

SD-WAN Orchestrator consists of roles with different set of privileges. As an Enterprise Super user, you can assign a pre-defined role to other Enterprise users. Role Customization allows you to customize the existing set of privileges for the Functional roles. The customization is applied to all the users available within the Enterprise.

You can customize only the Functional roles and not the Composite roles. When you customize a Functional role, the changes would impact the Composite roles that consist of the customized Functional role. For more information, see [Functional Roles](#).

Only an Operator super user can enable the Role Customization for an Enterprise super user. If the Role Customization option is not available for you, contact your Operator.

In the Enterprise portal, click **Role Customization**

You can perform the following operations:

- **Show Current Privileges** – Displays the current Functional role privileges. You can view the privileges of all the Functional roles and download them in CSV format.
- **New Package** – Enables to create a new package with customized role privileges. See [Create New Customized Package](#).
- **Reset to System Default** – Allows to reset the current role privileges to default settings. Only the customized privileges applied to the Functional roles in the Enterprise portal are reset to the default settings.

Click **Actions** to perform the following activities:

- **Upload Package** – Allows to upload a customized package. See [Upload Customized Package](#).
- **Clone Package** – Enables to create a copy of the selected package.

- **Modify Package** – Enables to edit the customization settings in the selected package. You can also click the link to the package to edit the settings.
- **Delete Package** – Removes the selected package. You cannot delete a package if it is already in use.
- **Apply Package** – Applies the customization available in the selected package to the existing Functional roles. This option modifies the role privileges only at the current level. If there are customizations available at the Enterprise level or a lower level for the same role, then the lower level takes precedence.

You can also click the Download Icon prior to the package name to download the package as a JSON file.

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

Create New Customized Package

You can create a customized package and apply the package to the existing Functional roles in the SD-WAN Orchestrator.

Procedure

- 1 In the Enterprise portal, click **Role Customization**.
- 2 Click **New Package**.

3 In the **Role Customization Package Editor** window, enter the following:

New Package

Name:

Description:

Scope:

Functional Roles	Privilege	Read	Create	Update	Delete
Network Enterprise Admin	Client Device	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Network Enterprise Support	Client User	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Network Enterprise Read Only	Diagnostics	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Security Enterprise Admin	Edge	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Security Enterprise Read Only	Customer	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
	Customer Alert	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
	Customer Authentication	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
	Customer Delegation	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
	Customer Event	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
	Customer Keys	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
	Customer PKI	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
	Customer Profile	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
	Enterprise Role	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
	Customer Token	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
	Customer User	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
	License	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
	Network Addressing	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
	Network Service	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

CSV Remove Privileges OK

- Enter a **Name** and a **Description** for the new custom package.
- In the **Roles** pane, select a Functional role and click **Remove Privileges** to customize the privileges for the selected role.

Note You can only add or remove Deny Privileges, that is take away privileges from the system default. You cannot grant additional privileges to a role using this option.

In the **Assign Privileges** window, select the features from the **Available Deny Privileges** and move them to the **Selected Deny Privileges** pane.

Assign Privileges to Network Enterprise Admin

Available Deny Privileges

Multi keyword search (comma separated)

- Update Profile Device VQM Settings
- Update Edge Overview Properties Name
- Update Edge Overview Properties Description
- Update Edge Overview Properties Enable Alerts
- Update Edge Overview Properties Serial Number
- Update Edge Overview Properties Activation Expiration
- Create Edge Overview Properties Activation Email
- Update Customer Alert
- Read Customer Delegation
- Read Customer User
- Read Edge Device Authentication Settings
- Read Edge Device BGP Settings

Selected Deny Privileges

Multi keyword search (comma separated)

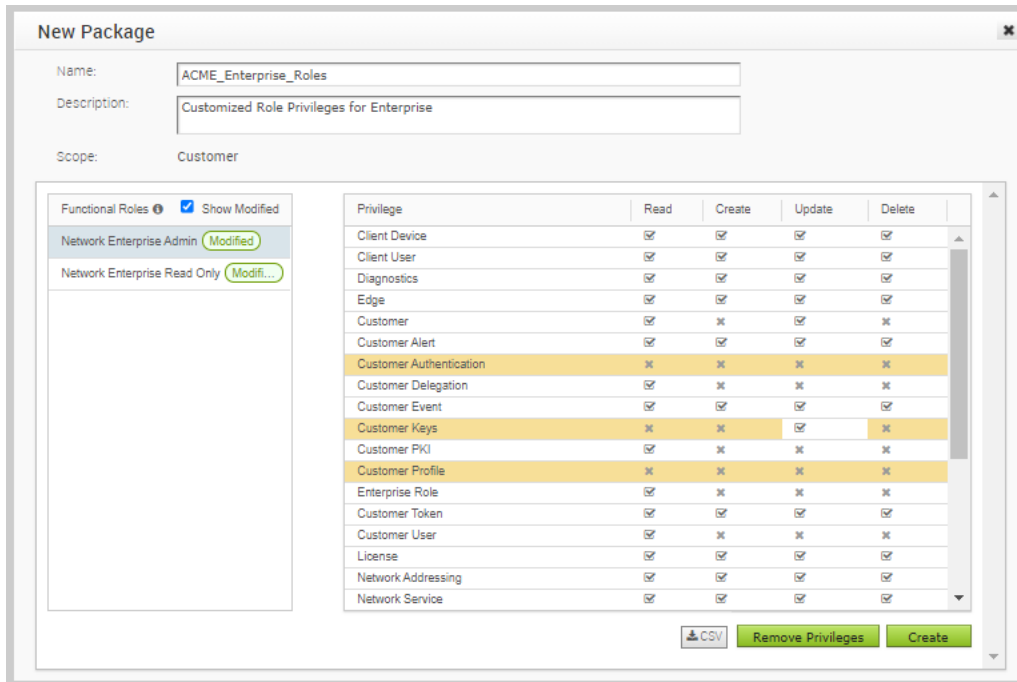
- Delete Customer Profile
- Update Customer Profile
- Create Customer Profile
- Read Customer Profile
- Delete Customer Keys
- Create Customer Keys
- Read Customer Keys
- Manage Customer Authentication
- Read Customer Authentication

OK Cancel

Note You can assign only **Deny** privileges to the Functional roles.

Click **OK**.

- 4 Repeat assigning privileges to the Functional roles in the **Role Customization Package Editor** window.
- 5 Select the **Show Modified** checkbox to filter and view the customized privileges. The changes to the privileges are highlighted in a different color.



- 6 Click **Create**. You can click **CSV** to download the Functional role privileges of selected role, in a CSV format.
- 7 The new package details are displayed in the **Role Customization Packages** window.



- 8 To edit the privileges, click the link to the package or select the package and click **Actions** > **Modify Package**. In the **Role Customization Package Editor** window that opens, add or remove Deny Privileges to the Functional roles in the package and click **OK**.

What to do next

Select the customized package and click **Actions** > **Apply Package** to apply the customization available in the selected package to the existing Functional roles across the SD-WAN Orchestrator.

You can edit the Deny privileges in an applied package whenever required. After modifying the privileges in the **Role Customization Package Editor** window, click **OK** to save and apply the changes to the Functional roles.

Note You can download the customized Functional role privileges as a JSON file and upload the customized package to another Orchestrator. For more information, see [Upload Customized Package](#).

Upload Customized Package

You can upload a package with customized role privileges assigned to different set of Functional roles in the SD-WAN Orchestrator.

You can download the already customized Functional role privileges as a package and upload the package to another Orchestrator.

Procedure

- 1 In the Enterprise portal, click **Role Customization**.
- 2 Click the Download Icon prior to a package name, which downloads the package as a JSON file.
- 3 Navigate to the Orchestrator to which you want to upload the customized package.
- 4 Click **Actions > Upload Package**.
- 5 Choose the JSON file you have downloaded, and the package is uploaded automatically.
- 6 The uploaded package is displayed in the **Role Customization Packages** window.

Monitor

Configure

Test & Troubleshoot

Administration

System Settings

Administrators

Role Customization

Edge Licensing

Role Customization Packages

Show Current Privileges...

New Package...

Reset to System Default...

Color

Reset View

Refresh

Display 2 Items.

0 selected

Actions

<input type="checkbox"/>	Name	InUse	Uploaded	Last Applied	Last Modified
<input type="checkbox"/>	<div><div><div><div></div></div><div>ACME_Enterprise_Roles</div></div><div>Customized role privileges for Enterprise ACME</div></div>	✖	Tue Mar 24, 16:52:54		Tue Mar 24, 16:52:54
<input type="checkbox"/>	<div><div><div><div></div></div><div>Role Customization Package</div></div><div>Tue Mar 24 2020 16:42:34 GMT+0530 (India Standard Time) created Tue Mar 24 2020 16:42:34 GMT+0530 (India Standard Time)</div></div>	✖	Tue Mar 24, 16:51:25		Tue Mar 24, 16:51:25

You can edit Deny privileges in an applied package whenever required. After modifying the privileges in the **Role Customization Package Editor** window, click **OK** to save and apply the changes to the Functional roles.

Monitor Role Customization Events

You can monitor the events related to changes in Role Customization.

In the Enterprise portal, click **Monitor > Events**.

To view the events related to Role Customization, you can use the filter option. Click the drop-down arrow next to the Search option and choose to filter by the Event column. The following events are available for Role Customization:

- Role customization package cloned
- Role customization package deleted
- Role customization package updated
- Role customization package uploaded
- Role customization package was applied
- All role customization packages were removed from the system

The following image shows some of the Role Customization events.

Time	Event	User	Se...	Message
Wed Mar 03, ...	All role customization packages were removed from the system	admin@test.com	Info	All custom roles at ENTERPRISE level (enterpriseld...
Wed Mar 03, ...	Role customization package was applied	admin@test.com	Info	ROLE_CUSTOMIZATION_PACKAGE (Admin User ...
Wed Mar 03, ...	Role customization package updated	admin@test.com	Info	ROLE_CUSTOMIZATION_PACKAGE (Role Custom...
Wed Mar 03, ...	Role customization package deleted	admin@test.com	Info	Test User Roles (version: 1614794000617 type: RO...
Wed Mar 03, ...	Role customization package uploaded	admin@test.com	Info	1614794011278 type (ROLE_CUSTOMIZATION_PA...
Wed Mar 03, ...	Role customization package cloned	admin@test.com	Info	Admin User Roles cloned to package Test User Roles
Wed Mar 03, ...	Role customization package uploaded	admin@test.com	Info	1614794000617 type (ROLE_CUSTOMIZATION_P...
Wed Mar 03, ...	Role customization package uploaded	admin@test.com	Info	1614793916169 type (ROLE_CUSTOMIZATION_P...
Wed Mar 03, ...	Browser enterprise Login	admin@test.com	Info	admin@test.com from [10.104.75.74]

List of Functional Role Privileges

The following table lists all the role privileges available in the Enterprise portal.

The columns in the table indicate the following:

- **Allow Privilege** – Do the roles have allow access?
- **Deny Privilege** – Do the roles have deny access?
- **Customizable** – Is the role privilege available for customization in the Role Customization window?

Navigation Path in the Enterprise Portal	Name of the Tab	Elements in the Tab	Name of the Role Privilege	Description	Allow Privilege	Deny Privilege	Customizable
Monitor > Edges > Select Edge	Overview						
		Top Sources	View Edge Sources	Grants ability to view Monitor Edge Sources tab	Yes	Yes	Yes
			Read Edge	Grants ability to view and manage Edge objects and their properties in general	Yes	No	No
			View User Identifiable Flow Stats	Grants ability to view potentially user identifiable flow source attributes	Yes	Yes	Yes
		Top Applications Top Categories Top Operating Systems Top Sources	Read Edge	Grants ability to view and manage Edge objects and their properties in general	Yes	No	No
			View Flow Stats	Grants ability to view collected flow statistics	Yes	Yes	Yes
	Sources		Read Edge	Grants ability to view and manage Edge objects and their properties in general	Yes	No	No
			View Edge Sources	Grants ability to view Monitor Edge Sources tab	Yes	Yes	Yes
		Devices	View User Identifiable Flow Stats	Grants ability to view potentially user identifiable flow source attributes	Yes	Yes	Yes

Navigation Path in the Enterprise Portal	Name of the Tab	Elements in the Tab	Name of the Role Privilege	Description	Allow Privilege	Deny Privilege	Customizable
			Create Client Device	Controls visibility to unique identifiers (IP or MAC address) of LAN-side client devices	Yes	No	No
			Read Client Device				
		Change Hostname	Update Client Device				
			Delete Client Device				
			Manage Client Device				
		Operating Systems	Create Client User	Controls visibility to potentially Personal Identifiable Information(PII) in flow statistics	Yes	No	No
			Read Client User				
			Update Client User				
			Delete Client User				
			Manage Client User				
	Applications Sources Destinations		Read Edge	Grants ability to view and manage Edge objects and their properties in general	Yes	No	No
			View Flow Stats	Grants ability to view collected flow statistics	Yes	Yes	Yes
	Events from this Edge		Read Customer Event	Grants ability to view customer level events	Yes	No	No
	Remote Actions		Read Remote Actions	Grants access to view and execute remote actions	No	Yes	Yes

Navigation Path in the Enterprise Portal	Name of the Tab	Elements in the Tab	Name of the Role Privilege	Description	Allow Privilege	Deny Privilege	Customizable
	Remote Actions Generate Diagnostic Bundle Remote Diagnostics		Read Diagnostics	Controls creation of and access to diagnostics bundles, both Edge and Gateway. Combine with Edge and Gateway privileges to control access to each type individually	Yes	Yes	Yes
	Generate Diagnostic Bundle		Create Diagnostic Bundle		No	Yes	Yes
	Remote Diagnostics		Read Remote Diagnostics	Privilege granting access to view and execute remote diagnostics	No	Yes	Yes
Monitor	Edges		Read Edge	Grants ability to view and manage Edge objects and their properties in general	Yes	No	No
		Edge Cluster	Read Edge Cluster	Controls the ability to create and configure Edge Clusters	No	Yes	Yes
	Network Services		Read Network Service	Grants ability to view and manage services with the Network Services configuration block	Yes	No	No
		Non SD-WAN Destinations via Gateway Non SD-WAN	Read Customer Event	Grants ability to view customer level events	Yes	No	No

Navigation Path in the Enterprise Portal	Name of the Tab	Elements in the Tab	Name of the Role Privilege	Description	Allow Privilege	Deny Privilege	Customizable
		Destinations via Edge					
		Non SD-WAN Destinations via Gateway	Read Non SD-WAN Destination via Gateway	Grants ability to view and manage Non SD-WAN Destinations via Gateway and Non SD-WAN Destinations via Edge	No	Yes	Yes
		Non SD-WAN Destinations via Edge					
		BGP Gateway Neighbor State	Read Network Service	Grants ability to view and manage services with the Network Services configuration block	Yes	No	No
		BGP Edge Neighbor State	Read Edge	Grants ability to view and manage Edge objects and their properties in general	Yes	No	No
		Edge VNFs	Read VNF Network Service	Grants ability to manage VNF Network Services	No	Yes	Yes
		Edge Cluster	Read Edge Cluster	Controls the ability to create and configure Edge Clusters	No	Yes	Yes
	Routing		Read Network Addressing	Grants ability to view and manage address block configuration in the legacy Network profile mode	Yes	No	No

Navigation Path in the Enterprise Portal	Name of the Tab	Elements in the Tab	Name of the Role Privilege	Description	Allow Privilege	Deny Privilege	Customizable
			Read Edge	Grants ability to view and manage Edge objects and their properties in general	Yes	No	No
			View Customer Routing	Grants ability to view the customer Routing	Yes	No	No
	Alerts		Create Customer Alert	Grants ability to view and manage customer alert configuration and generated alerts	Yes	No	No
			Read Customer Alert			Yes	Yes
			Update Customer Alert				
			Delete Customer Alert			No	No
			Manage Customer Alert				
	Events		Create Customer Event	Grants ability to view customer level events	Yes	No	No
			Read Customer Event				
			Update Customer Event				
			Delete Customer Event				
			Manage Customer Event				
	Reports		Update Customer	Grants ability to view and manage Customers,	Yes	Yes	Yes

Navigation Path in the Enterprise Portal	Name of the Tab	Elements in the Tab	Name of the Role Privilege	Description	Allow Privilege	Deny Privilege	Customizable
			Read Customer	from the Partner or Operator level		No	No
	Firewall	Firewall Logging	View Firewall Logs	Grants ability to view collected firewall logs	Yes	Yes	Yes
			Read Edge	Grants ability to view and manage Edge objects and their properties in general	Yes	No	No
			Read Customer Event	Grants ability to view customer level events	Yes	No	No
Configure > Edges > Select Edge	Edge Overview		Edge Overview	Controls ability to view or modify Edge overview page	No	Yes	Yes
		Properties	Create Edge Overview Properties	Controls ability to view or change items within the properties section of the Edge overview page	No	Yes	Yes
			Read Edge Overview Properties			No	No
			Update Edge Overview Properties			Yes	Yes
			Delete Edge Overview Properties				
		Name	Read Edge Overview Properties Name	Controls ability to view or change Edge name on the Edge overview page	No	Yes	Yes
			Update Edge Overview Properties Name				

Navigation Path in the Enterprise Portal	Name of the Tab	Elements in the Tab	Name of the Role Privilege	Description	Allow Privilege	Deny Privilege	Customizable
		Description	Read Edge Overview Properties Description	Controls ability to view or change Edge description on the Edge overview page	No	Yes	Yes
			Update Edge Overview Properties Description				
		Enable Alerts	Read Edge Overview Properties Enable Alerts	Controls ability to view or change Edge alert configuration on the Edge overview page	No	Yes	Yes
			Update Edge Overview Properties Enable Alerts				
		Authentication Mode	Read Edge Overview Properties Auth Mode	Controls ability to view or change Edge PKI configuration on the Edge overview page	No	Yes	Yes
			Update Edge Overview Properties Auth Mode				
			Read Customer PKI	Grants ability to view and manage enterprise PKI settings	Yes	No	No
			Update Customer PKI				
		Serial Number	Read Edge Overview Properties Serial Number	Controls ability to view or change Edge serial number, prior to activation, on the Edge overview page	No	Yes	Yes
			Update Edge Overview Properties Serial Number				
		Generate New Activation Key	Read Edge Overview Properties Activation Expiration	Controls ability to view or change the activation key expiration period on the Edge overview page	No	Yes	Yes

Navigation Path in the Enterprise Portal	Name of the Tab	Elements in the Tab	Name of the Role Privilege	Description	Allow Privilege	Deny Privilege	Customizable
			Update Edge Overview Properties Activation Expiration				
		Send Activation Email button	Create Edge Overview Properties Activation Email	Controls ability to generate an activation email on the Edge overview page	No	Yes	Yes
			Read Edge Overview Properties Activation Email				
		Local Credentials	Read Overview Properties Local Credentials	Grants ability to view and configure Edge local credentials	No	Yes	Yes
			Update Overview Properties Local Credentials				
		View	Read Edge	Grants ability to view and manage Edge objects and their properties in general	Yes	No	No
			Update Edge				
			Read Customer Keys	Grants ability to view and manage enterprise security keys such as Edge administrator credentials and IPSEC keys	Yes	Yes	Yes
			Update Customer Keys				
		License	Read License	Grants ability to view and manage Edge licensing	Yes	Yes	Yes
			Update License				

Navigation Path in the Enterprise Portal	Name of the Tab	Elements in the Tab	Name of the Role Privilege	Description	Allow Privilege	Deny Privilege	Customizable
		Profile	Create Edge Overview Profile	Controls visibility and control of Edges assigned profile on the Edge overview page	No	Yes	Yes
			Read Edge Overview Profile			No	No
			Update Edge Overview Profile			Yes	Yes
			Delete Edge Overview Profile				
			Assign Edge Profile	Grants ability to assign profiles to Edges	No	Yes	Yes
		RMA Reactivation	Create Edge	Grants ability to view and manage Edge objects and their properties in general	Yes	Yes	Yes
	Device						
		Authentication Settings	Create Edge Device Authentication Settings	Controls ability to view or change Edge Device Authentication Settings	No	Yes	Yes
			Read Edge Device Authentication Settings				
			Update Edge Device Authentication Settings				
			Delete Edge Device Authentication Settings				
		DNS Settings	Update Edge Device DNS Settings	Controls ability to view or change Edge Device DNS Settings	No	Yes	Yes

Navigation Path in the Enterprise Portal	Name of the Tab	Elements in the Tab	Name of the Role Privilege	Description	Allow Privilege	Deny Privilege	Customizable
		Netflow Settings	Create Edge Device Netflow Settings	Controls ability to view or change Edge Device Netflow Settings	No	Yes	Yes
			Read Edge Device Netflow Settings				
			Update Edge Device Netflow Settings				
			Delete Edge Device Netflow Settings				
		LAN-Side NAT Rules	Update Edge Device LAN-Side NAT Rules	Controls ability to view or change Edge Device LAN-Side NAT Rules	No	Yes	Yes
		Voice Quality Monitoring Settings	Read Edge Device VQM Settings	Controls ability to view or change Edge Device VQM Settings	No	Yes	Yes
			Update Edge Device VQM Settings				
		Syslog Settings	Read Edge Device Syslog Settings	Controls ability to view or change Edge Device Syslog Settings	No	Yes	Yes
			Update Edge Device Syslog Settings				
		Static Route Settings	Update Edge Device Static Route Settings	Controls ability to view or change Edge Device Static Route Settings	No	Yes	Yes
		ICMP Probes	Read Edge Device ICMP Probes	Controls ability to view or change Edge Device ICMP Probes	No	Yes	Yes
			Update Edge Device ICMP Probes				

Navigation Path in the Enterprise Portal	Name of the Tab	Elements in the Tab	Name of the Role Privilege	Description	Allow Privilege	Deny Privilege	Customizable
		ICMP Responders	Read Edge Device ICMP Responders	Controls ability to view or change Edge Device ICMP Responders	No	Yes	Yes
			Update Edge Device ICMP Responders				
		VRRP Settings	Update Edge Device VRRP Settings	Controls ability to view or change Edge Device VRRP Settings	No	Yes	Yes
		Cloud VPN	Read Edge Device Cloud VPN	Controls ability to view or change Edge Device Cloud VPN	No	Yes	Yes
			Update Edge Device Cloud VPN				
		BFD Rules	Update Edge Device BFD Rules	Controls ability to view or change Edge Device BFD Rules	No	Yes	Yes
		BGP Settings	Read Edge Device BGP Settings	Controls ability to view or change Edge Device BGP Settings	No	Yes	Yes
			Update Edge Device BGP Settings				
		Multicast Settings	Read Edge Device Multicast Settings	Controls ability to view or change Edge Device Multicast Settings	No	Yes	Yes
			Update Edge Device Multicast Settings				
		Cloud Security Service	Read Edge Device Cloud Security Service	Controls ability to view or change Edge Device Cloud Security Service	No	Yes	Yes
			Update Edge Device Cloud Security Service				

Navigation Path in the Enterprise Portal	Name of the Tab	Elements in the Tab	Name of the Role Privilege	Description	Allow Privilege	Deny Privilege	Customizable
		Gateway Handoff Assignment	Update Edge Device Gateway Handoff Assignment	Controls ability to view or change Edge Device Gateway Handoff Assignment	No	Yes	Yes
		High Availability	Create Edge Device High Availability	Controls ability to view or change Edge Device High Availability	No	Yes	Yes
			Read Edge Device High Availability				
			Update Edge Device High Availability				
			Delete Edge Device High Availability				
			Enable HA Standby Pair	Grants ability to configure standby HA	No	Yes	Yes
			Enable HA Cluster	Grants ability to configure HA Clustering	No	Yes	Yes
			Enable HA VRRP Pair	Grants ability to configure VRRP HA	No	Yes	Yes
		Configure VLAN	Read Edge Device Settings	Controls ability to view or change Edge Device Settings	No	Yes	Yes
		Management IP	Read Edge Device Management IP	Controls ability to view or change Edge Device Management IP	No	Yes	Yes
			Update Edge Device Management IP				
		Device Settings	Create Edge Device Settings	Controls ability to view or change Edge Device Settings	No	Yes	Yes

Navigation Path in the Enterprise Portal	Name of the Tab	Elements in the Tab	Name of the Role Privilege	Description	Allow Privilege	Deny Privilege	Customizable
			Read Edge Device Settings				
			Update Edge Device Settings				
			Delete Edge Device Settings				
		Interface Settings	Update Edge Device Interface Settings	Controls ability to view or change Edge Device Interface Settings	No	Yes	Yes
		WAN Settings	Update Edge Device WAN Settings	Controls ability to view or change Edge Device WAN Settings	No	Yes	Yes
		Security VNF	Update Edge Device Security VNF	Controls ability to view or change Edge Device Security VNF	No	Yes	Yes
		Wi-Fi Radio Settings	Create Edge Device Wi-Fi Settings	Controls ability to view or change Edge Device Wi-Fi Settings	No	Yes	Yes
			Read Edge Device Wi-Fi Settings				
			Update Edge Device Wi-Fi Settings				
			Delete Edge Device Wi-Fi Settings				
		Multi-Source QoS	Read Edge Device Cloud VPN QoS Settings	Controls ability to view or change Edge Device Cloud VPN QoS Settings	No	Yes	Yes

Navigation Path in the Enterprise Portal	Name of the Tab	Elements in the Tab	Name of the Role Privilege	Description	Allow Privilege	Deny Privilege	Customizable
			Update Edge Device Cloud VPN QoS Settings				
		TACACS Settings	Create Network Service	Grants ability to view and manage services with the Network Services configuration block	Yes	Yes	Yes
			Read Network Service			No	No
			Update Network Service			Yes	Yes
			Delete Network Service				
			Create Customer Keys	Grants ability to view and manage enterprise security keys such as Edge administrator credentials and IPSEC keys	Yes	Yes	Yes
			Read Customer Keys				
			Update Customer Keys				
			Delete Customer Keys				
			Manage Customer Keys			No	No
		L2 Settings	Update Edge Device L2 Settings	Controls ability to view or change Edge Device L2 Settings	No	Yes	Yes
		SNMP Settings	Create Edge Device SNMP Settings	Controls ability to view or change Edge Device SNMP Settings	No	Yes	Yes
			Read Edge Device SNMP Settings				

Navigation Path in the Enterprise Portal	Name of the Tab	Elements in the Tab	Name of the Role Privilege	Description	Allow Privilege	Deny Privilege	Customizable
			Update Edge Device SNMP Settings				
			Delete Edge Device SNMP Settings				
		NTP	Read Edge Device NTP Settings	Controls ability to view or change Edge Device NTP Settings	No	Yes	Yes
			Update Edge Device NTP Settings				
		Visibility Mode	Update Edge Device Config Visibility Mode	Controls ability to view or change Edge Device Config Visibility Mode	No	Yes	Yes
		Analytics Settings	Read Edge	Grants ability to view and manage Edge objects and their properties in general	Yes	No	No
			Update Edge				
	Business Policy		Edge Business Policy	Controls ability to view or change Edge business policy page	No	Yes	Yes
		SD-WAN Overlay Rate Limit	Read Edge Business Policy Rate Limit	Controls the ability to read and update the rate limiting business policy feature	No	Yes	Yes
			Update Edge Business Policy Rate Limit				
		SD-WAN Overlay Rate Limit SD-WAN Traffic Class and Weight Mapping	Read Edge	Grants ability to view and manage Edge objects and their properties in general	Yes	No	No

Navigation Path in the Enterprise Portal	Name of the Tab	Elements in the Tab	Name of the Role Privilege	Description	Allow Privilege	Deny Privilege	Customizable
			Read Customer Profile	Grants ability to view and edit enterprise configuration profiles	Yes	Yes	Yes
	Firewall		Edge Firewall	Controls ability to view or change Edge firewall page	No	Yes	Yes
		Firewall Logging Syslog Forwarding Stateful Firewall	Configure Edge Firewall Logging	Grants ability to configure Edges level firewall logging	No	Yes	Yes
		Firewall Logging	View Firewall Logs	Grants ability to view collected firewall logs	Yes	Yes	Yes
			Read Edge	Grants ability to view and manage Edge objects and their properties in general	Yes	No	No
		Syslog Forwarding	View Syslog Forwarding	Grants ability to see Syslog forwarding	No	Yes	Yes
			Read Edge	Grants ability to view and manage Edge objects and their properties in general	Yes	No	No
		Stateful Firewall Settings	Create Edge Firewall Edge Access	Privilege granting or denying visibility and control of an Edges Stateful Firewall Settings, Network & Flood Protection Settings and	No	Yes	Yes
		Network & Flood Protection Settings Edge	Read Edge Firewall Edge Access				
		Access	Update Edge Firewall Edge Access				

Navigation Path in the Enterprise Portal	Name of the Tab	Elements in the Tab	Name of the Role Privilege	Description	Allow Privilege	Deny Privilege	Customizable
			Delete Edge Firewall Edge Access	Edge Access on the Edge firewall page			
	Events from this Edge		Read Customer Event	Grants ability to view customer level events	Yes	No	No
	Remote Actions		Read Remote Actions	Privilege granting access to view and execute remote actions	No	Yes	Yes
	Remote Actions Generate Diagnostic Bundle Remote Diagnostics		Read Diagnostics	Controls creation of and access to diagnostics bundles, both Edge and Gateway. Combine with Edge and Gateway privileges to control access to each type individually	Yes	Yes	Yes
	Generate Diagnostic Bundle		Create Diagnostic Bundle		No	Yes	Yes
	Remote Diagnostics		Read Remote Diagnostics	Grants access to view and execute remote diagnostics	No	Yes	Yes
Configure > Profiles > Select Profile	Profile Overview		Profile Overview	Controls ability to view or change profile overview page	No	Yes	Yes
		Description	Create Profile Overview Description	Controls ability to view or change Profile Overview Description	No	Yes	Yes
			Read Profile Overview Description			No	No

Navigation Path in the Enterprise Portal	Name of the Tab	Elements in the Tab	Name of the Role Privilege	Description	Allow Privilege	Deny Privilege	Customizable
			Update Profile Overview Description			Yes	Yes
			Delete Profile Overview Description				
		Local Credentials	Read Overview Properties Local Credentials	Grants ability to view and configure Edge local credentials	No	Yes	Yes
			Update Overview Properties Local Credentials				
	Device						
		Authentication Settings	Create Profile Device Authentication Settings	Controls ability to view or change Profile Device Authentication Settings	No	Yes	Yes
			Read Profile Device Authentication Settings				
			Update Profile Device Authentication Settings				
			Delete Profile Device Authentication Settings				
		DNS Settings	Update Profile Device DNS Settings	Controls ability to view or change Profile Device DNS Settings	No	Yes	Yes
		Netflow Settings	Create Profile Device Netflow Settings	Controls ability to view or change Profile Device Netflow Settings	No	Yes	Yes

Navigation Path in the Enterprise Portal	Name of the Tab	Elements in the Tab	Name of the Role Privilege	Description	Allow Privilege	Deny Privilege	Customizable
			Read Profile Device Netflow Settings				
			Update Profile Device Netflow Settings				
			Delete Profile Device Netflow Settings				
		LAN-Side NAT Rules	Update Profile Device LAN-Side NAT Rules	Controls ability to view or change Profile Device LAN-Side NAT Rules	No	Yes	Yes
		Voice Quality Monitoring Settings	Read Profile Device VQM Settings	Controls ability to view or change Profile Device VQM Settings	No	Yes	Yes
			Update Profile Device VQM Settings				
		Syslog Settings	Read Profile Device Syslog Settings	Controls ability to view or change Profile Device Syslog Settings	No	Yes	Yes
			Update Profile Device Syslog Settings				
		Cloud VPN	Read Profile Device Cloud VPN	Controls ability to view or change Profile Device Cloud VPN	No	Yes	Yes
			Update Profile Device Cloud VPN				
		BFD Rules	Update Profile Device BFD Rules	Controls ability to view or change Profile Device BFD Rules	No	Yes	Yes
		OSPF Areas	Read Profile Device OSPF Settings	Controls ability to view or change Profile	No	Yes	Yes

Navigation Path in the Enterprise Portal	Name of the Tab	Elements in the Tab	Name of the Role Privilege	Description	Allow Privilege	Deny Privilege	Customizable
			Update Profile Device OSPF Settings	Device OSPF Settings			
		BGP Settings	Read Profile Device BGP Settings	Controls ability to view or change Profile Device BGP Settings	No	Yes	Yes
			Update Profile Device BGP Settings				
		Multicast Settings	Read Profile Device Multicast Settings	Controls ability to view or change Profile Device Multicast Settings	No	Yes	Yes
			Update Profile Device Multicast Settings				
		Cloud Security Service	Read Profile Device Cloud Security Service	Controls ability to view or change Profile Device Cloud Security Service	No	Yes	Yes
			Update Profile Device Cloud Security Service				
		Gateway Handoff Assignment	Update Profile Device Gateway Handoff Assignment	Controls ability to view or change Profile Device Gateway Handoff Assignment	No	Yes	Yes
		Configure VLAN	Read Profile Device Settings	Controls ability to view or change Profile Device Settings	No	Yes	Yes
		Management IP	Read Profile Device Management IP	Controls ability to view or change Profile Device Management IP	No	Yes	Yes
			Update Profile Device Management IP				

Navigation Path in the Enterprise Portal	Name of the Tab	Elements in the Tab	Name of the Role Privilege	Description	Allow Privilege	Deny Privilege	Customizable
		Device Settings	Create Profile Device Settings	Controls ability to view or change Profile Device Settings	No	Yes	Yes
			Read Profile Device Settings				
			Update Profile Device Settings				
			Delete Profile Device Settings				
		Interface Settings	Update Profile Device Interface Settings	Controls ability to view or change Profile Device Interface Settings	No	Yes	Yes
		Wi-Fi Radio Settings	Create Profile Device Wi-Fi Settings	Controls ability to view or change Profile Device Wi-Fi Settings	No	Yes	Yes
			Read Profile Device Wi-Fi Settings				
			Update Profile Device Wi-Fi Settings				
			Delete Profile Device Wi-Fi Settings				
		L2 Settings	Update Profile Device L2 Settings	Controls ability to view or change Profile Device L2 Settings	No	Yes	Yes
		Multi-Source QoS	Read Profile Device Cloud VPN QoS Settings	Controls ability to view or change Profile Device Cloud VPN QoS Settings	No	Yes	Yes
			Update Profile Device Cloud VPN QoS Settings				

Navigation Path in the Enterprise Portal	Name of the Tab	Elements in the Tab	Name of the Role Privilege	Description	Allow Privilege	Deny Privilege	Customizable
		SNMP Settings	Create Profile Device SNMP Settings	Controls ability to view or change Profile Device SNMP Settings	No	Yes	Yes
			Read Profile Device SNMP Settings				
			Update Profile Device SNMP Settings				
			Delete Profile Device SNMP Settings				
		NTP	Read Profile Device NTP Settings	Controls ability to view or change Profile Device NTP Settings	No	Yes	Yes
			Update Profile Device NTP Settings				
		Visibility Mode	Update Profile Device Config Visibility Mode	Controls ability to view or change Profile Device Config Visibility Mode	No	Yes	Yes
		Analytics Settings	Read Profile Device Analytics Settings	Controls ability to view or change Profile Device Analytics Settings	No	Yes	Yes
			Update Profile Device Analytics Settings				
			Create Profile Device Network Settings	Controls ability to view or change Profile Device Network Settings	No	Yes	Yes
			Read Profile Device Network Settings				
			Update Profile Device Network Settings				

Navigation Path in the Enterprise Portal	Name of the Tab	Elements in the Tab	Name of the Role Privilege	Description	Allow Privilege	Deny Privilege	Customizable
			Delete Profile Device Network Settings				
	Business Policy		Profile Business Policy	Controls ability to view or change profile business policy page	No	Yes	Yes
		SD-WAN Overlay Rate Limit	Read Profile Business Policy Rate Limit	Controls the ability to read and update the rate limiting business policy feature	No	Yes	Yes
			Update Profile Business Policy Rate Limit				
	Firewall		Profile Firewall	Controls ability to view or change profile firewall page	No	Yes	Yes
		Firewall Logging Syslog Forwarding Stateful Firewall	Configure Profile Firewall Logging	Grants ability to configure profile level firewall logging	No	Yes	Yes
		Firewall Logging	View Firewall Logs	Grants ability to view collected firewall logs	Yes	Yes	Yes
			Read Edge	Grants ability to view and manage Edge objects and their properties in general	Yes	No	No
		Syslog Forwarding	View Syslog Forwarding	Grants ability to see Syslog forwarding	No	Yes	Yes

Navigation Path in the Enterprise Portal	Name of the Tab	Elements in the Tab	Name of the Role Privilege	Description	Allow Privilege	Deny Privilege	Customizable
			Read Edge	Grants ability to view and manage Edge objects and their properties in general	Yes	No	No
		Stateful Firewall Settings	Create Edge Firewall Edge Access	Controls visibility and control of Stateful Firewall Settings, Network &	No	Yes	Yes
		Network & Flood Protection Settings Edge Access	Read Edge Firewall Edge Access	Flood Protection Settings, and Edge Access on the profile firewall page		No	No
			Update Edge Firewall Edge Access			Yes	Yes
			Delete Edge Firewall Edge Access				
Configure	Edges		Create Edge	Grants ability to view and manage Edge objects and their properties in general	Yes	Yes	Yes
			Read Edge			No	No
			Update Edge				
			Delete Edge			Yes	Yes
			Manage Edge			No	No
			Read Customer Profile	Grants ability to view and edit enterprise configuration profiles	Yes	Yes	Yes
	New Edge >	Authentication	Create Customer PKI	Grants ability to view and manage enterprise PKI settings	Yes	No	No
	Select Edge/Edges >	Local Credentials	Read Overview Properties Local Credentials	Grants ability to view and configure Edge local credentials	No	Yes	Yes

Navigation Path in the Enterprise Portal	Name of the Tab	Elements in the Tab	Name of the Role Privilege	Description	Allow Privilege	Deny Privilege	Customizable
			Update Overview Properties Local Credentials				
	Select Edge/Edges >	Assign Profile	Assign Edge Profile	Grants ability to assign profiles to Edges	No	Yes	Yes
	Select Edge/Edges >	Update Pre-Notifications	Update Edge Overview Properties Enable Alerts	Controls ability to view or change Edge alert configuration on the Edge overview page	No	Yes	Yes
	Select Edge/Edges >	Assign Edge License					
	Select Edge/Edges >	Update Customer Alerts					
		Edge Cluster	Read Edge Cluster	Grants ability to view Edge clusters	No	Yes	Yes
		Create Cloud Edge	Create DMZ Gateway	Grants ability to create DMZ Gateways	No	Yes	Yes
	Profiles		Create Customer Profile	Grants ability to view and edit enterprise configuration profiles	Yes	Yes	Yes
			Read Customer Profile				
			Update Customer Profile				
			Delete Customer Profile				
			Manage Customer Profile				
		Duplicate Profile	Duplicate Customer Profile	Grants ability to edit duplicate customer level profiles	No	Yes	Yes

Navigation Path in the Enterprise Portal	Name of the Tab	Elements in the Tab	Name of the Role Privilege	Description	Allow Privilege	Deny Privilege	Customizable
			Create Profile	Grants access to view and manage profiles at any level	No	Yes	Yes
			Read Profile				
			Update Profile				
			Delete Profile				
	Object Groups		Create Object Group	Grants ability to manage Object Group	Yes	Yes	Yes
			Read Object Group				
			Update Object Group				
			Delete Object Group				
			Manage Object Group			No	No
			Read Customer Profile	Grants ability to view and edit enterprise configuration profiles	Yes	Yes	Yes
	Segments/ Networks		Create Network Addressing	Grants ability to view and manage address block configuration in the legacy Network profile mode	Yes	Yes	Yes
			Read Network Addressing			No	No
			Update Network Addressing			Yes	Yes
			Delete Network Addressing			No	No
			Manage Network Addressing				
			Create Customer Segment	Grants ability to view and manage the creation of segments and their assignment to	No	Yes	Yes
			Read Customer Segment				

Navigation Path in the Enterprise Portal	Name of the Tab	Elements in the Tab	Name of the Role Privilege	Description	Allow Privilege	Deny Privilege	Customizable
			Update Customer Segment	configuration profiles			
			Delete Customer Segment				
	Overlay Flow Control		Create Overlay Flow Control	Grants ability to view and manage data and configuration presented on the Overlay Flow Control page	No	Yes	Yes
			Read Overlay Flow Control				
			Update Overlay Flow Control				
			Delete Overlay Flow Control				
			Read Customer Profile	Grants ability to view and edit enterprise configuration profiles	Yes	Yes	Yes
			Update Customer Profile				
	Network Services		Create Network Service	Grants ability to view and manage services with the Network Services configuration block	Yes	Yes	Yes
			Read Network Service			No	No
			Update Network Service			Yes	Yes
			Delete Network Service				
			Manage Network Service			No	No
			Create Customer Keys	Grants ability to view and manage enterprise	Yes	Yes	Yes

Navigation Path in the Enterprise Portal	Name of the Tab	Elements in the Tab	Name of the Role Privilege	Description	Allow Privilege	Deny Privilege	Customizable
			Read Customer Keys	security keys such as Edge administrator credentials and IPSEC keys			
			Update Customer Keys				
			Read Customer Profile	Grants ability to view and edit enterprise configuration profiles	Yes	Yes	Yes
		Edge Cluster	Create Edge Cluster	Controls the ability to create and configure Edge Clusters	No	Yes	Yes
			Read Edge Cluster				
			Update Edge Cluster				
			Delete Edge Cluster				
		Cloud VPN Hubs	Create VPN Hub Network Service	Grants ability to manage VPN Hubs as Network Services	No	Yes	Yes
			Read VPN Hub Network Service				
			Update VPN Hub Network Service				
			Delete VPN Hub Network Service				
		Non SD-WAN Destinations via Gateway	Create Non SD-WAN Destination via Gateway	Grants ability to view and manage Non SD-WAN Destinations via Gateway and Non SD-WAN Destinations via Edge	No	Yes	Yes
		Non SD-WAN Destinations via Edge	Read Non SD-WAN Destination via Gateway				
			Update Non SD-WAN Destination via Gateway				

Navigation Path in the Enterprise Portal	Name of the Tab	Elements in the Tab	Name of the Role Privilege	Description	Allow Privilege	Deny Privilege	Customizable
			Delete Non SD-WAN Destination via Gateway				
		Cloud Security Service	Create Cloud Security Service	Controls creation and configuration of third party cloud security services to which the traffic can be steered by business policy	No	Yes	Yes
			Read Cloud Security Service				
			Update Cloud Security Service				
			Delete Cloud Security Service				
		VNFs	Create VNF Network Service	Grants ability to manage VNF Network Services	No	Yes	Yes
			Read VNF Network Service				
			Update VNF Network Service				
			Delete VNF Network Service				
		VNF Licenses	Create VNF License Network Service	Grants ability to manage VNF licenses with Network Services	No	Yes	Yes
			Read VNF License Network Service				
			Update VNF License Network Service				
			Delete VNF License Network Service				

Navigation Path in the Enterprise Portal	Name of the Tab	Elements in the Tab	Name of the Role Privilege	Description	Allow Privilege	Deny Privilege	Customizable
		DNS Services	Create DNS Network Service	Controls the ability to create and configure DNS services for use in profiles	No	Yes	Yes
			Read DNS Network Service				
			Update DNS Network Service				
			Delete DNS Network Service				
		Private Network Names	Create Private Network Name Network Service	Grants ability to manage Private Network Name with Network Services	No	Yes	Yes
			Read Private Network Name Network Service				
			Update Private Network Name Network Service				
			Delete Private Network Name Network Service				
		Authentication Services	Create Authentication Service	Controls the creation and configuration of hosted 802.1x service providing LAN-side user authentication	No	Yes	Yes
			Read Authentication Service				
			Update Authentication Service				

Navigation Path in the Enterprise Portal	Name of the Tab	Elements in the Tab	Name of the Role Privilege	Description	Allow Privilege	Deny Privilege	Customizable
			Delete Authentication Service				
		TACACS Services	Create Network Service	Grants ability to view and manage services with the Network Services configuration block	Yes	Yes	Yes
			Read Network Service			No	No
			Update Network Service			Yes	Yes
			Delete Network Service				
			Create Customer Keys	Grants ability to view and manage enterprise security keys such as Edge administrator credentials and IPSEC keys	Yes	Yes	Yes
			Read Customer Keys				
			Update Customer Keys				
			Delete Customer Keys				
			Manage Customer Keys			No	No
		Cloud Subscriptions	Create Cloud Subscription Service	Grants ability to view and manage the configuration of access to IAAS providers, such as Azure, AWS and Google Cloud	No	Yes	Yes
			Read Cloud Subscription Service				
			Update Cloud Subscription Service				
			Delete Cloud Subscription Service				

Navigation Path in the Enterprise Portal	Name of the Tab	Elements in the Tab	Name of the Role Privilege	Description	Allow Privilege	Deny Privilege	Customizable
	Alerts & Notifications		Read Customer Alert Notification	Grants ability to view and manage customer alert configuration	No	Yes	Yes
			Create Customer Alert	Grants ability to view and manage customer alert configuration and generated alerts	Yes	No	No
			Read Customer Alert			Yes	Yes
			Update Customer Alert			No	No
			Delete Customer Alert				
			Manage Customer Alert				
		SMS Alert	Update Customer SMS Alert	Grants ability to configure SMS alerts at the customer level	No	Yes	Yes
	Customer		Update Enterprise	Grants ability to view and manage Customers, from the Partner or Operator level	Yes	Yes	Yes
		Other Settings	Read User Agreement	Privilege granting access to configure the customer user agreement feature	Yes	No	No
			Update User Agreement				

Navigation Path in the Enterprise Portal	Name of the Tab	Elements in the Tab	Name of the Role Privilege	Description	Allow Privilege	Deny Privilege	Customizable
Test & Troubleshoot			Read Diagnostics	Controls creation of and access to diagnostics bundles, both Edge and Gateway. Combine with Edge and Gateway privileges to control access to each type individually	Yes	Yes	Yes
	Remote Diagnostics		Create Remote Diagnostics	Grants access to view and execute remote diagnostics	No	No	No
			Read Remote Diagnostics			Yes	Yes
			Update Remote Diagnostics			No	No
			Delete Remote Diagnostics				
			Manage Remote Diagnostics			Yes	Yes
		Gateway	Remote Cloud Traffic Routing		No	Yes	Yes
		Reset USB Modem	Remote Reset USB Modem	Grants ability to execute the Edge USB modem reset remote action	No	Yes	Yes
		Scan for nearby Wi-Fi	Remote Scan for Wi-Fi Access Points	Grants ability to execute the Edge Wi-Fi scan remote action	No	Yes	Yes
		VPN Test	Remote VPN Test	Grants ability to execute the Edge VPN test remote action	No	Yes	Yes

Navigation Path in the Enterprise Portal	Name of the Tab	Elements in the Tab	Name of the Role Privilege	Description	Allow Privilege	Deny Privilege	Customizable
	Remote Actions		Create Remote Actions	Grants access to view and execute remote actions	No	Yes	Yes
			Read Remote Actions				
			Update Remote Actions				
			Delete Remote Actions				
		Select Edge > Shutdown button	Shutdown Edge	Grants ability to execute the Edge shutdown remote action	No	Yes	Yes
		Select Edge > Deactivate button	Deactivate Edge	Grants ability to execute the deactivate Edge remote action	No	Yes	Yes
	Diagnostic Bundles/ Packet Capture		Create Diagnostics	Controls creation of and access to diagnostics bundles, both Edge and Gateway. Combine with Edge and Gateway privileges to control access to each type individually	Yes	Yes	Yes
			Read Diagnostics				
			Update Diagnostics				
			Delete Diagnostics				
			Manage Diagnostics			No	No
		Request Diagnostic Bundle	Create Diagnostic Bundle	Grants ability to view and request Diagnostic bundles as part of remote diagnostics functionality	No	Yes	Yes
	Diagnostic Bundles/ Packet Capture	404 resource not found page *	Read Diagnostic Bundle				
			Update Diagnostic Bundle				

Navigation Path in the Enterprise Portal	Name of the Tab	Elements in the Tab	Name of the Role Privilege	Description	Allow Privilege	Deny Privilege	Customizable
		Delete Diagnostic Bundle	Delete Diagnostic Bundle				
		Request PCAP Bundle	Create PCAP Bundle	Grants ability to view and request PCAP bundles as part of remote diagnostics functionality	No	Yes	Yes
	Diagnostic Bundles/ Packet Capture	404 resource not found page *	Read PCAP Bundle				
			Update PCAP Bundle			No	No
			Delete PCAP Bundle			Yes	Yes
	Diagnostic Bundles/ Packet Capture	404 resource not found page *	Manage PCAP Bundle				
		Download Diagnostic Bundle	Download Edge Diagnostics	Grants ability to download Edge Diagnostics	No	Yes	Yes
Administration							
	System Settings		Read Customer Delegation	Grants ability to view and manage the delegation of privileges from the customer to Partners or the Operator	Yes	Yes	Yes
	General Information >	General Information	Read Customer General Information	Controls visibility and control of Customer General Information on the System Settings General Information page	No	Yes	Yes
			Update Customer General Information				
		Default Edge Authentication	Read Customer PKI	Grants ability to view and manage	Yes	No	No

Navigation Path in the Enterprise Portal	Name of the Tab	Elements in the Tab	Name of the Role Privilege	Description	Allow Privilege	Deny Privilege	Customizable
			Update Customer PKI	enterprise PKI settings			
		Edge Configuration	Read Customer Edge Settings	Controls visibility and control of Customer Edge Settings on the System Settings General Information page	No	Yes	Yes
			Update Customer Edge Settings				
		Privacy Settings	Read Customer Privacy Settings	Controls visibility and control of Customer Privacy Settings on the System Settings General Information page	No	Yes	Yes
			Update Customer Privacy Settings				
		Privacy Settings > Enforce PCI	Update Customer User	Grants ability to view and manage Customer administrators	Yes	Yes	Yes
		Contact Information	Read System Settings Contact Info	Controls visibility and control of System Settings Contact Info on the System Settings General Information page	No	Yes	Yes
			Update System Settings Contact Info				
	Authentication		Create Customer Authentication	Grants ability to view and manage customer authentication mode, for example SSO,	Yes	Yes	Yes
			Read Customer Authentication				

Navigation Path in the Enterprise Portal	Name of the Tab	Elements in the Tab	Name of the Role Privilege	Description	Allow Privilege	Deny Privilege	Customizable
			Update Customer Authentication	Radius or Native			
			Delete Customer Authentication				
			Manage Customer Authentication				
		API Tokens	Read Customer Token	Grants ability to view and manage authentication tokens at the Customer level	Yes	No	No
			Update Customer Token				
	Administrators		Create Customer User	Grants ability to view and manage Customer administrators	Yes	Yes	Yes
			Read Customer User				
			Update Customer User				
			Delete Customer User				
			Manage Customer User			No	No
	Select Enterprise User >	API Tokens	Create Customer Token	Grants ability to view and manage authentication tokens at the Customer level	Yes	No	No
			Read Customer Token				
			Update Customer Token				
			Delete Customer Token				

Navigation Path in the Enterprise Portal	Name of the Tab	Elements in the Tab	Name of the Role Privilege	Description	Allow Privilege	Deny Privilege	Customizable
			Manage Customer Token				
	Role Customization		Create Role Customization Package	Grants access to manage role customization packages	Yes	No	No
			Read Role Customization Package				
			Update Role Customization Package				
			Delete Role Customization Package				
			Manage Role Customization Package				
	Edge Licensing		Create License	Grants ability to view and manage Edge licensing	Yes	No	No
			Read License			Yes	Yes
			Update License			No	No
			Delete License				
			Manage License				
VeloCloud Support Access Role			Create Customer Delegation	Grants ability to view and manage the delegation of privileges from the customer to Partners or the Operator	Yes	Yes	Yes
			Read Customer Delegation				
			Update Customer Delegation				

Navigation Path in the Enterprise Portal	Name of the Tab	Elements in the Tab	Name of the Role Privilege	Description	Allow Privilege	Deny Privilege	Customizable
			Delete Customer Delegation				
			Manage Customer Delegation			No	No

* – When the corresponding user role privilege is denied, the Orchestrator window displays the *404 resource not found* error.

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 SD-WAN Orchestrator provides different types of licenses for deployed Edges. These license types account for POC enterprises 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 will 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-VC100M-PRE-HO-HG-L34S312P-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 New Edge** configuration screen includes an **Edge License** dropdown menu. This menu provides a list of available Edge licenses types which may be assigned to the newly created Edge and includes a search box for ease of locating the correct license.

For more 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 VMware 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:

- In the Enterprise portal, click **Configure > Edges**.
- To assign a license to each Edge, click the link to the Edge and select the license in the **Edge Overview** page. You can also select the Edge and click **Actions > Assign Edge License** to assign the license.
- To assign a license to multiple Edges, select the appropriate Edges, click **Actions > 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 VMware, then contact VMware SD-WAN 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 the incorrect Edge license type is chosen, the impact is that the activation report for that enterprise will be incorrect, and the license assignment will not align with the customer's purchases. These licensing inconsistencies would be 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 Enterprise portal, click **Administration > Edge Licensing**.



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

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:

Product	Description	Quantity
VC-510-HO-36-P	VMware SD-WAN Edge 510 Appliance, Deployment: Hosted Orchestrator for 3 years	11
VC-610-HO-36-P	VMware SD-WAN Edge 610 Appliance, Deployment: Hosted Orchestrator for 3 years	1

Product	Description	Quantity
VC100M-STD-HO-L34S1-36P	VMware SD-WAN 100 Mbps Standard Service Subscription for 3 years, Prepaid, Hosted Orchestrator, Basic Support Backline (L3-4)	11
VC350M-STD-HO-L34S1-36P	VMware SD-WAN 350 Mbps Standard Software Subscription for 3 year, Prepaid, Hosted Orchestrator, VMware Basic Support Backline(L3-4)	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 Enterprise portal, click **Configure > Edges**.
- 2 In the Edges screen, click **New Edge**.
- 3 In the **Provision New Edge** window, configure a new Edge and assign the license type.

- 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, click **Administration > Edge Licensing**.



<input type="checkbox"/>	Name	Term	Bandwidth	Edition	Region	Edges Assigned	Activated Ec
<input type="checkbox"/>	STANDARD 100 Mbps North America, Europe Middle East and Africa 36 Months	36 mon...	100 Mbps	Standard	North Ameri...	11	
<input type="checkbox"/>	STANDARD 350 Mbps North America, Europe Middle East and Africa 36 Months	36 mon...	350 Mbps	Standard	North Ameri...	1	

The above image shows a report of Edge licenses assigned to 12 Edges. You can click **Report** to generate a report of the licenses and the associated Edges in CSV format.

Configure High Availability on SD-WAN Edge

27

This section describes the high availability deployments and configuration supported on SD-WAN Edge.

Refer to the following topics:

Read the following topics next:

- [How SD-WAN Edge High Availability \(HA\) Works](#)
- [High Availability Deployment Models](#)
- [Split-Brain Condition](#)
- [Split-Brain Detection and Prevention](#)
- [Support for BGP Over HA Link](#)
- [Selection Criteria to Determine Active and Standby Status](#)
- [VLAN-tagged Traffic Over HA Link](#)
- [Configure High Availability \(HA\)](#)
- [HA Event Details](#)

How SD-WAN Edge High Availability (HA) Works

The high availability solution ensures continued traffic flow in case of failures. The SD-WAN Edge is the VMware data plane component that is deployed at an end user's branch location. SD-WAN Edge configured in High Availability (HA) mode are mirror images of each other and they show up on the SD-WAN Orchestrator as a single SD-WAN Edge.

In a high availability configuration, SD-WAN 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 SD-WAN 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.

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 SD-WAN 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 SD-WAN 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 SD-WAN Orchestrator generates the “High Availability Going Active” event and the Standby Edge takes over as Active.

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 SD-WAN Edge platform models: 510, 510N, 520, 520v, 540, 610, 610N, 620, 620N, 640, 640N, 680, 680N, 840, 2000, 3400, 3800, and Virtual Edge.

Note HA is supported only between identical SD-WAN Edge platform models. For more information on the Edge platform models, see <https://sdwan.vmware.com/get-started>.

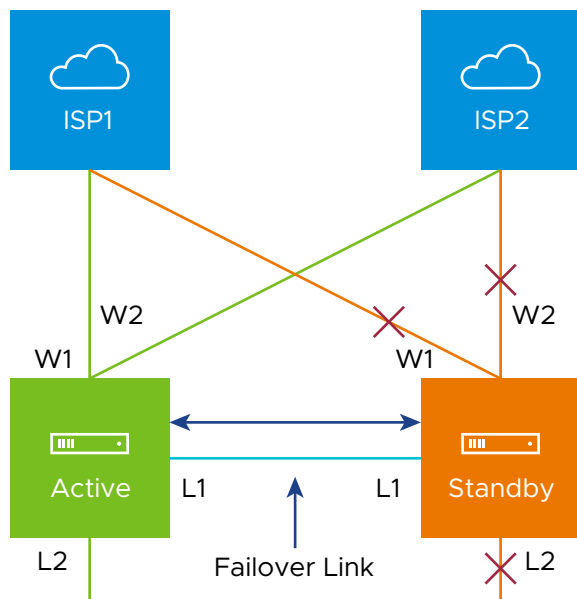
Note Mixing Wi-Fi capable and Non-Wi-Fi capable Edges in High Availability deployments is not supported. While the Edge models 510N, 610N, 620N, 640N, and 680N appear identical to their Wi-Fi capable counterparts, deploying a Wi-Fi capable Edge and a Non-Wi-Fi capable Edge of the same model (for example, an Edge 640 and an Edge 640N) as a High-Availability pair is not supported. Customers should ensure that the Edges deployed as a High Availability pair are of the same type: both Wi-Fi capable, or both Non-Wi-Fi capable.

Standard HA

This section describes Standard HA.

Topology Overview for Standard HA

The following figure shows a conceptual overview of Standard HA.



The Edges, one Active and one Standby, are connected by L1 ports to establish a failover link. The Standby SD-WAN 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, SD-WAN 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 SD-WAN 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 describe these two deployment types.

Deployment Type 1: HA using L2 switches

The following figure shows the network connections using only 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 SD-WAN Edges and is used for 'keep-alive' and communication between the SD-WAN Edges for HA support. The SD-WAN Edge's LAN connections are used to connect to the access layer L2 switches.

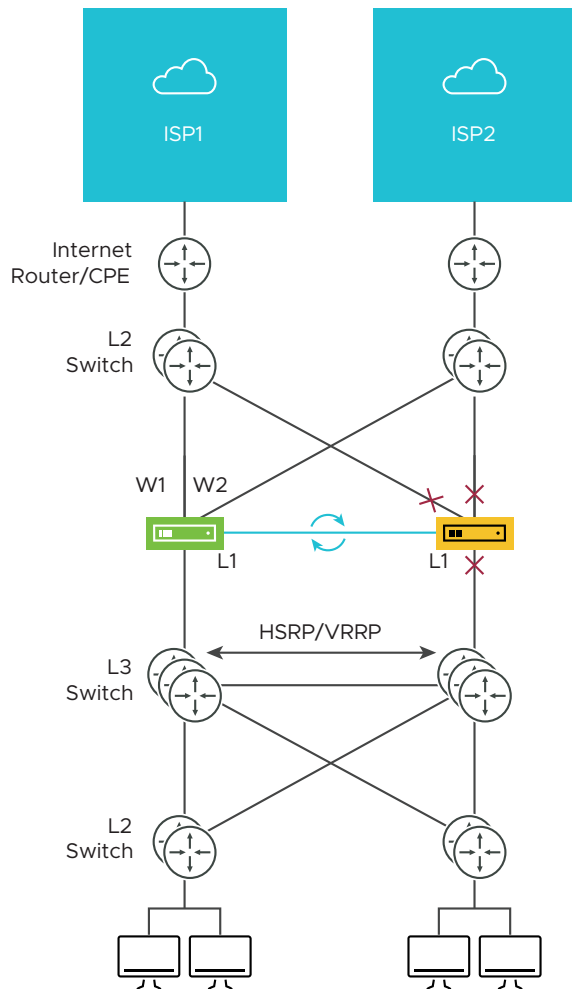
Considerations for HA Deployment using L2 switches

- 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 SD-WAN 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 SD-WAN 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.



The SD-WAN 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 SD-WAN Edge are connected to provide a failover link for HA support. The VMware 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 SD-WAN 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 SD-WAN Edges. The L2 switch must make the same ISP link available to both Edges.
- The Standby SD-WAN 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 SD-WAN 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.

Enhanced HA

This section describes Enhanced HA. The Enhanced HA eliminates the need for L2 Switches on WAN side of the Edges. 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.



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 SD-WAN 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. VMware SD-WAN supports LTE in 510 and 610 Edge models which have two SIM slots.

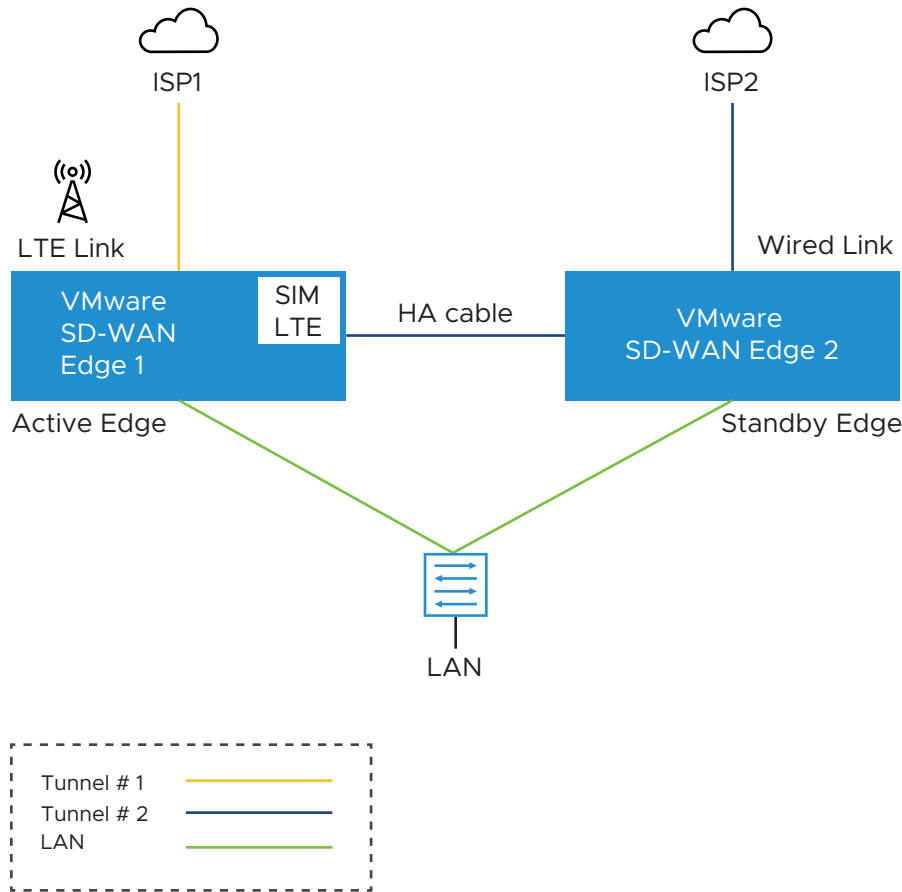
Starting with the 4.2 release, the LTE link/CELL interface is counted in the HA election. Internally, a lesser weight is provided for CELL links than wired links. So depending on the number of wired links connected to each Edge in the eHA pair, the Edge with the LTE link can either be the Active or the 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



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 (SD-WAN Edge 1) and one Standby (SD-WAN 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



The figure illustrates the topology of Enhanced HA support for LTE Interface on an Active Edge. In this example, the SD-WAN Edge 1 with one wired link and one LTE link acts as an Active Edge, and SD-WAN 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.

LTE Modem Information

This will fetch diagnostic information for the internal LTE modem.

Run

Interface

CELL1 ▾

Test Duration: 6.006 seconds

LTE CELL1

Modem Information

```
{
  "Manufacturer": "Sierra Wireless, Incorporated",
  "Model": "EM7511",
  "Model identifier": "353587100789907",
  "Firmware Revision": "SWI9X50C_01.07.02.00 6c91bc jenkins 2018/06/13 23 08 16",
  "Hardware Revision": "10001",
  "Supported capabilities": "gsm-umts, lte",
  "Current capabilities": "gsm-umts, lte",
  "own number": "NA",
  "state": "connected",
  "Failed reason": "--",
  "Power state": "on",
  "current_modes": "allowed 2g, 3g, 4g; preferred 4g",
  "imei": "353587100789907",
  "operator code": "310260",
  "operator name": "T-Mobile",
  "registration state": "home",
  "signal quality(%)": "52"
}
```

Connection Information

```
{
  "Bearer": "Available",
  "Connected": "yes",
  "Suspended": "no",
  "Interface": "wwan0",
  "APN": "",
  "IP type": "--",
  "user": "--",
  "password": "NA",
  "IP method": "static",
  "IP address": "100.232.152.201",
  "Gateway": "100.232.152.202",
  "DNS": "10.177.0.34",
  "MTU": "1430",
  "Stats Duration": "24359",
  "Rx bytes": "106396",
  "Tx bytes": "59484"
}
```

Location Information

```
{
  "Operator code": "310",
  "Operator name": "260",
  "Location area code": "FFFF",
  "tracking area code": "3A69",
  "cell id": "02CB0705"
}
```

Signal Information

```
{
  "Serving": {
    "EARFCN": "5035",
    "MCC": "310",
    "MNC": "260",
    "TAC": "14953",
    "CID": "02CB0705",
    "Bd": "12",
    "D": "2",
    "U": "2",
    "SNR": "4",
    "PCI": "334",
    "RSRQ": "11.8",
    "RSRP": "-107.4",
    "RSSI": "-81.6",
    "RXLV": "16"
  },
  "IntraFreq": {
    "PCI": "334",
    "RSRQ": "11.8",
    "RSRP": "-107.4",
    "RSSI": "-81.6",
    "RXLV": "16"
  }
}
```

Status Information

```
response: 'IGSTATUS:
Current Time: 24389
Reset Counter: 1
System mode: LTE
LTE band: B12
LTE Rx chan: 5035
LTE SSC1 state: NOT ASSIGNED
LTE SSC2 state: NOT ASSIGNED
LTE SSC3 state: NOT ASSIGNED
LTE SSC4 state: NOT ASSIGNED
EMM state: Registered
RRC state: RRC Idle
IMS reg state: No Srv

Temperature: 51
Mode: ONLINE
PS state: Attached
LTE bw: 5 MHz
LTE Tx chan: 23035

Normal Service

PCC RxM RSSI: -80
PCC RxM RSRP: -106
PCC RxD RSSI: -81
PCC RxD RSRP: -109
Tx Power: --
TAC: 3a69 (14953)
RSRQ (dB): -12.1
Cell ID: 02cb0705 (46860037)
SINR (dB): 4.2'
```

Debug Information

```
{
  "STATUS": "ERROR",
  "REASON": "Debug data not available"
}
```

- **Reset USB Modem** - Run this test on a selected Edge interface to reset an unworking USB modem connected to the given interface. Note that not all USB modems support this type of remote reset.

Reset USB Modem

This will attempt to reset an unworking USB modem connected to the given interface. Note that not all USB modems support this type of remote reset.

Interface

CELL1

Run

Test Duration: 55.115 seconds

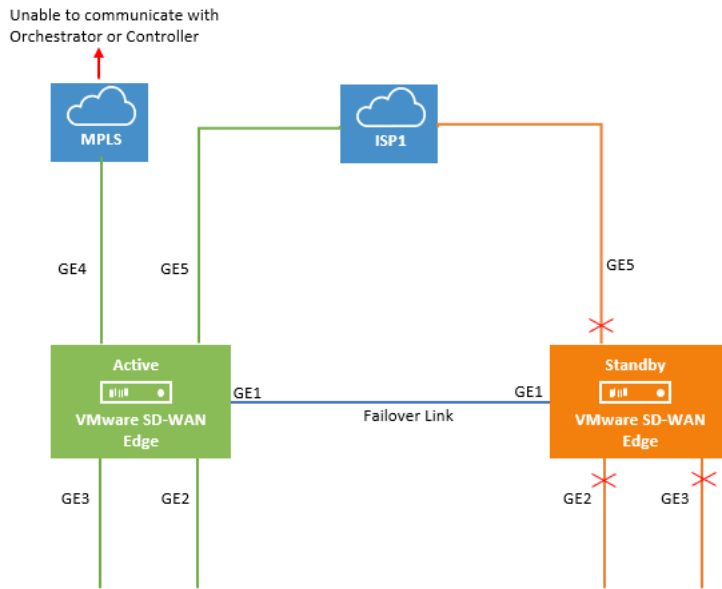
The restart command has been issued to the selected interface.

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.



In this topology, the Active and Standby Edges exchange heartbeat messages, synchronize configuration updates, and other information over the GE1 interface. Both SD-WAN 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.



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.

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 spanning tree loops on the LAN.

Typically, switches run the Spanning Tree Protocol to prevent loops 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](#).

Split-Brain Detection and Prevention

This section covers 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 Primary 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 Primary 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 Primary Gateway will always have tunnels only from the Active Edge.

As soon as the HA link fails, 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 failover from Active to Standby in a split-brain scenario is not the same as a normal failover where the Active Edge has gone down. A split-brain corrective failover may take a few extra milliseconds/seconds to converge.

Note When configuring WAN interface settings for an Edge, if you select "PPPoE" from the "Addressing Type" field, the Edge cannot send heartbeat packet by broadcast from the WAN interface which is configured as "PPPoE".



Support for BGP Over HA Link

In case the Edges switch to the enhanced HA option, the Active SD-WAN Edge will exchange BGP routes over the HA link. BGP on the Active Edge can now establish neighborhood with a peer connected only to the standby Edge's WAN link.

This will enable the Active Edge to learn routes from the WAN link(s) connected to the Standby Edge. The routing daemon on standby will not involve in any of the functionality. The standby Edge itself will just do a pass-through.

Note Routes are not synced between the active and the standby Edges. Therefore, in the above scenario, if there is a failover and a standby Edge becomes active, the BGP daemon on the newly active edge will establish a new neighborhood with the same BGP peer.

Selection Criteria to Determine Active and Standby Status

This section describes 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.

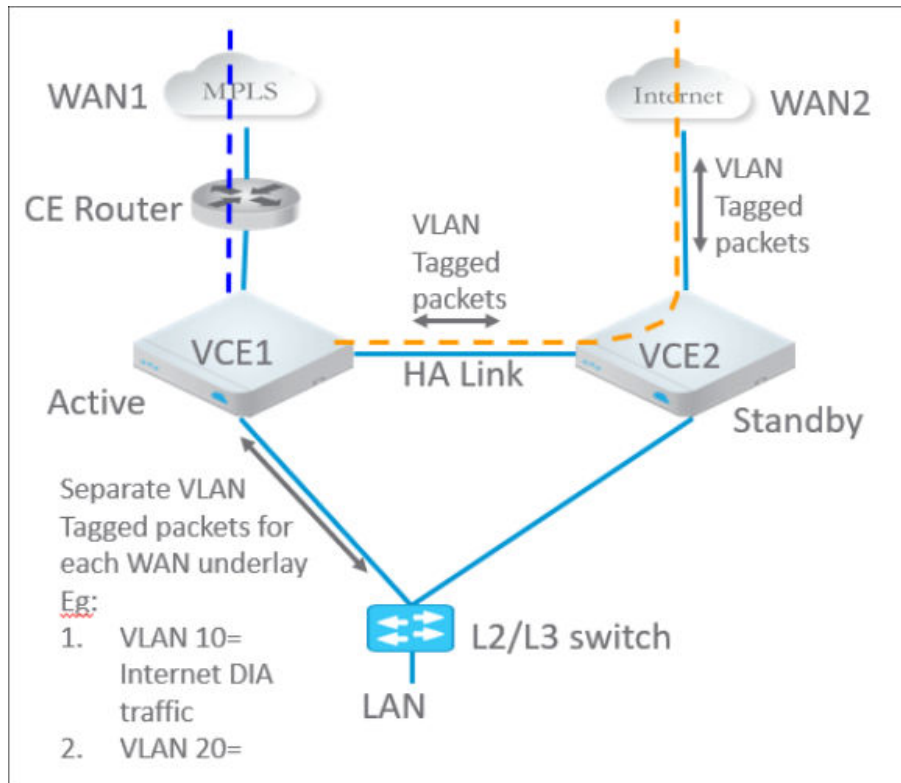
Note There is no preemption if the two Edges have the same number of LAN and WAN interfaces.

- 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.

VLAN-tagged Traffic Over HA Link

This section describes 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



Configure High Availability (HA)

To configure High Availability, configure the Active and Standby Edges.

Deploying High Availability on VMware ESXi

You can deploy the VMware 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 VMware 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.

Note For more information on **MAC address changes** and **Promiscuous mode operation**, refer to the topic <https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.security.doc/GUID-3507432E-AFEA-4B6B-B404-17A020575358.html>.

Limitations of VMware SD-WAN High Availability

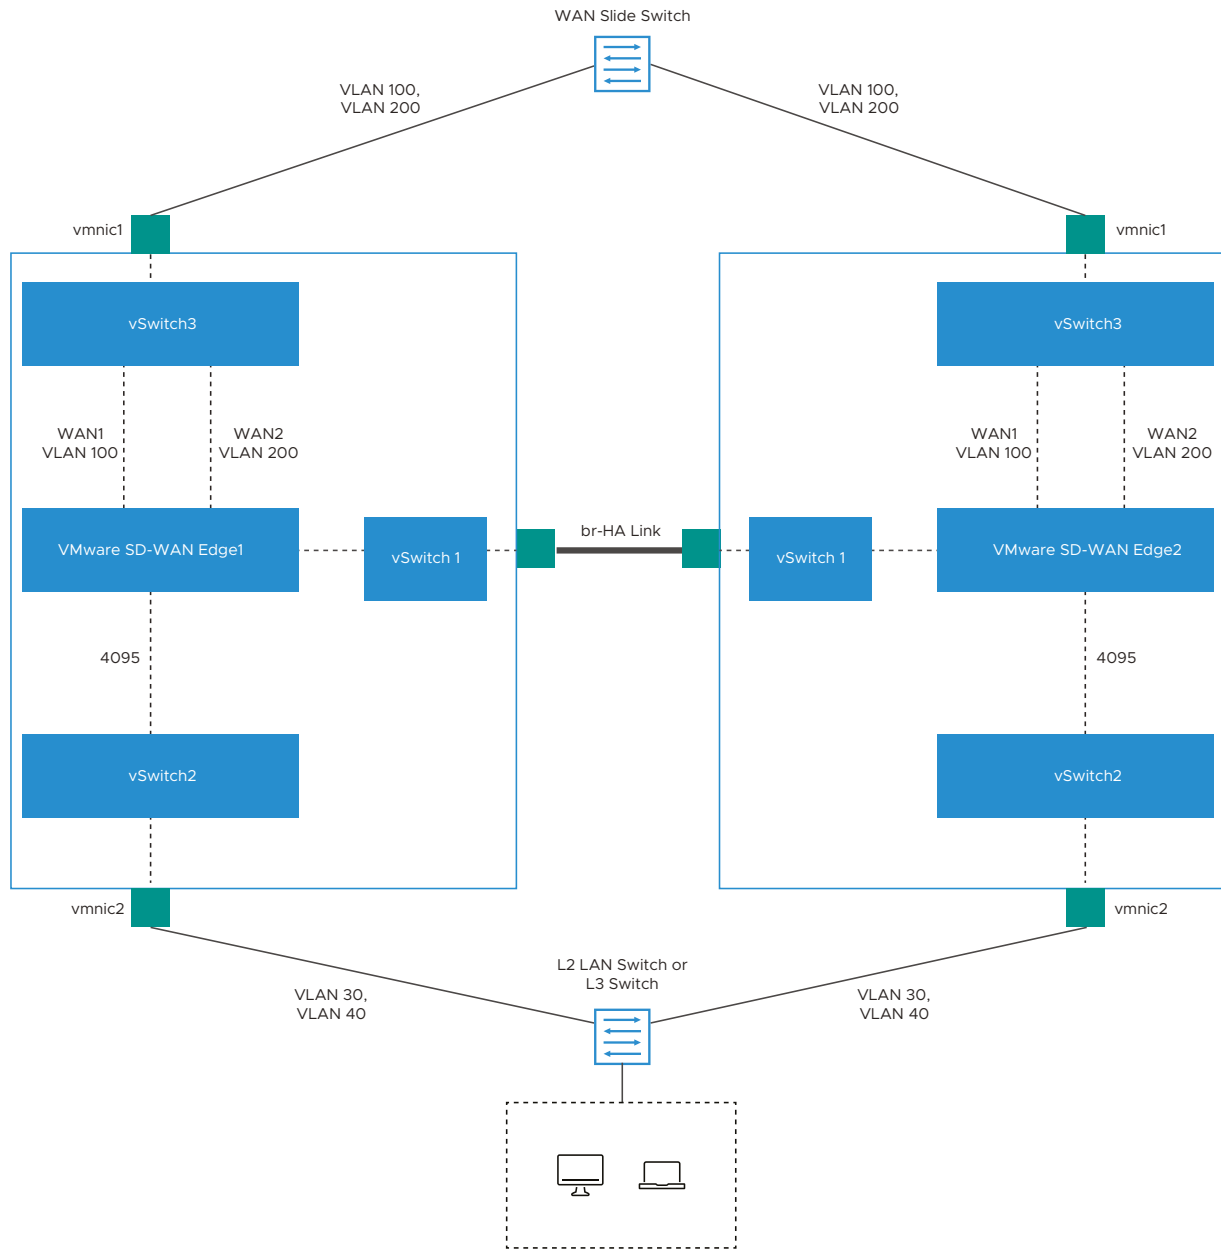
- There is no generic way of failure detection that will work on all the hardware, virtual, and uCPE platforms.

You can enable the Loss of Signal (LoS) detection to determine the HA Failover. For more information, see [HA LoS Detection on Routed Interfaces](#).

VMware 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.



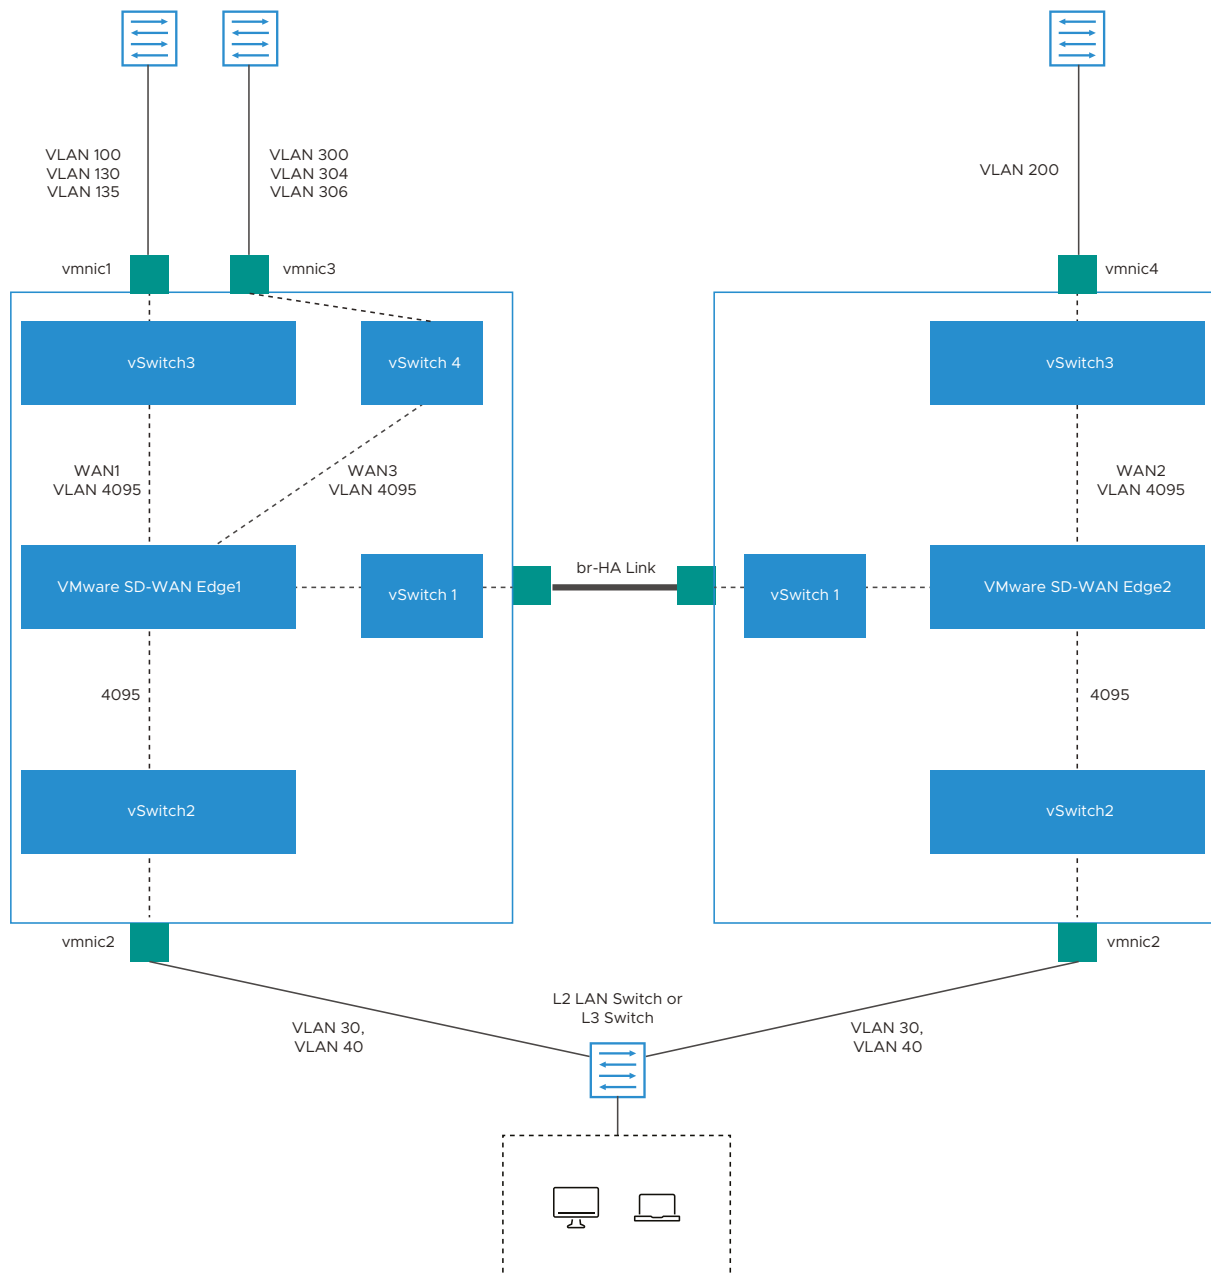
Topology 2: Enhanced HA with WAN Links

The following topology shows enhanced HA with three 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.



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 Enterprise portal, click **Configure > Edges**.
- 2 Click the Device Icon next to an Edge, or click the link to an Edge and then click the **Device** tab.
- 3 In the **Device** tab, scroll down to the **Interface Settings** section, which displays the Interfaces available in the selected Edge.
- 4 Click 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 click **Update**.

Virtual Edge

Interface GE3 ? x

☒ Override Interface

Interface Enabled ☒

Capability Routed

Segments All Segments

Addressing Type Static

IP Address 169.254.7.10

CIDR prefix 29

Gateway 169.254.7.9

WAN Overlay ☒ Auto-Detect Overlay unlock

OSPF ✗ OSPF not enabled for the selected Segment.

VNF Insertion ✗ VNF insertion is disallowed when an interface is configured for WAN overlays

Multicast Multicast is not enabled for the selected segment

RADIUS Authentication ⓘ ✗ WAN Overlay must be disabled to configure RADIUS Authentication.
Require User Authentication to access WAN

Advertise ☐

ICMP Echo Response ☒

NAT Direct Traffic ☒

Underlay Accounting ⓘ ☒

Trusted Source ⓘ ☐

Reverse Path Forwarding ⓘ Specific

VLAN

L2 Settings

Autonegotiate ☒

* MTU 1500

* Enable LoS Detection ☒

* ARP Probe Interval (in Seconds) 3

DHCP Server

Type Enabled Relay Disabled

Update GE3 Cancel

- 9 Click **Save Changes** in the **Devices** tab.

For more information on the other settings of the Interface, see [Configure Interface Settings](#).

To view the LoS detection events, see [Monitor Events for LoS Detection](#).

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, click **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>*

The following image shows the LoS events.

Events

Past 12 Hours

Tue Feb 2, 6:41

now

<

>

Search

Cols

Reset View

Refresh

CSV

Display 463 items

Time	Event	Segment	Edge	User	Severity	Message
Tue Feb 02, 19:01:29	High Availability R...		b1-edge1		Notice	Standby state ready for failover
Tue Feb 02, 19:01:17	High Availability G...		b1-edge1		Notice	Standby going active, Peer LAN interf
Tue Feb 02, 19:01:17	Interface LoS		b1-edge1		Alert	LoS detected on peer's interface GE5
Tue Feb 02, 19:01:17	Link alive		b1-edge1		Info	Link GE3 is no longer DEAD
Tue Feb 02, 19:01:17	Link alive		b1-edge1		Info	Link GE4 is no longer DEAD
Tue Feb 02, 19:01:15	Interface LoS		b1-edge1		Alert	LoS detected on interface GE5
Tue Feb 02, 18:59:06	Interface LoS		b1-edge1		Alert	LoS no longer seen on interface GE5
Tue Feb 02, 18:58:36	Interface LoS		b1-edge1		Alert	LoS detected on interface GE5

Unique MAC Address

Starting from 4.3.0 release, virtual Edges support a unique MAC address feature on a High Availability interface.

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.

Important On a customer enterprise using HA Edges and VMware 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.

For more information on MAC learning with vSphere Networking, see: [What is MAC Learning Policy](#).

Prerequisites

This section describes HA requirements that must be met before configuring a SD-WAN Edge as a Standby.

- The two SD-WAN Edges must be the same model.
- Only one SD-WAN Edge should be provisioned on the SD-WAN Orchestrator.
- The Standby SD-WAN Edge must not have an existing configuration on it.
- Ensure not to use 169.254.2.x for management interface.

Activate High Availability

You can activate High Availability (HA) on a pair of Edges to ensure redundancy.

- 1 In the Enterprise portal, click **Configure > Edges**.
- 2 Select the SD-WAN Edge from the list and click the **Device** tab.
- 3 Scroll down to the **High Availability** section and click **Active Standby Pair**.

High Availability

Type

- ☐ None
- ☒ Active Standby Pair
- ☐ Cluster
- ☐ VRRP with Third-Party Router

HA Interface

GE1

☒ Deploy with Unique LAN mac ⓘ

- 4 Click **Save Changes** at the top 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.

Note The above HA interfaces are selected automatically and you cannot configure an HA interface manually.

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** checkbox, 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, turn off High Availability as follows:

- 1 In the **High Availability** section, click **None**.
- 2 Click **Save Changes** at the top of the **Device** window.

Enable the High Availability again and then click the **Deploy with Unique LAN MAC** checkbox to activate or deactivate the option.

Wait for SD-WAN Edge to Assume Active

After the High Availability feature is enabled on the SD-WAN Orchestrator, wait for the existing SD-WAN Edge to assume an Active role, and wait for the SD-WAN Orchestrator Events to display **High Availability Going Active**.

i	Sun Jul 10, 23:00	High Availability Going Active	DC1 - Hub1	Notice	VeloCloud Edge going active, peer has not been detected
i	Sun Jul 10, 23:00	Edge service startup	DC1 - Hub1	Notice	VeloCloud edge service started
i	Sun Jul 10, 23:00	Edge online	DC1 - Hub1	Info	Management Daemon Started, version 2.1.0 build R21-
i	Sun Jul 10, 22:56	ENDPOINT_ACCEPTED_CERTIFICATE	DC1 - Hub1	Info	AE18A7B61185ABE827DBD8B98556C5AACA36C3ED
i	Sun Jul 10, 22:56	EDGE_OSPF_NSM	DC1 - Hub1	Notice	Edge NSM event: interface=172.31.2.1 nbr=172.31.2.2 router_id=172.31.2.2 status=Full
i	Sun Jul 10, 22:56	Link alive	DC1 - Hub1	Info	Link GE4 is no longer DEAD
i	Sun Jul 10, 22:56	Edge Interface Up	DC1 - Hub1	Info	Interface GE4 is up
i	Sun Jul 10, 22:56	Edge Interface Up	DC1 - Hub1	Info	Interface GE3 is up

Connect the Standby SD-WAN Edge to the Active Edge

- 1 Power on the Standby SD-WAN 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 SD-WAN Edge.
- 3 Wait for the Active SD-WAN Edge to detect and activate the standby SD-WAN Edge automatically. The SD-WAN Orchestrator Events displays **HA Standby Activated** when the SD-WAN Orchestrator successfully activates the standby SD-WAN Edge.

i	Fri Nov 18, 14:31:54	Edge service startup		Notice	VeloCloud edge service started
i	Fri Nov 18, 14:31:07	HA Standby Activated		Notice	Standby has been detected

The standby Edge will then begin to synchronize with the active SD-WAN Edge and reboot automatically during the process.

Note It may take up to 10 minutes for the Standby SD-WAN Edge to sync with the Active Edge and upgrade its software.

i	Fri Nov 18, 14:37:27	High Availability Ready		Notice	Standby state ready for failover
i	Fri Nov 18, 14:37:25	Edge service startup		Notice	VeloCloud edge service started
i	Fri Nov 18, 14:37:08	Edge online		Info	Management Daemon Started, version 2.2.1 build R221-20161109-GA
i	Fri Nov 18, 14:36:25	HA Peer State Unknown		Notice	Peer state unknown
i	Fri Nov 18, 14:34:59	Standby device software update started		Info	Begin HA Standby update with new software version
i	Fri Nov 18, 14:32:15	High Availability Ready		Notice	Standby state ready for failover
i	Fri Nov 18, 14:32:14	Edge service startup		Notice	VeloCloud edge service started

Connect LAN and WAN Interfaces on Standby SD-WAN Edge

Connect the LAN and WAN interfaces on the standby SD-WAN Edge mirroring the network connectivity on the Active Edge.

The SD-WAN Orchestrator Events will display **Standby device software update completed**. The **HA State** in the **Monitor > Edges** page appears green when ready.

Edge	Status	HA	Links	Gateways	Profile	Operator Profile
1 Bronze VCE	●	●	●	View	SF Branch Profile	Initial Operator Profile
2 DC1 - Hub1	●	●	●	View	DC1 Hub Profile	Hub Operator profile - no S...
3 DC2 - Hub1	●	●	●	View	DC2 Hub Profile	Hub Operator profile - no S...
4 SF1 - MPLS_Internet Branch	●	●	●	View	SF Branch Profile	Initial Operator Profile
5 SF2 - Dual Internet Branch	●	●	●	View	SF Branch Profile	Initial Operator Profile
6 Silver1 VCE	●	●	●	View	SF Branch Profile	Initial Operator Profile
7 Silver2 VCE	●	●	●	View	SF Branch Profile	Initial Operator Profile

Deactivate High Availability (HA)

This section covers 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 Enterprise portal, click **Configure > Edges**.
- 2 Select the SD-WAN Edge from the list and click the **Device** tab.
- 3 Scroll down to the **High Availability** section and click **None**.

High Availability

HA: None Segment Agnostic

High Availability is enabled at the Edge level. When using Active/Standby Pair HA, enable HA prior to connecting the Standby SD-WAN Edge. To learn more, please consult our [HA documentation](#).

Select Type

☒ None

☐ Active Standby Pair

☐ Cluster

☐ VRRP with 3rd party router

- 4 Click **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:

- 1 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.
- 2 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 repurpose 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.

HA Event Details

This section describes HA events.

HA Event	Description
HA_GOING_ACTIVE	A standby SD-WAN 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 SD-WAN 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 SD-WAN Edge no longer hears from the Standby SD-WAN Edge. For example, if the Standby SD-WAN Edge reboots, you will receive this message.
HA_READY	Means the Active SD-WAN Edge now hears from the Standby SD-WAN Edge. Once the Standby SD-WAN 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 SD-WAN Orchestrator is unable to verify the HA activation, it will generate this Event. Examples include: <ul style="list-style-type: none"> ■ the SD-WAN Orchestrator is unable to generate a certificate ■ the HA has been deactivated (rare)

VMware Virtual Edge Deployment

28

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 VMware Virtual Edge on KVM and VMware ESXi hypervisors.

Read the following topics next:

- [Deployment Prerequisites for VMware Virtual Edge](#)
- [Special Considerations for VMware Virtual Edge deployment](#)
- [Cloud-init Creation](#)
- [Install VMware Virtual Edge](#)

Deployment Prerequisites for VMware Virtual Edge

Describes the requirements for VMware 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.

	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

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

Recommended NIC Specifications

Hardware Manufacturer	Firmware Version	Host Driver for Ubuntu 16.04/18.04	Host Driver for ESXi 6.7
Dual Port Intel Corporation Ethernet Controller XL710 for 40GbE QSFP+	6.80	2.7.11	1.7.17
Dual Port Intel Corporation Ethernet Controller X710 for 10GbE SFP+	6.80	2.7.11	1.7.17
Quad Port Intel Corporation Ethernet Controller X710 for 10GbE SFP+	6.80	2.7.11	1.7.17

Supported Operating Systems

- Ubuntu 16.04
- VMware vSphere ESXi 6.7, and from version 4.3 and above VMware vSphere ESXi 6.7 and 7.0

Firewall/NAT Requirements

If the VMware Virtual Edge is deployed behind the Firewall and/or a NAT device, the following requirements apply:

- The Firewall must allow outbound traffic from the VMware Virtual Edge to TCP/443 (for communication with the SD-WAN 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 VMware Virtual Edge deployment](#).

Special Considerations for VMware Virtual Edge deployment

Describes the special considerations for VMware Virtual Edge deployment.

- The SD-WAN Edge is a latency-sensitive application. Refer to the [VMware 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 VMware and KVM, see:
 - KVM - [Enable SR-IOV on KVM](#)
 - VMware - [Enable SR-IOV on VMware](#)
 - 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`.

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 VMware 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 VMware 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 VMware 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: vedgel  
  
local-hostname: vedgel
```

- 2 Add the `network-interfaces` section, shown below, to specify the WAN configuration. By default, all SD-WAN 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: SD-WAN Orchestrator, Activation Code, and Ignore Certificates Errors.

Module	Description
<code>vco</code>	IP Address/URL of the SD-WAN Orchestrator.
<code>activation_code</code>	Activation code for the Virtual Edge. The activation code is generated while creating an Edge instance on the SD-WAN Orchestrator.
<code>vco_ignore_cert_errors</code>	Option to verify or ignore any certificate validity errors.

The activation code is generated while creating an Edge instance on the SD-WAN Orchestrator.

Important There is no default password in SD-WAN 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.

Install VMware Virtual Edge

You can install VMware 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.

Prerequisites

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. VMware recommends the following options:

- SR-IOV
- Linux Bridge
- OpenVSwitch Bridge

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

- [Enable SR-IOV on KVM](#)
- [Enable SR-IOV on VMware](#)

To install VMware Virtual Edge:

- On KVM, see [Install Virtual Edge on KVM](#).
- On VMware ESXi, see [Install Virtual Edge on VMware ESXi](#).

Enable SR-IOV on KVM

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

Prerequisites

This requires a specific NIC card. The following chipsets are certified by VMware to work with the SD-WAN Gateway and SD-WAN Edge.

- Intel 82599/82599ES
- Intel X710/XL710

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

Note SR-IOV mode is not supported if the KVM Virtual Edge is deployed with a High-Availability topology. For High-Availability deployments, ensure that SR-IOV is not enabled for that KVM Edge pair.

To enable SR-IOV on KVM:

- 1 Enable SR-IOV in BIOS. This will be dependent on your BIOS. Login to the BIOS console and look for SR-IOV Support/DMA. You can verify support on the prompt by checking that Intel has the correct CPU flag.

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

- 2 Add the options on Bboot (in /etc/default/grub).

```
GRUB_CMDLINE_LINUX="intel_iommu=on"
```

- a Run the following commands: `update-grub` and `update-initramfs -u`.
- b Reboot
- c Make sure iommu is enabled.

```
velocloud@KVMperf3:~$ dmesg | grep -i IOMMU
[ 0.000000] Command line: BOOT_IMAGE=/vmlinuz-3.13.0-107-generic root=/dev/mapper/qa--multiboot--002--vg-root ro intel_iommu=on splash quiet vt.handoff=7
[ 0.000000] Kernel command line: BOOT_IMAGE=/vmlinuz-3.13.0-107-generic root=/dev/mapper/qa--multiboot--002--vg-root ro intel_iommu=on splash quiet vt.handoff=7
[ 0.000000] Intel-IOMMU: enabled
...
velocloud@KVMperf3:~$
```

- 3 Based on the NIC chipset used, add a driver as follows:

- For the **Intel 82599/82599ES** cards in SR-IOV mode:

- 1 Download and install **ixgbe** driver from the [Intel](#) website.
- 2 Configure ixgbe config (tar and sudo make install).

```
velocloud@KVMperf1:~$ cat /etc/modprobe.d/ixgbe.conf
```

- 3 If the ixgbe config file does not exist, you have to create the file as follows.

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

- 4 Run the `update-initramfs -u` command and reboot the Server.
- 5 Use the `modinfo` command to verify if the installation is successful.

```
velocloud@KVMperf1:~$ modinfo ixgbe and ip link
filename: /lib/modules/4.4.0-62-generic/updates/drivers/net/ethernet/intel/ixgbe/ixgbe.ko
version: 5.0.4
```

```
license: GPL
description: Intel(R) 10GbE PCI Express Linux Network Driver
author: Intel Corporation, <linux.nics@intel.com>
srcversion: BA7E024DFE57A92C4F1DC93
```

- For the **Intel X710/XL710** cards in SR-IOV mode:

- 1 Download and install **i40e** driver from the [Intel](#) website.
- 2 Create the Virtual Functions (VFs).

```
echo 4 > /sys/class/net/device name/device/sriov_numvfs
```

- 3 To make the VFs persistent after a reboot, add the command from the previous step to the `/etc/rc.d/rc.local` file.
- 4 Deactivate the VF driver.

```
echo "blacklist i40evf" >> /etc/modprobe.d/blacklist.conf
```

- 5 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)
```

Install Virtual Edge on KVM

Describes 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 [Enable 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 VMware 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'>
  </virtualport>
  <vlan trunk='yes'>
    <tag id='33' nativeMode='untagged' />
    <tag id='200' />
    <tag id='201' />
    <tag id='202' />
  </vlan>
</network>
Bridge
<network>
  <name>passthrough</name>
  <model type='virtio' />
  <forward mode="bridge" />
</network>
<domain type='kvm'>
  <name>vedgel</name>
  <memory unit='KiB'>4194304</memory>
  <currentMemory unit='KiB'>4194304</currentMemory>
  <vcpu placement='static'>2</vcpu>
  <resource>
    <partition>/machine</partition>
  </resource>
  <os>
    <type arch='x86_64' machine='pc-i440fx-trusty'>hvm</type>
    <boot dev='hd' />
  </os>
  <features>
    <acpi />
    <apic />
    <pae />
  </features>
<!--
Set the CPU mode to host model to leverage all the available features on the host CPU
```

```

-->
<cpu mode='host-model'>
<model fallback='allow'/>
</cpu>
<clock offset='utc'/>
<on_poweroff>destroy</on_poweroff>
<on_reboot>restart</on_reboot>
<on_crash>restart</on_crash>
<devices>
<emulator>/usr/bin/kvm-spice</emulator>
<!--
Below is the location of the qcow2 disk image
-->
<disk type='file' device='disk'>
<driver name='qemu' type='qcow2'/>
<source file='/var/lib/libvirt/images/edge-VC_KVM_GUEST-x86_64-2.3.0-18- R23-20161114-GA-
updatable-ext4.qcow2'/>
<target dev='sda' bus='sata'/>
<address type='drive' controller='0' bus='0' target='0' unit='0'/>
</disk>
<!--
If using cloud-init to boot up virtual edge, attach the 2nd disk as CD-ROM
-->
<disk type='file' device='cdrom'>
<driver name='qemu' type='raw'/>
<source file='/home/vcadmin/cloud-init/vedgel/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>< tagid='#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>< tagid='#LAN2_VLAN#'></ vlan>

```

```

< aliasname='hostdev1' />
< addresstype='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
-->
< interfacetype='network'>
< modeltype='virtio' />
< sourcenetwork=WAN1 />
< vlan>< tagid='#hole2_vlan#' />< / vlan>
< aliasname=LAN1 />
< addresstype='pci' domain='0x0000' bus='0x00' slot='0x12' function='0x0' />
< / interface>
< interfacetype='network'>
< modeltype='virtio' />
< source network=LAN2 />
< vlan>< tag id='#LAN2_VLAN#' />< / vlan>
< aliasname='hostdev1' />
< addresstype='pci' domain='0x0000' bus='0x00' slot='0x13' function='0x0' />
< / interface>
< serial type='pty'>
< target port='0' />
< / serial>
< console type='pty'>
< target type='serial' port='0' />
< / console>
< input type='mouse' bus='ps2' />
< input type='keyboard' bus='ps2' />
< graphics type='vnc' port='-1' autoport='yes' listen='127.0.0.1'>
< listen type='address' address='127.0.0.1' />
< / graphics>
< sound model='ich6'>
< address type='pci' domain='0x0000' bus='0x00' slot='0x04' function='0x0' />
< / sound>
< video>
< model type='cirrus' vram='9216' heads='1' />
< address type='pci' domain='0x0000' bus='0x00' slot='0x02' function='0x0' />
< / video>
< memballoon model='virtio'>
< address type='pci' domain='0x0000' bus='0x00' slot='0x06' function='0x0' />
< / memballoon>
< / devices>
< / domain>

```

- 3 Save the domain XML file that defines the VM (for example, `vedge1.xml` created in step 2).
- 4 Launch the VM by performing the following steps:
 - a Create VM.

```
virsh define vedge1.xml
```

- b Start VM.

```
virsh start vedgel
```

Note `vedgel` is the name of the VM defined in the `<name>` element of the domain XML file. Replace `vedgel` 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 SD-WAN 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 SD-WAN Orchestrator using the activation key. The SD-WAN Orchestrator IP address and the activation key have been defined in the Cloud-init file.

Enable SR-IOV on VMware

Enabling SR-IOV on VMware is an optional configuration.

Prerequisites

This requires a specific NIC card. The following chipsets are certified by VMware to work with the SD-WAN Gateway.

- Intel 82599/82599ES
- Intel X710/XL710

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

To enable SR-IOV on VMware:

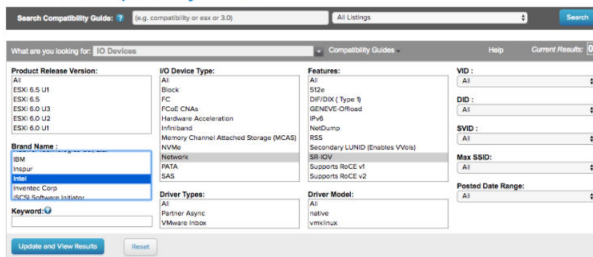
- 1 Make sure that your NIC card supports SR-IOV. Check the VMware Hardware Compatibility List (HCL) at <https://www.vmware.com/resources/compatibility/search.php?deviceCategory=io>

Brand Name: Intel

I/O Device Type: Network

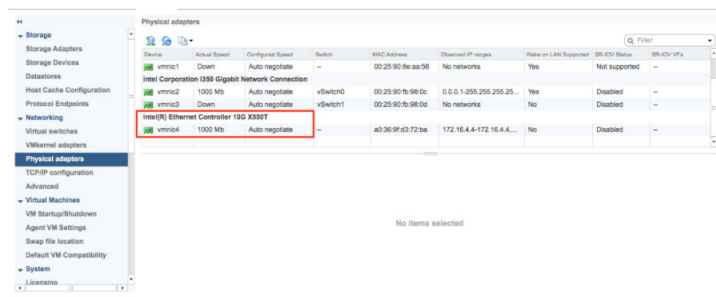
Features: SR-IOV

VMware Compatibility Guide

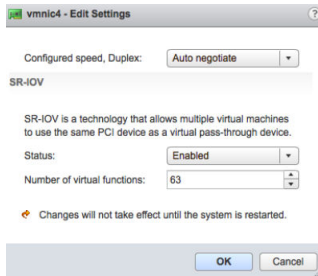


The following VMware KB article provides details of how to enable SR-IOV on the supported NIC: <https://kb.vmware.com/s/article/2038739>

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



- If SR-IOV is successfully enabled, the number of Virtual Functions (VFs) will show under the particular NIC after ESXi reboots.

Device	Actual Speed	Configured Speed	Switch	MAC Address	Observed IP ranges	Wake on LAN Supported	SR-IOV Status	SR-IOV VFs
vmnic4	1000 Mb	Auto negotiate	—	80:36:9f:d3:72:be	172.16.4.4-172.16.4.4	No	Enabled	63 (61 currently)
vmnic2	1000 Mb	Auto negotiate	vSwitch0	00:25:90:fb:98:0c	0.0.0.1-255.255.255.25...	Yes	Disabled	—
vmnic3	Down	Auto negotiate	vSwitch1	00:25:90:fb:98:0c	No networks	No	Disabled	—
vmnic0	Down	Auto negotiate	—	00:25:90:8e:aa:54	No networks	Yes	Not supported	—

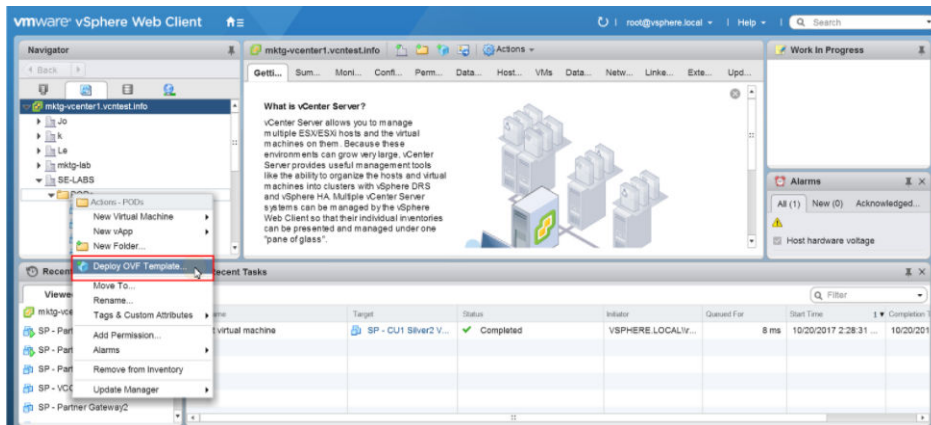
Install Virtual Edge on VMware ESXi

Describes how to install Virtual Edge on VMware ESXi.


If you decide to use SR-IOV mode, enable SR-IOV on VMware. For steps, see [Enable SR-IOV on VMware](#).

To install Virtual Edge on VMware ESXi:

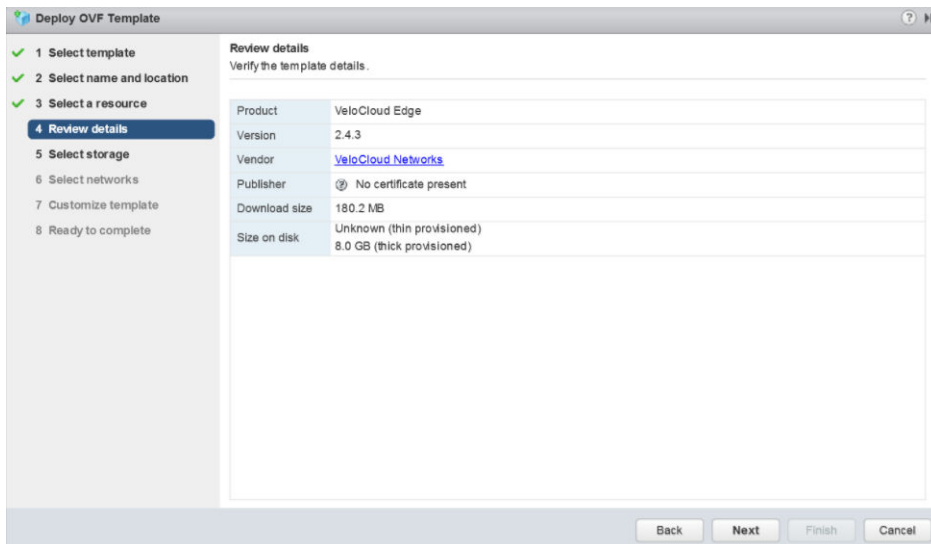
- Use the vSphere client to deploy an OVF template, and then select the Edge OVA file.



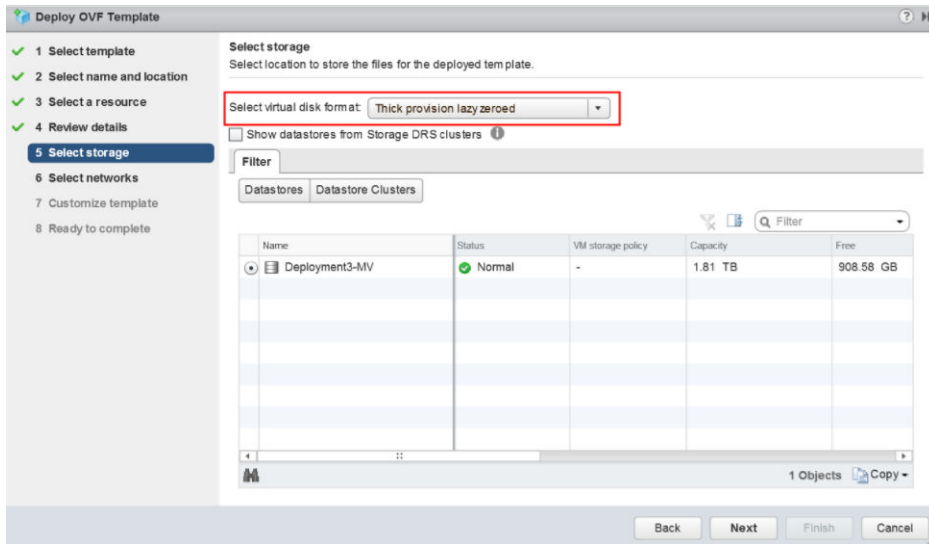
- Select an OVF template from an URL or Local file.



- 3 Select a name and location of the virtual machine.
- 4 Select a resource.
- 5 Verify the template details.

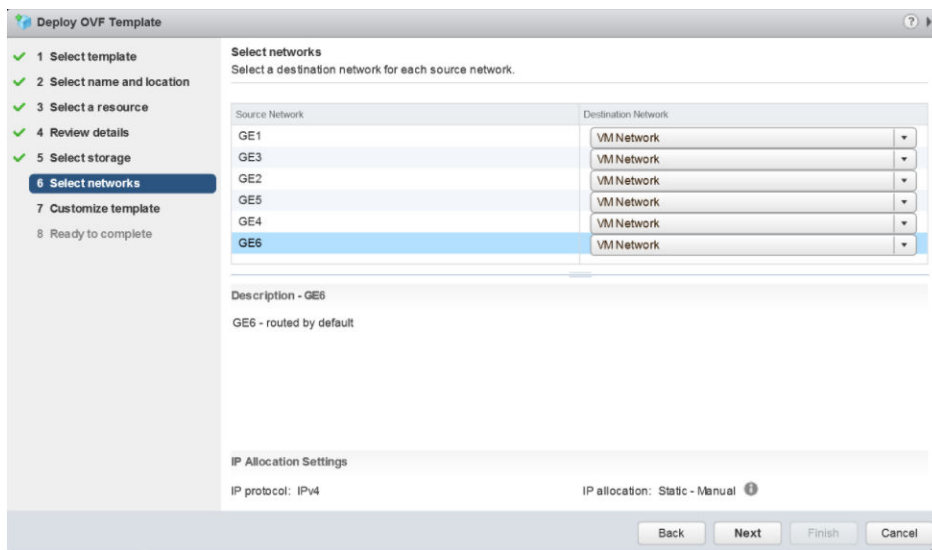


- 6 Select the storage location to store the files for the deployment template.




- 7 Configure the networks for each of the interfaces.

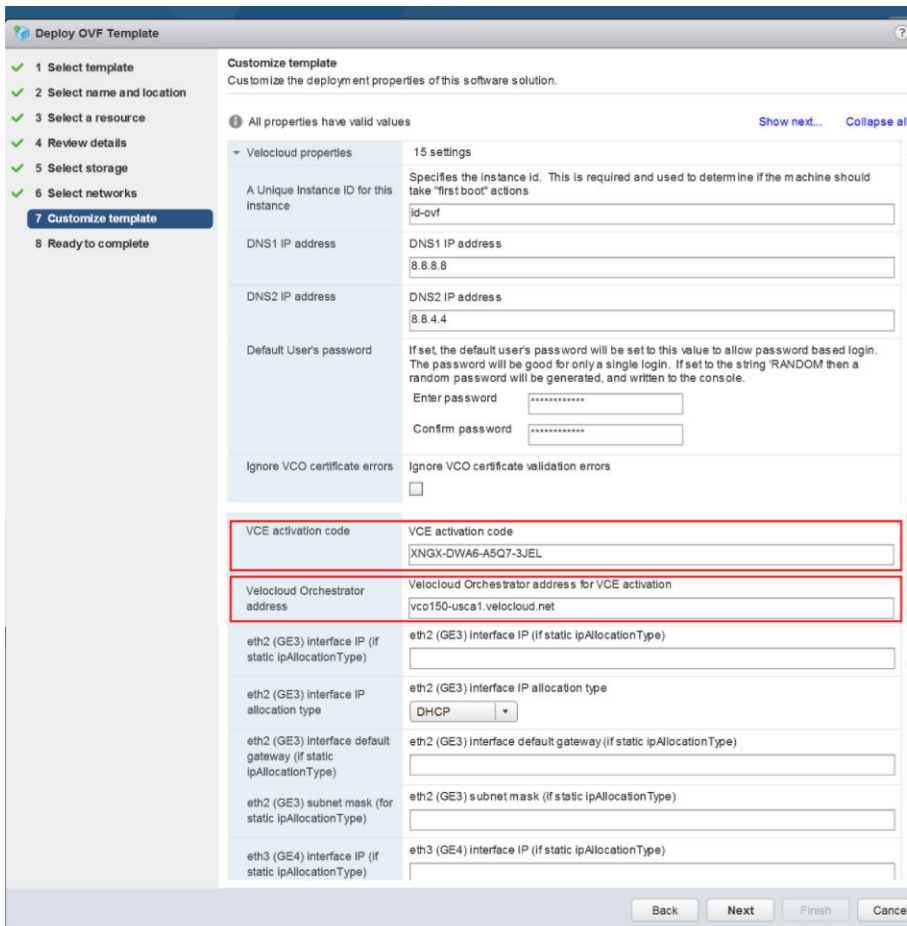
Note Skip this step if you are using a cloud-init file to provision the Virtual Edge on ESXi.

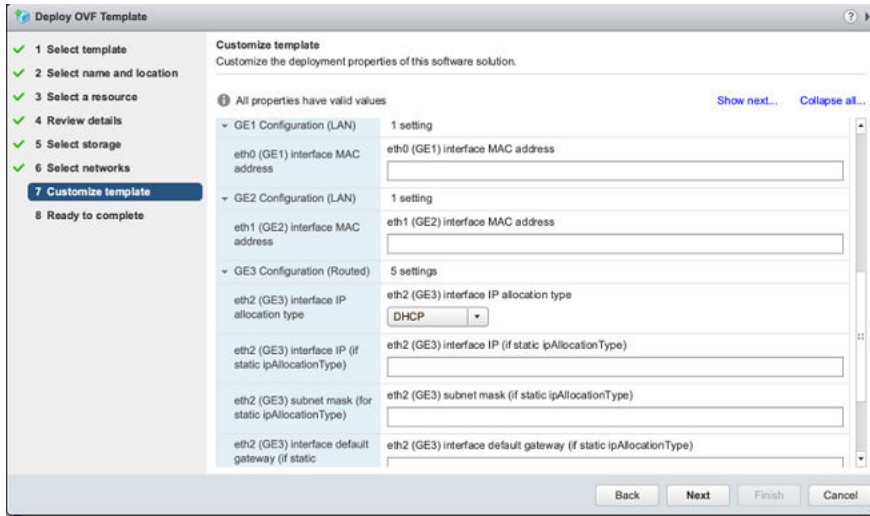


- 8 Customize the template by specifying the deployment properties. The following image highlights:
- From the SD-WAN Orchestrator UI, retrieve the URL/IP Address. You will need this address for Step c below.
 - Create a new Virtual Edge for the Enterprise. Once the Edge is created, copy the Activation Key. You will need the Activation Key for Step c" below.

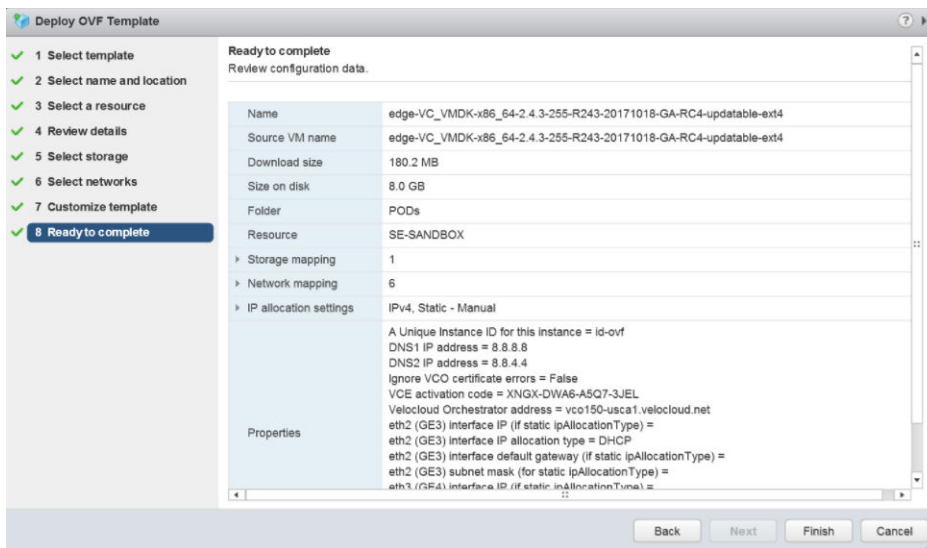


- c On the customize template page shown in the image below, type in the Activation Code that you retrieved in Step b above, and the SD-WAN Orchestrator URL/IP Address retrieved in Step a above, into the corresponding fields.

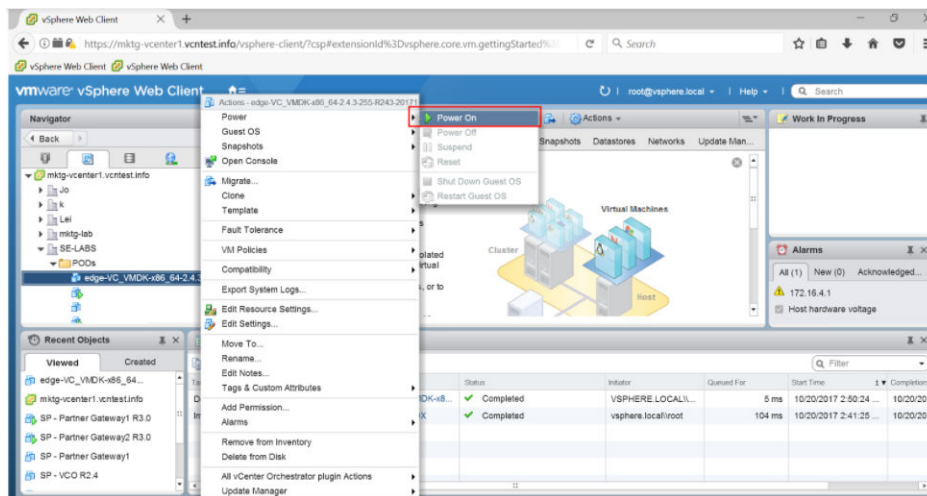




9 Review the configuration data.



10 Power on the Virtual Edge.



Once the Edge powers up, it will establish connectivity to the SD-WAN Orchestrator.

Read the following topics next:

- [Enterprise-Level Orchestrator Alerts and Events](#)
- [Supported VMware SD-WAN Edge Events for Syslogs](#)

Enterprise-Level Orchestrator Alerts and Events

Describes a summary of alerts and events generated within the VMware SD-WAN Orchestrator at the Enterprise level.

The document provides details about all Enterprise-level Orchestrator events. Although these events are stored within the SD-WAN Orchestrator and displayed on the Orchestrator UI, most of them are generated by either a VMware SD-WAN Edge or a VMware SD-WAN 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:

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 VMware SD-WAN component generating the notification can be one of the following: <ul style="list-style-type: none">■ SD-WAN Orchestrator■ SD-WAN Edge (MGD)■ SD-WAN Edge (EDGED)■ SD-WAN Edge (PROCMON)
GENERATED WHEN	Technical reason(s) and circumstances under which this event is generated.

Column name	Details
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

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
EDGE_UP	Edge Up	ALERT	SD-WAN Orchestrator	Edge comes back after losing connectivity with the SD-WAN Orchestrator through heartbeats. 2 consecutive heartbeats by an Edge causes the SD-WAN Orchestrator to change its status to EDGE_UP. The SD-WAN Orchestrator runs a monitor every 15 seconds that will update the status of all Edges.		
EDGE_DOWN	Edge Down	ALERT	SD-WAN Orchestrator	Edge loses connectivity with the SD-WAN Orchestrator and fails performing 2 or more consecutive heartbeats. The SD-WAN Orchestrator runs a monitor every 15 seconds that will update the status of all Edges.		

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
LINK_UP	Link Up	ALERT	SD-WAN Orchestrator	A WAN Link returns to a normal functioning state.		
LINK_DOWN	Link Down	ALERT	SD-WAN 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	SD-WAN 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	SD-WAN Orchestrator	An HA Edge fails-over to its standby.		
EDGE_SERVICE_DOWN	Edge Service Down	ALERT	SD-WAN Orchestrator	The Edge service running on the SD-WAN Edge may be down. This may indicate Edge device failure or failure of network connectivity.		
EDGE_CSS_TUNNEL_UP	Edge CSS Tunnel Up	ALERT	SD-WAN Orchestrator	A Cloud Security Service tunnel from Edge is UP.		

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
EDGE_CSS_TUNNEL_DOWN	Edge CSS Tunnel Down	ALERT	SD-WAN Orchestrator	A Cloud Security Service tunnel from Edge is DOWN.		
NVS_FROM_EDGE_TUNNEL_DOWN	NVS From Edge Tunnel Down	ALERT	SD-WAN Orchestrator	A NSD via Edge tunnel is DOWN.		
NVS_FROM_EDGE_TUNNEL_UP	NVS From Edge Tunnel Up	ALERT	SD-WAN Orchestrator	A NSD via Edge tunnel is UP.		
VNF_VM_DEPLOYED	VNF VM Deployed	ALERT	SD-WAN Orchestrator	An Edge VNF virtual machine gets deployed on to the Edge.		
VNF_VM_POWERED_ON	VNF VM Powered ON	ALERT	SD-WAN 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	SD-WAN Orchestrator	An Edge VNF virtual machine is powered off.		
VNF_VM_DEPLOYED_AND_POWERED_OFF	VNF VM Deployed and Powered OFF	ALERT	SD-WAN Orchestrator	An Edge VNF virtual machine gets deployed on to the Edge and is immediately powered on.		
VNF_VM_DELETED	VNF VM Deleted	ALERT	SD-WAN Orchestrator	An Edge VNF virtual machine is removed from the Edge.		
VNF_VM_ERROR	VNF VM error	ALERT	SD-WAN Orchestrator	An error occurs during deployment of an Edge VNF virtual machine.		

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
VNF_INSERTION_ENABLED	VNF insertion enabled	ALERT	SD-WAN Orchestrator	Insertion of an Edge VNF virtual machine is enabled on the Edge.		
VNF_INSERTION_DISABLED	VNF insertion disabled	ALERT	SD-WAN Orchestrator	Insertion of an Edge VNF virtual machine is deactivated on the Edge.		
VNF_IMAGE_DOWNLOAD_IN_PROGRESS	VNF Image Download In Progress	ALERT	SD-WAN Orchestrator	An Edge VNF virtual machine image download is in progress.		
VNF_IMAGE_DOWNLOAD_COMPLETED	VNF Image Download Completed	ALERT	SD-WAN Orchestrator	An Edge VNF virtual machine image download is completed.		
VNF_IMAGE_DOWNLOAD_FAILED	VNF Image Download Failed	ALERT	SD-WAN 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	SD-WAN Orchestrator	A BFD session has been established to Edge neighbor.		
EDGE_BFD_NEIGHBOR_DOWN	Edge BFD neighbor unavailable	INFO	SD-WAN Orchestrator	A BFD session to Edge neighbor is not established.		
PKI_PROMOTION	Endpoint PKI mode promoted	INFO	SD-WAN 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	SD-WAN Orchestrator	Edge certificate revocation occurs intentionally or due to an expired certificate (The latter should rarely happen, given Edge certificates automatically renew after 30 days into the 90 day period).		
CERTIFICATE_RENEWAL	Certificate renewal request	INFO	SD-WAN Orchestrator	Edge certificate automatically renews after 30 days into the 90 day period.		
UPDATE_EDGE_IMAGE_MANAGEMENT	Update Edge image management	INFO	SD-WAN Orchestrator	Activates/deactivates management of Edge software images for a customer.		
SET_EDGE_SOFTWARE	Updated Edge software image	INFO	SD-WAN 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.		

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
UNSET_EDGE_SOFTWARE	Unset overridden Edge software image	INFO	SD-WAN 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	SD-WAN Orchestrator	A new operator profile has been associated with this enterprise.		
REMOVE_OPERATOR_PROFILE	Removed operator profile	INFO	SD-WAN Orchestrator	An existing operator profile has been removed from this enterprise.		
ADD_SOFTWARE_IMAGE	Added software image	INFO	SD-WAN 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	SD-WAN Orchestrator	List of operator profiles associated with the Enterprise has been modified.		
MODIFY_ASSIGNED_SOFTWARE_IMAGE_LIST	Modified the assigned software image list	INFO	SD-WAN Orchestrator	List of software images associated with the Enterprise has been modified.		

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
CLOUD_SECURITY_ENABLED	Cloud Security enabled	INFO	SD-WAN Orchestrator	Cloud Security is activated in enterprise's profile or Edge-specific profile		
CLOUD_SECURITY_DISABLED	Cloud Security disabled	INFO	SD-WAN Orchestrator	Cloud Security is deactivated in enterprise's profile		
CLOUD_SECURITY_PROVIDER_DELETED	Cloud security provider deleted	INFO	SD-WAN Orchestrator	Cloud Security provider associated with an enterprise's profile has been deleted.		
CLOUD_SECURITY_TUNNELING_PROTOCOL_CHANGE	Cloud Security Tunneling Protocol Change	INFO	SD-WAN 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	SD-WAN Orchestrator	Cloud Security provider associated with an Edge-specific profile has been added.		
CLOUD_SECURITY_PROVIDER_REMOVED	CLOUD_SECURITY_PROVIDER_REMOVED	INFO	SD-WAN Orchestrator	Cloud Security provider associated with an Edge-specific profile has been removed.		
CLOUD_SECURITY_OVERRIDE_ENABLED	CLOUD_SECURITY_OVERRIDE_ENABLED	INFO	SD-WAN Orchestrator	Cloud Security override has been activated in an Edge-specific profile.		

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
CLOUD_SECURITY_OVERRIDE_DISABLED	CLOUD_SECURITY_OVERRIDE_DISABLED	INFO	SD-WAN Orchestrator	Cloud Security override has been deactivated in an Edge-specific profile.		
CREATE_CLOUD_SERVICE_SITE	Cloud Security Service site creation enqueued	INFO	SD-WAN 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	SD-WAN 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	SD-WAN 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	SD-WAN Orchestrator	An API automation job for Cloud Security Service Zscaler Sub Location as been enqueued.		
EDGE_NVS_TUNNEL_UP	Edge Direct IPsec tunnel up	INFO	SD-WAN Orchestrator	A Cloud Security Service tunnel or NSD via Edge tunnel is up.		

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
EDGE_NVS_TUNNEL_DOWN	Edge Direct IPsec tunnel down	INFO	SD-WAN Orchestrator	A Cloud Security Service tunnel or NSD via Edge tunnel is down.		
DIAGNOSTIC_REQUEST	New diagnostic bundle request	INFO	SD-WAN Orchestrator	A new Edge diagnostic bundle is requested by an enterprise or an operator user.		
EDGE_DIRECT_SITE_DELETED	Edge direct site deleted	INFO	SD-WAN Orchestrator	A NSD via Edge tunnel has been deleted.		
EDGE_DIRECT_TUNNELS_DISABLED	Edge direct tunnels disabled	INFO	SD-WAN Orchestrator	NSD via Edge deactivated in profile device settings.		
EDGE_DIRECT_TUNNELS_ENABLED	Edge direct tunnels enabled	INFO	SD-WAN Orchestrator	NSD via Edge enabled in profile device settings.		
EDGE_DIRECT_TUNNEL_PROVIDER_DELETED	Edge direct tunnel provider deleted	INFO	SD-WAN 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	SD-WAN 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	SD-WAN Orchestrator	An API automation job to update a NSD via Edge tunnel has been enqueued.		

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
DELETE_NVS_FROM_EDGE_SITE	NSD via Edge site deletion enqueued	INFO	SD-WAN 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	SD-WAN 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	SD-WAN Orchestrator	An enterprise has successfully delegated access to operator to manager its users.		
ENTERPRISE_DISABLE_OPERATOR_ACCESS	User management access revoked from operator	INFO	SD-WAN Orchestrator	An enterprise revokes access that was previously delegated to operator to manage its entities.		
ENTERPRISE_ENABLE_OPERATOR_ACCESS	Access delegated to operator	INFO	SD-WAN Orchestrator	An enterprise has successfully delegated access to operator to manager its entities.		
ENTERPRISE_ENABLE_PROXY_ACCESS	Access revoked from operator	INFO	SD-WAN Orchestrator	An enterprise has successfully delegated access to partner to manager its entities.		

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
ENTERPRISE_DISABLE_PROXY_ACCESS	Access delegated to partner	INFO	SD-WAN 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	SD-WAN 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	SD-WAN Orchestrator	Edge to Edge VPN associated with an Edge device or its corresponding profile has been enabled.		
VPN_DISABLE	Cloud VPN disabled	INFO	SD-WAN Orchestrator	Cloud VPN settings associated with an Edge device or its corresponding profile has been deactivated.		
VPN_ENABLE	Cloud VPN enabled	INFO	SD-WAN Orchestrator	When cloud VPN settings associated with an Edge device or its corresponding profile has been enabled.		

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
VPN_UPDATE	Cloud VPN updated	INFO	SD-WAN 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	SD-WAN Orchestrator	A remote action is performed on an online Edge.		
RECURRING_REPORT_ERROR	Recurring report error	ERROR	SD-WAN Orchestrator	When recurring report fails.		
CREATE_COMPOSITE_ROLE	Composite Role Created	INFO	SD-WAN Orchestrator	When a composite role is created by an Enterprise, Partner, or Operator.	4.5	
UPDATE_COMPOSITE_ROLE	Composite Role Updated	INFO	SD-WAN Orchestrator	When a composite role is updated by an Enterprise, Partner, or Operator.	4.5	
DELETE_COMPOSITE_ROLE	Composite Role Deleted	INFO	SD-WAN 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	SD-WAN Orchestrator	When sublocation configuration of Edge deviceSettings are modified.	4.5	

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
ENQUEUE_UPDATE_ZSCALER_SUBLOCATION	Zscaler Sub Location update enqueued	INFO	SD-WAN Orchestrator	When sublocation configuration of Edge deviceSettings are modified.	4.5	
ENQUEUE_DELETE_ZSCALER_SUBLOCATION	Zscaler Sub Location deletion enqueued	INFO	SD-WAN Orchestrator	When sublocation configuration of Edge deviceSettings are modified.	4.5	
CREATE_ZSCALER_SUBLOCATION	Zscaler Sub Location object created	INFO	SD-WAN Orchestrator	When sublocation configuration of Edge deviceSettings are modified.	4.5	
UPDATE_ZSCALER_SUBLOCATION	Zscaler Sub Location object updated	INFO	SD-WAN Orchestrator	When sublocation configuration of Edge deviceSettings are modified.	4.5	
DELETE_ZSCALER_SUBLOCATION	Zscaler Sub Location object deleted	INFO	SD-WAN Orchestrator	When sublocation configuration of Edge deviceSettings are modified.	4.5	
ENQUEUE_UPDATE_ZSCALER_LOCATION	Zscaler Location update enqueued	INFO	SD-WAN Orchestrator	When location configuration of Edge deviceSettings are modified.	4.5	
CREATE_ZSCALER_LOCATION	Zscaler Location object created	INFO	SD-WAN Orchestrator	When location configuration of Edge deviceSettings are modified.	4.5	
UPDATE_ZSCALER_LOCATION	Zscaler Location object updated	INFO	SD-WAN Orchestrator	When location configuration of Edge deviceSettings are modified.	4.5	

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
DELETE_ZSCALER_LOCATION	Zscaler Location Object deleted	INFO	SD-WAN Orchestrator	When location configuration of Edge deviceSettings are modified.	4.5	
EDGE_BGP_NEIGHBOR_UP	BGP session established to Edge neighbor	INFO	SD-WAN Edge	Generated when a BGP peer establishes tunnel with an SD-WAN Edge.		
EDGE_BGP_NEIGHBOR_DOWN	Edge BGP neighbor unavailable	INFO	SD-WAN Edge	Generated when the Edge's BGP peer loses tunnel with the Edge.		
GATEWAY_BGP_NEIGHBOR_UP	BGP session established to Gateway neighbor	INFO	SD-WAN Gateway	When a BGP peer establishes tunnel with a Gateway.		
GATEWAY_BGP_NEIGHBOR_DOWN	Gateway BGP neighbor unavailable	INFO	SD-WAN Gateway	When a Gateway's BGP peer loses tunnel with a Gateway.		
VRF_MAX_LIMIT_EXCEEDED	VMware SD-WAN Partner Gateway: Maximum rules in a routemap limit hit for enterprise <enterprise-name>	WARNING	SD-WAN Gateway	Maximum inbound route map config limit reached.		
VRF_ROUTE_MAP_RULES_MAX_LIMIT_HIT	VMware SD-WAN Partner Gateway: Maximum rules in a routemap limit hit for enterprise <enterprise-name>	WARNING	SD-WAN Gateway	Maximum outbound route map config limit reached.		
VRF_LIMIT_EXCEEDED	VMware SD-WAN gateway: Maximum VRF limit(1000) reached	ALERT	SD-WAN Gateway	Maximum VRF limit reached for Partner Gateway.		
GATEWAY_STARTUP	VMware SD-WAN gateway service started	INFO	SD-WAN Gateway	Gateway daemon has started.		

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
ZSCALER_MONITOR_DISABLED	Zscaler monitor disabled	CRITICAL	SD-WAN Edge/SD-WAN Gateway (PROCمون)	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	SD-WAN Edge/SD-WAN Gateway (PROCمون)	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	SD-WAN Edge/SD-WAN Gateway (PROCمون)	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	SD-WAN Edge/SD-WAN Gateway (PROCمون)	Generated when procmon starts the services.	4.5	
EDGE_SERVICES_STOPPED/ GATEWAY_SERVICES_STOPPED	Edge/Gateway Services Stopped	INFO	SD-WAN Edge/SD-WAN Gateway (PROCمون)	Generated when procmon stops all the services.	4.5	
EDGE_SERVICES_RESTARTED/ GATEWAY_SERVICES_RESTARTED	Edge/Gateway Services Restarted	INFO	SD-WAN Edge/SD-WAN Gateway (PROCمون)	Generated when procmon restarts all the services.	4.5	
EDGE_SERVICES_TERMINATED/ GATEWAY_SERVICES_TERMINATED	Edge/Gateway Services terminated	INFO	SD-WAN Edge/SD-WAN Gateway (PROCمون)	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	SD-WAN 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	SD-WAN Gateway (PROCMON)	Generated by vc_procm on Gateway when MGD gets stopped.	4.4	
GATEWAY_NAT_SERVICE_FAILED	Service natd failed with error, restarting	ERROR	SD-WAN Gateway (PROCMON)	Generated by vc_procm on Gateway when natd daemon gets stopped.	4.4	
EDGE_DNSMASQ_FAILED	dnsmasq FAILED to start up	ERROR	SD-WAN Edge (PROCMON)	Generated when dnsmasq daemon failed to start up.	4.4	
EDGE_SSH_LOGIN	sshd accepted connection	INFO	SD-WAN 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	SD-WAN 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	SD-WAN Edge (PROCMON)	LED service deactivated.		
EDGE_LED_SERVICE_FAILED	Edge front-panel LED service failed	ERROR	SD-WAN Edge (PROCMON)	LED service failed.		
EDGE_MGD_SERVICE_DISABLED	Management service disabled	CRITICAL	SD-WAN Edge (PROCMON)	Management service is unable to activate for too many failures.		

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
EDGE_MGD_SERVICE_FAILED	Management service failed	ERROR	SD-WAN Edge (PROCMON)	Management service failed.		
EDGE_SERVICE_DISABLED	Edge dataplane service disabled	WARNING/ CRITICAL	SD-WAN Edge (PROCMON)	Edge Dataplane service is deactivated.		
EDGE_SERVICE_ENABLED	Edge dataplane service enabled	WARNING	SD-WAN Edge (PROCMON)	Edge Dataplane service is activated by user from local UI.		
EDGE_SERVICE_FAILED	Edge dataplane service failed	ERROR	SD-WAN Edge (PROCMON)	Edge Dataplane service failed.		
EDGE_VNFD_SERVICE_DISABLED		WARNING	SD-WAN Edge (PROCMON)	Edge VNFD service deactivated.		
EDGE_VNFD_SERVICE_FAILED		ERROR	SD-WAN Edge (PROCMON)	Edge VNFD service failed.		
EDGE_DOT1X_SERVICE_DISABLED	Edge 802.1x service disabled	WARNING, CRITICAL	SD-WAN Edge (PROCMON)	SD-WAN Edge 802.1x service is deactivated.		
EDGE_DOT1X_SERVICE_FAILED	Edge 802.1x service failed	ERROR	SD-WAN Edge (PROCMON)	SD-WAN Edge 802.1x service failed.		
EDGE_NYANSA_SYSLOG_SERVICE_FAILED		ERROR	SD-WAN Edge (PROCMON)	Nyansa Syslog service failed.		
EDGE_NYANSA_SYSLOG_SERVICE_DISABLED		WARNING	SD-WAN Edge (PROCMON)	Nyansa Syslog service deactivated.		
EDGE_NYANSA_AMOND_SERVICE_FAILED		ERROR	SD-WAN Edge (PROCMON)	Nyansa Amond service failed.		
EDGE_NYANSA_AMOND_SERVICE_DISABLED		WARNING	SD-WAN Edge (PROCMON)	Nyansa Amond service deactivated		

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
EDGE_NYANSA_SNMP_TRAPD_SERVICE_FAILED		ERROR	SD-WAN Edge (PROCMON)	Nyansa SNMP Trapd service failed.		
EDGE_NYANSA_SNMP_TRAPD_SERVICE_DISABLED		WARNING	SD-WAN Edge (PROCMON)	Nyansa SNMP Trapd service deactivated.		
EDGE_NYANSA_SNMP_READER_SERVICE_FAILED		ERROR	SD-WAN Edge (PROCMON)	Nyansa SNMP Reader service failed.		
EDGE_NYANSA_SNMP_READER_SERVICE_DISABLED		WARNING	SD-WAN Edge (PROCMON)	Nyansa SNMP Reader service deactivated.		
EDGE_USB_PORTS_ENABLED/ GATEWAY_USB_PORTS_ENABLED	Edge/Gateway USB ports Enabled	INFO	SD-WAN Edge/SD-WAN Gateway (MGD)	Generated when USB ports is activated.	4.5	
EDGE_USB_PORTS_DISABLED/ GATEWAY_USB_PORTS_DISABLED	Edge/Gateway USB ports Disabled	INFO	SD-WAN Edge/SD-WAN 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	SD-WAN Edge/SD-WAN Gateway (MGD)	Generated when procmon activates USB ports failure.	4.5	

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
EDGE_USB_PORTS_DISABLE_FAILURE/ GATEWAY_USB_PORTS_DISABLE_FAILURE	Edge/Gateway USB ports Disable Failure	CRITICAL	SD-WAN Edge/SD-WAN Gateway (MGD)	Generated when procmon deactivates USB ports failure.	4.5	
VNF_VM_EVENT	VNF VM Event	INFO	SD-WAN 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	SD-WAN Edge (MGD)	VNF insertion is activated or deactivated. Event detail will help distinguish the type.		
VNF_IMAGE_DOWNLOAD_EVENT	VNF image download event	INFO	SD-WAN Edge (MGD)	VNF download is in progress, completed, or failed. Event detail will help distinguish the type.		
MGD_START	Online	INFO	SD-WAN Edge (MGD)	Management daemon on Edge has started.		
MGD_EXITING	Shutting Down	INFO	SD-WAN Edge (MGD)	Management service on a SD-WAN Edge is shutting down for a restart.		

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
MGD_SET_CERT_SUCCESS	Set Certificate Successful	INFO	SD-WAN Edge (MGD)	New PKI certificate for Orchestrator communication is installed successfully on a SD-WAN Edge.		
MGD_SET_CERT_FAIL	Set Certificate Failed	ERROR	SD-WAN Edge (MGD)	Installation of a new PKI certificate for Orchestrator communication on a SD-WAN Edge has failed.		
MGD_CONFIG_APPLIED	Configuration Applied	INFO	SD-WAN Edge (MGD)	Configuration change made on the Orchestrator has been pushed to SD-WAN Edge and is successfully applied.		
MGD_CONFIG_PENDING	New configuration pending	INFO	SD-WAN Edge (MGD)	New configuration is pending application (This event is currently NOT generated anywhere)		
MGD_CONFIG_ROLLBACK	Bad configuration rolled back	CRITICAL	SD-WAN Edge (MGD)	Configuration policy sent from the Orchestrator had to be rolled back because it destabilized the SD-WAN Edge.		

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
MGD_CONF_FAILED	Failed to apply configuration	ERROR	SD-WAN Edge (MGD)	Edge failed to apply a configuration change made on the Orchestrator.		
MGD_CONF_UPDATE_INVALID	Invalid software update configuration	WARNING	SD-WAN Edge (MGD)	Edge has been assigned an Operator Profile with an invalid software image that the Edge cannot use.		
MGD_DEVICE_CONFIG_WARNING		WARNING	SD-WAN Edge (MGD)	Inconsistent device settings are detected. MGD continues with warnings.		
MGD_DEVICE_CONFIG_ERROR		ERROR	SD-WAN Edge (MGD)	Invalid device settings are detected by MGD.		
MGD_SWUP_IGNORED_UPDATE	Software update ignored	INFO	SD-WAN Edge (MGD)	Software update is ignored at the activation time, because SD-WAN Edge is already running that version.		
MGD_SWUP_INVALID_SWUPDATE	Invalid software update	WARNING	SD-WAN Edge (MGD)	Software update package received from the Orchestrator is invalid.		
MGD_SWUP_DOWNLOAD_FAILED	Software download failed	ERROR	SD-WAN Edge (MGD)	Download of an Edge software update image has failed.		

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
MGD_SWUP_UNPACK_FAILED	Software update unpack failed	ERROR	SD-WAN Edge (MGD)	Edge has failed to unpack the downloaded software update package.		
MGD_SWUP_INSTALL_FAILED	Software update install failed	ERROR	SD-WAN Edge (MGD)	Edge software update installation failed.		
MGD_SWUP_INSTALLED	Software update	INFO	SD-WAN Edge (MGD)	Software update was successfully downloaded and installed.		
MGD_SWUP_REBOOT	Restart after software update	INFO	SD-WAN Edge (MGD)	Edge is being rebooted after a software update.		
MGD_SWUP_STANDBY_UPDATE_START	Standby device software update started	INFO	SD-WAN 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 SD-WAN Orchestrator.		
MGD_SWUP_STANDBY_UPDATE_FAILED	Standby device software update failed	ERROR	SD-WAN 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		

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
MGD_SWUP_STANDBY_UPDATED	Standby device software update completed	INFO	SD-WAN Edge (MGD)	When Active Edge detects standby comes up with expected image version		
MGD_VCO_ADDR_RES_OLV_FAILED	Cannot resolve Orchestrator address	WARNING	SD-WAN Edge (MGD)	DNS resolution of the Orchestrator address failed.		
MGD_DIAG_REBOOT	User-initiated restart	INFO	SD-WAN Edge (MGD)	Edge is rebooted by a Remote Action from the Orchestrator.		
MGD_DIAG_RESTART	Services restarted	INFO	SD-WAN Edge (MGD)	Data plane service on the SSD-WAN Edge is restarted by a Remote Action from the Orchestrator.		
MGD_SHUT_DOWN	Powered off	INFO	SD-WAN Edge (MGD)	Edge diagnostic shutdown based on user request.		
MGD_HARD_RESET	Reset to factory defaults	INFO	SD-WAN Edge (MGD)	Edge is restored to its factory-default software and configuration.		
MGD_DEACTIVATED	Deactivated	INFO	SD-WAN Edge (MGD)	Edge is deactivated based on user request by mgd.		
MGD_NETWORK_SETTINGS_UPDATED	Network settings updated	INFO	SD-WAN Edge (MGD)	Network settings are applied to a SD-WAN Edge.		
MGD_NETWORK_MGMT_IF_BROKEN	Management Network incorrectly set up	ALERT	SD-WAN Edge (MGD)	Management network is set up incorrectly.		

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
MGD_NETWORK_MGMT_IF_FIXED	Network was restarted twice to fix Management Network inconsistency	WARNING	SD-WAN Edge (MGD)	Network is restarted twice to fix the Management Network inconsistency.		
MGD_INVALID_VCO_ADDRESS	Unable to heartbeat to new VCO %(newprimary)s, keep talking to old VCO %(oldprimary)s	WARNING	SD-WAN Edge (MGD)	Invalid address for Orchestrator was sent in a management plane policy update and was ignored.		
MGD_ACTIVATION_PARTIAL	Activation incomplete	INFO	SD-WAN Edge (MGD)	Edge is activated partially, but a software update failed.		
MGD_ACTIVATION_SUCCESS	Activated	INFO	SD-WAN Edge (MGD)	Edge has been activated successfully.		
MGD_ACTIVATION_ERROR	Activation failed	ERROR	SD-WAN 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	SD-WAN Edge (MGD)	Standby Edge send this event when HA is deactivated.		
EDGE_INTERFACE_DOWN	Edge Interface Down	INFO	SD-WAN Edge (MGD)	Generated by hotplug scripts when the interface is down.		
EDGE_INTERFACE_UP	Edge Interface Up	INFO	SD-WAN Edge (MGD)	Generated by hotplug scripts when the interface is up.		

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
EDGE_KERNEL_PANIC		ALERT	SD-WAN 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.		
EDGE_OSPF_NSM	Edge OSPF NSM Event	INFO	SD-WAN Edge (EDGED)	Edge send this event when OSPF neighbor state changes.		
IP_SLA_PROBE	IP SLA Probe	INFO	SD-WAN Edge (EDGED)	Edge generates when IPSLA state changes.		
IP_SLA_RESPONDER	IP SLA Responder	ALERT, INFO	SD-WAN Edge (EDGED)	When IPSLA responder state changes from up to down and vice versa.		
ALL_CSS_DOWN	ALL_CSS_DOWN	ALERT	SD-WAN Edge (EDGED)	When all CSS paths go down.		
CSS_UP	CSS_UP	ALERT	SD-WAN Edge (EDGED)	When atleast one CSS path is up.		

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
LINK_MTU	Link MTU detected	INFO	SD-WAN 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	SD-WAN 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	SD-WAN Edge (EDGED)	Peer is unusable.		Deprecated

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
PEER_USABLE	Peer usable	INFO	SD-WAN Edge (EDGED)	Peer is usable.		Deprecated
BW_UNMEASURABLE	Error measuring bandwidth	ALERT	SD-WAN Edge (EDGED)	Bandwidth measurement failed to the Primary Gateway. Reattempt at measurement in 30minutes. Reasons include a link experiencing 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.		
SLOW_START_CAP_MET	Bandwidth measured exceeds the slow start cap. Moving to burst mode.	NOTICE	SD-WAN 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	SD-WAN Edge (EDGED)	BFD configured with incorrect local address.		

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
FLOOD_ATTACK_DETECTED		INFO	SD-WAN Edge (EDGED)	Generated when a malicious host floods the SD-WAN Edge with new connections.		
LINK_ALIVE	Link alive	INFO	SD-WAN Edge (EDGED)	When link state (link_fsm) becomes alive.		
LINK_DEAD	Link dead	ALERT	SD-WAN Edge (EDGED)	When link state (link_fsm) becomes dead.		
LINK_USABLE	Link usable	INFO	SD-WAN Edge (EDGED)	When link state (link_fsm) becomes usable.		
LINK_UNUSABLE	Link unusable	ALERT	SD-WAN Edge (EDGED)	When link state (link_fsm) becomes unusable.		
VPN_DATA_CENTER_STATUS	VPN Tunnel state change	INFO, ERROR	SD-WAN Edge (EDGED)	VPN Tunnel state change.		
INTERFACE_CONFIG_ERROR	Interface config error	ALERT	SD-WAN Edge (EDGED)			
HA_STANDBY_ACTIVATED	HA Standby Activated	INFO	SD-WAN Edge (EDGED)	When active Edge detects standby peer send this event to SD-WAN Orchestrator to activate standby Edge.		
HA_INTF_STATE_CHANGED	HA Interface State Changed	ALERT	SD-WAN Edge (EDGED)	HA interface went down/up.		
HA_GOING_ACTIVE	High Availability Going Active	INFO	SD-WAN Edge (EDGED)	Standby Edge transition to Active Edge after detecting no heartbeat for more than 700ms.		

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
HA_FAILED	High Availability Peer State Unknown	INFO	SD-WAN Edge (EDGED)	Active Edge detects no heartbeat or activity from standby Edge for more than 700 milliseconds.		
HA_READY	High Availability Ready	INFO	SD-WAN Edge (EDGED)	Active Edge detects activated standby peer.		
MGD_UNREACHABLE	Management Proxy unreachable	EMERGENCY	SD-WAN Edge (EDGED)	Data plane process could not communicate to the management plane proxy.		
VRRP_INTOMASTER_STATE	VRRP HA updated to primary state	INFO	SD-WAN Edge (EDGED)	VRRP get into primary state		
VRRP_OUTOFMASTER_STATE	VRRP HA updated out of primary state	INFO	SD-WAN Edge (EDGED)	VRRP get out of primary state.		
VRRP_FAILINFO	VRRP failed	INFO	SD-WAN Edge (EDGED)	VRRP failed.		
EDGE_HEALTH_ALERT	Edge Health Alert	EMERGENCY	SD-WAN Edge (EDGED)	Data plane is unable to allocate necessary resources for packet processing.		
EDGE_STARTUP	Edge service startup	INFO	SD-WAN Edge (EDGED)	Edge is running in mgmt-only mode.		
EDGE_DHCP_BAD_OPTION	Invalid DHCP Option	WARNING	SD-WAN Edge (EDGED)	SD-WAN Edge is configured with an invalid DHCP option.		

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
EDGE_NEW_USER	New client user seen	INFO	SD-WAN Edge (EDGED)	New or updated client user detected on a given MAC address.		
EDGE_NEW_DEVICE	New client device seen	INFO	SD-WAN Edge (EDGED)	A new device is detected during DHCP.		
INVALID_JSON		CRITICAL	SD-WAN Edge (EDGED)	The Edged received invalid json data from the mgd.		
QOS_OVERRIDE	QoS override	INFO	SD-WAN 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	SD-WAN Edge (EDGED)	Edge L2 loop is detected.		
EDGE_TUNNEL_CAPACITY_WARNING	Edge Tunnel CAP warning	WARNING	SD-WAN Edge (EDGED)	Edge has reached its maximum tunnel capacity.		
LOS_DETECTED	LoS no longer seen on interface <iface-name>/ LoS detected on interface <iface-name>	ALERT	SD-WAN Edge (EDGED)	Loss of Signal state changed on the interface in HA setup.	4.4	
EDGE_LOCAL_UI_LOGIN	Edge Local UI Login	INFO	SD-WAN Edge	LOCAL UI login is successful for a user.		

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
EDGE_MEMORY_USAGE_ERROR	Memory Usage Critical	ERROR	SD-WAN 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.		

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
EDGE_MEMORY_USAGE_WARNING	Memory Usage Warning	WARNING	SD-WAN 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	SD-WAN Edge	User initiates a Edge service restart.		
EDGE_REBOOTING	User-initiated Edge reboot	WARNING	SD-WAN Edge	User initiates a Edge reboot.		
EDGE_HARD_RESET	User-initiated Edge hard reset	WARNING	SD-WAN Edge	Edge hard reset		
EDGE_DEACTIVATED	Edge deactivated	WARNING	SD-WAN Edge	SD-WAN 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	SD-WAN Edge	SD-WAN Edge login via console port.		
EDGE_COMMAND	Edge Command	INFO	SD-WAN Edge	Generated by a SD-WAN Edge during remote diagnostics when executing Edge commands.		

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
EDGE_BIOS_UPDATED	Edge BIOS updated	INFO	SD-WAN Edge	Generated by 12-upgrade-bios.sh script when SD-WAN Edge BIOS is successfully updated.		
EDGE_BIOS_UPDATE_FAILED	Edge BIOS update failed	ERROR	SD-WAN Edge	Generated by 12-upgrade-bios.sh script when SD-WAN Edge BIOS update failed.		
IPV6_ADDR_DELETED	Deleted IPv6 address <v6addr> on interface/sub-interface <iface/subiface name>	INFO	SD-WAN Edge/SD-WAN 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 <iface-name>	INFO	SD-WAN Edge	When IPv6 address is added on interface.	4.4	
IPV6_ADDR_DEPRECATED	Deprecated IPv6 address <v6-addr> on interface <iface-name>	INFO	SD-WAN Edge	When IPv6 address gets deprecated on an interface.	4.4	
IPV6_ADDR_PREFERRED	Preferred IPv6 address <v6-addr> on interface <iface-name>	INFO	SD-WAN 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	SD-WAN 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	SD-WAN Edge	When IPv6 NDP DAD is failed.	4.4	
EDGE_SHUTTING_DOWN	Edge is shutting down - must be restarted by power-cycling	WARNING	SD-WAN Edge (LUA Backend)	When Edge is shutting down.	4.4	

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
BIOS_PHY_RESET_CMOS_SET	BIOS - Phy reset CMOS bit is set/ BIOS - Phy reset CMOS bit can't be set	WARNING	SD-WAN 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	SD-WAN Edge	When CPLD Firmware is being updated during software upgrade.	4.4	
EVDSL_INTERFACE_UP_EVENT	Contains json string with evdslModem name, status, serial number	INFO	SD-WAN Edge	Generated when EVDSL interface moves to Up state.	4.5	
EVDSL_INTERFACE_DOWN_EVENT	contains json string with evdslModem name, status, serial number	INFO	SD-WAN Edge	Generated when EVDSL interface moves to Down state.	4.5	
NAT_PORT_ASSIGN_FAILURE	NAT Ports exhausted from <src_ip> to <dst_ip>:<dport>	WARNING	SD-WAN Edge/SD-WAN Gateway	Generated when NAT port allocation range is exhausted.	4.5	
IPv6_MAX_DAD_FAILURE	IPv6 <link local / RA > stable secret address generation failed on interface <iface name> after multiple DAD failures	ALERT	SD-WAN Edge	Generated when we fail to generate stateless IPv6 address after multiple DAD failures.	4.5	
IPv6_ADDRESS_GENERATION_FAILURE	IPv6 <link local / RA> stable secret address generation failed on interface <iface name> after generating multiple invalid addresses	ALERT	SD-WAN 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	SD-WAN Edge	Generated for invalid static route.	4.5	

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
INVALID OSPF_ROUTE	Rejected invalid routes <route-prefix>/0 flag <route flags in hex>	ALERT	SD-WAN Edge	Generated for invalid OSPF routes.	4.5	
INVALID_BGP_ROUTE	Rejected invalid routes <route-prefix>/0 flag <route flags in hex>	ALERT	SD-WAN Edge	Generated for invalid BGP routes.	4.5	
INVALID_REMOTE_OSPF_ROUTE	Rejected invalid routes <route-prefix>/0 flag <route flags in hex>	ALERT	SD-WAN 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	SD-WAN Edge	Generated for invalid remote BGP route.	4.5	
INVALID_OVERLAY_ROUTE	Rejected invalid routes <route-prefix>/0 flag <route flags in hex>	ALERT	SD-WAN Edge	Generated for invalid Overlay route.	4.5	
INVALID_ROUTE	Rejected invalid routes <route-prefix>/0 flag <route flags in hex>	ALERT	SD-WAN Edge	Generated for invalid routes.	4.5	
EDGE_BFDv6_CONFIG	Incorrect local address <IP address>. IP Address not present	INFO	SD-WAN Edge	Generated when invalid IPv6 BFD configuration is received.	4.5	
EDGE_USB_DEVICE_INSERTED	Edge USB device inserted	ALERT	SD-WAN Edge	Generated when USB device is inserted.	4.5	
EDGE_USB_DEVICE_REMOVED	Edge USB device removed	ALERT	SD-WAN 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	SD-WAN Edge	Generated when WiFi card at a port is no longer usable.	4.5	

Supported VMware SD-WAN Edge Events for Syslogs

The following table describes all the possible VMware SD-WAN Edge events that could be exported to syslog collectors.

Events	Severity	Description
BW_UNMEASURABLE	ALERT	Generated by a SD-WAN Edge when the path bandwidth is unmeasurable.
EDGE_BIOS_UPDATE_FAILED	ERROR	Generated by 12-upgrade-bios.sh script when SD-WAN Edge BIOS is updated.
EDGE_BIOS_UPDATED	INFO	Generated by 12-upgrade-bios.sh script when SD-WAN Edge BIOS update failed.
EDGE_CONSOLE_LOGIN	INFO	Generated by a SD-WAN Edge during login via console port.
EDGE_DEACTIVATED	WARNING	Generated when a SD-WAN 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 SD-WAN Edge is configured with an invalid DHCP option.
EDGE_DISK_IO_ERROR	WARNING	Generated by a SD-WAN Edge when the Disk IO error has occurred during upgrade/downgrade.
EDGE_DISK_READONLY	CRITICAL	Generated by a SD-WAN 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 SD-WAN Edge 802.1x service is deactivated.
EDGE_DOT1X_SERVICE_FAILED	ERROR	Generated by vc_procmon when the SD-WAN Edge 802.1x service failed.
EDGE_HARD_RESET	WARNING	Generated when user has initiated SD-WAN Edge hard reset.
EDGE_HEALTH_ALERT	EMERGENCY	Generated by the SD-WAN 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.

Events	Severity	Description
EDGE_KERNEL_PANIC	ALERT	Generated by a SD-WAN 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 SD-WAN EdgeL2 loop is detected.
EDGE_LED_SERVICE_DISABLED	WARNING, CRITICAL	Generated by vc_procmon when the SD-WAN Edge LED service is deactivated.
EDGE_LED_SERVICE_FAILED	ERROR	Generated by vc_procmon when the SD-WAN 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 SD-WAN 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 SD-WAN 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.
EDGE_NEW_USER	INFO	Generated when a new client user is added.

Events	Severity	Description
EDGE_OSPF_NSM	INFO	Generated by the SD-WAN Edge when the OSPF Neighbor state Machine (NSM) state occurred.
EDGE_REBOOTING	WARNING	Generated when a user has initiated SD-WAN Edge reboot.
EDGE_RESTARTING	WARNING	Generated when a user has initiated SD-WAN Edge service restart.
EDGE_SERVICE_DISABLED	WARNING	Generated when the SD-WAN Edge data plane service is deactivated.
EDGE_SERVICE_ENABLED	WARNING	Generated when the SD-WAN Edge data plane service is enabled.
EDGE_SERVICE_FAILED	ERROR	Generated when the SD-WAN Edge data plane service failed.
EDGE_SHUTTING_DOWN	WARNING	Generated when a SD-WAN Edge is shutting down.
EDGE_STARTUP	INFO	Generated when a SD-WAN Edge is running in mgmt-only mode.
EDGE_SSH_LOGI	INFO	Generated by a SD-WAN Edge during login via SSH protocol.
EDGE_TUNNEL_CAP_WARNING	WARNING	Generated when a SD-WAN Edge has reached its maximum tunnel capacity.
EDGE_USB_PORTS_ENABLED	INFO	Generated when USB ports are enabled on a SD-WAN Edge.
EDGE_USB_PORTS_DISABLED	INFO	Generated when USB ports are deactivated on a SD-WAN Edge.
EDGE_USB_PORTS_ENABLE_FAILURE	CRITICAL	Generated by a SD-WAN Edge when the enable operation for its USB ports fails.
EDGE_USB_PORTS_DISABLE_FAILURE	CRITICAL	Generated by a SD-WAN Edge when the deactivate operation for its USB ports fails.
EDGE_USB_DEVICE_REMOVED	ALERT	Generated by a SD-WAN Edge when a device is removed from its USB port.
EDGE_USB_DEVICE_INSERTED	ALERT	Generated by a SD-WAN 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 SD-WAN Edge with new connections.

Events	Severity	Description
GATEWAY_SERVICE_STATE_UPDATE D		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 SD-WAN Edge.
INVALID_JSON	CRITICAL	Generated when a SD-WAN 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.
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.

Events	Severity	Description
MGD_ACTIVATION_ERROR	ERROR	Generated when a SD-WAN 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 SD-WAN Edge is activated partially, but a software update failed.
MGD_ACTIVATION_SUCCESS	INFO	Generated when a SD-WAN Edge has been activated successfully.
MGD_CONF_APPLIED	INFO	Generated when a configuration change made on the Orchestrator has been pushed to SD-WAN Edge and is successfully applied.
MGD_CONF_FAILED	INFO	Generated when the SD-WAN 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 SD-WAN Edge.
MGD_CONF_UPDATE_INVALID	INFO	Generated when a SD-WAN Edge has been assigned an Operator Profile with an invalid software image that the Edge cannot use.
MGD_DEACTIVATED	INFO	Generated when a SD-WAN 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 SD-WAN Edge is rebooted by a Remote Action from the Orchestrator.
MGD_DIAG_RESTART	INFO	Generated when the data plane service on the SD-WAN Edge is restarted by a Remote Action from the Orchestrator.
MGD_EMERG_REBOOT	CRITICAL	Generated when a SD-WAN Edge is rebooted to recover from stuck processes by vc_procmon.
MGD_ENTER_LIVE_MODE	DEBUG	Generated when the management service on a SD-WAN Edge is entering the LIVE mode.
MGD_EXIT_LIVE_MODE	DEBUG	Generated when the management service on a SD-WAN Edge is exiting the LIVE mode.

Events	Severity	Description
MGD_EXITING	INFO	Generated when the management service on a SD-WAN Edge is shutting down for a restart.
MGD_EXTEND_LIVE_MODE	DEBUG	Generated by a SD-WAN Edge when Live mode is extended.
MGD_FLOW_STATS_PUSH_FAILED	DEBUG	Generated by a SD-WAN Edge when Flow stats pushed to Orchestrator failed.
MGD_FLOW_STATS_PUSH_SUCCEEDED	DEBUG	Generated by a SD-WAN Edge when Flow stats pushed to Orchestrator succeeded.
MGD_FLOW_STATS_QUEUED	INFO	Generated by a SD-WAN Edge when Flow stats pushed to Orchestrator is queued.
MGD_HARD_RESET	INFO	Generated when a SD-WAN Edge is restored to its factory-default software and configuration.
MGD_HEALTH_STATS_PUSH_FAILED	DEBUG	Generated by a SD-WAN Edge when Health stats pushed to Orchestrator failed.
MGD_HEALTH_STATS_PUSH_SUCCEEDED	DEBUG	Generated by a SD-WAN Edge when Health stats pushed to Orchestrator succeeded.
MGD_HEALTH_STATS_QUEUED	INFO	Generated by a SD-WAN Edge when Health stats pushed to Orchestrator is queued.
MGD_HEARTBEAT	INFO	Generated by a SD-WAN Edge when Heartbeat is generated to Orchestrator.
MGD_HEARTBEAT_FAILURE	INFO	Generated by a SD-WAN Edge when generated Heartbeat to Orchestrator failed.
MGD_HEARTBEAT_SUCCESS	INFO	Generated by a SD-WAN 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 SD-WAN Edge when Link stats pushed to Orchestrator failed.
MGD_LINK_STATS_PUSH_SUCCEEDED	DEBUG	Generated by a SD-WAN Edge when Link stats pushed to Orchestrator succeeded.

Events	Severity	Description
MGD_LINK_STATS_QUEUED	INFO	Generated by a SD-WAN Edge when Link stats pushed to Orchestrator is queued.
MGD_LIVE_ACTION_FAILED	DEBUG	Generated by a SD-WAN Edge when Live Action failed.
MGD_LIVE_ACTION_REQUEST	DEBUG	Generated by a SD-WAN Edge when Live Action is requested.
MGD_LIVE_ACTION_SUCCEEDED	DEBUG	Generated by a SD-WAN 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_UPDATE D	INFO	Generated when new network settings are applied to a SD-WAN Edge.
MGD_SET_CERT_FAIL	ERROR	Generated when the installation of a new PKI certificate for Orchestrator communication on a SD-WAN Edge has failed.
MGD_SET_CERT_SUCCESS	INFO	Generated when a new PKI certificate for Orchestrator communication is installed successfully on a SD-WAN Edge.
MGD_SHUTDOWN	INFO	Generated when the SD-WAN Edge diagnostic shutdown based on user request.
MGD_START	INFO	Generated when the management daemon on the SD-WAN 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 SD-WAN 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.

Events	Severity	Description
MGD_SWUP_INVALID_SWUPDATE	WARNING	Generated when a software update package received from the Orchestrator is invalid.
MGD_SWUP_REBOOT	INFO	Generated when the SD-WAN 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.
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).

Events	Severity	Description
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 SD-WAN 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.