

Using SD-WAN Orchestrator - 9.7.2

October 27, 2026

Copyright and Trademarks

© Copyright 2026 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein. Aruba Networks and the Aruba logo are registered trademarks of Aruba Networks, Inc. Third-party trademarks mentioned are the property of their respective owners. To view the end-user software agreement, go to: HPE Networking EULA

Support

For product and technical support, contact support at either of the following:

1.800.943.4526 (toll-free in USA and Canada)
+1.408.941.4300
www.silver-peak.com/support

We are dedicated to continually improving our products and documentation. If you have suggestions or feedback for our documentation, send an e-mail to sp-techpubs@hpe.com.

Table of Contents

Using SD-WAN Orchestrator - 9.7.2 **29**

What's New **30**
- Orchestrator 9.7.2 . . . 30
  - HPE Mist Integration for EdgeConnect . . . 30

Getting Started **31**
- Supported Browsers . . . 31
- Guidelines for Creating Passwords . . . 31
- Overview of SD-WAN Prerequisites . . . 31

SASE Copilot **34**
- Availability . . . 34
- Curated Workflows . . . 35
  - Workflow Categories . . . 35
  - How Do I Launch SASE Copilot? . . . 35
- Enable SASE Copilot . . . 36
  - Before You Begin . . . 36
  - Launch SASE Copilot . . . 36
  - Re-enable SASE Copilot . . . 37
- Permissions and Roles . . . 37
  - Read & Write and Administrator Roles . . . 37
  - Read-Only Roles . . . 38
  - Operational and Important Behavior . . . 38
- Chat Management . . . 38
  - Display Mode . . . 38
  - Clear Context . . . 38
  - Chat Refresh . . . 38
  - Response Feedback . . . 39

Light or Dark Theme **40**

Menu Options **41**
- Monitoring . . . 41
- Configuration . . . 41
- Administration . . . 41
- Orchestrator . . . 42
- Support . . . 42
- Monitoring . . . 42
  - Monitoring > Summary . . . 42
    - Dashboard . . . 42

Topology . . . 43
Health Map . . . 46
Alarms Tab . . . 48
Standard . . . 48
Correlated . . . 49
Filtering Alarms . . . 49
Disable Alarms . . . 50
Customize Alarms . . . 50
Alarm Severity . . . 50
Alarm Recipients . . . 51
Correlation Settings . . . 52
Additional Alarm Indications . . . 52
Export Alarm Descriptions . . . 53
List of Alarms . . . 53
EdgeConnect Appliance Alarms . . . 53
Tunnels . . . 53
Software . . . 56
Equipment . . . 73
Threshold Crossing Alerts . . . 85
Orchestrator Alarms . . . 88
Tunnels . . . 88
Software . . . 97
Equipment . . . 126
Route Next Hops . . . 126
Route Label Status . . . 127
Route Labels Dialog Box . . . 128
Add a Route Label . . . 128
Edit a Route Label . . . 128
Delete a Route Label . . . 128
Appliance Power Consumption . . . 129
Monitoring > Security . . . 130
Firewall Drops . . . 130
Protection Profile Top Talkers . . . 131
Protection Profile Trends . . . 132
Flow Baselines . . . 135
Flow Baseline Trends . . . 137
Internal IP Reputation . . . 138
Secure Web Services Top Talkers . . . 140
Security Log Viewer . . . 141
View Real-Time Security Logs . . . 142
View Stored Security Logs . . . 142
Configure Maximum Line Limit for Security Logs . . . 143
Filter Security Logs . . . 144
Security Log Attributes . . . 146
Monitoring > Performance . . . 150
Availability . . . 151
Availability Tables . . . 151
Configure Color Codes for Availability Thresholds . . . 154

AppExpress Summary Tab . . . 154
QoE Trends . . . 157
Ping QoE Trends . . . 157
About AppExpress . . . 158
Prerequisites . . . 158
Determining the Best Path for an Application Flow . . . 159
Transport Types . . . 161
IP SLA Summary . . . 161
Time Range . . . 161
Search . . . 162
Table Columns . . . 162
Expandable Rows . . . 163
State Indicators . . . 163
Internet Breakout Trends . . . 163
Internet Breakout Modes . . . 164
Waterfall . . . 164
Balanced . . . 166
Non-Default to Default Segment Internet Breakout . . . 166
Example Scenario . . . 166
Traffic Flow Diagram . . . 167
Configuration Steps . . . 167
Best Practices . . . 167
Troubleshooting . . . 167
Additional Operational Notes . . . 167
Application Performance Summary . . . 169
Application Performance Trends . . . 171
Monitoring > Reporting . . . 171
Schedule and Run Reports . . . 171
View Reports . . . 175
Sample Report . . . 177
Scheduled and Historical Jobs . . . 177
Metered Usage Reports . . . 178
Monitoring > Bandwidth . . . 178
Overlay Interface Transport . . . 178
Key Information . . . 178
Interface Bandwidth Summary . . . 180
Interpreting the Interface Summary Report . . . 181
Interface Bandwidth Trends . . . 183
Application Bandwidth Summary . . . 184
Application Bandwidth Pie Charts . . . 185
Application Bandwidth Trends . . . 186
User Bandwidth Summary . . . 187
User Bandwidth Pie Charts . . . 188
User Bandwidth Trends . . . 189
Top Talkers . . . 190
Top Domains . . . 192
Top Countries . . . 192
Top Ports . . . 193

Traffic Behavior . . . 194
Appliance Bandwidth Summary . . . 195
Appliance Bandwidth Trends . . . 196
Appliance Packet Counts . . . 197
Tunnel Bandwidth Summary . . . 197
Show Underlays . . . 198
Traceroute . . . 198
Live View . . . 199
Tunnel Bandwidth Pie Charts . . . 200
Key Information . . . 200
Tunnel Bandwidth Trends . . . 201
Tunnel Packet Counts . . . 202
Flows . . . 203
Segment and Zone Filters . . . 206
Filter Field . . . 207
Reset or Reclassify Flows . . . 207
Export . . . 207
Flow Detail . . . 207
Inbound/Outbound Reduction % . . . 207
Additional Information about Flows . . . 208
Flows in Drop State . . . 208
ICMP/UDP Flows . . . 208
TCP Non-Accelerated Flows . . . 208
TCP Accelerated Flows . . . 208
Outbound and Inbound . . . 209
Appliance Flow Counts . . . 210
Appliance Flow Trends . . . 211
Tunnel Flows . . . 211
DSCP Bandwidth Summary . . . 212
DSCP Bandwidth Pie Charts . . . 213
DSCP Bandwidth Trends . . . 213
Traffic Class Bandwidth Summary . . . 214
Traffic Class Bandwidth Pie Charts . . . 214
Shaper Bandwidth Summary . . . 215
Shaper Bandwidth Trends . . . 216
WAN Optimization . . . 217
Change WAN Optimization Configuration . . . 219
Monitoring > Tunnel Health . . . 220
Live View . . . 220
Loss Summary . . . 221
Loss Trends . . . 222
Jitter Summary . . . 223
Jitter Trends . . . 224
Jitter Trend Interpretation Guidelines . . . 225
Latency Summary . . . 226
Latency Trends . . . 228
Out of Order Packets Summary . . . 229
Out of Order Packets Trends . . . 230

MOS Summary . . . 231
MOS Trends . . . 232
Tunnels Summary . . . 233
Configuration . . . 236
Configuration > Overlays & Security . . . 236
Business Intent Overlays . . . 236
Overview . . . 237
SD-WAN Traffic to Internal Subnets . . . 237
Building SD-WAN Using These Interfaces . . . 239
Service Level Objective (SLO) . . . 239
Link Bonding Policy . . . 240
QoS and Optimization . . . 242
Breakout Traffic to Internet and Cloud Services . . . 243
Hub Versus Branch Breakout Settings . . . 243
Preferred Policy Order and Available Policies . . . 243
Break Out Locally Using These Interfaces, Available Interfaces, and Link Selection . . . 244
Apply Overlays . . . 248
Interface Labels . . . 248
Manage Labels . . . 249
Create a Label . . . 249
Edit a Label . . . 250
Delete a Label . . . 250
Route Labels . . . 251
Route Labels Dialog Box . . . 251
Add a Route Label . . . 252
Edit a Route Label . . . 252
Delete a Route Label . . . 252
Hubs . . . 252
Deployment Profiles . . . 253
Map Labels to Interfaces . . . 253
LAN-side Configuration: Segments and Firewall Zones . . . 254
LAN–side Configuration: DHCP and Router Advertisements . . . 254
V4 . . . 254
V6 . . . 257
Add a Router Advertisement Prefix . . . 258
WAN–side Configuration . . . 259
A More Comprehensive Guide to Basic Deployments . . . 262
Bridge Mode . . . 262
Router Mode . . . 263
Server Mode . . . 267
Deployment - EdgeHA . . . 267
Enable EdgeHA Mode . . . 268
IPSec over UDP Tunnel Configuration . . . 268
VRRP Configuration . . . 268
LAN-side Monitoring . . . 269

Internal Subnets . . . 269
Add Internal Subnets . . . 269
Add or Replace Internal Subnets in Bulk . . . 270
Consider Non-default Routes as Internal Subnets . . . 270
Global Firewall Policies . . . 270
Global Firewall Policies Tab . . . 271
Wildcard-based Prefix Matching Rules . . . 272
Firewall Zone Definition . . . 273
Appliance Security Policies . . . 273
Appliance Security Policies Tab . . . 274
Wildcard-based Prefix Matching Rules . . . 275
Firewall Protection Profiles . . . 276
Baseline Learning, Auto Rate Limit, and Smart Burst . . . 276
Auto Rate Limit . . . 276
Smart Burst . . . 277
Enable Baseline Learning . . . 277
SYN Cookie, Smart SYN Cookie, and Dynamic IP Reputation Management . . . 279
Smart SYN Cookie . . . 279
Dynamic IP Reputation Management . . . 280
Enable Internal IP Reputation . . . 280
Create a Firewall Protection Profile . . . 281
Set Firewall Protection Profile Thresholds . . . 284
Add Profile Mappings . . . 287
Add Firewall Protection Profile to a Template Group . . . 287
View DoS Threshold Alarms . . . 287
Schedule IPSec Key Rotation . . . 287
Failure Handling and Orchestrator Reachability . . . 288
Schedule IPSec Key Rotation Dialog Box . . . 288
Monthly Schedule Rotation . . . 289
Inbound Port Forwarding . . . 290
When Translate Mode Is Disabled . . . 290
When Translate Mode Is Enabled . . . 291
Advanced Security Settings . . . 292
Intrusion Detection/Prevention System . . . 295
Prerequisites for IDS/IPS . . . 297
Apply IDS/IPS on Appliances . . . 298
Manage Signature Profiles . . . 298
Update Signatures on Appliances . . . 299
View Signature History . . . 300
Specify Traffic to Be Inspected . . . 302
Advanced Reporting and Analytics . . . 302
Signature Profiles . . . 303
Create a Signature Profile . . . 305
Roles . . . 305
Prerequisites . . . 306
Add Roles . . . 306
Import CSV . . . 306

Trust Store . . . 307
Enable/Disable Auto Sync with Cloud Portal . . . 308
View Built-In Certificate Details . . . 309
Add a Custom CA Certificate to the Trust Store . . . 309
Deploy/Withdraw a Certificate . . . 309
Manage Certificates in Orchestrator Versions Earlier than 9.7.0 . . . 310
End Entity Certificates Tab . . . 311
Adding End Entity Certificates . . . 312
End Entity Certificate Validation at the Time of Upload . . . 313
Certificate Expiry Checking . . . 315
Prepare the Trust Store . . . 315
EST Servers . . . 316
Add an EST Server Profile . . . 316
Appliance End Entity Profiles . . . 318
Add an Appliance End Entity Profile . . . 318
Manually Obtain a Signed End Entity Certificate . . . 322
Create a Certificate Signing Request (CSR) . . . 322
Send the CSR to the Certificate Authority (CA) . . . 324
Obtain the Signed Certificate From the CA . . . 325
Upload the Signed Certificate to the End Entity Certificate Tab . . . 325
Use an End Entity Certificate or Profile for a Service . . . 327
End Entity Certificates Dialog Box . . . 327
Clients Table . . . 328
Network Access Control (NAC) . . . 330
Active Stations . . . 331
Cached Stations . . . 332
Deny List Stations . . . 332
Stations History . . . 333
Enable or Edit Network Access Control (NAC) . . . 333
802.1x/MAC Authentication Profiles . . . 334
802.1x Authentication Profile Fields . . . 334
MAC Authentication Profile Fields . . . 335
Server . . . 336
Server Fields . . . 336
Server Groups Fields . . . 337
AAA Profile . . . 337
Apply Policies . . . 339
Delete a Policy . . . 340
Secure Web Services . . . 340
Secure Web Services Dialog Box . . . 342
Cipher Settings . . . 343
Service Impacts of Activating Cipher Profiles . . . 345
Preconfigured Cipher Profiles . . . 347
Cipher Settings Tab . . . 348
Edit a Cloned Cipher Profile . . . 350
TLS Settings . . . 350
Validate TLS Settings on Self-Hosted/On-Prem Orchestrators . . . 352
Validate TLS Settings on Orchestrator-as-a-Service . . . 352

Validate TLS Settings on ECOS . . . 352
Tunnels (ECOS) Settings . . . 353
Avoid Dropped Tunnels When Changing Cipher Profile Tunnel Settings . . . 355
Implement Tunnel Setting Changes for Orchestrated Underlay Tunnels for ALL Labels . . . 356
Orchestrated IPSec UDP Underlay Tunnels . . . 356
Orchestrated IKE-based IPSec Underlay Tunnels . . . 356
Implement Tunnel Setting Changes for Orchestrated IKE-based, IPSec Passthrough Tunnels to Select SSE Services (HPE SSE, Zscaler, Netskope, etc.) . . . 357
Implement Tunnel Setting Changes for Unified Fabric, HPE ANW Central Tunnels . . . 359
Implement Tunnel Setting Changes for Manually Constructed, EdgeConnect-to-EdgeConnect Tunnels . . . 359
Implement Tunnel Setting Changes for Manually Constructed, EdgeConnect-to-Third Party Passthrough Tunnels 360
Troubleshoot Dropped Tunnels When Changing Cipher Profile Tunnel Settings . . . 361
Alarms . . . 361
Tunnels Tab Status . . . 361
Tunnel Settings Error Messages . . . 361
SSH (ECOS) Settings . . . 362
Configure SSH on ECOS . . . 364
Certificate Settings . . . 364
Configure Certificates on Orchestrator and ECOS . . . 365
SNMP (ECOS) Settings . . . 365
Configure SNMP on ECOS . . . 365
NTP (ECOS) Settings . . . 366
Configure NTP on ECOS . . . 366
Set Up NTP Servers Using the Orchestrator Template . . . 366
Configure Secure NTP via the CLI . . . 366
Secure NTP Parameters . . . 366
Secure NTP Procedure . . . 366
Manual NTP Server Setup . . . 367
Manually Change Appliance Time or Date . . . 367
Cluster Settings . . . 367
Configure Clusters on Orchestrator . . . 368
SSL WAN Optimization Certificates . . . 368
SSL WAN Optimization Certificates Dialog Box . . . 368
SSL WAN Optimization CA Certificates Tab . . . 369
SSL WAN Optimization CA Certificates Edit Row . . . 369
Discovered Appliances . . . 369
Preconfigure Appliances . . . 370
Appliance Wizard . . . 372
EC-Enterprise Licenses . . . 375
Assign a License to an Appliance . . . 376

Licenses . . . 376
EC-Enterprise Licenses . . . 377
Licensing Subscription Tiers . . . 377
Assign a License to an Appliance . . . 379
EC-Metered Licenses . . . 379
Assign a License to an Appliance . . . 380
Bandwidth Usage Report . . . 381
Feature License Usage Report . . . 381
EC-Metered Licenses . . . 381
Assign a License to an Appliance . . . 382
Bandwidth Usage Report . . . 383
Feature License Usage Report . . . 383
Cloud Portal . . . 383
Secondary Accounts . . . 384
Configuration > Networking . . . 385
Deployment Tab . . . 385
Deployment Dialog Box . . . 387
Enable EdgeHA . . . 387
IPSec over UDP Tunnel Configuration . . . 388
LAN-side High Availability . . . 388
LAN-side Monitoring . . . 388
Map Labels to Interfaces . . . 388
LAN and WAN Interface Configuration: Segment, Firewall Zone, and Role . . . 389
LAN–side Configuration: DHCP and Router Advertisements . . . 389
V4 . . . 390
V6 . . . 392
Add a Router Advertisement Prefix . . . 393
WAN–side Configuration . . . 394
Interfaces Tab . . . 396
Terminology . . . 398
PoE (Power Over Ethernet) . . . 398
Interfaces Dialog Box . . . 399
Hardware . . . 399
Dynamic . . . 400
PoE Configuration Dialog Box . . . 400
Automatic Bandwidth Detection . . . 401
Prerequisites . . . 401
Running a Speed Test . . . 401
Speed Test Results . . . 402
Automatic Bandwidth Detection — Scope and Limitations . . . 402
VRRP Tab . . . 404
VRRP Configuration Considerations . . . 404
VRRP Edit Row . . . 404
VRRP Settings . . . 404
DNS Proxy . . . 407
Configure DNS Proxy Policies . . . 407
Branch NAT . . . 408

NAT Rules and Pools . . . 409
NAT Rules . . . 409
NAT Pools . . . 411
PPPoE Tab . . . 412
USB LTE Tab . . . 414
Understanding RSSI Values . . . 414
Technical Tips and Best Practices . . . 415
Enable or Edit a USB LTE Modem . . . 415
Loopback Interfaces . . . 416
Loopback Orchestration . . . 417
Create a Loopback Interface . . . 419
Change the Subnet IP for a Loopback Pool . . . 420
Reclaim Deleted Loopback IP Addresses . . . 421
Virtual Tunnel Interfaces (VTI) . . . 421
VTI Dialog Box . . . 421
Add a VTI . . . 422
Edit a VTI . . . 424
Delete a VTI . . . 425
DHCP Manager . . . 425
Address pool . . . 425
Lease . . . 426
DHCP lease statuses . . . 426
Scope . . . 427
Reservation . . . 427
Service . . . 427
Profiler . . . 428
Failover state . . . 428
DHCP failover state fields . . . 428
DHCP failover state descriptions . . . 429
Relay . . . 429
Failover Configuration . . . 430
Monitoring and Alarms . . . 430
Release and Compatibility Notes . . . 430
DHCP settings Integration . . . 430
DHCP Pseudo LAN Client Isolation . . . 431
ARP Visibility . . . 431
Logging and Troubleshooting Notes . . . 431
Known Limitations . . . 431
DHCP Server Defaults . . . 431
DHCP Settings / Router Advertisements . . . 432
V4 . . . 432
V6 . . . 435
Add a Router Advertisement Prefix . . . 436
DHCP Leases . . . 437
DHCP Lease Statuses . . . 437
DHCP Failover . . . 438
Preparation . . . 438
Procedure . . . 438

DHCP Failover Settings Examples . . . 441
Single DHCP Failover Group Example . . . 442
Two DHCP Failover Groups Example . . . 442
DHCP Failover Fundamentals . . . 443
DHCP Failover Primary or Secondary Roles . . . 444
DHCP Active and Backup Server Behavior . . . 444
DHCP Database Synchronization . . . 444
FAQs . . . 445
DHCP Failover State . . . 445
DHCP Failover State Fields . . . 446
DHCP Failover State Descriptions . . . 446
Link Aggregation . . . 447
Configure Link Aggregation . . . 449
Add or Modify a Channel Group . . . 449
Delete a Channel Group . . . 450
Apply Your Changes . . . 450
Cluster Profiles . . . 451
Create a Cluster Profile . . . 451
Clusters . . . 452
Cluster Config Dialog Box . . . 453
Add an Appliance to a Cluster . . . 454
Applying or Removing a Cluster Profile . . . 455
Bridge Groups . . . 455
Configure Bridge Groups . . . 456
Add or Modify a Bridge Group . . . 457
Delete a Bridge Group . . . 457
Regions . . . 457
Regional Routing . . . 458
View Status . . . 459
Edit Regions . . . 459
Regional Routing . . . 460
Routing Segmentation . . . 460
Segment Configuration . . . 461
Overlays & Breakout Policies for Segments . . . 461
Firewall Zone Policies . . . 461
Inter-Segment Routing & DNAT Exceptions . . . 463
Add a Rule . . . 463
Delete a Rule . . . 464
Inter-Segment SNAT Exceptions . . . 464
Loopback . . . 464
Appliances . . . 464
Comment . . . 465
Delete a Segment . . . 465
Management Services . . . 465
Management Services Dialog Box . . . 466
Inter-Segment Routing and DNAT Exceptions . . . 469
Add a Rule . . . 469
Delete a Rule . . . 470

Inter-Segment SNAT Exceptions . . . 470
BGP Tab . . . 471
BGP Information . . . 474
Add Peer . . . 476
BGP Inbound and Outbound Route Redistribution Maps . . . 477
Prefix Match Criteria . . . 478
GE Clause . . . 478
LE Clause . . . 479
Combining LE and GE Clauses . . . 479
Exact Match . . . 480
BGP ASN Global Pool . . . 483
Routes Tab . . . 483
Route Maps . . . 483
System-Generated Default Routes . . . 488
Characteristics of system-generated routes . . . 488
About metric values . . . 488
Interpretation guidance . . . 488
Edit or Add Routes . . . 489
Add Routes . . . 491
Segment / VRF Pruning . . . 493
Import Subnets . . . 494
Export Subnets . . . 494
SD-WAN Fabric Route Redistribution Maps . . . 494
Prefix Match Criteria . . . 495
GE Clause . . . 495
LE Clause . . . 496
Combining LE and GE Clauses . . . 497
Exact Match . . . 497
OSPF Tab . . . 498
OSPF Edit Row . . . 500
Add Area . . . 500
Add Interface . . . 501
OSPF Route Redistribution Maps . . . 502
Prefix Match Criteria . . . 502
GE Clause . . . 503
LE Clause . . . 503
Combining LE and GE Clauses . . . 504
Exact Match . . . 504
BFD Tab . . . 506
BFD Dialog Box . . . 507
Multicast . . . 508
Multicast and Segments . . . 508
Multicast Configuration Workflow Overview . . . 509
Multicast Tab Views . . . 509
Segment Filter . . . 510
Multicast Dialog Box . . . 510
Peer Priority Tab . . . 512
Peer Priority Edit Row . . . 513

Admin Distance Tab . . . 514
Admin Distance Edit Row . . . 515
Management Routes Tab . . . 515
VXLAN Tab . . . 516
Prerequisites . . . 516
Common Settings for all VNIs . . . 517
VNI Mappings . . . 517
Role to GPID Mapping . . . 518
VXLAN Statistics Reporting . . . 518
VTEP Details . . . 518
ARP Tab . . . 519
ARP, Routes, and VRFs . . . 519
ARP Table Columns . . . 519
Adding a Static ARP Entry . . . 520
Deleting ARP Entries . . . 521
Tunnels Tab . . . 521
Troubleshooting . . . 528
Tunnels Dialog Box . . . 529
About Authentication in IPSec Tunnels . . . 529
Add or Modify a Manually Created Underlay Tunnel . . . 530
Add or Modify a Manually Created Passthrough Tunnel . . . 539
Delete a Tunnel . . . 544
Use Passthrough Tunnels . . . 544
IPSec Suite B Presets . . . 544
Schedule Auto MTU Discovery . . . 545
Tunnel Troubleshooting . . . 546
Tunnel Exception . . . 546
Configuration > Policies . . . 547
Route Policies Tab . . . 547
Route Policies Dialog box . . . 547
Priority . . . 548
Match Criteria . . . 548
Source or Destination . . . 549
Wildcard-based Prefix Matching Rules . . . 550
QoS Policies Tab . . . 550
QoS Policies Dialog Box . . . 550
Handle and Mark DSCP Packets . . . 551
Apply DSCP Markings to Optimized (Tunnelized) Traffic . . . 551
Apply DSCP Markings to Passthrough Traffic . . . 552
Priority . . . 554
Match Criteria . . . 554
Source or Destination . . . 555
Wildcard-based Prefix Matching Rules . . . 556
Schedule QoS Map Activation . . . 556
Optimization Policies Tab . . . 557
Optimization Policies Dialog Box . . . 557
Priority . . . 557
Match Criteria . . . 557

Source or Destination . . . 558
Wildcard-based Prefix Matching Rules . . . 559
Set Actions . . . 559
TCP Acceleration . . . 560
Access Lists Tab . . . 564
Match Criteria . . . 565
Wildcard-based Prefix Matching Rules . . . 566
Access Lists Edit Row . . . 567
Address Groups . . . 567
Add an Address Group . . . 568
Add a Rule to an Address Group . . . 569
Delete an Address Group . . . 570
Export Address Groups . . . 570
Import Address Groups . . . 571
View a Single Address Group . . . 572
Edit or Delete a Rule . . . 572
Using Address Groups in Match Criteria . . . 572
Requirement and Behavior of Excluded IPs in Nested Groups . . . 573
Address Group Formats . . . 574
Service Groups . . . 574
Add a Service Group . . . 575
Add a Rule to a Service Group . . . 576
Delete a Service Group . . . 577
Export Service Groups . . . 577
Import Service Groups . . . 578
View a Single Service Group . . . 579
Edit or Delete a Rule . . . 579
Using Service Groups in Match Criteria . . . 580
Shaper Tab . . . 580
Application Definitions Tab . . . 582
Application Definition Dialog Box . . . 583
Add an Application Definition . . . 583
Enable AppExpress for an Application . . . 585
Edit an Application Definition . . . 587
Disable an Application Definition . . . 587
Delete a User-created Application Definition . . . 587
Application Groups Tab . . . 588
Apply AppExpress Groups Tab . . . 589
AppExpress Groups Tab . . . 589
Add or Edit AppExpress Group Dialog Box . . . 589
Threshold Crossing Alerts Tab . . . 591
ON by Default . . . 592
OFF by Default . . . 592
Threshold Crossing Alerts Edit Row . . . 593
IP SLA . . . 594
IP SLA Tab . . . 594
IP SLA Monitor Use Cases . . . 595
Example #1 – Ping via Interface . . . 595

Example #2 – HTTP/HTTPS via Interface . . . 596
Example #3 – Monitor Interface . . . 596
Example #4 – Monitor Interface (WAN0) to Ensure High Availability . . . 597
Example #5 – Monitor VRRP . . . 597
Example #6 – Revoke Static Route Using IP SLA . . . 598
IP SLA Rule . . . 599
Configuration > Templates . . . 604
Templates Overview . . . 604
Modifying a Template . . . 605
Template Groups . . . 606
System Template . . . 606
Disable Stateful+SNAT Processing for Selected LAN-side Subnets . 610
Auth/Radius/TACACS+ Template . . . 611
Authentication and Authorization . . . 611
Appliance-based User Database . . . 611
RADIUS . . . 611
TACACS+ . . . 612
What Is Recommended . . . 612
Network Access Control (NAC) Template . . . 612
802.1x/MAC Authentication Profiles . . . 613
802.1x Authentication Profile Fields . . . 613
MAC Authentication Profile Fields . . . 614
Server . . . 615
Servers Fields . . . 615
Server Group Fields . . . 616
AAA Profile . . . 616
Apply Policies . . . 617
Delete a Policy . . . 617
Merge / Replace . . . 617
Appliance IP Allow List Template . . . 618
Appliance IP Allow List Template Settings . . . 618
Flow Export Template . . . 618
Firewall Protection Profiles Template . . . 619
Enable Baseline Learning . . . 619
Enable Internal IP Reputation . . . 621
Add New Profiles . . . 621
Set Firewall Protection Profile Thresholds . . . 624
Add Profile Mappings . . . 626
Logging Template . . . 627
Minimum Severity Levels . . . 627
Configure Remote Logging . . . 628
Add a Client Certificate . . . 629
Banner Messages Template . . . 630
HTTPS Certificate Template . . . 631
Password Settings Template . . . 632
User Management Template . . . 633
ECOS Local User Accounts . . . 633

Command Line Interface Privileges . . . 634
DNS Template . . . 634
Date/Time Setting . . . 634
Data Collection . . . 635
SNMP Template . . . 635
SNMP V1/V2 . . . 635
SNMP v3 . . . 636
Trap Receivers . . . 636
SSL WAN Optimization Certificates Template . . . 637
SSL WAN Optimization CA Certificates Template . . . 638
Multicast Template . . . 638
VRRP Template . . . 639
Peer Priority Template . . . 640
Route Redistribution Maps Template . . . 641
Routes Template . . . 642
BGP Template . . . 644
OSPF Template . . . 645
BFD Template . . . 647
VXLAN Template . . . 648
Common Settings for all VNIs . . . 648
VNI Mappings . . . 649
Role to GPID Mapping . . . 649
Admin Distance Template . . . 649
Access Lists Template . . . 650
Priority . . . 652
Match Criteria . . . 652
Source or Destination . . . 652
Wildcard-based Prefix Matching Rules . . . 652
Route Policies Template . . . 653
Why? . . . 653
Priority . . . 654
Match Criteria . . . 654
Source or Destination . . . 655
Wildcard-based Prefix Matching Rules . . . 655
Set Actions Fields . . . 656
Where the Appliance Directs Traffic . . . 656
How Traffic Is Managed If a Tunnel Is Down . . . 656
QoS Policies Template . . . 656
Priority . . . 657
Match Criteria . . . 657
Source or Destination . . . 658
Wildcard-based Prefix Matching Rules . . . 659
Handle and Mark DSCP Packets . . . 659
Apply DSCP Markings to Optimized (Tunnelized) Traffic . . . 659
Apply DSCP Markings to Pass-through Traffic . . . 661
Optimization Policies Template . . . 662
Priority . . . 663
Match Criteria . . . 663

Source or Destination . . . 664
Wildcard-based Prefix Matching Rules . . . 664
Set Actions Fields . . . 665
TCP Acceleration Options . . . 666
Threshold Crossing Alerts Template . . . 669
ON by Default . . . 671
OFF by Default . . . 671
TCA Metrics . . . 671
Security Policies Template . . . 672
Implicit Drop Logging . . . 673
Template . . . 673
Wildcard-based Prefix Matching Rules . . . 673
DNS Proxy Policies . . . 674
Shaper Template . . . 674
Add / Delete Shaper . . . 675
Enable Interface Shaper . . . 675
Shaper Options . . . 675
Shaper Configuration . . . 676
Per-interface Shaping . . . 677
Dynamic Rate Control . . . 677
Management Services Template . . . 678
CLI Template . . . 678
Session Management Template . . . 679
Web Proxy . . . 680
Secure Web Services Template . . . 680
Apply Template Groups . . . 681
Best Practices When Applying Template Groups . . . 682
Apply Template Groups Behavior . . . 684
Multiple Template Groups Example . . . 685
Template Options . . . 688
Configuration > Cloud Services . . . 692
AWS Network Manager . . . 692
Prerequisites for AWS Transit Gateway Network Manager . . . 692
Create a User Profile in AWS . . . 692
Create Transit Gateways . . . 694
Create a Network Manager . . . 696
Orchestrator Configuration . . . 696
Subscription . . . 696
Interface Labels . . . 697
Tunnel Settings . . . 697
VTI Subnet Pool . . . 697
Segment & Zone Association . . . 697
AWS Resources . . . 698
Appliance Association . . . 698
Verification . . . 698
Route Tables and Static Routes . . . 699
Peering . . . 700

Azure Network Manager . . . 701
Microsoft Azure Prerequisites . . . 702
Orchestrator Prerequisites . . . 702
Orchestrator Configuration . . . 703
Cloud Gateways . . . 703
Subscription . . . 703
LAN interface labels . . . 703
Azure resources . . . 703
Appliance association . . . 704
Pause Orchestration . . . 704
Deploy Route Servers . . . 704
Deploy Load Balancers . . . 705
On-prem Gateways . . . 707
Subscription . . . 707
Storage URL . . . 707
WAN Interface Labels . . . 708
Tunnel Settings . . . 708
VTI Subnet Pool . . . 708
Zone/Role . . . 708
Virtual WAN Association . . . 708
Pause Orchestration . . . 708
Verification . . . 709
HPE SSE . . . 709
Secure Web Gateway . . . 709
Configure Secure Web Gateway . . . 710
Subscription . . . 710
Interface Labels . . . 711
Tunnel Settings . . . 712
IP SLA . . . 712
Sub-Locations . . . 712
HPE SSE POP Override . . . 713
HPE SSE Association . . . 714
Pause Orchestration . . . 714
Redirect HPE SSE Application Traffic via Break Out Locally . . . 714
Set Match Criteria . . . 714
Verify HPE SSE Deployment . . . 715
SSE Connector . . . 715
Configure SSE Connector . . . 716
Subscription . . . 717
SSE Connector Interface Orchestration . . . 718
Add SSE Connector Interface Pools . . . 719
SSE Connector Interfaces . . . 721
SSE Connector Management . . . 722
Connector Images Table . . . 722
Target Appliances Table . . . 722
Download and Install an SSE Connector Image . . . 723
Uninstall a Connector Image . . . 724
EdgeHA Functionality and Connector Traffic Flow . . . 725

Troubleshoot SSE Connector . . . 725
Microsoft Office 365 . . . 726
Zscaler Internet Access . . . 727
Configure Zscaler . . . 728
Subscription . . . 728
Interface Labels . . . 729
Tunnel Settings . . . 730
Service Edge Override . . . 730
IP SLA . . . 730
Country / Timezone . . . 731
Gateway Options . . . 731
Configure Gateway Options . . . 732
Configure Bandwidth Control . . . 732
Zscaler Association . . . 733
Pause Orchestration . . . 733
Using Zscaler for Breakout Traffic . . . 734
Verify Zscaler Deployment . . . 734
Netskope . . . 734
Configure Netskope . . . 735
Subscription . . . 735
Interface Labels . . . 736
Tunnel Settings . . . 736
Netskope Service Edge Override . . . 737
IP SLA . . . 737
Netskope Association . . . 737
Pause Orchestration . . . 737
Using Netskope for Breakout Traffic . . . 738
Verify Netskope Deployment . . . 738
Service Orchestration . . . 738
Prerequisites . . . 739
Set Up a New Service . . . 739
Remote Endpoint Configuration . . . 739
Add Endpoints One at a Time . . . 739
Add Endpoints in Bulk . . . 741
Bulk Edits . . . 741
Interface Labels . . . 742
Tunnel Settings . . . 742
IP SLA Settings . . . 746
Pause Orchestration (Optional) . . . 747
BIO Breakout . . . 747
Remote Endpoint Association . . . 747
Add Tunnel Local Identifiers to the Third-Party Service Provider . . . 748
Verification . . . 748
Deploy Cloud Hubs . . . 748
Cloud Hubs in AWS . . . 749
Create or Modify an AWS Account . . . 750
Deploy a New EC-V . . . 750
Remove an EC-V . . . 750

AWS Accounts . . . 751
AWS Account Configuration . . . 751
Create a Policy with Required Permissions . . . 751
Attach Policy to the Orchestrator IAM User Account . . . 752
Download Orchestrator IAM User Account Credentials . . . 752
Create a Key Pair to Assign to EC-Vs . . . 752
Subscribe to EdgeConnect SD-WAN Product on the AWS Marketplace Portal . . . 752
Add the AWS Account to Orchestrator . . . 753
AWS Deployment Configuration . . . 753
Deploy Cloud Hubs in Azure . . . 755
Create or modify an Azure subscription . . . 757
Deploy a new EC-V . . . 757
Remove an EC-V . . . 757
Azure subscriptions . . . 757
Add new Azure subscription . . . 757
Edit an existing Azure subscription . . . 758
Deploy cloud hubs in Transit VNet . . . 758
Configure the Azure subscription . . . 758
Accept Azure Marketplace image terms . . . 758
Create a new app registration . . . 759
Create a new resource group . . . 759
Create a custom role . . . 760
Assign the custom role to the resource group . . . 761
Add Azure subscription details to Orchestrator . . . 761
Deploy EC-V into a transit VNet from Orchestrator . . . 762
Deploy Cloud Hubs in Virtual WAN Hub . . . 766
Configure the Azure subscription . . . 766
Create a new app registration . . . 766
Create a new resource group . . . 767
Create a custom role . . . 767
Assign the custom role to the resource group . . . 768
Add Azure subscription details on Orchestrator . . . 768
Deploy Managed NVA . . . 769
Cloud Hubs in GCP . . . 771
Add or Modify a GCP Account . . . 772
Deploy a New EC-V . . . 772
Manage an EC-V . . . 772
GCP Accounts . . . 772
GCP Account Configuration . . . 773
Create a GCP Project . . . 773
Enable GCP APIs . . . 774
Enable Compute Engine API . . . 774
Enable Google Cloud Resource Manager API . . . 774
Create a Custom Role . . . 774
Create a Custom Role Using Google Cloud Shell (Recommended) . . . 774
Create a Custom Role Using Google Cloud Console . . . 775
Create a GCP Service Account . . . 775

Create a Service Account Key Pair . . . 776
Add the GCP Account to Orchestrator . . . 776
GCP Deployment Configuration . . . 776
Administration . . . 779
Administration > General Settings . . . 779
Appliance Users . . . 779
Appliance Users Dialog Box . . . 780
Auth/RADIUS/TACACS+ Tab . . . 780
Authentication and Authorization . . . 780
RADIUS and TACACS+ . . . 781
Auth/RADIUS/TACACS+ Edit Row . . . 781
Authentication Order . . . 781
Authorization Information . . . 782
Authentication and Authorization . . . 782
RADIUS and TACACS+ . . . 782
Use Appropriate RADIUS Configuration Options . . . 782
Appliance IP Allow List . . . 783
Appliance IP Allow List Settings . . . 783
Appliance Password Settings . . . 784
Date/Time Tab . . . 784
Date/Time Settings Dialog Box . . . 785
DNS (Domain Name Servers) Tab . . . 785
DNS (Domain Name Servers) Edit Row . . . 786
SNMP . . . 786
Modify SNMP Configuration . . . 787
SNMP V1/V2 . . . 787
SNMP V3 . . . 787
Trap Receivers . . . 788
Flow Export Tab . . . 789
Custom Information Elements . . . 789
Flow Export Edit Row . . . 793
Logging Tab . . . 794
Severity Levels . . . 795
Remote Logging . . . 795
Logging Dialog Box . . . 796
Log Settings . . . 796
Log Facilities Configuration . . . 797
Remote Log Receivers . . . 797
Banners Tab . . . 798
Banners Edit Row . . . 798
HTTPS Certificate Tab . . . 798
HTTPS Certificate Dialog Box . . . 799
Orchestrator Reachabililty Tab . . . 800
Custom Appliance Tags . . . 801
Administration > Software . . . 802
System Information . . . 802
Disable Stateful+SNAT Processing for Selected LAN-side Subnets . . . 808
Software Versions . . . 809

Upgrade Appliance Software . . . 809
Appliance Configuration Backup . . . 811
View Configuration History . . . 813
Appliance Restore and Reboot . . . 813
Restore an Appliance Configuration Backup . . . 815
View or Copy an Appliance Configuration Backup File . . . 816
Download an Appliance Configuration Backup File . . . 816
Remove Appliance from Orchestrator . . . 816
Remove Appliance from Orchestrator and Account . . . 817
Administration > Tools . . . 817
Synchronize Appliance Configuration . . . 817
Broadcast CLI Commands . . . 818
Link Integrity Test . . . 819
TCPPERF Version 1.4.8 . . . 820
Disk Management . . . 826
Erase Network Memory . . . 827
Reboot or Shut Down an Appliance . . . 828
Behavior During Reboot . . . 829
Schedule an Appliance Reboot . . . 829
Behavior During Reboot . . . 830
Reachability Tab . . . 831
Appliances/Orchestrator . . . 831
Select Fast WebSocket Failover Mode . . . 833
Appliances/HPE ANW Central . . . 835
Active Sessions Tab . . . 835
Orchestrator . . . 836
Orchestrator > Orchestrator Server . . . 836
View Orchestrator Server Information . . . 836
Restart, Reboot, or Shutdown . . . 837
Orchestrator High Availability . . . 837
Set Up Email Notifications . . . 839
Configure Orchestrator HA Cluster Reachability . . . 840
Promote Backup Orchestrator to Primary . . . 841
Decommission an Orchestrator Instance . . . 842
Orchestrator Users . . . 842
Enforce Password Criteria and Set Password Requirements . . . 843
Add a User . . . 847
Multi-Factor Authentication . . . 847
Configuring Multi-Factor Authentication Through an Application 848
Configuring Multi-Factor Authentication Through Email . . . 849
Using Multi-Factor Authentication . . . 849
Modify a User . . . 850
Role Based Access Control . . . 851
Roles . . . 852
Appliance Access . . . 853
Assign Roles and Appliance Access . . . 855

Authentication . . . 856
Configure a RADIUS or TACACS+ Server . . . 856
Authenticate Using RADIUS or TACACS+ . . . 857
Configure an OAuth Server . . . 857
Prerequisites . . . 857
Register Orchestrator as an App . . . 858
Configure OAuth Server Properties in Orchestrator . . . 858
Configure a JWT Server . . . 860
Configure a SAML Server . . . 863
SAML and Orchestrator Configuration . . . 863
API Key . . . 869
Create, Change, or Reset Password . . . 870
Create or Change Password . . . 870
Reset Password . . . 871
Cloud Portal . . . 871
Secondary Accounts . . . 872
Air-Gap Tab . . . 872
Before You Begin . . . 873
Enable Air-Gap Mode . . . 873
Register Air-Gap to Orchestrator . . . 873
Upload Air-Gap Files . . . 874
Audit Logs . . . 874
Orchestration Settings . . . 875
Maintenance Mode . . . 877
Set Maintenance Mode Using the Menu Available from the Appliance Tree . . . 877
Set Maintenance Mode Using the Orchestrator Menu . . . 877
Tunnel Settings Tab . . . 878
IPSec Suite B Presets . . . 878
General Tab . . . 879
IKE Tab . . . 883
IPSec Tab . . . 885
Orchestrator Blueprint Export . . . 887
Brand Customization . . . 888
Orchestrator > Software & Setup . . . 888
Upgrade Orchestrator Software . . . 888
Upgrade On-Prem Orchestrator . . . 889
Upgrade via HTTP . . . 889
Upgrade via SCP . . . 889
Upgrade Orchestrator in the Cloud . . . 890
Check for Orchestrator and Appliance Software Updates . . . 890
Back Up on Demand . . . 891
Schedule Orchestrator Backup . . . 892
User-Configured Orchestrator Backup (On-Premises Only) . . . 892
Daily EdgeConnect Backup (All Orchestrators) . . . 892
Scheduling Details . . . 892
Schedule Stats Collector Backup . . . 894
SMTP Server Settings . . . 896

Proxy Configuration . . . 897
Web Proxy . . . 897
Configure Orchestrator as a Proxy for EdgeConnect Onboarding . . 897
Onboarding Appliances Using Orchestrator as a Proxy . . . 897
Orchestrator HTTPS Certificate . . . 898
Timezone for Scheduled Jobs . . . 899
Orchestrator Advanced Properties . . . 900
IP Allow List . . . 900
Orchestrator Setup . . . 901
Availability Time Settings . . . 903
Statistics Retention . . . 904
DoS Stats . . . 905
Stats Configuration . . . 907
Stats Collector Configuration . . . 909
Features that Require Distributed Stats Collector . . . 910
Prerequisites . . . 910
Enable Distributed Stats Collector . . . 911
Before You Begin . . . 911
Create a VM for the Stats Collector . . . 911
Configure the VM as a Stats Collector . . . 912
Configure the Stats Collector with a Public Key . . . 912
Copy the Public Key File from Orchestrator to the Stats Collector 912
Confirm the New Public Key File . . . 913
Create and Install an End Entity Certificate . . . 914
Add the Root CA Certificate for the Certificate Authority . . . 914
Create and Send the CSR in Orchestrator and Upload the Signed Certificate in Orchestrator . . . 915
Create and send the CSR in Stats Collector and upload the signed certificate in Stats Collector . . . 915
Configure the Stats Collector Feature . . . 917
Add a Stats Collector . . . 917
Delete a Stats Collector . . . 918
Associate Appliances with a Stats Collector . . . 918
Associate Appliances with the Predefined Local Stats Collector 919
Enable the Distributed Stats Collector . . . 919
Discontinue Legacy Stats Collection . . . 920
Back Up and Restore Stats Collector . . . 920
Back Up the Stats Collector . . . 920
Restore the Stats Collector from the CLI . . . 921
Notification Banner . . . 921
Change Stats Collector Log Level . . . 921
Orchestrator > HPE Integration Services . . . 922
HPE ANW Central Site Mapping . . . 922
Prerequisites . . . 922
Create HPE Aruba Networking Central Sites in Bulk . . . 923
Create an HPE Aruba Networking Central Account in Orchestrator . 923
Edit EdgeConnect to HPE Aruba Networking Central Site Mapping . 925
Check for Site List Updates . . . 925

ClearPass Policy Manager . . . 926
Manage ClearPass Policy Manager Accounts . . . 927
View ClearPass Policy Manager Accounts . . . 927
Add a ClearPass Policy Manager Server . . . 927
Edit a ClearPass Policy Manager Server . . . 928
Pause ClearPass Policy Manager Integration . . . 928
Mist . . . 928
Prerequisites . . . 929
Onboard EdgeConnect in Mist . . . 929
Add the Registration Code to Orchestrator . . . 929
Use Automation to Add the Registration Code to Orchestrator 930
Manually Add the Registration Code to Orchestrator . . . 930
Enable WAN Assurance in Orchestrator . . . 931
Assign the EdgeConnect Appliance to a Site in Mist . . . 931
Disable WAN Assurance in Orchestrator . . . 931
Release a WAN Edge Device from a Mist Organization . . . 932
Monitor Mist Connections in Orchestrator . . . 932
HPE Aruba Networking Central . . . 933
Overview of the Unified Fabric . . . 933
Prerequisites . . . 934
Enable HPE Aruba Networking Central . . . 934
Support . . . 935
Support > Technical Assistance . . . 935
Tech Support - Appliances . . . 935
Tech Support - Orchestrator . . . 936
Take Action with Files . . . 936
Log In to the Support Portal . . . 936
Monitor Transfer Progress . . . 937
Packet Capture . . . 938
Upload Local Files . . . 939
Create a Support Case . . . 940
Partition Management . . . 941
Orchestrator Remote Log Receivers . . . 941
Supported Receiver Types . . . 941
Add a Remote Log Receiver . . . 942
Configure an HTTP Receiver . . . 942
Configure an HTTPS Receiver . . . 943
Configure a KAFKA Receiver . . . 943
Configure a SYSLOG Receiver . . . 944
Configure a WEBSOCKET Receiver . . . 945
Event Message Retention . . . 946
Remote Log Messages . . . 946
JSON Format . . . 946
JSON Alarm Message Format . . . 946
JSON Audit Log Message Format . . . 947
JSON Data Key for Orchestrator and Appliance Alarms . . . 947
JSON Data Key for Audit Logs . . . 948

RFC5424 Syslog Format . . . 949
RFC5424 Syslog Message Details . . . 949
Structured Data . . . 950
RFC5424 Structured Data Element for Orchestrator Alarms . . . 950
RFC5424 Structured Data Element for Appliance Alarms . . . 951
RFC5424 Structured Data Element for Audit Logs . . . 951
RFC5424 Meta Structured Data Element . . . 952
Routing Peer Table . . . 952
RMA Wizard . . . 953
Run the RMA Wizard . . . 954
Add a Backup Appliance . . . 955
Upgrade and Downgrade . . . 955
Support > User Documentation . . . 955
Alarm Descriptions . . . 955
Built-in Policies . . . 956
Support > Reporting . . . 957
Dropped Packet Trends . . . 957
Appliance Memory Trends . . . 958
System Performance . . . 959
System PoE State . . . 960
Appliance CPU Usage . . . 961
Appliance Hardware Debug . . . 962
Orchestrator Debug . . . 962
IPSec UDP Status . . . 962
Unverified Emails . . . 963
Orchestrator Debug Log Level . . . 964
Global Log Level . . . 964
Customized Log Levels . . . 964
Notable Customized Loggers . . . 965
Live Tail Logger . . . 965

Using SD-WAN Orchestrator - 9.7.2

This guide contains information about how to get started with Orchestrator and how to use Orchestrator to manage your HPE Networking EdgeConnect products.

- This guide documents the latest Orchestrator release.
- This guide may contain updates not included in the help that ships with the product.
- PDF versions of the latest and previous versions of the Orchestrator User Guide can be found in User Guides.

What's New

This page provides brief descriptions of new features in the recent Orchestrator release.

Orchestrator 9.7.2

The following feature is included in the user guide for Orchestrator 9.7.2.

HPE Mist Integration for EdgeConnect

SASE Orchestrator release 9.7.2 introduces integration with HPE Mist, enabling administrators to onboard EdgeConnect appliances into the Mist cloud as supported WAN Edge devices. EdgeConnect appliances securely establish telemetry channels to stream device health, performance statistics, and event alerts to HPE Mist. The integration enables administrators to use Mist capabilities such as Insights, Service Level Experiences (SLEs), Security Events, Marvis AI, and Premium Analytics.

See Mist.

Getting Started

Orchestrator enables you to globally monitor performance and manage EdgeConnect (EC) appliances, whether you are configuring a WAN Optimization network (NX, VX, or VRX appliances) or an SD-WAN network (EC or EC-V appliances).

On this page:

- Supported Browsers
- Guidelines for Creating Passwords
- Overview of SD-WAN Prerequisites

Supported Browsers

Orchestrator and the Appliance Web user interfaces support the following browsers:

- Google Chrome (recommended)
- Microsoft Edge
- Mozilla Foxfire
- Opera
- Safari

We recommend that you use the latest version available for your browser.

Guidelines for Creating Passwords

See Create, Change, or Reset Password for password requirements.

Overview of SD-WAN Prerequisites

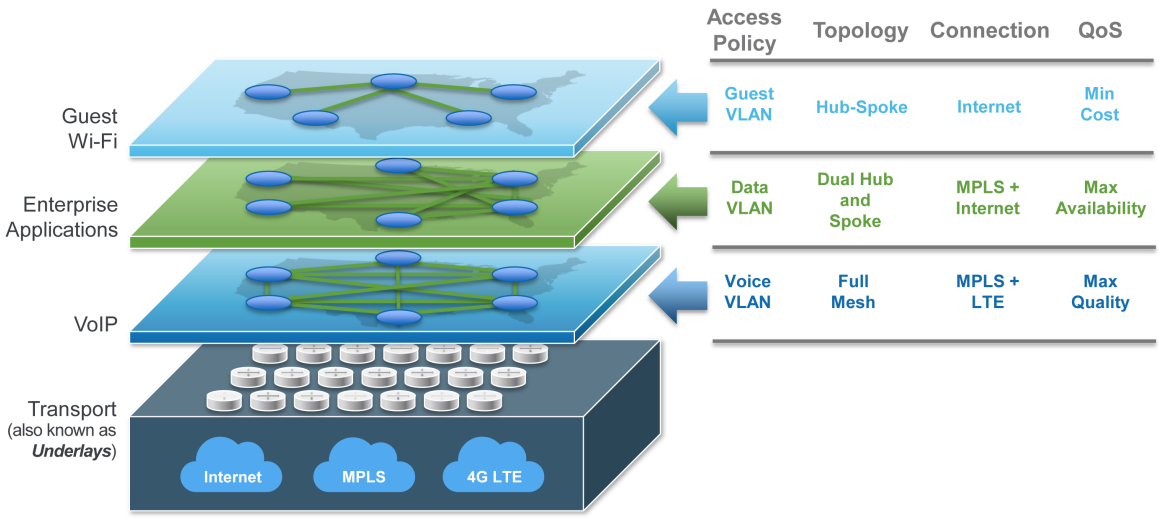
With Orchestrator, you create virtual network overlays to apply business intent to network segments. Provisioning a device is managed by applying profiles.

- Interface Labels associate each interface with a use.

  - **LAN** labels refer to traffic type, such as **VoIP**, **data**, or **replication**.
  - **WAN** labels refer to the service or connection type, such as **MPLS**, **internet**, or **Verizon**.

- **Deployment Profiles** configure the interfaces and map the labels to them, to characterize the appliance.
- **Business Intent Overlays** use the Labels specified in Deployment Profiles to define how traffic is routed and optimized between sites. These overlays can specify preferred paths and can link bonding policies based on **application**, **VLAN**, or **subnet**, independent of the brand and physical routing attributes of the underlay.

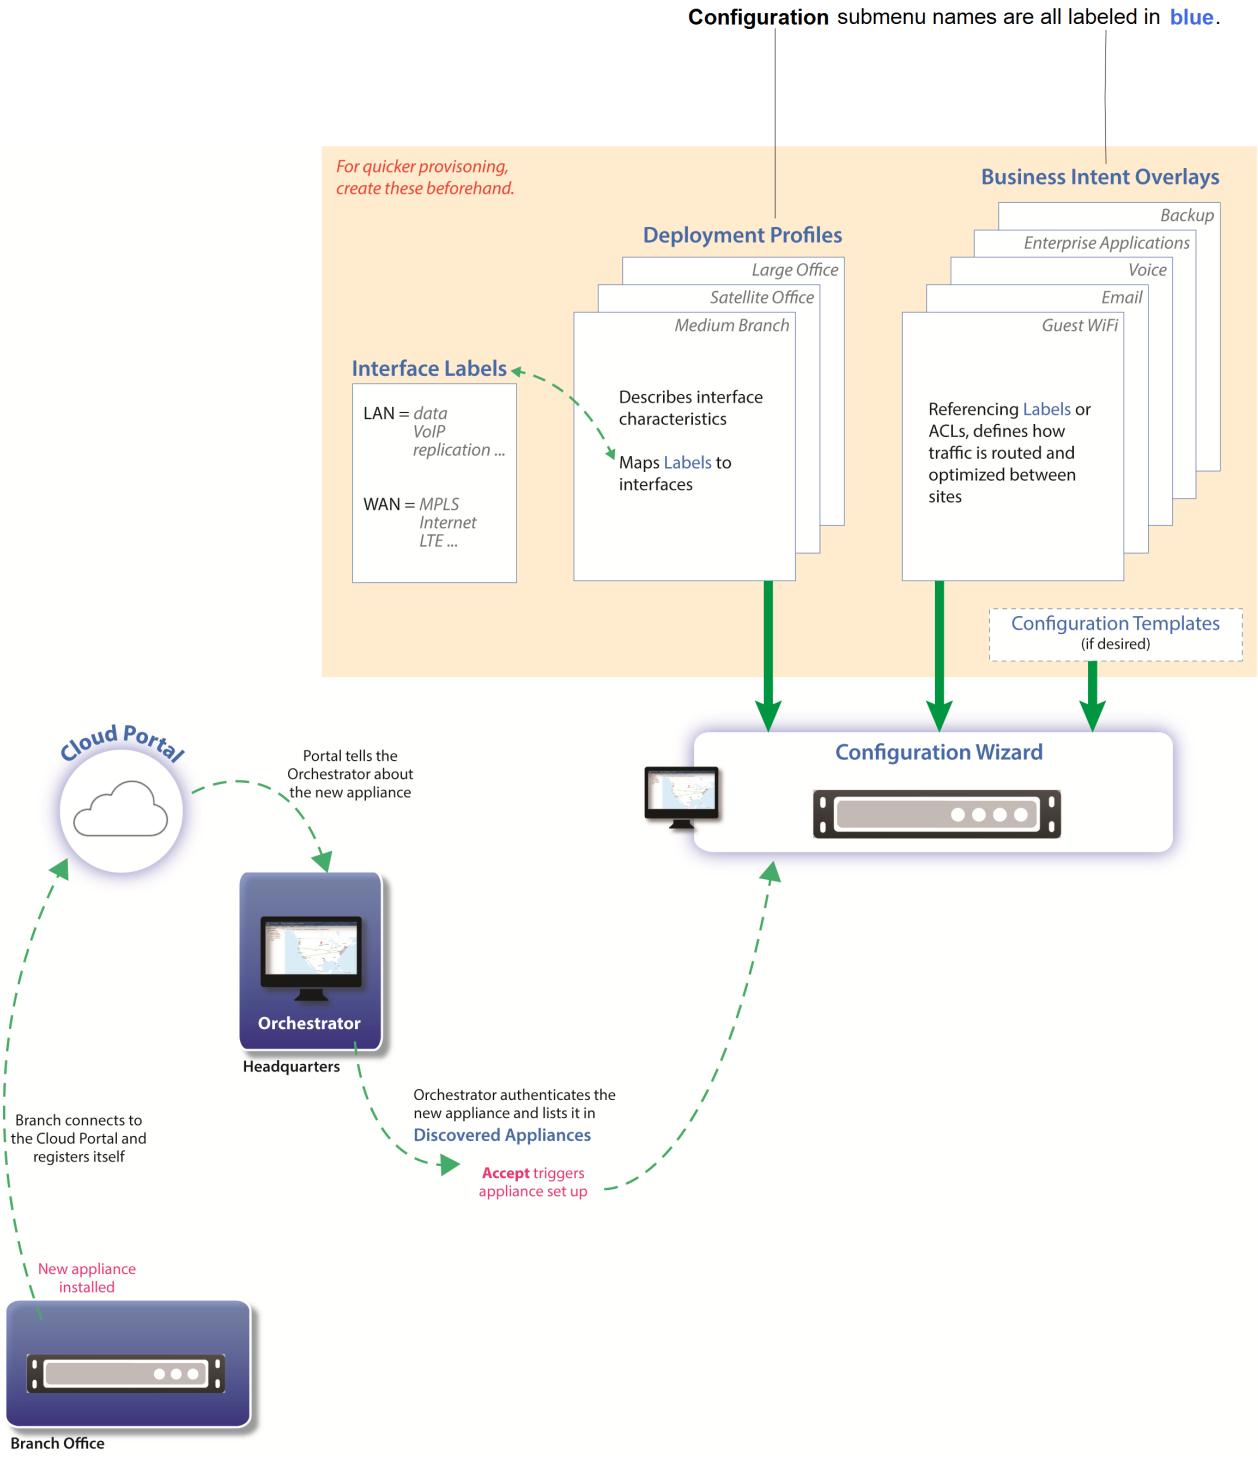
  This diagram shows the basic architecture and capabilities of **Overlays**.



Including a new appliance into the SD-WAN fabric consists of two basic steps:

1. **Registration and discovery.** After you **Accept** the discovered appliance, the **Configuration Wizard** opens.
2. **Provisioning.** Because the wizard prompts you to select profiles, it is easier to create these ahead of time.

The following figure shows the process of installing and provisioning an appliance for SD-WAN.


Configuration submenu names are all labeled in blue.
For quicker provisoning, create these beforehand.
Business Intent Overlays
Deployment Profiles
Backup
Enterprise Applications
Voice
Email
Guest WiFi
Large Office
Satellite Office
Medium Branch
Interface Labels
LAN = data VoIP replication ...
WAN = MPLS Internet LTE ...
Describes interface characteristics
Maps Labels to interfaces
Referencing Labels or ACLs, defines how traffic is routed and optimized between sites
Configuration Templates
(if desired)
Configuration Wizard
Cloud Portal
Portal tells the Orchestrator about the new appliance
Orchestrator
Headquarters
Orchestrator authenticates the new appliance and lists it in Discovered Appliances
Accept triggers appliance set up
Branch connects to the Cloud Portal and registers itself
New appliance installed
Branch Office

SASE Copilot

SASE Copilot provides AI-assisted search and contextual help within Orchestrator, enabling users to understand their SD-WAN and SASE environment, interpret system state, and navigate product documentation. The capability is designed to augment, not replace, existing Orchestrator workflows.

NOTE: SASE Copilot is delivered as a cloud service and is available only on cloud-hosted Orchestrator. It is not available on on-premises (self-hosted) Orchestrator deployments.

IMPORTANT: SASE Copilot answers only from its indexed knowledge base. If information is not available, it may be unable to provide a complete answer or may return a partial response. SASE Copilot is powered by AI, so mistakes are possible. Review all output carefully before use.

When using SASE Copilot:

- Expect strong results for documentation-backed “what is” and “explain” questions, and for the curated workflows described in Curated Workflows.
- Open-ended operational questions may be answered only in part, because the curated workflows and prompts cover a subset of everything the system can do. This coverage continues to expand over time.
- Do not rely on it as a replacement for advanced operational runbooks.
- Treat answers as contextual assistance, validated against official documentation when performing changes.

Availability

SASE Copilot is delivered as a cloud service and is updated on its own cadence, separate from the Orchestrator software release cycle. New SASE Copilot capabilities can become available between Orchestrator releases. The documentation is also updated regularly.

To discover the current capabilities at any time, open the SASE Copilot panel and ask a general question, such as:

> What are your capabilities?

SASE Copilot responds with a summary of what it can currently help you do, reflecting the most recent updates to the service.

Curated Workflows

SASE Copilot uses a set of curated internal workflows to interpret and respond to user questions. Each workflow is optimized for a specific class of questions and determines how the system retrieves information and generates responses.

For most questions, workflow selection is handled automatically. However, there are two exceptions: hop-by-hop path analysis and tunnel troubleshooting. These workflows are not launched from the SASE Copilot panel; instead, you launch them from a specific Orchestrator page by clicking the SASE Copilot icon (<img src](./Media/sase-copilot-icon.svg" alt="SASE Copilot icon" width="18" height="18" style="vertical-align:middle">). See How Do I Launch SASE Copilot?

Workflow Categories

Examples of curated workflow categories include the following:

- **Documentation workflow**—Answers questions that map directly to published SD-WAN and Orchestrator documentation.
- **Network health workflow**—Addresses questions related to the interpretation of network state, health indicators, and system behavior as reported by Orchestrator.
- **License summary workflow**—Provides structured, summary-style responses related to licensing information.
- **Path or hop-by-hop analysis workflow**—Supports structured visualization and analysis of traffic paths, including LAN and WAN routing, overlays, security, and per-hop performance details. This workflow is launched from the **Flows** page by clicking the SASE Copilot icon. See How Do I Launch SASE Copilot?
- **Tunnel troubleshooting workflow**—Supports structured diagnostic analysis of tunnel state, errors, and performance. This workflow is launched from the **Tunnels** page by clicking the SASE Copilot icon. See How Do I Launch SASE Copilot?

How Do I Launch SASE Copilot?

You can access the SASE Copilot icon (<img src](./Media/sase-copilot-icon.svg" alt="SASE Copilot icon" width="18" height="18" style="vertical-align:middle">) from two places in the Orchestrator interface:

- To the right of the search box at the top of the screen, where it opens the SASE Copilot panel.
- Within specific monitoring and configuration grids, where it launches a workflow scoped to the row you select.

The hop-by-hop path analysis and tunnel troubleshooting workflows are launched **only** from this icon on their respective pages. They are not available from the Copilot panel prompt suggestions and cannot be initiated as free-text questions.

Workflow	Where to find the icon	How to launch
Hop-by-hop path analysis	Flows tab, in the AI Insights column	Select the flow you want to analyze, then click the SASE Copilot icon in that row.
Tunnel troubleshooting	Tunnels tab, in the Troubleshooting column of the underlay tunnels grid	Select the underlay tunnel you want to troubleshoot, then click the SASE Copilot icon in that row.

NOTE: The SASE Copilot icon appears in these tables only when SASE Copilot is enabled, and you are signed in with Read & Write permissions or as an administrator.

Enable SASE Copilot

SASE Copilot is disabled by default. An administrator must explicitly enable the feature before it can be used. Enabling SASE Copilot allows Orchestrator to send user queries and related context to an AI service to generate responses within Orchestrator.

Before You Begin

- You must be signed in to Orchestrator with Read & Write permissions or as an administrator.
- You must review and accept the SASE Copilot Terms of Use during the enablement process.

Launch SASE Copilot

1. Sign in to Orchestrator with Read & Write permissions or as an administrator.
2. Click the SASE Copilot icon (<img src](./Media/sase-copilot-icon.svg" alt="SASE Copilot icon" width="18" height="18" style="vertical-align:middle">) to the right of the search box at the top of the screen.

   The SASE Copilot panel opens.

   NOTE: The Copilot interface may be visible even when the feature is disabled, but access is restricted until enablement is complete.
3. When prompted, review the **Terms of Use**.

   The Terms of Use covers five sections: Scope and Acceptance, Use of AI Technology, User Responsibilities, Safety and Security, and Modification and Termination. A **View Data Privacy Policy** link is provided. You must review the entire document before the accept button becomes available.

4. Click **Accept Terms and Enable SASE Copilot**.

   This action accepts the Terms of Use and enables SASE Copilot for all authorized users on this Orchestrator instance.

After you accept the Terms of Use, SASE Copilot becomes available to authorized users.

NOTE: If you choose not to enable SASE Copilot at this time, the feature remains disabled, and no AI queries are processed.

Re-enable SASE Copilot

Disabling SASE Copilot does not prevent it from being re-enabled later. Only users with Read & Write permissions or an administrator can enable or disable SASE Copilot. If SASE Copilot is disabled or was not accepted during the initial prompt, you can re-enable it using the procedure below.

To re-enable SASE Copilot:

1. Sign in to Orchestrator with Read & Write permissions or as an administrator.
2. Navigate to **Orchestration > SASE Copilot**.
3. On the SASE Copilot page, locate the **Enable SASE Copilot** toggle switch.
4. Set the toggle to **ON**.

After enabling, SASE Copilot is available for use by authorized users.

Permissions and Roles

Access to SASE Copilot is controlled by Orchestrator roles. Only users with Read & Write permissions or administrators can use SASE Copilot. Read-only users cannot access SASE Copilot.

Read & Write and Administrator Roles

Users with Read & Write permissions or administrators:

- Can enable or disable SASE Copilot.
- Can review and accept the Terms of Use during the enablement process.
- Can use SASE Copilot after it is enabled.
- Can access chat history, clear context, and submit response feedback.

Read-Only Roles

Users with a read-only role:

- Cannot access SASE Copilot.
- May see the SASE Copilot interface entry point, but all queries and chat history requests are blocked.
- Receive a permissions error if they attempt to use the feature.

Operational and Important Behavior

- SASE Copilot operates in read-only mode with respect to Orchestrator configuration.
- The AI service does not make configuration changes or apply actions automatically.
- Enabling SASE Copilot does not grant additional privileges beyond a user's existing role.

Chat Management

When SASE Copilot is open, users with Read & Write permissions or administrators have access to the following controls in the Copilot panel header.

Display Mode

You can display SASE Copilot in two modes: docked and undocked.

- **Docked mode**—SASE Copilot opens as a panel on the right side of the Orchestrator interface.
- **Undocked mode**—SASE Copilot opens as a floating window. To switch to undocked mode, click the **Undock** icon in the panel header.

Orchestrator stores your display mode preference in the browser, allowing it to persist across sessions.

Clear Context

Click **Clear context** to close the current chat session and start a new one. Use this option to begin a fresh conversation without context carried over from previous questions.

Chat Refresh

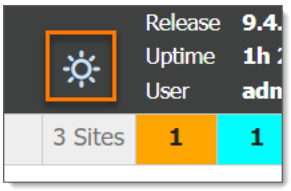
Click **Chat refresh** to reload the conversation history from the server. Use this option if the conversation appears out of sync or to restore chat history after navigating away.

Response Feedback

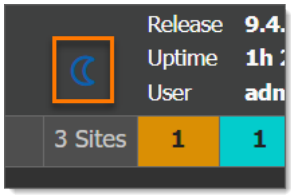
After SASE Copilot returns a response, you can submit feedback using the thumbs-up or thumbs-down icons. You can add an optional text comment before submitting your feedback. Feedback is used to improve response quality.

Light or Dark Theme

The Orchestrator user interface provides a light theme (light background with dark text) and a dark theme (dark background with light text). To toggle to the dark theme (dark background with light text), click the sun icon in the upper-right corner of the Orchestrator application.



Click the moon icon to toggle back to the light theme.



Orchestrator saves theme selections by user and browser.

Menu Options

All of Orchestrator's monitoring and configuration options are organized into five main menu groups, or tabs, located at the top of the main screen, including the following:

- Monitoring
- Configuration
- Administration
- Orchestrator
- Support

Monitoring

The options under the **Monitoring** tab focus on reports related to performance, traffic, and appliance status. Additionally, Threshold Crossing Alerts are helpful in monitoring your network.

Configuration

The options under the **Configuration** tab focus on how to configure Orchestrator. The options available under this menu are organized as follows:

- Overlays & Security
- Networking
- Templates & Policies
- Cloud Services

Administration

The options under the **Administration** tab are related to appliance administration. They include general settings, software management, and tools for troubleshooting and maintenance.

Orchestrator

The options under the **Orchestrator** tab are used for managing Orchestrator itself. These options do not relate to managing appliances.

Support

The options under the **Support** tab can be used when working with Support to facilitate opening a case or providing Support with data and reports needed to troubleshoot network issues.

Monitoring

The options under **Monitoring** focus on performance, traffic, and appliance status. Additionally, Threshold Crossing Alerts are helpful in monitoring your network.

Categories include the following:

- Summary
- Reporting
- Bandwidth
- Tunnel Health

Monitoring > Summary

The options under **Monitoring > Summary** focus on Orchestrator monitoring features, such as the Dashboard, which provides a unified display of your network; the Topology Tab, which provides a visual display of your network; the Health Map, which provides a high-level view of your network's health; and the Alarms Tab, which provides details about both appliance and Orchestrator alarms.

The Alarms topic also provides a detailed list of alarms related to EdgeConnect appliances and Orchestrator.

Dashboard

Monitoring > Summary > Dashboard

The Dashboard integrates information from multiple components—or widgets—into a unified display for monitoring your network. It displays appliance license information, topology, health map, top talkers, top domains, and so forth, on one tab. The collection of widgets are customizable and persist for each user account.

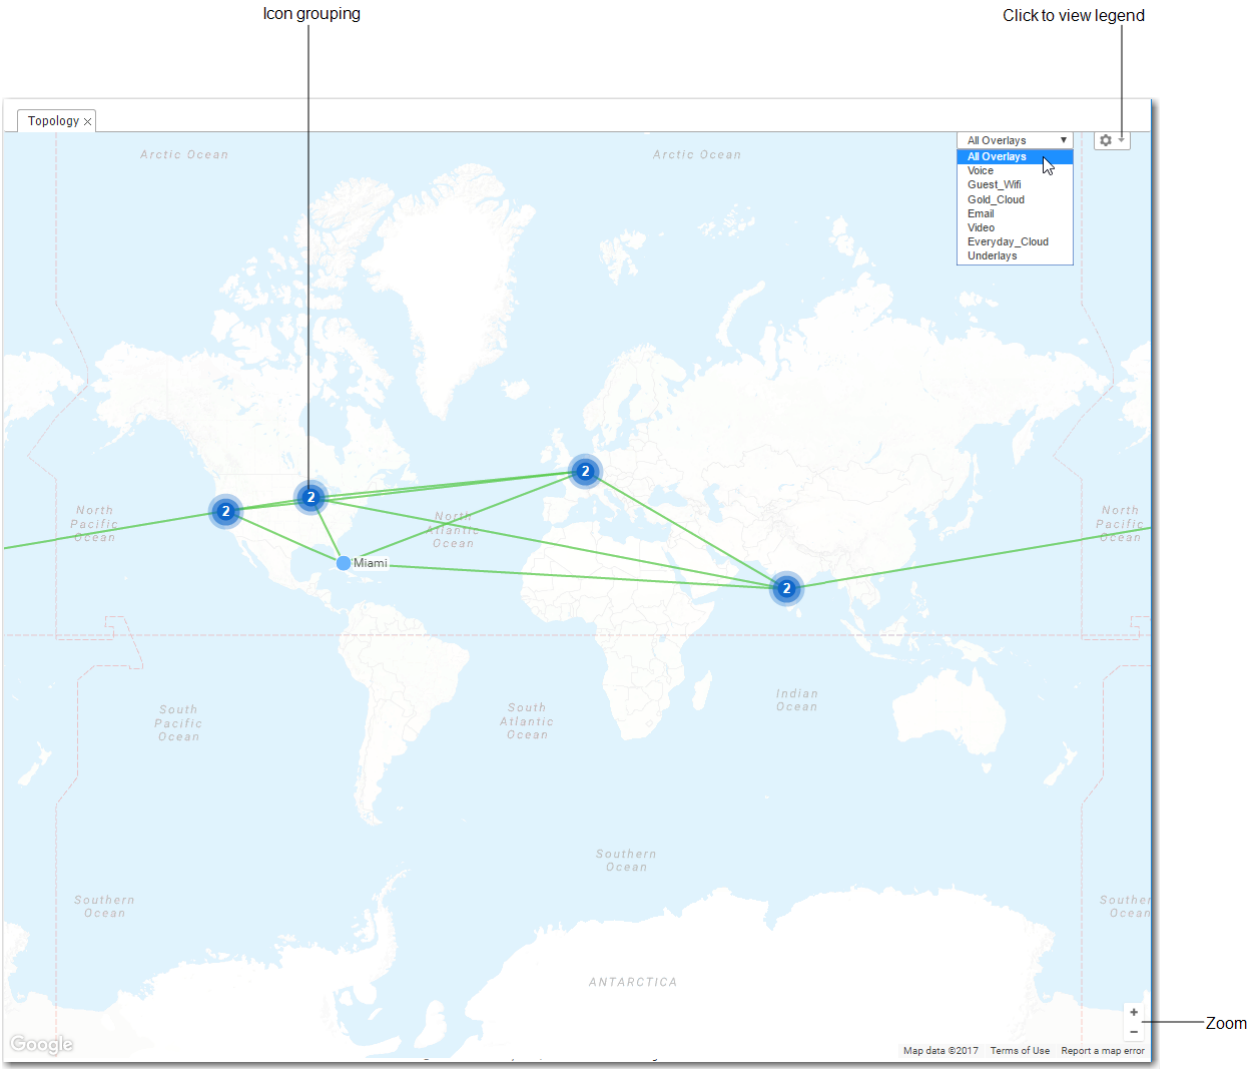
- Click **Settings** [ ] to select the widgets you want to show or hide.
- To move widgets, drag them by title.
- To access more detail in its corresponding tab, click a widget's title.
- To filter on various widgets, select **Src** or **Dest**, **Overlay** or **Underlay**, or **Inbound** or **Outbound**. The filter varies depending on the widget you are selecting.
- You can choose and change the grouping variable for Overlay-Transport and Overlay-Interface by clicking **Flip**.
- The **Appliance Licenses** widget displays an inventory by appliance model, as well as license type, availability, and usage.
- To search for appliances in the tree, enter an appliance name and the tag will be displayed above the tree.
- To filter collections of appliances, select **Show Tags** and select from among the tag options.

Topology

Monitoring > *Summary* > *Topology*

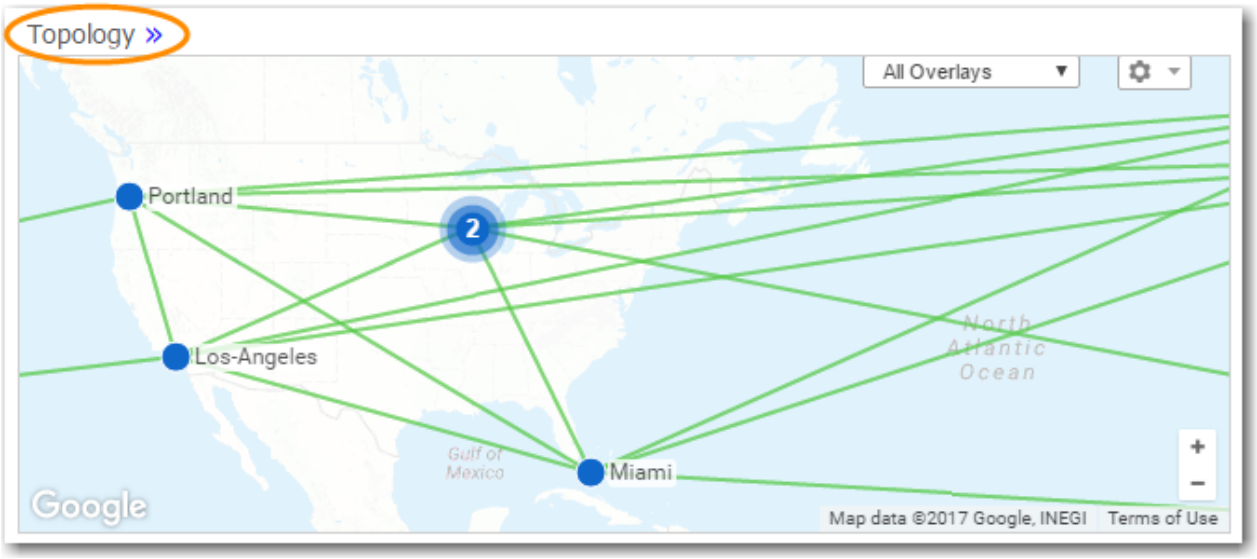
The Topology tab provides a visual summary of your Silver Peak network.

When configuring a software-defined WAN (**SD-WAN**), you can view **All Overlays**, individual **Business Intent Overlays** (BIOs), or the single and bonded **Underlay** tunnels that support them.

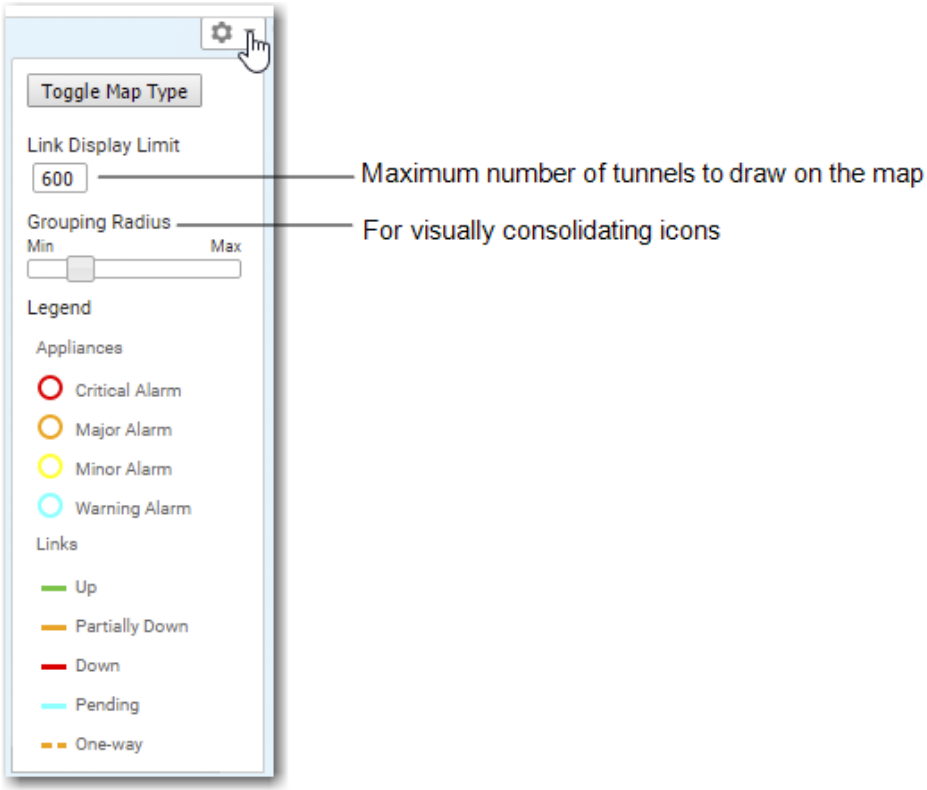


You can access it under **Monitoring** in the menu bar, or by clicking the widget title on the **Dashboard** tab.

Topology widget on Dashboard tab

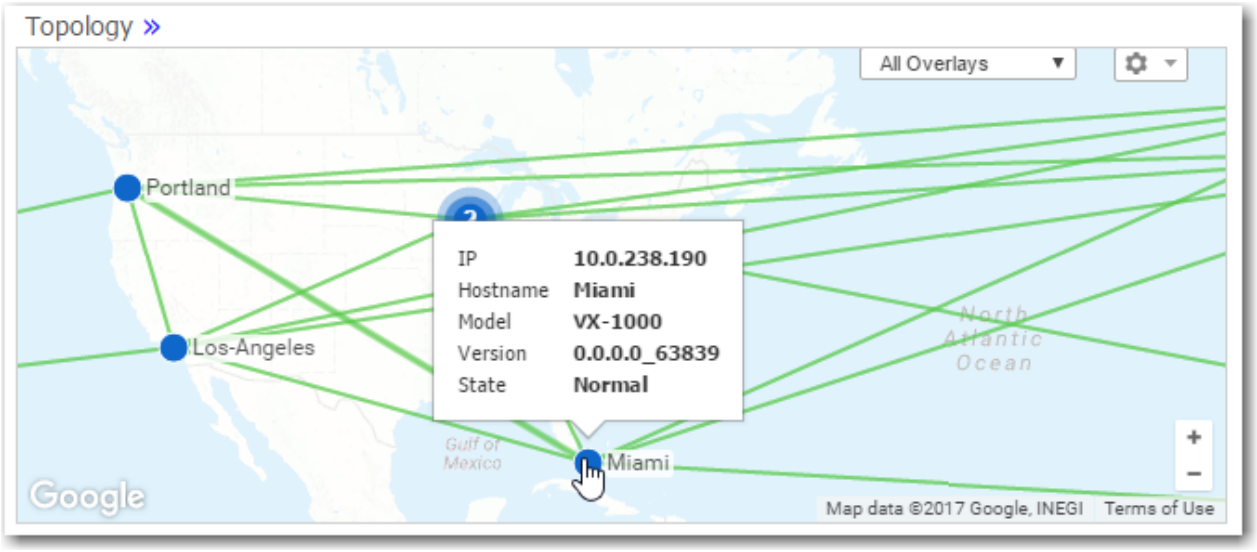


- The Legend details the appliances' management and operational states.

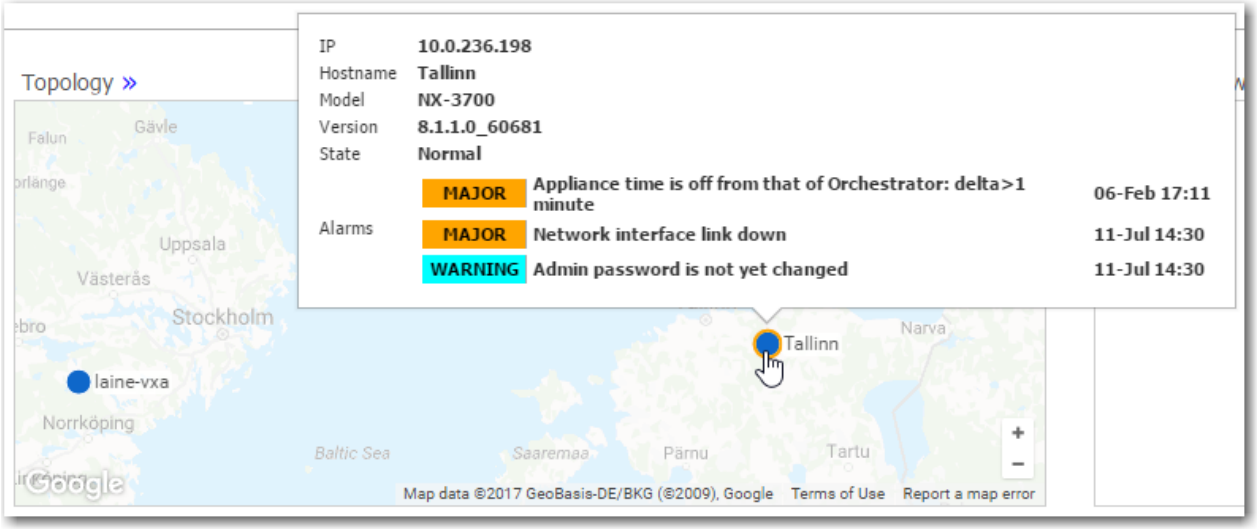


- The **Topology** map can dynamically geolocate an appliance when you enter a location (City, State, Country) in an appliance Configuration Wizard, or when you modify the appliance by right-clicking to access its contextual menu.
- The map tile renders to support variable detail at different zoom levels.

- You can use icon grouping to visually consolidate adjacent appliances. The status bubbles up, and you can configure relative grouping distance in the map's legend. The grouping is also a function of how far you zoom in or out.
- Rolling over an individual appliance's icon displays basic system information.



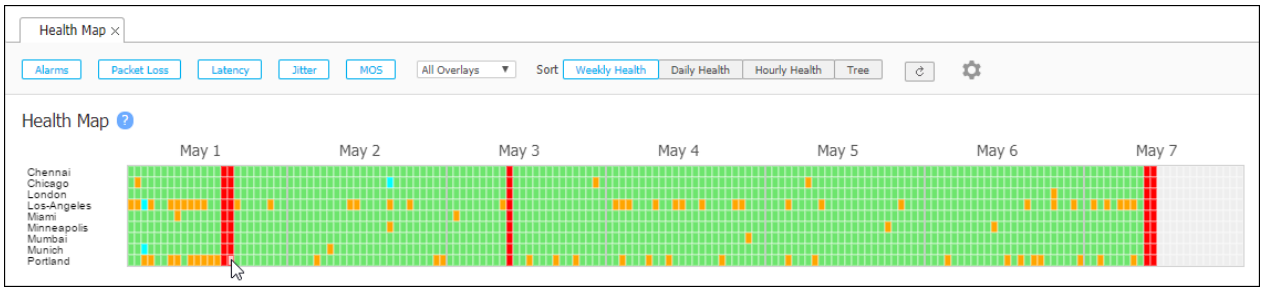
When the icon is encircled by a ring, indicating an alarm, those also display.



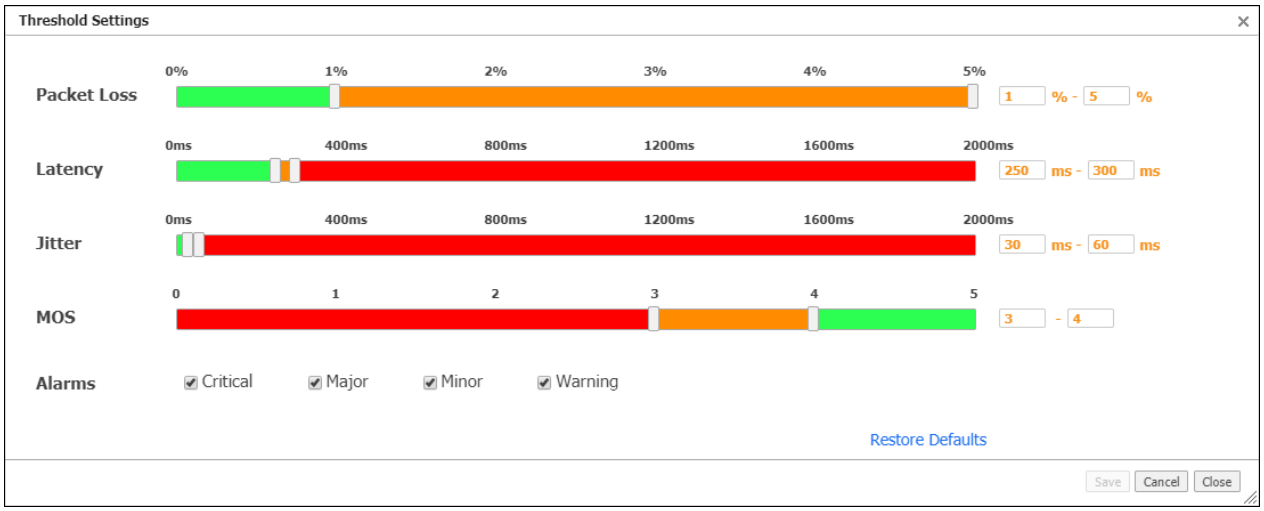
Health Map

Monitoring > Summary > Health Map

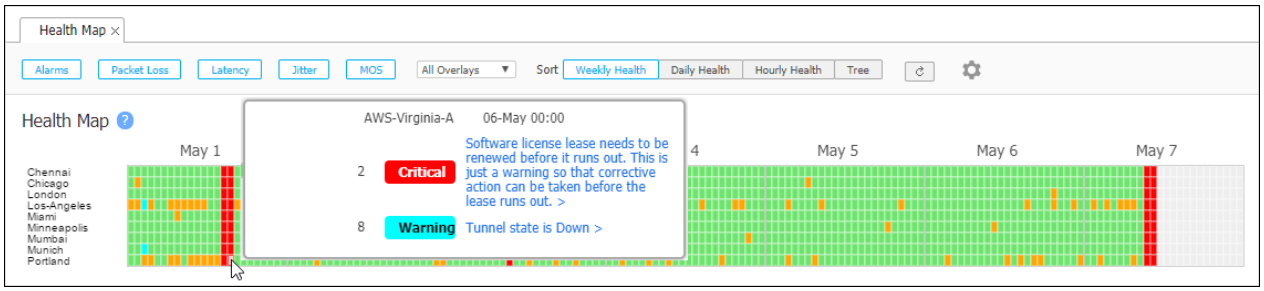
The **Health Map** provides a high-level view of your network's health, based on real-time measurements of network conditions between appliances.



- View filters are available for alarms, packet loss, latency, jitter, MOS (mean opinion score), and Business Intent Overlay.
- The health map can be sorted by weekly, daily, hourly health, or tree (by group, and then alphabetical by hostname).
- Each block represents one hour and uses color coding to display the most severe event among the selected filters. Color codes correspond to alarm severity and thresholds.
  - **Green** – Normal operation.
  - **Red** – Critical. Steps must be taken immediately in order to restore the affected service.
  - **Orange** – Major. Steps must be taken as soon as possible because the affected service has degraded drastically.
  - **Yellow** – Minor. A problem that does not yet affect service, but could if the problem is not corrected.
  - **Aqua** – Warning. A potential problem that could affect service.
  - **Grey** – No data available.
- Thresholds can be configured by clicking on the gear icon .



- Clicking a color block displays a pop-up with specifics about that event, what value triggered it, and any additional threshold breach for that appliance during the same hour.



- While filter and sort order customizations persist for each user account, threshold settings apply globally.
- Threshold settings are not retroactive. Setting new thresholds does not redisplay historical data based on newly edited values.
- Deleting an appliance deletes its data.
- If you are using overlays, note the following:
  - You can view each overlay's health individually.
  - If you remove an individual overlay, its data is not recoverable. However, its historical data remains included in **All Overlays**.

Alarms Tab

Monitoring > *Summary* > *Alarms*

An alarm is a system-generated notification that is triggered by an event, error, or condition within Orchestrator or an appliance. The table on the Alarms tab is displayed in one of two modes, which is toggled by clicking **Standard** or **Correlated**.

- **Standard mode**: the table displays a list of appliance and Orchestrator alarms that indicate conditions that may require human intervention. Alarms are triggered by conditions that may come and go without management involvement.
- **Correlated mode**: when correlation is enabled, the table displays a list of "Root Causes". Each root cause can be expanded to display the alarms that were triggered by conditions that either resulted in or are related to the root cause. Correlation is enabled from the **Correlation Settings** dialog box.

Standard

Click **Standard** to place the Alarms table in Standard viewing mode. In this mode, each entry is an appliance or Orchestrator alarm that indicates a condition that may require human intervention. Because alarms are conditions, they can come and go without management involvement. While merely acknowledging most alarms does not clear them, some alarm conditions are set up to self-clear when you acknowledge them. For example, if you remove a hard disk drive, it generates an alarm; after you replace it and it finishes rebuilding itself, the alarm clears.

Information that each alarm entry provides includes the Host Name, Alarm Time, Condition Source, an Alarm Description, and a recommended action.

Correlated

Click **Correlated** to place the Alarms table in Correlated viewing mode. In this mode, each entry is a root cause and can be expanded to display the alarms that generated the root cause. Each displayed alarm shown in Correlated mode is also shown in Standard mode.

Intelligent Alarm Correlation aggregates multiple alarms with the same root cause into a single, correlated alarm. Alarm Correlation reduces alarm noise by intelligently grouping related alarms that are triggered by a single underlying event. Correlation simplifies root cause monitoring and analysis.

Large SD WAN deployments often generate many alarms from a single event that can obscure root causes and increase diagnostic and remedial efforts. A single event can also generate alarms at multiple locations. Alarm Correlation can help determine if the trigger is a local or network failure, the root location of the alarm, the cause of the alarm, and the nature of the failure, such as an downed interface or unreachable next hop.

When Alarm Correlation is enabled, Orchestrator evaluates inbound alarms and determines their relationship to existing root cause conditions. Related alarms are grouped under a correlated alarm; the original alarms are retained for troubleshooting context. The Correlated table displays a list of root causes, each of which can be expanded to display the alarms that Orchestrator determines is related to the cause. The Standard table displays a sortable list of all alarms generated by the selected appliances.

Orchestrator automatically performs Alarm Correlation using predefined correlation logic. Manual alarm grouping is not required or supported. Orchestrator Correlation interval is 5 minutes, which is not configurable.

Filtering Alarms

Unless otherwise stated, you can filter the alarms listed in each Alarms table mode as follows:

- Time: **1h**, **4hr**, **1d**, **7d**, or **Custom**. Custom enables you to specify a range of dates in the Range fields.
- **Active:** All uncleared alarms. Acknowledged alarms go to the bottom of this list.
- **History:** Filtered to show only cleared alarms.
- **All:** All uncleared and cleared alarms (not available in Correlated mode).

NOTE: Orchestrator keeps a history of alarms for 7 days. If you use the on-prem version of Orchestrator, you can configure it to keep a history for more than 7 days. If you use Orchestrator-as-a-service, this cannot be changed.

The Alarms tab also includes the following functionality:

- **Alarm Emails ON** and **Alarm Emails Paused:** Enable or disable to receive an email if an alarm is on or paused.

- **Alarm Email Recipients:** Each configured recipient can receive emails about either appliance alarms or Orchestrator alarms. Click **Add Recipient** in the Alarm Recipients dialog box. Select the appropriate type of alarm from the **Alarm Type** drop-down list, and then select the check boxes (**Critical**, **Major**, **Minor**, **Warning**) for which you want to receive emails. Click **Save** or **Reload**.
- **Wait to Send Emails:** You can customize the amount of time you want the system to wait to send you an email notifying you of an alarm. Click this button to open the Wait to Send Emails dialog box, and then enter the number of minutes you want the system to wait. Click **Save**.
- **Export:** You can export a CSV file of your alarms.
- **Ack**, **Acked By**, and **Acked Time:** These columns in the Alarms table indicate whether an acknowledgment has been received between devices.
  - **Acked By:** Name of the appliance that did the acknowledgment.
  - **Acked Time:** Time when the acknowledgment was received by the appliance.

Disable Alarms

You can specify the alarms you want to disable by clicking **Customize / Disable Alarms**, which opens the Alarm Information dialog box.

To disable alarms:

1. Click **Disable All Alarms on Specific Appliances**.
2. Enter the name of the appliance that has the alarms you want disabled.
3. Click **Disable Alarms**.
4. Click **Save**.

Customize Alarms

Complete the following steps to customize a pre-existing alarm.

1. Select the edit icon next to the selected appliance in the Alarm Information dialog box.
2. Choose **Enable/Disable**.
3. If enabling, specify the **Custom Severity** by choosing from the list: **None**, **CRITICAL**, **MAJOR**, **MINOR**, **WARNING**.
4. If disabling, the following message appears: *You are about to disable this alarm. Click **Save**.

Alarm Severity

Orchestrator has four severity levels for alarms:

- **Critical** (red): Critical alarms are service-affecting and require immediate attention. They reflect conditions that adversely affect an appliance or indicate the loss of a broad category of service.

- **Major** (orange): While service-affecting, major alarms are less severe than critical alarms. They reflect conditions that should be addressed in the next 24 hours. An example would be an alarm caused by an unexpected traffic class error.
- **Minor** (yellow): Minor alarms are not service-affecting and can be addressed at any time. Examples include alarms caused by a user who has not changed their account's default password, a degraded disk, or a software version mismatch.
- **Warning** (blue): Warning alarms are not service-affecting. They warn of conditions that could become problems over time—for example, an alarm caused by IP SLA being down.

Alarm Recipients

To add alarm recipients to receive email messages that notify them of alarms within the network:

1. Click **Alarm Email Recipients**.
2. Click **Add Recipient**.
3. Enter the following information:
   - In the **Alarm Type** field, select **Orchestrator** for Orchestrator alarms or **Appliance** for appliance-generated alarms.
   - To indicate the severity of alarms to send to recipients, select the **Critical**, **Major**, **Minor**, and **Warning** check boxes as appropriate. Major and Minor are selected by default.
   - In the **Format** field, select the format to use for the alarm notifications. By default, alarms are **HTML Formatted**. However, you can choose **Plain Text** or **Both**.
   - In the **Email Addresses** field, enter email addresses for recipients you want to receive alarm notifications.
   - In the **Groups** field, select the appropriate groups from the drop-down list. The listed groups reflect the groups configured in the appliance tree.
   - To include detailed diagnostic or debugging information in alarm notification email messages, select the **Debug** check box.

     IMPORTANT: It is recommended that you use the Debug option cautiously and disable it after your troubleshooting is finished to avoid unnecessary logging or performance overhead. Because debugging information can contain sensitive system information, consider limiting access to authorized personnel only.
4. Click **Save**.

Plain Text alarms are emailed as pipe-separated data. Recipients can create a script to parse the message and read the fields.

Example:

```
Hostname|Alarm_Status|Time|Alarm_ID|Type_ID|Source|Severity|Description|
    Recommended_action
Orchestrator|1|1526341365000|94|6815775|orchestrator|MINOR|Backup configuration not
    set|
Orchestrator|1|1526341362000|93|6815762|orchestrator|MAJOR|Orchestrator is using the
    default SMTP settings|
```

Alarm ID is the auto-incremented, primary key in the database. Alarm Status indicates 0 for cleared or 1 for raised.

Correlation Settings

The **Correlation settings** dialog box controls Correlation alarm content received by the selected remote server. This dialog box is accessed by clicking the Correlation settings text in the title bar.

To enable alarm aggregation and Orchestration correlation, select the **Correlation** check box.

To control the Alarm and Correlation information sent to the remote server, set the **Alarm notification** dropdown to the desired setting. The following table displays the sent content for each valid alarm notification setting.

Alarm Notification Content Settings

Correlation	Alarm Notification Setting	Remote server receives the following:
Disabled	Standard	Orchestrator Alarms, All appliance alarms
Enabled	Standard	Orchestrator Alarms, Appliance independent alarms, Appliance aggregated alarms
Enabled	Correlated	Orchestrator Alarms, Appliance independent alarms, Root cause alarms (identified through correlation)
Enabled	Correlated and constituents	Orchestrator Alarms, Appliance independent alarms, Appliance aggregated alarms, Root cause alarms (identified through correlation)

Additional Alarm Indications

- A cumulative (Orchestrator + appliances) alarm summary always displays at the right side of the header. Clicking it opens a top-level summary and access to the Alarms tab.
- Appliances are color coded to indicate their severest alarms on the Topology tab and in the navigation pane.
- Threshold crossing alerts are related to alarms. They are preemptive, user-configurable thresholds that declare a Major alarm when crossed. For more information about their configuration and use, see Threshold Crossing Alerts Template and Threshold Crossing Alerts Tab.

Export Alarm Descriptions

Orchestrator enables you to export to a CSV file a full list of alarms you could potentially receive. This file includes a variety of details about the listed alarms, including alarm descriptions and recommended actions. For details, see Alarm Descriptions.

To automatically export the CSV file, go to **Support > User Documentation > Alarm Descriptions**.

List of Alarms

This topic provides lists of alarms related to EdgeConnect appliances and Orchestrator.

NOTE: The tables in this topic use the decimal numeral system for Alarm ID. You can convert these numbers to the hexadecimal numeral system if you have applications that can do their own filtering, such as SNMP.

EdgeConnect Appliance Alarms

Appliances can raise alarms based on issues that occur with tunnels, software, equipment, and Threshold Crossing Alerts (TCAs). TCAs are visible on the appliance but are managed by Orchestrator.

Tunnels

System Type 0 (Appliance); Source Type 1 (Tunnel)

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
65537: CRITICAL	Tunnel is down. **Recommended Action:** Tunnel peer is unreachable. Check tunnel configuration. Verify Local & Remote IPs, Admin up, and peer's Mode matches. Check network connectivity. **NOTE:** Down-in progress overlay tunnels are not counted as down tunnels, as the appliance is in the process of bringing the tunnel up, and they are not hard down.	Appliance	TRUE	TRUE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
65539: CRITICAL	Tunnel protocol version mismatch. **Recommended Action:** Tunnel peers are running incompatible software versions. Normal during a software upgrade. Run the same or compatible software releases among the tunnel peers.	Appliance	TRUE	TRUE
65542: CRITICAL	Tunnel peer type mismatch. [For VX-Xpress only] **Recommended Action:** VX-Xpress appliance can only peer with another VX-Xpress appliance. Create a tunnel to another VX-Xpress appliance.	Appliance	TRUE	FALSE
65543: CRITICAL	Duplicate license detected. **Recommended Action:** Duplicate serial numbers detected. Install unique license on all virtual appliances. To check and/or change license: In Appliance Manager: Administration > Basic Settings > License & Registration In Orchestrator: Configuration > Overlays & Security > Licensing > Licenses	Appliance	TRUE	TRUE
65544: CRITICAL	Tunnel has invalid source IP address. **Recommended Action:** Delete the tunnel and re-create it with a valid IP address.	Appliance	TRUE	TRUE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
65545: CRITICAL	Tunnel received an unmatched GRE packet. **Recommended Action:** Check for tunnel encapsulation mismatch. On the Tunnels page, go to specified tunnel and verify both tunnel peers are using the same encapsulation method.	Appliance	TRUE	TRUE
65548: CRITICAL	Tunnel is using a not allowed cipher. **Recommended Action:** Check tunnel configuration and ensure that only allowed ciphers are used.	Appliance	TRUE	TRUE
65536: MAJOR	Tunnel is misconfigured. **Recommended Action:** System ID is not valid. Was appliance registration completed?	Appliance	TRUE	TRUE
65546: MAJOR	Tunnel is in reduced functionality. **Recommended Action:** Tunnel peers are not running the same release of software. This results in reduced functionality. Run the same or compatible software releases among the tunnel peers.	Appliance	TRUE	TRUE
65547: MAJOR	Tunnel UDP port conflicts with cluster port. **Recommended Action:** [Deprecated alarm] Choose another number for UDP Destination Port on local and remote appliances if using the same interface for UDP tunnel and flow redirection.	Appliance	TRUE	TRUE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
65541: MINOR	Tunnel software version mismatch. **Recommended Action:** Tunnel peers are not running the same release of software, but the releases are completely compatible. Normal during an upgrade. Run the same software version to eliminate the alarm.	Appliance	TRUE	TRUE
131072: MINOR	Tunnel has unexpected traffic class error. **Recommended Action:** [Deprecated alarm]	Appliance	TRUE	TRUE

Software

System Type 0 (Appliance); Source Type 4 (Software)

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
262147: CRITICAL	The licensing for this virtual appliance has expired. [For VX series only] **Recommended Action:** Enter a new license key for the appliance. **NOTE:** The VX appliances are a family of virtual appliances, comprised of the VX-n000 software, an appropriately paired hypervisor and server, and a valid software license.	Appliance	TRUE	FALSE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
262148: CRITICAL	There is no license installed on this virtual appliance. [For VX series only] **Recommended Action:** Enter a valid license key for the appliance. **NOTE:** The VX appliances are a family of virtual appliances, comprised of the VX-n000 software, an appropriately paired hypervisor and server, and a valid software license.	Appliance	TRUE	FALSE
262156: CRITICAL	Invalid virtual appliance license. [For VX series only] **Recommended Action:** Enter a valid license key for the appliance.	Appliance	TRUE	FALSE
262165: CRITICAL	Software license token needs to be renewed. **Recommended Action:** Software will automatically renew the license lease as long as it has HTTPS connectivity to the internet.	Appliance	FALSE	TRUE
262166: CRITICAL	Software capability token expired. **Recommended Action:** Portal (30-day token) expired; no communication with Portal in 30 days. You must have HTTPS connectivity to internet to renew the license lease.	Appliance	TRUE	FALSE
262171: CRITICAL	Invalid account registration information. **Recommended Action:** Check that the account registration information is correct.	Appliance	TRUE	FALSE
262172: CRITICAL	EC Base license not granted. **Recommended Action:** Contact Support to obtain additional EdgeConnect licenses. If you have licenses, approve this appliance from your Orchestrator.	Appliance	TRUE	FALSE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
262173: CRITICAL	Orchestrator is unreachable. **Recommended Action:** Either the primary or backup Orchestrator (or both) cannot be reached using HTTPS. This connectivity is required for appliances to be managed by Orchestrator.	Appliance	TRUE	FALSE
262175: CRITICAL	Cloud Portal host name cannot be resolved. **Recommended Action:** Check if appliance has been configured with a reachable DNS server. If there is no DNS server configured, appliance tries to use built-in DNS servers on the Internet to resolve the portal hostname.	Appliance	TRUE	FALSE
262176: CRITICAL	EC Plus license not granted. **Recommended Action:** Contact Support to obtain additional licenses.	Appliance	TRUE	FALSE
262177: CRITICAL	EC WAN Optimization license not granted. **Recommended Action:** Contact Support to obtain additional licenses.	Appliance	TRUE	FALSE
262178: CRITICAL	Appliance has not been approved by Orchestrator. **Recommended Action:** Approve the appliance from your Orchestrator.	Appliance	TRUE	FALSE
262179: CRITICAL	Software licensing error. **Recommended Action:** Failing to get token from Portal. Contact Support.	Appliance	TRUE	FALSE
262180: CRITICAL	No public IP address detected on an interface behind Internet. **Recommended Action:** [Deprecated alarm] Connect the interface to Internet.	Appliance	TRUE	FALSE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
262182: CRITICAL	DHCP server misconfiguration. **Recommended Action:** DHCP server configuration contains invalid entry that prevented it from running. Check log file for error and verify your configuration.	Appliance	TRUE	FALSE
262185: CRITICAL	Unable to resolve Orchestrator DNS name. **Recommended Action:** Could not resolve one or more Orchestrator DNS names. Check DNS server configuration.	Appliance	TRUE	TRUE
262209: CRITICAL	Invalid account registration information. **Recommended Action:** Check that the account registration information is correct.	Appliance	TRUE	FALSE
262210: CRITICAL	Cloud Portal not connected. **Recommended Action:** Check the Cloud Portal connection.	Appliance	TRUE	FALSE
262211: CRITICAL	Database errors. **Recommended Action:** If the issue is not resolved, contact Technical Support.	Appliance	TRUE	FALSE
262213: CRITICAL	Config DB load partially failed. **Recommended Action:** Check configuration and apply again.	Appliance	TRUE	TRUE
262220: CRITICAL	End entity certificate has expired. **Recommended Action:** Renew the end entity certificate that has expired.	Appliance	TRUE	TRUE
262224: CRITICAL	One or more peers cannot support the current number of IPv4 routes. **Recommended Action:** Some routes are being dropped. Upgrade all appliances to a compatible ECOS version to receive all routes.	Appliance	TRUE	TRUE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
262225: CRITICAL	One or more peers cannot support the current number of IPv6 routes. **Recommended Action:** Some routes are being dropped. Upgrade all appliances to a compatible ECOS version to receive all routes.	Appliance	TRUE	TRUE
262229: CRITICAL	End entity certificate est enrollment failed. **Recommended Action:** Check that est server enrollment is working correctly.	Appliance	TRUE	TRUE
262230: CRITICAL	HPE End entity certificate procurement timed out. **Recommended Action:** Hpe_cert procurement failed. Check connectivity to Portal, and then disable and enable hpe_cert.	Appliance	TRUE	TRUE
262231: CRITICAL	HPE End entity certificate validation failed. **Recommended Action:** Hpe_cert validation or save failed. It is possible that the trust store does not have RDA root CA certificate. Add it, and then disable and enable hpe_cert.	Appliance	TRUE	TRUE
262232: CRITICAL	HPE End entity certificate expired or revoked. **Recommended Action:** HPE certificate expired or revoked. Disable and enable hpe_cert to reissue.	Appliance	TRUE	TRUE
262243: CRITICAL	Secure Web Services database load failed. **Recommended Action:** Secure Web Services Database load failure. The Secure Web Services–based categorization of the flows shall continue as per the older database if already loaded.	Appliance	TRUE	TRUE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
262244: CRITICAL	Secure Web Services threat intelligence server not reachable. **Recommended Action:** Secure Web Services external server is not reachable. Check the connectivity.	Appliance	TRUE	TRUE
262245: CRITICAL	Secure Web Services not enough disk space on the appliance. **Recommended Action:** Secure Web Services database files cannot be downloaded to the appliance. Check the disk space availability in the /var/opt/tms.	Appliance	TRUE	TRUE
262247: CRITICAL	Unable to resolve proxy server DNS name. **Recommended Action:** Could not resolve the proxy server DNS name. Check the DNS server configuration.	Appliance	TRUE	FALSE
262248: CRITICAL	Proxy server is unreachable. **Recommended Action:** Portal and Orchestrator cannot be reached using HTTPS via Proxy. Check Proxy configuration.	Appliance	TRUE	FALSE
262144: MAJOR	Software upgrade process has failed. **Recommended Action:** If the software upgrade repeatedly fails, open a support case.	Appliance	FALSE	TRUE
262145: MAJOR	System is low on resources. **Recommended Action:** The appliance is running low on resources (memory). If this alarm persists, open a support case.	Appliance	TRUE	FALSE
262146: MAJOR	Significant change in time of day has occurred, and might compromise statistics. Please contact TAC. **Recommended Action:** [Deprecated alarm] Appliance time changed. Appliance statistics could be missing for an extended interval. Contact Support.	Appliance	FALSE	TRUE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
262149: MAJOR	A disk self-test has been performed. You must reboot the appliance after the test has been completed. **Recommended Action:** Reboot the appliance. Traffic will not be optimized until this is performed.	Appliance	TRUE	FALSE
262154: MAJOR	Software license will expire in 15 days. [For VX series only] **Recommended Action:** Enter new license key to avoid loss of optimization or potential traffic disruption.	Appliance	FALSE	FALSE
262157: MAJOR	Dual wan-next-hop topology is no longer supported. **Recommended Action:** Reconfigure appliance as single bridge with one next hop, or as dual bridge with two IP addresses and two next hops.	Appliance	TRUE	TRUE
262160: MAJOR	Major inconsistency among tunnel traffic class settings found during upgrade. **Recommended Action:** [Deprecated alarm] Review the WAN shaper traffic class settings.	Appliance	TRUE	TRUE
262161: MAJOR	Tunnel IP header disable setting was discarded during upgrade. **Recommended Action:** [Deprecated alarm] Review the optimization map header compression settings.	Appliance	TRUE	TRUE
262163: MAJOR	A peer name has been specified in the configuration matching no existing remote peer. **Recommended Action:** Correct route-map entry or build tunnel. A route policy peer hostname might have changed.	Appliance	TRUE	TRUE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
262168: MAJOR	Cloud Portal WebSocket is down. **Recommended Action:** Appliance cannot connect to the Cloud Portal using HTTPS WebSockets.	Appliance	FALSE	TRUE
262174: MAJOR	Cloud Portal is unreachable for licensing. **Recommended Action:** Appliance cannot connect to the Cloud Portal using HTTPS WebSockets. Verify the connectivity between the appliance and the portal. This connectivity is needed for licensing.	Appliance	TRUE	FALSE
262184: MAJOR	IPv4 subnet table is full. **Recommended Action:** IPv4 subnet table has reached its maximum allowable size. Additional subnets will not be added unless others are removed.	Appliance	TRUE	TRUE
262188: MAJOR	A BGP peer session is not in Established state. **Recommended Action:** A BGP peer session is Down. Verify BGP neighbor, ASN, or next-hop IP address is configured correctly.	Appliance	TRUE	TRUE
262190: MAJOR	An OSPF neighbor session is no longer in Full or Two-Way state. **Recommended Action:** An OSPF neighbor session is Down. Verify whether OSPF neighbor connectivity still exists on this interface.	Appliance	TRUE	TRUE
262193: MAJOR	DHCP server failover my state communications interrupted. **Recommended Action:** DHCP server failover my state communications interrupted. Check for partner reachability and verify your configuration.	Appliance	TRUE	TRUE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
262197: MAJOR	Excessive route advertisement updates detected. **Recommended Action:** Verify proper configuration or route filtering of the route indicated.	Appliance	TRUE	TRUE
262201: MAJOR	EC Feature License not granted. **Recommended Action:** Contact Support.	Appliance	TRUE	FALSE
262205: MAJOR	ACL Groups File Handling Failed. **Recommended Action:** Verify that valid IPs/ports are used for configuration. If the issue persists, contact Support with support logs.	Appliance	FALSE	TRUE
262206: MAJOR	ACL Groups Config Memory Limit Exceeded. **Recommended Action:** To free up memory, reduce the group name lengths used in the Address/Service groups configuration and try again.	Appliance	FALSE	TRUE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
262207: MAJOR	ACL rule has invalid syntax. **Recommended Action:** Check ACL rules syntax and apply again. Syntax rules: Even when using a range or a wildcard, the IPv4 address must be specified in the 4-octet format, separated by the dot notation. For example, A.B.C.D. Range is specified using a single dash. For example, 128-129. Wildcard is specified as an asterisk (*). Range and Wildcard can both be used in the same address, but an octet can only contain one or the other. For example, 10.136-137.*.64-95. A wildcard can only be used to define an entire octet. For example, 10.13*.*.64-95 is not supported. Use 10.130-139.*.64-95 to specify this range. The same rules apply to IPv6 addressing. CIDR notation and (Range or Wildcard) are mutually exclusive in the same address. For example, 192.168.0.1-127/24 is not supported. Use either 192.168.0.0/24 or 192.168.0.1-127. After fixing the syntax, the alarm does not clear automatically. Clear it manually.	Appliance	TRUE	TRUE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
262208: MAJOR	AVC Attributes File Handling Failed. **Recommended Action:** Contact Support.	Appliance	FALSE	TRUE
262214: MAJOR	Maximum allowable OSPF neighbors has been exceeded. **Recommended Action:** OSPF maximum allowable neighbors exceeded. Check how many OSPF neighbors are active on each OSPF interface.	Appliance	TRUE	TRUE
262215: MAJOR	An LACP port is not in sync with its neighbor. **Recommended Action:** Verify its status and whether the aggregator is configured correctly.	Appliance	TRUE	TRUE
262216: MAJOR	An LACP aggregator has no operational ports. **Recommended Action:** Verify whether it is configured correctly and the status of its ports.	Appliance	TRUE	TRUE
262218: MAJOR	One or more DoS thresholds exceeded limit. **Recommended Action:** Check appliance Firewall Protection Profiles status for thresholds state and firewall log for details of DoS events.	Appliance	TRUE	TRUE
262221: MAJOR	End entity certificate is nearing expiry. **Recommended Action:** Renew the end entity certificate that is expiring soon.	Appliance	TRUE	TRUE
262227: MAJOR	IPv6 subnet table is full. **Recommended Action:** IPv6 subnet table has reached its maximum allowable size. Additional subnets will not be added unless others are removed.	Appliance	TRUE	TRUE
262242: MAJOR	Route label is down. **Recommended Action:** Route label and associated routes are down due to IP SLA status.	Appliance	TRUE	TRUE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
262170: MINOR	Performance is limited by max WAN Optimization bandwidth. **Recommended Action:** Recommend subscribing to more WAN Optimization bandwidth.	Appliance	FALSE	TRUE
262183: MINOR	IPv4 subnet table reached High water mark. **Recommended Action:** IPv4 subnet table has reached its maximum level for adding BGP/OSPF-learned routes. Only local subnets added beyond this number.	Appliance	TRUE	TRUE
262194: MINOR	Secure shell challenge-response succeeded. **Recommended Action:** Secure shell authentication succeeded. No action required if authorized personnel are trying to access secure shell.	Appliance	TRUE	TRUE
262195: MINOR	Secure shell challenge-response failed. **Recommended Action:** Secure shell authentication failed. Verify if authorized personnel are trying to access secure shell.	Appliance	TRUE	TRUE
262199: MINOR	DSCP label is unassigned. **Recommended Action:** Label is not assigned to interface.	Appliance	TRUE	TRUE
262200: MINOR	Peer interface admin or oper or nh reachability is down. **Recommended Action:** Peer interface admin or operational or next hop reachability status changed.	Appliance	TRUE	TRUE
262226: MINOR	IPv6 subnet table reached High water mark. **Recommended Action:** IPv6 subnet table has reached its maximum level for adding BGP/OSPF-learned routes. Only local subnets added beyond this number.	Appliance	TRUE	TRUE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
262228: MINOR	DNS Proxy model limit. **Recommended Action:** DNS proxy model limit reached. DNS for express Apps may no longer take the same path as express App traffic.	Appliance	TRUE	TRUE
262246: MINOR	VRRP instance removed since its interface attributes has changed. **Recommended Action:** Attributes of the interface associated with VRRP instance have changed. Reconfigure VRRP.	Appliance	TRUE	TRUE
262150: WARNING	The SSL private key is invalid. **Recommended Action:** The key is not an RSA standard key that meets the minimum requirement of 1024 bits. Regenerate a key that meets this minimum requirement.	Appliance	TRUE	FALSE
262151: WARNING	The SSL certificate is not yet valid. **Recommended Action:** The SSL certificate has a future start date. It will correct itself when the future date becomes current. Otherwise, install a certificate that is current.	Appliance	TRUE	FALSE
262152: WARNING	The SSL certificate has expired. **Recommended Action:** Reinstall a valid SSL certificate that is current.	Appliance	TRUE	FALSE
262153: WARNING	The NTP server is unreachable. **Recommended Action:** Check the appliance's NTP server IP and version configuration. Can appliance reach the NTP server? Is UDP port 123 open between the appliance's mgmt0 IP and the NTP server?	Appliance	TRUE	FALSE
262155: WARNING	Software license will expire in 45 days. [For VX series only] **Recommended Action:** Enter a new license key to avoid loss of optimization or potential traffic disruption.	Appliance	FALSE	TRUE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
262158: WARNING	Setting default system wan-next-hop to VLAN next-hop no longer necessary. **Recommended Action:** [Deprecated alarm] Use the VLAN IP address as tunnel source endpoints instead of bvi0.	Appliance	FALSE	TRUE
262159: WARNING	Minor inconsistency among tunnel traffic class settings found during upgrade. **Recommended Action:** [Deprecated alarm] Review the WAN shaper traffic class settings.	Appliance	FALSE	TRUE
262162: WARNING	A very large range has been configured for a local subnet. **Recommended Action:** Confirm that you intended to configure such a large local subnet (/8 or larger).	Appliance	FALSE	TRUE
262164: WARNING	Interface shaper max bandwidth exceeds system max bandwidth. **Recommended Action:** Review the interface shaper max bandwidth settings. Make sure it does not exceed system max bandwidth.	Appliance	TRUE	TRUE
262167: WARNING	Cloud Portal is unreachable. **Recommended Action:** Appliance cannot connect to the Cloud Portal using HTTPS. This connectivity is needed for internet applications classification.	Appliance	FALSE	TRUE
262169: WARNING	SaaS application is no longer supported. **Recommended Action:** SaaS application is no longer supported.	Appliance	FALSE	TRUE
262181: WARNING	Admin password is not yet changed. **Recommended Action:** Change admin password.	Appliance	FALSE	TRUE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
262186: WARNING	Built-in CA certificate was invalid and it has been deleted internally. **Recommended Action:** Built-in CA Certificate was invalid, and a new one has been automatically generated. Install the built-in CA certificate on clients as needed.	Appliance	FALSE	TRUE
262187: WARNING	CA Bundle was invalid and it has been deleted internally. **Recommended Action:** CA Certificate Bundle is invalid and will be fixed automatically by Cloud Portal in a couple of hours, or contact Support.	Appliance	FALSE	TRUE
262189: WARNING	An IP SLA monitor is in the Down state. **Recommended Action:** An IP SLA monitor has reported Down status. Check and correct the source of the failure.	Appliance	TRUE	TRUE
262191: WARNING	DNS proxy process is in Down state. **Recommended Action:** DNS proxy is in down state.	Appliance	TRUE	TRUE
262192: WARNING	EC Licensing Warning. **Recommended Action:** Check your EC license.	Appliance	FALSE	TRUE
262196: WARNING	An IP SLA monitor is not installed. **Recommended Action:** An IP SLA monitor is not installed. Check and fix the source of the failure.	Appliance	TRUE	TRUE
262198: WARNING	CPU utilization threshold exceeded. **Recommended Action:** CPU utilization reached almost 100%. Ignore if it is intended. Otherwise, take action to reduce CPU utilization.	Appliance	TRUE	TRUE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
262202: WARNING	Stats collection slow or incomplete. **Recommended Action:** In Orchestrator, go to **Orchestrator > Software & Setup > Setup > Stats Collector Configuration** and look for the following issues: 1. Stats collection paused due to low disk space. 2. Stats collection failing because the Stats Collector is unreachable. 3. Too many appliances assigned to a single Stats Collector.	Appliance	FALSE	TRUE
262203: WARNING	Unable to resolve Stats Collector DNS name. **Recommended Action:** Could not resolve Stats Collector DNS name. Check DNS server configuration.	Appliance	FALSE	TRUE
262204: WARNING	Stats Collector is unreachable. **Recommended Action:** Appliance cannot connect to Stats Collector using HTTPS. This connectivity is required for Appliance to upload stats.	Appliance	FALSE	TRUE
262212: WARNING	BFD process is unresponsive and maybe crashed. **Recommended Action:** BFD process is unresponsive and might have crashed. Open a support case.	Appliance	TRUE	TRUE
262217: WARNING	DNS resolution for one of the IPSLA targets failed. **Recommended Action:** Check DNS configuration and/or verify the domain names in IP SLA configuration.	Appliance	TRUE	TRUE
262219: WARNING	The SSL certificate will expire in less than 30 days. **Recommended Action:** Renew a valid SSL certificate before it expires.	Appliance	TRUE	FALSE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
262222: WARNING	Trusted certificate has expired. **Recommended Action:** Replace the expired certificate in the trust store.	Appliance	TRUE	TRUE
262223: WARNING	Trusted certificate is nearing expiry. **Recommended Action:** Replace the trusted certificate that is expiring soon.	Appliance	TRUE	TRUE
262233: WARNING	Authentication delay detected. **Recommended Action:** Check the connectivity to ensure seamless authentication.	Appliance	TRUE	TRUE
262234: WARNING	Authentication failures detected. **Recommended Action:** Check auth-server connectivity/supplicant credentials.	Appliance	TRUE	TRUE
262235: WARNING	Misconfigurations detected. **Recommended Action:** Check the configurations and apply again.	Appliance	TRUE	TRUE
262236: WARNING	RADIUS servers unreachable. **Recommended Action:** Check the server configuration and reachability.	Appliance	TRUE	TRUE
262237: WARNING	Bad authenticator detected. **Recommended Action:** Check the server-side configurations.	Appliance	TRUE	TRUE
262238: WARNING	RADIUS server misconfigurations detected. **Recommended Action:** Check the RADIUS server configurations.	Appliance	TRUE	TRUE
262239: WARNING	No loopback interface for the selected labels detected. **Recommended Action:** Failed to configure AppExpress group because there is no loopback interface for the selected labels.	Appliance	TRUE	TRUE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
262240: WARNING	Auto-update of baseline for threshold is skipped since new baseline value is greater than “Baseline Freeze %” of configured Max Threshold. **Recommended Action:** Maximum threshold needs to be revised.	Appliance	TRUE	TRUE
262241: WARNING	Appliance reserve flow capacity is less than ‘Minimum reserve capacity limit’. So baseline raise is skipped. **Recommended Action:** Consider adjusting the Baseline Settings immediately and plan for a capacity upgrade as flow utilization nears its limit.	Appliance	TRUE	TRUE

Equipment

System Type 0 (Appliance); Source Type 3 (Equipment)

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
196608: CRITICAL	RAID array is degraded. **Recommended Action:** Initiate a self-service RMA for the degraded disk from the Support Portal.	Appliance	FALSE	TRUE
196611: CRITICAL	Fan failure detected. **Recommended Action:** Proceed to the Online RMA tool of the HPE Networking Support Portal to request a hardware replacement.	Appliance	FALSE	FALSE
196612: CRITICAL	System bypass mode. **Recommended Action:** Normal with factory default configuration, during reboot, and if user has put the appliance in bypass mode. Check the system bypass configuration.	Appliance	FALSE	FALSE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
196613: CRITICAL	LAN/WAN fail-to-wire card failure. **Recommended Action:** Proceed to the Online RMA tool of the HPE Networking Support Portal to request a hardware replacement.	Appliance	FALSE	FALSE
196614: CRITICAL	LAN/WAN fail-to-wire card relay failure. **Recommended Action:** Open a support case for assistance.	Appliance	FALSE	FALSE
196615: CRITICAL	Encryption card hardware failure. **Recommended Action:** [Deprecated alarm] Proceed to the Online RMA tool of the HPE Networking Support Portal to request a hardware replacement.	Appliance	FALSE	FALSE
196628: CRITICAL	WAN next-hop router discovered on a LAN port (box is in backwards). **Recommended Action:** Check WAN next hop IP address. Check lan0 and wan0 cabling (in-line mode only). If not resolved, contact Support.	Appliance	TRUE	FALSE
196641: CRITICAL	NIC failure. **Recommended Action:** Proceed to the Online RMA tool of the HPE Networking Support Portal to request a hardware replacement.	Appliance	FALSE	FALSE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
196644: CRITICAL	Insufficient configured memory size for this virtual appliance. [For VX series only] **Recommended Action:** Assign more memory to the virtual machine and restart the appliance. Traffic will not be optimized until this is resolved.	Appliance	TRUE	FALSE
196645: CRITICAL	Insufficient configured processor count for this virtual appliance. [For VX series only] **Recommended Action:** Assign more processors to the virtual machine and restart the appliance. Traffic will not be optimized until this is resolved.	Appliance	TRUE	FALSE
196646: CRITICAL	Insufficient configured disk storage for this virtual appliance. [For VX series only] **Recommended Action:** Assign more storage to the virtual machine and restart the appliance. Traffic will not be optimized until this is resolved.	Appliance	TRUE	FALSE
196649: CRITICAL	Bridge loop is detected. **Recommended Action:** Make sure bridge ports are connected to different virtual switches and restart the appliance. Traffic will not be optimized until this is resolved.	Appliance	TRUE	FALSE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
196650: CRITICAL	Network interface is unassigned. **Recommended Action:** Assign the network interface to an existing MAC address, and then restart the appliance. Or, if the network interface is not being used, set its admin state to down.	Appliance	TRUE	FALSE
196651: CRITICAL	Bridge creation failed. **Recommended Action:** Check log messages for more details on the failure.	Appliance	TRUE	FALSE
196657: CRITICAL	Cell interface is down. **Recommended Action:** Device could be out of range of cellular network. Can try: admin down and admin back up the device. Last resort: check alternative values for APN.	Appliance	TRUE	TRUE
196658: CRITICAL	LTE modem unreachable. **Recommended Action:** LTE modem is unplugged or not visible to system. Suggested action: check LTE modem connection to the system. Unplug and plug LTE modem back in and reboot the system.	Appliance	TRUE	TRUE
196659: CRITICAL	Cellular interface cannot establish a call with the carrier. **Recommended Action:** Make sure APN is set to either default or a carrier-specified APN.	Appliance	TRUE	TRUE
196665: CRITICAL	Equipment Ambient temperature is high. **Recommended Action:** Make sure that the room temperature is normal and wait for the appliance to recover.	Appliance	TRUE	TRUE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
196666: CRITICAL	Temperature at a node of PCB is high. **Recommended Action:** Make sure that the room temperature is normal and wait for it to recover.	Appliance	TRUE	TRUE
196667: CRITICAL	Equipment CPU temperature is high. **Recommended Action:** Make sure that the room temperature is normal and wait for it to recover.	Appliance	TRUE	TRUE
196668: CRITICAL	FAN is running with low duty cycle. **Recommended Action:** Proceed to the Online RMA tool of the HPE Networking Support Portal to request a hardware replacement.	Appliance	TRUE	TRUE
196671: CRITICAL	An uncorrected memory corruption occurred. **Recommended Action:** Proceed to the Online RMA tool of HPE Networking Support Portal to request a hardware replacement.	Appliance	TRUE	TRUE
196609: MAJOR	Disk is failed. **Recommended Action:** Proceed to the Online RMA tool of the HPE Networking Support Portal to request a hardware replacement.	Appliance	FALSE	FALSE
196617: MAJOR	Network interface link down. **Recommended Action:** Is the system in bypass mode? Check cables and interface admin status on the router.	Appliance	TRUE	TRUE
196618: MAJOR	Management interface link down. **Recommended Action:** Check cables and interface admin status on the router.	Appliance	TRUE	TRUE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
196619: MAJOR	Interface is half duplex. **Recommended Action:** Check speed/duplex settings on the router/switch port.	Appliance	TRUE	TRUE
196620: MAJOR	Interface speed is 10 Mbps. **Recommended Action:** Check speed/duplex settings. Use a 100/1000 Mbps port on the router/switch.	Appliance	TRUE	TRUE
196621: MAJOR	Config DB disk full. **Recommended Action:** [Deprecated alarm]	Appliance	TRUE	FALSE
196622: MAJOR	Operating System disk full. **Recommended Action:** [Deprecated alarm]	Appliance	TRUE	FALSE
196623: MAJOR	File System disk full. **Recommended Action:** [Deprecated alarm]	Appliance	TRUE	FALSE
196624: MAJOR	Datapath internal loopback test failed. **Recommended Action:** [Deprecated alarm]	Appliance	TRUE	FALSE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
196625: MAJOR	Next-hop unreachable. **Recommended Action:** Check cables on EdgeConnect appliance and router. Check IP/mask on EdgeConnect appliance and router. Next hop should be only a single IP hop away. To troubleshoot, use: *show cdp neighbor*, *show arp*, and: *ping/ping6 -I <appliance IP> <next-hop IP>* Packets are sent with ttl=1, so ensure that the next hop IP has no intermediate routers. **NOTE:** If there is either a LAN Next-Hop Unreachable or WAN Next-Hop Unreachable alarm, resolve the alarm(s) immediately by configuring the gateway(s) to respond to ICMP pings from the EdgeConnect appliance IP Address.	Appliance	TRUE	FALSE
196626: MAJOR	VRRP instance is down. **Recommended Action:** Check the interface. Is the link down?	Appliance	TRUE	TRUE
196629: MAJOR	Disk is not-in-service. **Recommended Action:** Check whether the disks are properly seated. Contact Support for further assistance.	Appliance	FALSE	FALSE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
196631: MAJOR	Disk has been removed by operator. **Recommended Action:** Normal during disk replacement. Insert the disk using Appliance Manager or Orchestrator. Contact Support if insertion fails.	Appliance	FALSE	FALSE
196632: MAJOR	LAN/WAN interfaces have different admin states. [For Bridge mode only] **Recommended Action:** Check interface admin configuration (lan0/wan0, lan1/wan1). Applicable only to in-line mode.	Appliance	TRUE	TRUE
196633: MAJOR	LAN/WAN interfaces have different link carrier states. [For Bridge mode only] **Recommended Action:** Check interface configured speed settings and current values (lan0/wan0, lan1/wan1). Applicable only to in-line mode.	Appliance	TRUE	TRUE
196634: MAJOR	LAN/WAN interface has been shut down due to link propagation of paired interface. [For Bridge mode only] **Recommended Action:** Check cables and connectivity. For example, if lan0 is shut down, check why wan0 is down. Applicable only to in-line mode.	Appliance	TRUE	TRUE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
196636: MAJOR	Cluster peer unreachable. **Recommended Action:** Check Cluster configuration on all applicable appliances and check L3/L4 connectivity between the peers. Open TCP and UDP ports 4164 between the cluster peer IPs if they are blocked.	Appliance	TRUE	FALSE
196637: MAJOR	Bonding members have different speed/duplex. **Recommended Action:** Check interface speed/duplex settings and negotiated values on wan0/wan1 and lan0/lan1 ether-channel groups.	Appliance	TRUE	TRUE
196638: MAJOR	WCCP adjacency(ies) down. **Recommended Action:** Cannot establish WCCP neighbor. Check WCCP configuration on appliance and router. Verify reachability. Enable debugging on router: `debug ip wccp packet`	Appliance	TRUE	TRUE
196639: MAJOR	WCCP assignment table mismatch. **Recommended Action:** Check the WCCP mask/hash assignment configuration on all EdgeConnect appliances and ensure that they match.	Appliance	TRUE	TRUE
196640: MAJOR	Power supply not connected, not powered, or failed. [EC-M, EC-L, and EC-XL only (dual supplies)] **Recommended Action:** Connect to a power outlet. Check power cable connectivity.	Appliance	FALSE	FALSE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
196642: MAJOR	LAN next-hop unreachable. **Recommended Action:** Check Appliance configuration: LAN side next hop IP, Appliance IP/Mask, VLAN IP/mask, and VLAN ID. **NOTE:** If there is either a LAN Next-Hop Unreachable or WAN Next-Hop Unreachable alarm, resolve the alarm(s) immediately by configuring the gateway(s) to respond to ICMP pings from the EdgeConnect appliance IP Address.	Appliance	TRUE	FALSE
196643: MAJOR	Unexpected system restart. **Recommended Action:** The appliance rebooted unexpectedly. Power issues? Was the appliance shut down ungracefully? Contact Support if the shutdown was not planned.	Appliance	FALSE	TRUE
196647: MAJOR	Interfaces have different MTUs. [For Bridge mode only: lan0/wan0] **Recommended Action:** Check interface MTU settings on lan0/wan0 (pairwise) on dual bridge mode and lan0/lan1/wan0/wan1… on single bridge mode.	Appliance	TRUE	TRUE
196648: MAJOR	Interfaces have different MTUs. [For Bridge mode only: lan1/wan1] **Recommended Action:** Check interface MTU settings on lan1/wan1 or tlan1/twan1 interfaces.	Appliance	TRUE	TRUE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
196652: MAJOR	System optimization disabled. **Recommended Action:** [Deprecated alarm] Turn on system optimization.	Appliance	TRUE	FALSE
196656: MAJOR	HASync peer is down. **Recommended Action:** Check HA link connectivity.	Appliance	TRUE	FALSE
196662: MAJOR	USB disabled due to Over Current. **Recommended Action:** USB port disabled due to overcurrent. Connected device might be pulling too much or in charging device mode.	Appliance	FALSE	FALSE
196664: MAJOR	PoE fault detected. **Recommended Action:** See the Alarm Logs for detected reason.	Appliance	TRUE	TRUE
196610: MINOR	Disk is degraded. **Recommended Action:** Wait for disk to recover. If it does not recover, contact Support.	Appliance	FALSE	FALSE
196630: MINOR	Disk is rebuilding. **Recommended Action:** Normal. Wait for the disk to rebuild. If it does not rebuild, contact Support.	Appliance	FALSE	FALSE
196635: MINOR	Disk SMART threshold exceeded. **Recommended Action:** SSD Smart alert detected. Contact Support.	Appliance	FALSE	TRUE
196653: MINOR	Non-optimal configured memory size for this virtual appliance. [For VX series only] **Recommended Action:** Assign more memory to the virtual machine and restart the appliance. Traffic will be sub-optimal until this is resolved.	Appliance	TRUE	TRUE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
196654: MINOR	Non-optimal configured processor count for this virtual appliance. [For VX series only] **Recommended Action:** Assign more processors to the virtual machine and restart the appliance. Traffic will be sub-optimal until this is resolved.	Appliance	TRUE	TRUE
196655: MINOR	Non-optimal configured disk storage for this virtual appliance. [For VX series only] **Recommended Action:** Assign more storage to the virtual machine and restart the appliance. Traffic will be sub-optimal until this is resolved.	Appliance	TRUE	TRUE
196670: MINOR	A corrected memory corruption occurred. [For VX series only] **Recommended Action:** An occasional few of these are normal. If frequent, contact HPE Networking Support.	Appliance	TRUE	TRUE
196616: WARNING	Network interface admin down. **Recommended Action:** Check the EdgeConnect interface configuration.	Appliance	TRUE	TRUE
196627: WARNING	VRRP state changed from Master to Backup. **Recommended Action:** VRRP state has changed from Master to Backup. Check VRRP Master for uptime and connectivity.	Appliance	TRUE	TRUE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
196660: WARNING	Unsupported USB device. **Recommended Action:** Detected USB device recognized, but not supported on this gateway. Unplug the device and reboot the system.	Appliance	FALSE	FALSE
196661: WARNING	Unrecognized USB device. **Recommended Action:** Detected USB device not recognized by gateway. Unplug the device and reboot the system.	Appliance	FALSE	FALSE
196663: WARNING	Activity detected only on disabled combo port. **Recommended Action:** Check cables or change active combo port setting.	Appliance	TRUE	TRUE
196672: WARNING	Secure Boot is unexpectedly off. **Recommended Action:** Contact HPE Networking Support.	Appliance	TRUE	TRUE

Threshold Crossing Alerts

System Type 0 (Appliance); Source Type 5 (Threshold Crossing Alerts)

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
327681: WARNING	WAN Tx throughput threshold exceeded. **Recommended Action:** User configured. Check bandwidth reports for tunnel bandwidth.	Appliance	FALSE	TRUE
327682: WARNING	LAN Rx throughput threshold exceeded. [LAN Rx outbound] **Recommended Action:** User configured. Check bandwidth reports.	Appliance	FALSE	TRUE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
327683: WARNING	Optimized flows count threshold exceeded. **Recommended Action:** User configured. Check flow and real-time connection reports.	Appliance	FALSE	TRUE
327684: WARNING	Total flows count threshold exceeded. **Recommended Action:** User configured. Check flow and real-time connection reports.	Appliance	FALSE	TRUE
327685: WARNING	File system utilization threshold exceeded. **Recommended Action:** Disk is almost full. In Appliance Manager, go to **Support > Debug Files**, and then delete the old tcpdumps, snapshots, sysdumps, and show-tech files.	Appliance	FALSE	TRUE
327686: WARNING	Latency threshold exceeded. **Recommended Action:** User Configured. Check Latency reports. If latency is too high, check routing between the appliances and QoS policy on upstream routers. Check tunnel DSCP marking. If latency persists, contact Internet Service Provider (ISP) and Support.	Appliance	FALSE	TRUE
327687: WARNING	Pre-FEC loss threshold exceeded. **Recommended Action:** User configured. Check Loss reports. Check for loss between EdgeConnect appliances (interface counters on upstream routers). Use a network bandwidth measurement tool, such as iPerf, to measure loss. Contact Internet Service Provider (ISP).	Appliance	FALSE	TRUE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
327688: WARNING	Post-FEC loss threshold exceeded. **Recommended Action:** User configured. Check Loss reports. Check for loss between EdgeConnect appliances (interface counters on upstream routers). Use a network bandwidth measurement tool, such as iPerf, to measure loss. Enable/Adjust Forward Error Correction (FEC). Contact ISP (Internet Service Provider).	Appliance	FALSE	TRUE
327689: WARNING	Out of order packets threshold exceeded. **Recommended Action:** User configured. Check Out-of-Order Packets Reports. Normal in a network with multiple paths and different QoS queues. Normal in a dual-homed router or four port in-line configuration. Contact Support if out-of-order packets are not 100% corrected.	Appliance	FALSE	TRUE
327690: WARNING	Corrected out of order packets threshold exceeded. **Recommended Action:** User configured. Check Out-of-Order Packets Reports. Normal in a network with multiple paths and different QoS queues. Normal in a dual-homed router or four port in-line configuration. Contact Support if out-of-order packets are not 100% corrected.	Appliance	FALSE	TRUE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
327691: WARNING	Bandwidth utilization threshold exceeded. **Recommended Action:** User configured. Check bandwidth reports for tunnel bandwidth utilization.	Appliance	FALSE	TRUE
327692: WARNING	Low reduction threshold exceeded. **Recommended Action:** User configured. Check bandwidth reports for dedupe. Check if the traffic is pre-compressed or encrypted.	Appliance	FALSE	TRUE
327693: WARNING	Appliance flow limit threshold exceeded. **Recommended Action:** If this condition persists, a larger appliance will be necessary to fully optimize all flows.	Appliance	FALSE	TRUE

Orchestrator Alarms

Orchestrator can raise alarms based on issues with tunnels, software, and equipment.

Tunnels

System Type 100 (Orchestrator); Source Type 1 (Tunnel)

Alarm ID: Severity	Alarm Text	Source	Service Affect-ing	Clearable
6619136: CRITICAL	Interfaces with duplicate IP exists: {0}. **Recommended Action:** No overlays will be applied to appliances with duplicate IP address.	/orchestrator/interfaces	TRUE	FALSE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
6619137: CRITICAL	Interfaces with no public IP exists: {0}. **Recommended Action:** NAT flag is turned on for this interface. Refer to the Deployment page. If this interface is not behind NAT, remove the NAT flag. No tunnels will be built to or from the interface with missing public IP.	/orchestrator/interfaces	TRUE	FALSE
6619138: CRITICAL	Failed to apply overlays. **Recommended Action:** One or more steps failed while applying overlays. Refer to the Audit Logs for more details.	/orchestrator/orchestration/overlays	TRUE	FALSE
6619139: CRITICAL	ACL used in an overlay is not defined on the appliance. acl name: {0} overlay name: {1}. **Recommended Action:** ACLs can be created on the appliance by applying ACL templates.	/orchestrator/orchestration/overlays	TRUE	FALSE
6619140: CRITICAL	Interfaces with duplicate wan/vti label exists: {0}. **Recommended Action:** Assign unique labels to all WAN/VTI interfaces. No overlays will be applied to appliances with duplicate WAN/VTI labels.	/orchestrator/interfaces	TRUE	FALSE
6619141: CRITICAL	Interfaces with duplicate public IP exists: {0}. **Recommended Action:** No overlays will be applied to appliances with duplicate IP addresses.	/orchestrator/interfaces	TRUE	FALSE
6619142: CRITICAL	Failed to apply tunnel group. **Recommended Action:** Refer to the Audit Logs for more details.	/orchestrator/orchestration/tunnelgroups	TRUE	FALSE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
6619143: CRITICAL	Interface has bad IP address: {0}. **Recommended Action:** No overlay tunnels will be built using this interface.	/orchestrator/interfaces	TRUE	FALSE
6619146: CRITICAL	Cannot build tunnel with src IP {0} and dest IP {1}. IP versions mismatch. **Recommended Action:** Make sure the tunnel source and destination IP address are both IPv4 or are both IPv6 addresses.	/orchestrator/orchestration/tunnels	TRUE	TRUE
6619147: CRITICAL	Failed to apply labels. **Recommended Action:** Refer to the Audit Logs for more details.	/orchestrator/orchestration	TRUE	FALSE
6619148: CRITICAL	Failed to apply internal subnets. **Recommended Action:** Refer to the Audit Logs for more details.	/orchestrator/orchestration	TRUE	FALSE
6619149: CRITICAL	Failed to apply application classification data to appliance. **Recommended Action:** Make sure the appliance can connect to Orchestrator. Refer to the Audit Logs for more details.	/orchestrator/orchestration/applications	TRUE	FALSE
6619150: CRITICAL	Appliance has the same IPSec UDP port as the other HA peer, overlays will not be applied to this appliance. HA Peer: {0}. **Recommended Action:** You can change the IPSec UDP Port of an appliance by editing the System Information from the System Information tab.	/orchestrator/orchestration/tunnels/ha	TRUE	FALSE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
6619151: CRITICAL	Both Overlay Manager and Tunnel Group manager are ENABLED. **Recommended Action:** You can enable one or the other. Turn one of them OFF.	/orchestration	TRUE	FALSE
6619152: CRITICAL	Overlay {0} has no hub defined. No tunnels will be built between appliances that are part of this overlay. **Recommended Action:** To add a Hub to an Overlay, either (1) apply the Overlay to a Hub appliance or (2) go to the Hubs tab and make an appliance that is currently in the Overlay a Hub.	/orchestration/overlays	TRUE	FALSE
6619153: CRITICAL	Overlay {0} has no WAN ports defined. No tunnels will be built between appliance that are part of this overlay. **Recommended Action:** At least one WAN port needs to be defined in an overlay.	/orchestration/overlays	TRUE	FALSE
6619154: CRITICAL	Tunnel Group {0} has no hub defined. No tunnels will be built between appliances that are part of this tunnel group. **Recommended Action:** To add a Hub to a Tunnel Group, either (1) apply the Tunnel Group to a Hub appliance or (2) go to the Hubs tab and make an appliance that is currently in the Tunnel Group a Hub.	/orchestration/tunnelgroups	TRUE	FALSE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
6619155: CRITICAL	Tunnel Group {0} has no interfaces defined. No tunnels will be built between appliances that are part of this tunnel group. **Recommended Action:** Go to Tunnel Groups to configure interfaces.	/orchestration/tunnelgroups	TRUE	FALSE
6619156: CRITICAL	Failed to apply templates. **Recommended Action:** One or more templates failed to apply. Refer to the Audit Logs for more details.	/orchestrator/orchestration/templates	TRUE	FALSE
6619157: CRITICAL	Failed to apply port forwarding rules to appliance. **Recommended Action:** Make sure appliance can connect to Orchestrator. Refer to the Audit Logs for more details.	/orchestrator/orchestration	TRUE	FALSE
6619158: CRITICAL	Overlay {0} is using local breakout without any interfaces selected. **Recommended Action:** Go to Business Intent Overlays to configure local breakout interface.	/orchestration/overlays	TRUE	FALSE
6619159: CRITICAL	Maximum number of tunnels exceeded. {0}. **Recommended Action:** Configure Overlays to create fewer tunnels.	/orchestrator/orchestration/tunnels	FALSE	FALSE
6619160: CRITICAL	At least one region is missing a hub appliance. **Recommended Action:** To add a Hub to an Overlay, either (1) apply the Overlay to a Hub appliance or (2) go to the Hubs tab and make an appliance that is currently in the Overlay a Hub.	/orchestration/overlays	TRUE	FALSE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
6619161: CRITICAL	Appliance has a duplicate hostname with another appliance in the network. No overlays will be built to this appliance. **Recommended Action:** Change the hostname of one of the appliances.	/orchestrator/orchestration/tunnels	FALSE	FALSE
6619166: CRITICAL	Orchestration failed. {0}. **Recommended Action:** Go to the Audit Logs for more details.	/orchestrator/orchestration	TRUE	FALSE
6619168: CRITICAL	Duplicate IPSec UDP Port {0} detected on the following appliances [{1}] that belong to site {2}. Appliances sharing the same site name must have unique IPSec UDP port numbers. Orchestration for these appliances will be suspended until this is addressed. **Recommended Action:** You can change the IPSec UDP Port of an appliance on the System Information tab.	/orchestrator/orchestration/tunnels	TRUE	FALSE
6619173: CRITICAL	Failed to apply traffic behavior data to appliance. **Recommended Action:** Make sure appliance can connect to Orchestrator. Refer to the Audit Logs for more details.	/orchestrator/orchestration	TRUE	FALSE
6619174: CRITICAL	Appliance does not have geo location information. Zscaler Service Edges cannot be auto discovered. **Recommended Action:** Update appliance location in Configuration Wizard.	/orchestrator/orchestration/zscaler	TRUE	FALSE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
6619175: CRITICAL	Only IPSec UDP tunnel mode is supported on EdgeHA devices. **Recommended Action:** Check Tunnel Settings.	/orchestration	TRUE	TRUE
6619177: CRITICAL	Edge HA peer {0} has tunnels with source port {1}. Orchestration will be skipped for this appliance until the conflicting tunnels are deleted. **Recommended Action:** Wait for the conflicting tunnels to tear down.	/orchestrator/orchestration/tunnels	TRUE	FALSE
6619178: CRITICAL	This appliance and appliances [{0}] have been manually configured with same IPSec UDP port {1}. Orchestration will be paused on all the conflicting appliances until unique ports are assigned to the appliance. **Recommended Action:** You can change the IPSec UDP Port of an appliance on the System Information tab.	/orchestrator/orchestration/tunnels	TRUE	FALSE
6619189: CRITICAL	{0} exceeded threshold {1} by {2}% ({3}) at {4}. **Recommended Action:** Check internal LAN.	/orchestrator/internaldrops	FALSE	TRUE
6619144: MINOR	Appliance does not have any wan labels required for an overlay. No tunnels will be built on this appliance for the overlay. Overlay name: {0} wan labels: {1}. **Recommended Action:** Assign at least one WAN label selected for this overlay in the deployment configuration of the appliance.	/orchestrator/orchestration/overlays	TRUE	TRUE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
6619145: MINOR	Appliance missing lan label {1} for traffic access policy of overlay. No traffic on this appliance will be routed to the overlay. Overlay name: {0}. **Recommended Action:** Assign a LAN port the required LAN label selected for this overlay in the deployment configuration of the appliance. If this appliance is in server mode, you should use an ACL instead of selecting a LAN label in the overlay configuration.	/orchestrator/orchestration/overlays	TRUE	TRUE
6619163: WARNING	Mesh Overlay - {0} has no hub defined. No tunnels will be built for Hub & Spoke Interface label {1} for this overlay. **Recommended Action:** To add a Hub to an Overlay, either (1) apply the Overlay to a Hub appliance or (2) go to the Hubs tab and make an appliance that is currently in the Overlay a Hub.	/orchestration/overlays	TRUE	FALSE
6619164: WARNING	Mesh Tunnel Group - {0} has no hub defined. No tunnels will be built for Hub & Spoke Interface label {1} for this group. **Recommended Action:** To add a Hub to a Tunnel Group, either (1) apply the Tunnel Group to a Hub appliance or (2) go to the Hubs tab and make an appliance that is currently in the Tunnel Group a Hub.	/orchestration/tunnelgroups	TRUE	FALSE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
6619165: WARNING	Appliance does not have public IP for wan label {0} auto discovered Zscaler Service Edges may not be correct. **Recommended Action:** On the Deployment page, toggle the Not Behind NAT flag to NAT, or create a Service Edge Override on the Zscaler tab.	/orchestrator/orchestration zscaler	FALSE	FALSE
6619167: WARNING	At least one hub is not part of any overlay. **Recommended Action:** To add a Hub to an Overlay, either (1) apply the Overlay to a Hub appliance or (2) go to the Hubs tab and make an appliance that is currently in the Overlay a Hub.	/orchestration/overlays	TRUE	FALSE
6619169: WARNING	At least one appliance is associated with a region and regional routing is currently disabled. **Recommended Action:** If regional routing is desired and has been authorized, go to Regional Routing and enable the feature.	/orchestration	TRUE	TRUE
6619170: WARNING	{0} is only supported on inline router deployment mode appliances. **Recommended Action:** Go to Deployment and choose “Router” mode.	/orchestrator	FALSE	TRUE
6619171: WARNING	Regional routing is enabled but {0} not associated with any region, so no tunnel will be built. **Recommended Action:** Associate the appliance with a region or disable regional routing.	/orchestration	TRUE	TRUE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
6619176: WARNING	{0} Appliance {1} configured Bandwidth for {2} interface is below min threshold for the number of configured underlay tunnels. **Recommended Action:** On Deployment page, update Inbound/outbound interface Bandwidth to be above the minimum Bandwidth. (Check the “interfaceBandwidth CheckPerTunnelOverhead” property under Advanced Properties.)	/orchestrator/ interfaceBandwidt	TRUE	FALSE

Software

System Type 100 (Orchestrator); Source Type 4 (Software)

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
6815744: CRITICAL	Orchestrator detected possible cloned appliances - cloned: {0} clone: {1}. **Recommended Action:** [Deprecated alarm]	/orchestrator/ discovery/ clone	FALSE	TRUE
6815745: CRITICAL	Appliance backup failed: {0}. **Recommended Action:** Check Orchestrator reachability, and then go to Backup Now and retry the backup. Open a support case if this alarm persists more than 24 hours.	/orchestrator/ system/backup	FALSE	TRUE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
6815746: CRITICAL	Appliances with the same serial numbers exist: {0}. **Recommended Action:** This alarm might be caused by applying identical license keys on multiple appliances. If the appliances are cloned, contact Support for the steps to correctly clone the appliance.	/orchestrator/system	TRUE	TRUE
6815748: CRITICAL	Orchestrator cannot reach this appliance. **Recommended Action:** Check appliance reachability to Orchestrator and HPE Aruba Networking EdgeConnect Cloud Portal.	/orchestrator/connectivity	TRUE	TRUE
6815749: CRITICAL	Appliance version not supported: {0}. **Recommended Action:** Upgrade appliance to minimum supported version.	/orchestrator/system	TRUE	TRUE
6815752: CRITICAL	Appliance is configured with labels to build IPSec UDP tunnels, but the appliance version does not support IPSec UDP tunnels. **Recommended Action:** You can change the tunnel modes for labels in the Overlay Manager Settings.	/orchestrator/orchestration/tunnels	TRUE	FALSE
6815754: CRITICAL	Orchestrator requires a validated portal account name and key (or for NX/VX/VRX products, a valid license key). **Recommended Action:** Go to Licensing to provide the required information.	/license	TRUE	FALSE
6815757: CRITICAL	Orchestrator portal account or license expired on {0 date}. **Recommended Action:** Go to Licensing to provide the required information.	/license	TRUE	FALSE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
6815758: CRITICAL	Orchestrator cannot connect to HPE Aruba Networking EdgeConnect Cloud Portal using HTTPS. **Recommended Action:** Check portal connection and refer to the Audit Logs for more information.	/portal connectivity	TRUE	TRUE
6815760: CRITICAL	Orchestrator cannot register with HPE Aruba Networking EdgeConnect Cloud Portal using the credentials provided. **Recommended Action:** Go to the Licensing tab and provide the required account information.	/portal/ registration	TRUE	TRUE
6815766: CRITICAL	CPX license expired on {0 date}. **Recommended Action:** [Deprecated alarm] Renew your license to avoid service interruption.	/portal/license/ cpx	TRUE	TRUE
6815770: CRITICAL	Your EdgeConnect account has licenses that expired on {0 date}. EdgeConnect devices in your network will stop passing traffic. **Recommended Action:** Renew your license to avoid service interruption.	/portal/license/ ec	TRUE	TRUE
6815774: CRITICAL	SaaS license expired on {0 date}. **Recommended Action:** Renew your license to avoid service interruption.	/portal/license/ saas	TRUE	TRUE
6815778: CRITICAL	Discovered appliances list contains cloned appliances. Clones: {0}. **Recommended Action:** Contact Support for steps to correctly clone the appliance.	/discovery/ clone	FALSE	TRUE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
6815779: CRITICAL	Orchestrator backup failed. **Recommended Action:** Go to Historical Jobs for details.	/system/backup	FALSE	TRUE
6815780: CRITICAL	Orchestrator failed to get update from portal for application definition data. **Recommended Action:** Check portal connection and refer to the Audit Logs for more information.	/orchestration/applications	FALSE	TRUE
6815782: CRITICAL	Orchestrator failed to get update from portal for traffic behavior data. **Recommended Action:** Check portal connection and refer to the Audit Logs for more information.	/orchestration/applications	FALSE	TRUE
6815790: CRITICAL	Orchestrator is not registered with HPE Aruba Networking EdgeConnect Cloud Portal. **Recommended Action:** Use your previous Orchestrator to approve this one. If you do not have another Orchestrator, contact Support for assistance.	/portal/registration	TRUE	TRUE
6815791: CRITICAL	Failed to connect to Zscaler. **Recommended Action:** Check Zscaler subscription.	/orchestration	TRUE	TRUE
6815792: CRITICAL	Your Orchestrator service will expire on {0 date}. **Recommended Action:** Contact the HPE Aruba Networking Sales team to order an extension.	/portal/license/cloudorch	FALSE	TRUE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
6815796: CRITICAL	Your WAN Optimization expired on {0 date}. EdgeConnect devices in your network will stop using WAN Optimization. **Recommended Action:** Renew your WAN Optimization license to avoid service interruption.	/portal/license/ec	TRUE	TRUE
6815800: CRITICAL	Cannot get Azure data. Details : {0}. **Recommended Action:** Check Azure subscription. Go to the Audit Logs for more details.	/orchestration/azure	TRUE	TRUE
6815801: CRITICAL	Cannot download Azure configuration. Details : {0}. **Recommended Action:** Check Azure subscription.	/orchestration/azure	TRUE	TRUE
6815802: CRITICAL	Cannot create IPSEC Tunnels for Azure VPN Site - {0} for appliance {1} and label {2}. **Recommended Action:** Associate Hub to Azure VPN Site using the Azure Portal. If Hub is already associated, wait for deployment to complete to start Azure Orchestration.	/orchestrator/orchestration/azure	TRUE	TRUE
6815803: CRITICAL	Cannot Orchestrate association to Azure VWan. **Recommended Action:** Use the VTI IP Pool dialog box to configure the subnet pool.	/orchestration/azure	TRUE	TRUE
6815804: CRITICAL	Cannot connect to Azure. Details : {0}. **Recommended Action:** Check Azure subscription. Go to the Audit Logs for more details.	/orchestration/azure	TRUE	TRUE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
6815808: CRITICAL	Appliance was manually added to Orchestrator. **Recommended Action:** Remove the appliance from Orchestrator, discover and approve it.	/orchestrator/system	TRUE	TRUE
6815815: CRITICAL	Custom bonding policy or secondary WAN interface configured in overlay, but appliance does not support this feature. **Recommended Action:** Check overlay configuration or upgrade appliance.	/orchestrator/orchestration/overlays	TRUE	FALSE
6815817: CRITICAL	Invalid IPSec UDP Key Material lifetime. Lifetime must be greater than rotation period. **Recommended Action:** Change IPSec UDP Key Material lifetime.	/ikeless	TRUE	FALSE
6815820: CRITICAL	Cannot connect to AWS. Details : {0}. **Recommended Action:** Check AWS subscription. Go to the Audit Logs for more details.	/orchestration/awstgnm	TRUE	TRUE
6815822: CRITICAL	Cannot Orchestrate association to AWS Transit Gateway. **Recommended Action:** Use AWS VTI Subnet Pool Dialog to configure the subnet pool.	/orchestration/awstgnm	TRUE	TRUE
6815825: CRITICAL	Azure VWAN has duplicate ASN in the network. Details : {0}. **Recommended Action:** Use Azure Portal to assign unique ASNs to the VPN Sites.	/orchestration/azure	TRUE	TRUE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
6815826: CRITICAL	Orchestrator cannot register with HPE Aruba Networking EdgeConnect Cloud Portal. **Recommended Action:** Go to the HPE Aruba Networking EdgeConnect Cloud Portal and provide the portal host information. If the issue is not resolved, contact Support.	/portal/ registration	TRUE	TRUE
6815827: CRITICAL	Routing Segmentation is only supported in inline-router mode. **Recommended Action:** Check the deployment mode.	/orchestrator/ routing Segmentation	TRUE	TRUE
6815828: CRITICAL	Cannot Orchestrate association to AWS Transit Gateway. No primary interface configured for AWS TGNM Integration. **Recommended Action:** Configure the primary interfaces using Interface Label dialog in AWS Network Manager tab.	/orchestration/ awstgnm	TRUE	TRUE
6815829: CRITICAL	Cannot Orchestrate association to Azure VWAN. {0}. **Recommended Action:** Configure the interfaces using Interface Label dialog in Microsoft Azure Virtual WAN tab to proceed with the integration.	/orchestration/ azure	TRUE	TRUE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
6815830: CRITICAL	Cannot create AWS Customer Gateway for Appliance {0} and label {1}. Reason - No valid interface public ip address found. **Recommended Action:** Ensure interface has public IP address. Refer to deployment page. No AWS Customer Gateway will be created with missing public IP.	/orchestrator/orchestration/awstgnm	TRUE	TRUE
6815838: CRITICAL	New IPSec UDP Key Material will be activated at {0}. **Recommended Action:** Ensure the appliance is reachable from Orchestrator.	/orchestrator/ikeless	TRUE	TRUE
6815839: CRITICAL	Orchestrator is unable to rotate IPSec UDP key material. **Recommended Action:** Ensure the appliance is reachable from Orchestrator.	/orchestrator/ikeless	TRUE	TRUE
6815841: CRITICAL	Unable to activate IPSec UDP Key Material. IPSec UDP tunnels will go down without activation. **Recommended Action:** Go to Audit Logs for more details.	/orchestrator/ikeless	TRUE	TRUE
6815842: CRITICAL	Orchestration of the appliance is in-progress for more than 24 hrs. **Recommended Action:** Reboot Orchestrator.	/orchestrator/orchestration/hung	TRUE	TRUE
6815843: CRITICAL	Orchestrator needs to restart in order to reflect changes to custom cert settings. **Recommended Action:** Restart Orchestrator for new custom cert settings to take effect.	/orchestrator/customCerts	TRUE	FALSE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
6815846: CRITICAL	Cannot Orchestrate association of Remote Endpoints for {0} service. No primary/backup interfaces configured. **Recommended Action:** Configure the primary/backup interfaces using Interface Label dialog in Service Orchestration tab.	/orchestration/service Orchestration	TRUE	TRUE
6815850: CRITICAL	Unable to establish connection with stats collector. **Recommended Action:** Unable to establish connection with stats collector. Check the status of Stats collector. It may not be running.	/connectivity/statsCollector	TRUE	TRUE
6815856: CRITICAL	Cannot connect to HPE ANW Central. Details : {0}. **Recommended Action:** Check HPE ANW Central subscription. Go to the Audit Logs for more details.	HPEAruba Networking Central	TRUE	TRUE
6815857: CRITICAL	Orchestrator cannot register with the HPE Aruba Networking EdgeConnect Cloud Portal due to a network issue. **Recommended Action:** Check the HPE Aruba Networking EdgeConnect Cloud Portal connection and view the Audit Logs for more information.	/portal/connectivity	TRUE	TRUE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
6815860: CRITICAL	Cannot create Zscaler gre tunnels for Appliance {0} and label {1}. Reason - No valid interface public ip address found. **Recommended Action:** Ensure the interface has public IP address. Refer to the Deployment page. No Zscaler GRE tunnels will be created with missing public IP.	/orchestrator/orchestration/zscaler	TRUE	TRUE
6815861: CRITICAL	Unable to orchestrate Zscaler sub-location. The sum of all the sub-location bandwidths cannot be greater than location bandwidth. **Recommended Action:** Refer to the Audit Logs for more details : {0}. Reduce the sub-location bandwidth to be less than the location bandwidth.	/orchestrator/orchestration/zscaler	TRUE	TRUE
6815862: CRITICAL	Orchestrator cannot register secondary accounts with HPE Aruba Networking EdgeConnect Cloud Portal using the credentials provided. **Recommended Action:** Check secondary account information by going to HPE Aruba Networking EdgeConnect Cloud Portal.	/portal/registration	TRUE	TRUE
6815863: CRITICAL	Secondary account(s) are not registered with HPE Aruba Networking EdgeConnect Cloud Portal. Details: {0}. **Recommended Action:** Check secondary account information by going to HPE Aruba Networking EdgeConnect Cloud Portal.	/portal/registration	TRUE	TRUE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
6815867: CRITICAL	Your EdgeConnect account {0} has licenses that expired on {1 date}. EdgeConnect devices in your network that are licensed with this account will stop passing traffic {2 date}. **Recommended Action:** Renew your license to avoid service interruption.	/portal/license/ec	TRUE	TRUE
6815871: CRITICAL	Your WAN Optimization for account {0} expired on {1 date}. EdgeConnect devices in your network that are licensed with this account will stop using WAN Optimization on {2 date}. **Recommended Action:** Renew your WAN Optimization license to avoid service interruption.	/portal/license/ec	TRUE	TRUE
6815873: CRITICAL	Failed to connect to Netskope. **Recommended Action:** Check Netskope subscription.	/orchestration	TRUE	TRUE
6815876: CRITICAL	Cannot create Netskope gre tunnels for Appliance {0} and label {1}. Reason - No valid interface public ip address found. **Recommended Action:** Ensure interface has public IP address. Refer to the Deployment page. No Netskope GRE tunnels will be created with missing public IP.	/orchestrator/orchestration/netskope	TRUE	TRUE
6815877: CRITICAL	Appliance does not have geo location information. Netskope POPs cannot be auto discovered. No third-party tunnels will be built on this appliance. **Recommended Action:** Update appliance location in Configuration Wizard.	/orchestrator/orchestration/netskope	TRUE	FALSE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
6815886: CRITICAL	Primary orchestrator is down. **Recommended Action:** [Deprecated alarm]	/orchestrator-HA	TRUE	TRUE
6815891: CRITICAL	Your AAS license expired on {0 date}. **Recommended Action:** Renew your license. Go to Audit Logs for more details.	portal/license/orchestrator	TRUE	TRUE
6815899: CRITICAL	Failed to connect to HPE SSE Service. {0}. **Recommended Action:** Check HPE SSE subscription token key and expiration.	/orchestration/HPE SSE	TRUE	TRUE
6815914: CRITICAL	Unable to connect to telemetry database. {0}. **Recommended Action:** Check telemetry database connection. Check Layer 3 and 4 (TCP port 8123) connectivity and authentication.	/orchestration/clickhouse	TRUE	TRUE
6815918: CRITICAL	Cannot connect to Azure LAN. Details : {0}. **Recommended Action:** Check Azure LAN subscription. Go to Audit Logs for more details.	/orchestration/azure	TRUE	TRUE
6815747: MAJOR	Appliance time is off from that of Orchestrator: {0}. **Recommended Action:** Check NTP settings for server consistency between appliances and Orchestrator.	/orchestrator/system/time	FALSE	TRUE
6815750: MAJOR	Appliance is needed to reboot. **Recommended Action:** You can reboot the appliance under Appliance Reboot / Shutdown.	/orchestrator/system	FALSE	TRUE
6815751: MAJOR	Appliance configuration changes have not been saved. **Recommended Action:** [Deprecated alarm]	/orchestrator/system	FALSE	TRUE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
6815756: MAJOR	Orchestrator portal account or license will expire on {0 date}. **Recommended Action:** Go to Licensing to provide the required information.	/license	FALSE	TRUE
6815761: MAJOR	Orchestrator does not have a set email address for alarm delivery. **Recommended Action:** Go to Alarms to configure email recipient(s).	/email/alarm	FALSE	FALSE
6815762: MAJOR	Orchestrator is using the default SMTP settings. **Recommended Action:** Go to SMTP Server Settings to configure SMTP server.	/email/smtp	FALSE	FALSE
6815776: MAJOR	Orchestrator SMTP configuration is not set. If you do not set Orchestrator SMTP configuration, Orchestrator cannot notify users about alarms, send reports, or send password recovery emails. **Recommended Action:** Go to SMTP Server Settings to configure SMTP server.	/email/smtp	FALSE	FALSE
6815777: MAJOR	Failed to deliver an email. **Recommended Action:** Check SMTP Server Settings.	/email/smtp	FALSE	FALSE
6815784: MAJOR	HPE Aruba Networking diagnostic remote access has been enabled from {0} to {1}. **Recommended Action:** You can disable this in the Remote Access Settings.	/system/support	FALSE	FALSE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
6815787: MAJOR	Failed to apply appliance preconfiguration. **Recommended Action:** Applying preconfiguration to an appliance failed. Refer to the Preconfiguration tab and the Audit Logs for more details.	/orchestrator/preconfiguration	FALSE	TRUE
6815788: MAJOR	Changes done on the appliance will not be auto saved. Enable Auto Save in Orchestration Settings. **Recommended Action:** Enable Auto Save in Orchestration Settings.	/orchestration	TRUE	FALSE
6815807: MAJOR	Duplicate ASNs found in the network for appliances - {0} with asn - {1}. **Recommended Action:** Assign unique ASNs for appliances using the Orchestrator BGP menu.	/orchestrator/orchestration/bgp	TRUE	TRUE
6815809: MAJOR	Invalid ASN found in the network for appliance - {0} with asn - {1}. **Recommended Action:** Assign unique ASN for the appliance using Orchestrator BGP menu.	/orchestrator/orchestration/bgp	TRUE	TRUE
6815831: MAJOR	Invalid ASN found in the network for appliance - {0} with asn - {1}. Amazon EC2 supports all 2-byte ASN numbers in the range of 1 - 65534, with the exception of 7224, which is reserved in the us-east-1 Region, and 9059, which is reserved in the eu-west-1 Region. **Recommended Action:** Assign unique ASN for the appliance using Orchestrator BGP menu.	/orchestrator/orchestration/aws_tgnm	TRUE	TRUE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
6815837: MAJOR	Unable to assign ASN for the appliance. All ASNs from ASN Range are reserved. **Recommended Action:** Use BGP ASN Global Pool dialog to increase the scope of ASN Range.	/orchestrator/orchestration/bgp	FALSE	TRUE
6815844: MAJOR	Source Interface is not configured in the IP SLA configuration on Zscaler Internet Access. **Recommended Action:** Configure the Source Interface for the Zscaler IP SLA configuration.	/orchestration/zscaler/ipsla	TRUE	TRUE
6815847: MAJOR	Failed to initialize connection with ClearPass Policy Manager. {0}. **Recommended Action:** Check ClearPass Policy Manager server parameters.	/orchestration/clearPass	TRUE	TRUE
6815848: MAJOR	Failed to connect with ClearPass Policy Manager service endpoints. {0}. **Recommended Action:** Check the Audit Logs for details.	/orchestration/clearPass	TRUE	TRUE
6815859: MAJOR	This appliance does not support Stats Collection. **Recommended Action:** Upgrade the appliance.	/orchestrator/statsCollector	TRUE	TRUE
6815878: MAJOR	Source Interface is not configured in the IP SLA configuration on Netskope. **Recommended Action:** Configure the Source Interface for the Netskope IP SLA configuration.	/orchestration/netskope/ipsla	TRUE	TRUE
6815887: MAJOR	Backup orchestrator is down. **Recommended Action:** Check backup orchestrator status.	/orchestrator-HA	TRUE	TRUE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
6815892: MAJOR	Out of compliance AAS licensing: {0}. **Recommended Action:** Go to the Audit Logs for more details.	/portal/ license/ orchestrator	FALSE	TRUE
6815898: MAJOR	Appliance does not currently support Air-Gap mode. **Recommended Action:** Upgrade the appliance to ECOS 9.4.2.0 or above.	/orchestrator/ airGap	TRUE	FALSE
6815903: MAJOR	Missing AWS vpc and route table association for the appliance. Unable to automate AWS lan integration. **Recommended Action:** In Orchestrator, go to **Orchestrator > Cloud Services > AWS Network Manager**, click **AWS Resources**, and then configure AWS VPC and AWS Route Table for the appliance.	/orchestrator/ orchestration/ aws_tgnm	FALSE	TRUE
6815904: MAJOR	Missing AWS transit gateway cidr block for [{0}]. Unable to automate AWS lan integration for appliance {1}. **Recommended Action:** Use the AWS VPC console to configure the transit gateway CIDR block and force refresh in Orchestrator - AWS Network Manager Subscription UI.	/orchestrator/ orchestration/ aws_tgnm	FALSE	TRUE
6815905: MAJOR	Missing edge connect segment association for label {0}. Unable to automate AWS lan integration for appliance {1}. **Recommended Action:** In Orchestrator, go to **Orchestrator > Cloud Services > AWS Network Manager**, click **Tunnel Settings**, and then associate segments to LAN labels.	/orchestration/ aws_tgnm	TRUE	TRUE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
6815906: MAJOR	Missing edge connect segment{0} association to aws tgw route table. Unable to automate AWS lan integration for appliance {1}. **Recommended Action:** In Orchestrator, go to **Orchestrator > Cloud Services > AWS Network Manager**, click **Segment & Zone/Role Association**, and then associate EdgeConnect segments to AWS transit gateway route tables.	/orchestration/aws_tgnm	TRUE	TRUE
6815907: MAJOR	Unable to orchestrate appliance[{0}] associated to aws transit gateway[{1}] due to aws region mismatch [Error Code: InvalidTransitGatewayID.NotFound].Details - edge connect region [{2}] and transit gateway region [{3}]. **Recommended Action:** Associate appliance to transit gateway within the same region as EdgeConnect using Orchestrator AWS Network Manager Association UI.	/orchestrator/orchestration/aws_tgnm	FALSE	TRUE
6815908: MAJOR	Unable to orchestrate appliance[{0}] associated to aws core network edge[{1}] due to aws region mismatch [Error Code: InvalidEdgeLocation.NotFound].Details - edge connect region [{2}] and core network edge location [{3}]. **Recommended Action:** Associate appliance to core network edge location within the same region as EdgeConnect using Orchestrator AWS Network Manager Association UI.	/orchestrator/orchestration/aws_tgnm	FALSE	TRUE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
6815909: MAJOR	HPE SSE Tunnels has exceeded SSE tenant quota. **Recommended Action:** Increase the HPE SSE tenant size if more tunnels are needed.	/orchestration/HPE SSE	TRUE	TRUE
6815913: MAJOR	Your SAML Identity Provider certificate for {0} has expired. **Recommended Action:** Get the latest certificate from your Identity Provider and update it in Orchestrator SAML configuration.	/orchestrator/samlIdpCerts	FALSE	TRUE
6815919: MAJOR	Cipher profile changes could not be applied. {0}. **Recommended Action:** Validate your active cipher profile against the appliance and make sure it is supported.	/orchestrator/orchestration/cipherProfile	TRUE	TRUE
6815753: MINOR	There were {0} failed attempts to login over last 5 minutes. **Recommended Action:** Go to the Audit Logs for details on failed login attempts. Review usernames for validity and adherence to password policies.	/authentication	FALSE	TRUE
6815775: MINOR	Backup configuration not set. **Recommended Action:** Go to Backup to schedule backup or back up now.	/system/backup	FALSE	TRUE
6815835: MINOR	Appliance does not have any wan labels required for Azure VWAN Orchestration. No third party tunnels will be built on this appliance. **Recommended Action:** Use appliance deployment to assign at least one WAN label matching Azure VWAN Interface Label list.	/orchestrator/orchestration/azure	TRUE	TRUE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
6815836: MINOR	Appliance does not have any wan labels required for AWS TGNM Orchestration. No third party tunnels will be built on this appliance. **Recommended Action:** Use appliance deployment to assign at least one WAN label matching AWS Interface Label list.	/orchestrator/ orchestration/ aws_tgnm	TRUE	TRUE
6815845: MINOR	Appliance does not have any wan labels required for {0} Orchestration. No third party tunnels will be built on this appliance. **Recommended Action:** Use appliance deployment to assign at least one wan label matching Third party service Interface Label list.	/orchestrator/ orchestration/ service Orchestration	TRUE	TRUE
6815854: MINOR	HA Sync Communication is not enabled for {0} appliances. **Recommended Action:** Follow release documents to enable this feature.	/orchestration/ deployment	FALSE	TRUE
6815911: MINOR	The data of one or more tables might be leaking into their defaultPartitions. **Recommended Action:** Go to the Partition Management tab and verify that all rows are eventually moved out from the defaultPartition (i.e., Rows=0). Contact support if the issue persists.	/orchestrator/ Partition Management	TRUE	TRUE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
6815916: MINOR	Appliance does not have any lan labels required for Azure LAN Orchestration. **Recommended Action:** Use appliance deployment to assign at least one LAN label matching Azure cloud gateways LAN Interface Label list.	/orchestrator/orchestration/azure	TRUE	TRUE
6815755: WARNING	Orchestrator portal account or license will expire on {0 date}. **Recommended Action:** Go to Licensing to provide the required information.	/license	FALSE	TRUE
6815759: WARNING	Orchestrator cannot connect to HPE Aruba Networking EdgeConnect Cloud Portal using HTTPS. **Recommended Action:** Check portal connection and refer to the Audit Logs for more information.	/portal/connectivity	TRUE	TRUE
6815763: WARNING	CPX license will expire on {0 date}. **Recommended Action:** [Deprecated alarm] Renew your license to avoid service interruption.	/portal/license/cpx	FALSE	TRUE
6815764: WARNING	CPX license will expire on {0 date}. **Recommended Action:** [Deprecated alarm] Renew your license to avoid service interruption.	/portal/license/cpx	FALSE	TRUE
6815765: WARNING	CPX license will expire on {0 date}. **Recommended Action:** [Deprecated alarm] Renew your license to avoid service interruption.	/portal/license/cpx	FALSE	TRUE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
6815767: WARNING	Your EdgeConnect account has licenses that will expire in {0} day(s). EdgeConnect devices in your network will stop passing traffic on {1 date}. **Recommended Action:** Renew your license to avoid service interruption.	/portal/license/ec	FALSE	TRUE
6815768: WARNING	Your EdgeConnect account has licenses that will expire in {0} day(s). EdgeConnect devices in your network will stop passing traffic on {1 date}. **Recommended Action:** Renew your license to avoid service interruption.	/portal/license/ec	FALSE	TRUE
6815769: WARNING	Your EdgeConnect account has licenses that will expire in {0} day(s). EdgeConnect devices in your network will stop passing traffic on {1 date}. **Recommended Action:** Renew your license to avoid service interruption.	/portal/license/ec	FALSE	TRUE
6815771: WARNING	SaaS license will expire on {0 date}. **Recommended Action:** Renew your license to avoid service interruption.	/portal/license/saas	FALSE	TRUE
6815772: WARNING	SaaS license will expire on {0 date}. **Recommended Action:** Renew your license to avoid service interruption.	/portal/license/saas	FALSE	TRUE
6815773: WARNING	SaaS license will expire on {0 date}. **Recommended Action:** Renew your license to avoid service interruption.	/portal/license/saas	FALSE	TRUE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
6815783: WARNING	Orchestrator deployment size has exceeded the recommended level of {0} appliances. **Recommended Action:** Contact Support to increase the allocation of cloud resources.	/system/database	FALSE	TRUE
6815785: WARNING	Some appliances are paused from orchestration. **Recommended Action:** Go to Pause Orchestration List to see detail.	/orchestration	TRUE	FALSE
6815786: WARNING	Apply Overlays is currently disabled. **Recommended Action:** Enable Apply Overlays in Orchestration Settings.	/orchestration	TRUE	FALSE
6815789: WARNING	Apply Templates is currently disabled. **Recommended Action:** Enable Apply Templates in Orchestration Settings.	/orchestration/templates	TRUE	FALSE
6815793: WARNING	Your WAN Optimization will expire in {0} day(s). EdgeConnect devices in your network will stop using WAN Optimization on {1 date}. **Recommended Action:** Renew your WAN Optimization license to avoid service interruption.	/portal/license/ec	FALSE	TRUE
6815794: WARNING	Your WAN Optimization will expire in {0} day(s). EdgeConnect devices in your network will stop using WAN Optimization on {1 date}. **Recommended Action:** Renew your WAN Optimization license to avoid service interruption.	/portal/license/ec	FALSE	TRUE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
6815795: WARNING	Your WAN Optimization will expire in {0} day(s). EdgeConnect devices in your network will stop using WAN Optimization on {1 date}. **Recommended Action:** Renew your WAN Optimization license to avoid service interruption.	/portal/license/ec	FALSE	TRUE
6815798: WARNING	Paused stats collection for some of the appliances. **Recommended Action:** Go to Statistics Retention to resume stats collection.	/orchestration	FALSE	TRUE
6815806: WARNING	Stats Collection is paused and will resume after Orchestrator backup is completed. **Recommended Action:** Go to Scheduled & Historical Jobs to view the status of the backup process. Open a support case if this alarm persists more than 12 hours.	/orchestration/backup	FALSE	TRUE
6815811: WARNING	Zscaler Internet Access orchestration is paused. **Recommended Action:** Go to Zscaler Internet Access to resume orchestration.	/orchestration/zscaler	FALSE	FALSE
6815812: WARNING	Microsoft Azure Virtual WAN orchestration is paused. **Recommended Action:** Go to Microsoft Azure Virtual WAN to resume orchestration.	/orchestration/azure	FALSE	FALSE
6815813: WARNING	Could not allocate IPs from Loopback Pool {0}. **Recommended Action:** Change Loopback pool with enough IPs in Loopback Orchestration tag.	/orchestrator/orchestration	FALSE	TRUE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
6815814: WARNING	Loss and Latency metrics are available for IPSLA monitors with appliance version later than {0}. **Recommended Action:** Upgrade appliance to enable loss and latency metrics for IP SLA.	/orchestrator/orchestration/ipsla	FALSE	TRUE
6815816: WARNING	Best internet breakout is configured in overlay, but appliance does not support this feature (Deprecated). **Recommended Action:** Check the overlay configuration or upgrade the appliance.	/orchestrator/orchestration/overlays	TRUE	FALSE
6815818: WARNING	Shell access settings are different on the appliance than on Orchestrator. **Recommended Action:** Reconcile shell access setting. Matching Orchestrator policy with appliance setting is recommended.	/orchestrator/orchestration/shellAccessSetting	FALSE	TRUE
6815819: WARNING	Connection not established for websocket receiver: {0}. **Recommended Action:** Check WebSocket receiver configuration.	remoteLogWebSocket	FALSE	TRUE
6815821: WARNING	AWS Transit Gateway Network Manager orchestration is paused. **Recommended Action:** Go to AWS Network Manager to resume orchestration.	/orchestration/aws_tgnm	FALSE	FALSE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
6815823: WARNING	A new maintenance alert was received from HPE Aruba Networking. **Recommended Action:** Review the maintenance alert, then open a support case if you have any questions regarding the maintenance activity.	/portal/ SilverPeak Maintenance	TRUE	TRUE
6815824: WARNING	Stats Collection is lagging behind. **Recommended Action:** Open a support case if this alarm persists more than 24 hours.	/orchestrator/ statistics	FALSE	TRUE
6815832: WARNING	Inter-Segment Routing & DNAT rules have duplicate ip address with existing Inter-Segment Routing & DNAT Exceptions rules. **Recommended Action:** Check the Inter-Segment Routing & DNAT rules and solve the duplicate IP address.	/orchestrator/ routing Segmentation	TRUE	TRUE
6815840: WARNING	This appliance does not support Routing Segmentation. **Recommended Action:** Upgrade the appliance.	/orchestrator/ routing Segmentation	TRUE	TRUE
6815849: WARNING	ClearPass Policy Manager session paused. **Recommended Action:** Go to ClearPass Policy Manager to resume orchestration.	/orchestration/ clearPass	FALSE	FALSE
6815851: WARNING	Stats Collection is paused and will resume after backup is completed. **Recommended Action:** [Deprecated alarm]	/orchestration/ backup	FALSE	TRUE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
6815855: WARNING	Some of the HPE ANW Central sites don't have lat and lon. Site Name : {0}. **Recommended Action:** Orchestrator cannot derive a lat and lon from the site address information configured in HPE ANW Central.	/HPEAruba Networking Central/ HPEAruba Networking Central	FALSE	TRUE
6815864: WARNING	Your EdgeConnect account {0} has licenses that will expire in {1} day(s). EdgeConnect devices in your network that are licensed with this account will stop passing traffic on {2 date}. **Recommended Action:** Renew your license to avoid service interruption.	/portal/license/ ec	FALSE	TRUE
6815865: WARNING	Your EdgeConnect account {0} has licenses that will expire in {1} day(s). EdgeConnect devices in your network that are licensed with this account will stop passing traffic on {2 date}. **Recommended Action:** Renew your license to avoid service interruption.	/portal/license/ ec	FALSE	TRUE
6815866: WARNING	Your EdgeConnect account {0} has licenses that will expire in {1} day(s). EdgeConnect devices in your network that are licensed with this account will stop passing traffic on {2 date}. **Recommended Action:** Renew your license to avoid service interruption.	/portal/license/ ec	FALSE	TRUE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
6815868: WARNING	Your WAN Optimization for account {0} will expire in {1} day(s). EdgeConnect devices in your network that are licensed with this account will stop using WAN Optimization on {2 date}. **Recommended Action:** Renew your WAN Optimization license to avoid service interruption.	/portal/license/ec	FALSE	TRUE
6815869: WARNING	Your WAN Optimization for account {0} will expire in {1} day(s). EdgeConnect devices in your network that are licensed with this account will stop using WAN Optimization on {2 date}. **Recommended Action:** Renew your WAN Optimization license to avoid service interruption.	/portal/license/ec	FALSE	TRUE
6815870: WARNING	Your WAN Optimization for account {0} will expire in {1} day(s). EdgeConnect devices in your network that are licensed with this account will stop using WAN Optimization on {2 date}. **Recommended Action:** Renew your WAN Optimization license to avoid service interruption.	/portal/license/ec	FALSE	TRUE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
6815872: WARNING	EdgeHA peer {0} is configured on a different account. EdgeHA Group reporting may be inaccurate unless the appliances are on the same metered account. (Deprecated) **Recommended Action:** [Deprecated alarm] Check the accounts that the EdgeHA peers are registered to in the Licenses tab.	/orchestrator/system/deployment/ha	FALSE	TRUE
6815874: WARNING	Netskope orchestration is paused. **Recommended Action:** Go to Netskope to resume orchestration.	/orchestration/netskope	FALSE	FALSE
6815888: WARNING	Your AAS license expires on {0 date} (in {1} days). **Recommended Action:** [Deprecated alarm]	/portal/license/orchestrator	FALSE	TRUE
6815889: WARNING	Your AAS license expires on {0 date} (in {1} days). **Recommended Action:** [Deprecated alarm]	/portal/license/orchestrator	FALSE	TRUE
6815890: WARNING	Your AAS license expires on {0 date} (in {1} days). **Recommended Action:** [Deprecated alarm]	/portal/license/orchestrator	FALSE	TRUE
6815894: WARNING	Non Edge HA Profile active on Edge HA Site/Cluster ({0}). **Recommended Action:** Map the Edge HA Profile to this site/cluster.	/orchestration/clusterManager	TRUE	FALSE
6815895: WARNING	Edge HA Profile active on Non Edge HA Site/Cluster ({0}). **Recommended Action:** Map the Non Edge HA Profile to this site/cluster.	/orchestration/clusterManager	TRUE	FALSE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
6815896: WARNING	Interface ({0}) not available in appliance ({1}). **Recommended Action:** Configure deployment or profile to continue.	/orchestration/ clusterManager	TRUE	FALSE
6815900: WARNING	HPE SSE orchestration is paused. **Recommended Action:** Go to HPE SSE to resume orchestration.	/orchestrator/ HPE SSE	FALSE	FALSE
6815901: WARNING	Unable to orchestrate edgeHA appliances. EdgeHA {0} site name is not configured. **Recommended Action:** Configure same site name for both EdgeHA appliances.	/orchestrator/ HPE SSE/ha	FALSE	TRUE
6815902: WARNING	Unable to orchestrate edgeHA appliances. EdgeHA peer {0} has a different site name from its peer. **Recommended Action:** Configure same site name for both EdgeHA appliances.	/orchestrator/ HPE SSE/ha	FALSE	TRUE
6815910: WARNING	One or more certificates have expired or are due to expire within the next 60 days. **Recommended Action:** Go to Custom CA Certificate Trust Store to update the certificate.	/orchestrator/ customCerts	TRUE	FALSE
6815912: WARNING	Your SAML Identity Provider certificate for {0} will expire on {1}. **Recommended Action:** Get the latest certificate from your Identity Provider and update it in Orchestrator SAML configuration.	/orchestrator/ samlIdpCerts	FALSE	TRUE

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
6815915: WARNING	Azure Network Manager orchestration is paused. **Recommended Action:** Go to Azure Network Manager Cloud Gateways to resume orchestration.	/orchestration/azureLan	FALSE	FALSE
6815917: WARNING	Either labels are not configured in appliance or WAN side interfaces / labels are not permitted. **Recommended Action:** Configure deployment or profile to continue.	/orchestration/clusterManager	TRUE	FALSE

Equipment

System Type 100 (Orchestrator); Source Type 3 (Equipment)

Alarm ID: Severity	Alarm Text	Source	Service Affecting	Clearable
6750234: CRITICAL	Failed to get database connection. Details: {0}. **Recommended Action:** Reserve required Memory and CPU.	/system/resource	TRUE	TRUE
6750209: MAJOR	Disk partition {0} is dangerously full - {1}% used. **Recommended Action:** Go to Server Information to see detailed disk usage.	/system/disk	FALSE	FALSE
6750208: WARNING	Disk partition {0} is more than {1}% used. **Recommended Action:** Go to Server Information to see detailed disk usage.	/system/disk	FALSE	FALSE

Route Next Hops

Monitoring > Summary > Route Next Hops

The Route Next Hops tab displays the state of each management, WAN, and LAN next hop. Next hop status is derived from Layer 2 reachability using Address Resolution Protocol (ARP). When a next hop is deemed unreachable, all associated routes for the next hop will go down.

Filter	Description
Appliance	Appliance name.
Next Hop IP	IP address of the router to which the appliance sends datapath (or management) traffic.
Interface	Logical port associated with the Next Hop IP.
Source	Direction of the next hop router, relative to the appliance.
State	One of four possible states: **Initializing**, **Reachable**, **Unreachable**, **Test disabled** (when appliance is in Bypass mode).
Uptime	Amount of time that the next hop router has been reachable.

Route Label Status

Monitoring > *Summary* > *Route Label Status*

The Route Label Status tab provides an at-a-glance view of the status of all route labels. It shows whether each route label is UP or DOWN, based on the status of the associated IP SLA rules. This tab helps users quickly determine the reachability of critical routes and applications.

The IP SLA Condition on the route label determines whether the label should be marked DOWN if *all* IP SLA rules are DOWN or if *any* IP SLA rules are DOWN.

IP SLA Condition	Description
ALL	If *all* IP SLA rules associated with a route label are in the DOWN state, all routes assigned to that label are marked DOWN and are removed from the routing table.
ANY	If *any* IP SLA rule associated with a route label is in the DOWN state, all routes assigned to that label are marked DOWN and are removed from the routing table.

- Route labels do not impact routes received from BGP, OSPF, or fabric peers.
- Route labels can only be applied to locally configured static routes.
- You can apply one label per static route.

To create a route label, click **Edit Labels** in the tab header. The Route Labels dialog box opens.

To create a route, click **Routes** in the tab header.

To create an IP SLA rule, click **IP SLA** in the tab header.

To refresh the status, click the **Refresh** icon.

NOTE: The default label called “Internet” is a built-in route label used to manage the status of routes related to internet connectivity. It is typically associated with a default route (0.0.0.0) to dynamically control its availability based on the status of IP SLA rules.

Route Labels Dialog Box

Add, edit, or delete a route label from this dialog box. You can create up to 127 route labels.

Add a Route Label

1. Click **Add Label**. A new row is added to the table.
2. Click in the new cell in the Label column and then enter a name for the route label. Route Label names can be up to 15 characters and contain alphanumeric characters, dashes, and underscores.
3. Press **Tab** or click in the new cell in the IP SLA Condition column, and then select **ALL** or **ANY**.

   By default, if *all* IP SLA rules associated with a route label are in the DOWN state, the label will go DOWN, causing all routes assigned to it to also be marked DOWN and be removed from the routing table. If you set the IP SLA Condition to ANY, the route label and assigned routes go DOWN if *any* linked IP SLA rule associated with the label is in the DOWN state.
4. (Optional) Click in the Comment field and enter a comment.
5. Click **Save**.

Edit a Route Label

1. Click **Edit Labels**.

   The Edit Labels dialog box opens.
2. Click in the cell to edit the label, IP SLA Condition, or comment.
3. Click **Save**.

Delete a Route Label

1. Click **Edit Labels**.

   The Edit Labels dialog box opens.
2. Click the delete icon (**X**) in the last column of the route label you want to delete.
3. Click **Save**.

NOTICE: If you delete a route label, the route label is removed from IP SLA rules and static routes currently using that route label.

Appliance Power Consumption

Monitoring > Summary > Appliance Power Consumption

The Appliance Power Consumption tab displays the power consumed by each selected EdgeConnect appliance that supports this feature.

The graph at the top of the tab displays an aggregated view of throughput bytes and power during a specified time range for the selected appliances. Click a predefined time period (**1h**, **4h**, **1d**, **7d**) to display statistics over the last hour, four hours, day, or seven days. Click **Custom** and set your own custom time range to display statistics for that time period.

The table at the bottom of the tab displays a list of selected appliances and provides the following information for each appliance:

- Power Delivered by the Power Adapter (Voltage, Current and Power)
- Power Delivered by PSU (Voltage, Current and Power)
- Serial Number of the PSU
- Firmware Version of the PSU
- Temperature of the PSU
- The PSU system connection status
- The PSU power connection status

Field	Description
Appliance	Name of the appliance
Model	The model number of the appliance (for example, EC-XL, EC10106, EC10104)
Power Supply Type	The type of power supply installed in the appliance (Modular - Internal or External power adapter)
Last Week energy (kWh)	The total power consumed by the appliance over the last seven days, measured in kilowatt-hours
Last 24h energy (kWh)	The total power consumed by the appliance over the last 24 hours, measured in kilowatt-hours
Power (W)	The current instantaneous power draw of the appliance, measured in watts
Connection Status (PS 1 / PS 2)	Indicates whether the power supply unit is connected and active. Green indicates **On**; Gray indicates **Off**. Appliances with an external power adapter display a single connection status indicator.
Input (V) (PS 1 / PS 2)	The input voltage for the power supply unit. For modular internal power supplies, this value reflects AC input voltage. For external power adapters, this value reflects DC input voltage.
Current (A) (PS 1 / PS 2)	The current draw for the power supply unit, measured in amperes.

Field	Description
Temperature (C) (PS 1 / PS 2)	The operating temperature of the power supply unit, measured in degrees Celsius.

The method of obtaining data depends on the type of power supply on the appliance: - **EC-10104, EC-10106, and EC-10108**: Uses an external power adapter. Current and voltage of the power output from the adapter is measured. - **EC-10150, EC-XL, EC-SP, EC-L, EC-M, and EC-S**: Uses a modular power supply. Current and voltage of the power output from the Power supply is measured. In addition to this, serial number, firmware version, and power connectivity status to the power supply is reported. - **EC-XS and EC-US**: Uses an external power adapter. These appliances do not have a provision to measure power consumption.

Monitoring > Security

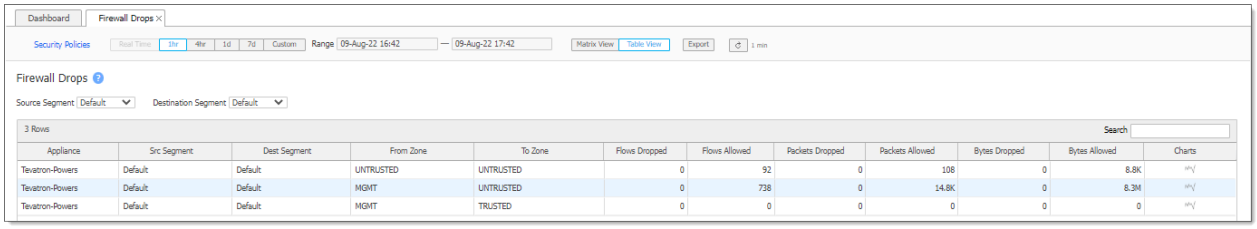
The options under **Monitoring > Security** focus on reports related to firewall drops, firewall protection profiles, and flow baselines.

Firewall Drops

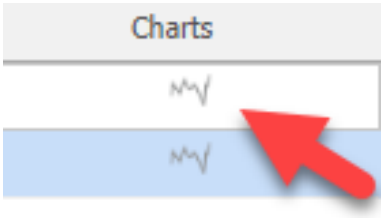
Monitoring > Security > Firewall Drops

You can use the Firewall Drops tab to view the statistics on various flows, packets, and bytes dropped or allowed by a zone-based firewall for a given time range.

- You can select a range of time (in hours and days) to view the firewall drops. You can also select to view in Matrix or Table view.
- Select **Export** to export the report to an excel spreadsheet.



- If segmentation is enabled, you can specify the **Source Segment** and the **Destination Segment** to search for the flows, packets, and firewall drops in that segment.
- In the Charts column, you can select the chart icon.
- In this pop-up, you can see packets, and bytes dropped or allowed by a zone-based firewall for a given time range.



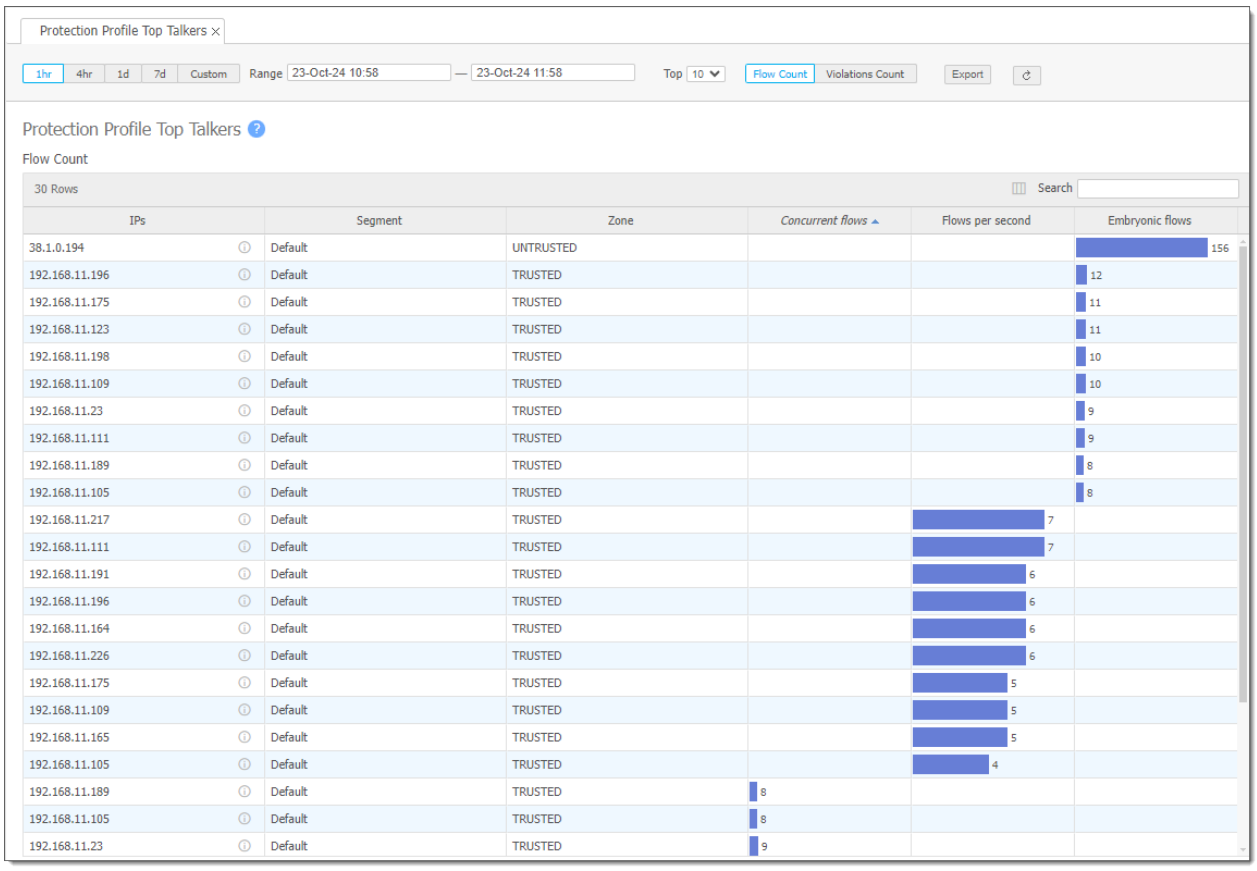
Protection Profile Top Talkers

Monitoring > Security > Protection Profile Top Talkers

The Protection Profile Top Talkers tab shows historical statistics related to Denial of Service (DoS) for the selected appliance. This chart is only available if the selected appliance has a firewall protection profile applied.

The Protection Profile Top Talkers chart will show 10, 25, or 50 IPs that produce the top concurrent flow, flows per second, or embryonic flows. You can also select to show the top IPs that cross the min/max thresholds. Use this data to monitor network behavior based on DoS activity to identify how the device is operating with the current DDoS settings and determine how to tune the settings to obtain the expected behavior.

TIP: Click the information icon in the IPs column to display details about the IP address.



You can customize the chart settings using the controls at the top of the tab, as follows:

Option	Description
Time period	Click a predefined time period (**1h**, **4h**, **1d**, **7d**) to display statistics over the last hour, four hours, day, or seven days. Click **Custom** and set your own custom time range to display statistics for that time period.
Top	Select **10**, **25**, or **50** to view the top 10, 25, or 50 IPs that produce the most concurrent flow, flows per second, or embryonic flows.
Flow Count	Click this option to display the concurrent flow counts, embryonic flow counts, and flows per second. The DoS feature monitors the Protection Profile Top Talkers on a per host basis. You can sort the Concurrent flows, Flows per second, and Embryonic flows columns in ascending or descending order. Click the column heading to toggle the sort order.
Violations Count	Click this option to display the minimum and maximum threshold violation counts. The chart shows the maximum and minimum threshold violations for the segment and zone of each IP.
Export	Click **Export** to save the contents of the Protection Profile Top Talkers table to a CSV file.
Refresh	Click the **Refresh** button to fetch data again for the selected time period.

Protection Profile Trends

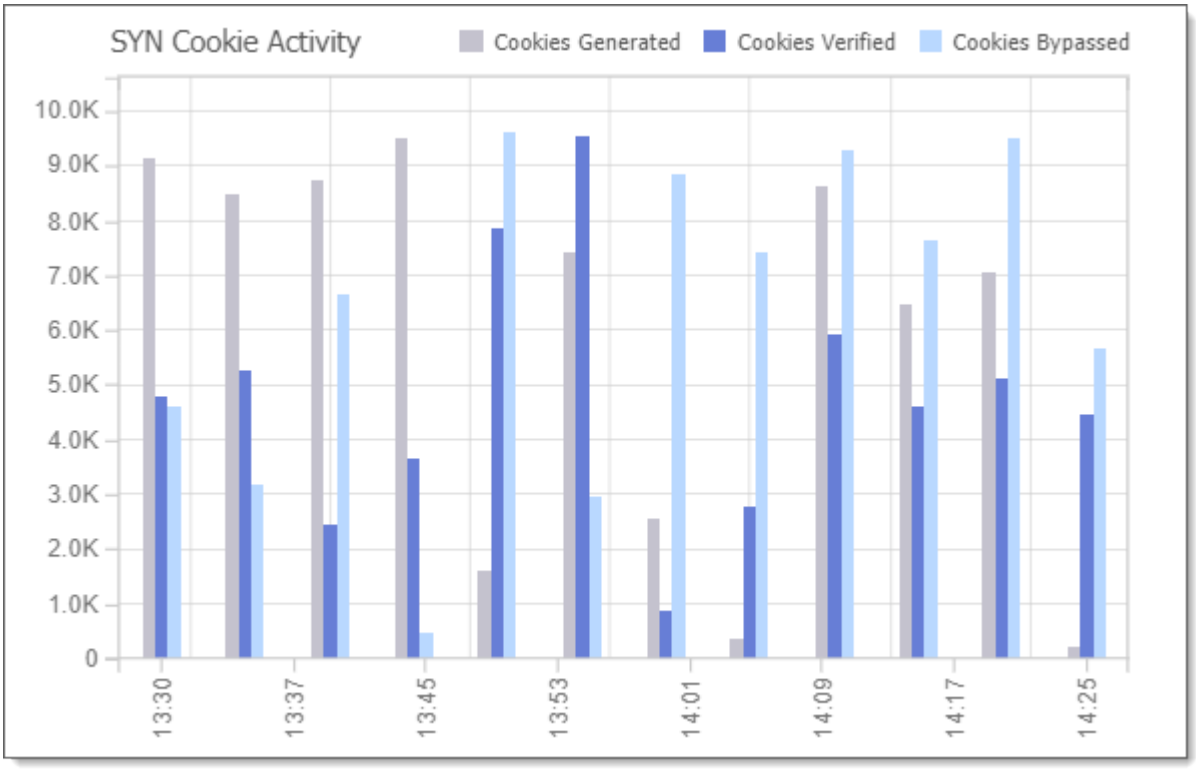
Monitoring > Security > Protection Profile Trends

The Protection Profile Trends tab provides a view of historical and real-time statistics related to Firewall Protection Profile thresholds. These thresholds are configured and mapped to specific zones and segments as defined by your setup. These charts are only available if the appliance has a Firewall Protection Profile applied to a zone/segment pair. These charts offer insights into statistics corresponding to Firewall Protection Profile thresholds configured for the specific zone and segment combinations.

The following filters allow you to select the data corresponding to the configured Firewall Protection Profile thresholds that are set using the Classification, Metric, and IP Protocol settings in the Firewall Protection Profile, which are then applied to the selected segment and zone pairs. By using these filters effectively, you can gain insights into the Firewall Protection Profile threshold statistics mapped to the zone/segment pairs of the EdgeConnect appliance.

- **Segment:** Select a Segment and Zone combination to see the statistics related to the threshold configured and mapped to the selected segment and zone combination in the Firewall Protection Profile mapping. If a Firewall Protection Profile is not mapped to a segment and zone combination, charts for this segment and zone will not be shown.
- **Zone:** Select a Segment and Zone combination to see the statistics related to the threshold configured and mapped to the selected segment and zone combination in the Firewall Protection Profile mapping. If a Firewall Protection Profile is not mapped to a segment and zone combination, charts for this segment and zone will not be shown.
- **Metric:** Select **Flows per second**, **Concurrent flows**, **Embryonic flows**, or **Embryonic Two-Way flows** to see the corresponding statistics. If a threshold for the selected metric is present in the Firewall Protection Profile, relevant drops and violation count charts are shown. If the threshold for the selected metric is not present, only the peak chart is shown.

  If you select Embryonic Two-Way flows, a SYN Cookie Activity chart is shown with the other charts.

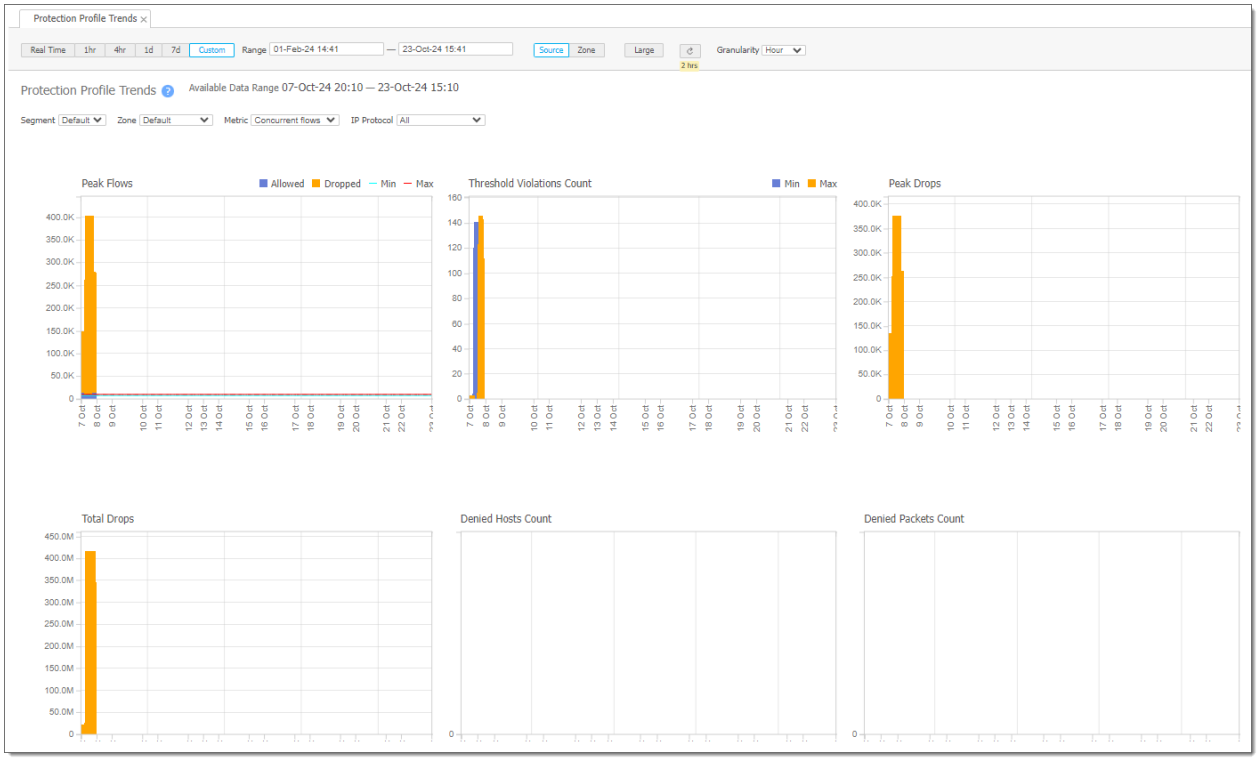


The SYN Cookie Activity chart includes these statistics:

  - Cookies Generated: The number of SYN cookie challenges generated.
  - Cookies Verified: The number of SYN cookie challenges that received a verified response, which allowed TCP session initiation with the server.

NOTE: Unlike other Firewall Protection Profile settings, Smart SYN Cookie threshold actions are not applied to fabric traffic. The Cookies Bypassed counter represents the number of flows that either originate from trusted IP addresses or are identified as fabric traffic.

- **IP Protocol:** Select **All**, **TCP**, **UDP**, or **Others** to see the corresponding statistics. If a threshold for the selected protocol is present in the Firewall Protection Profile, relevant drops and violation count charts are shown. If the threshold for the selected protocol is not present, only the peak chart is shown.



You can customize the chart settings using the controls at the top of the tab, as follows:

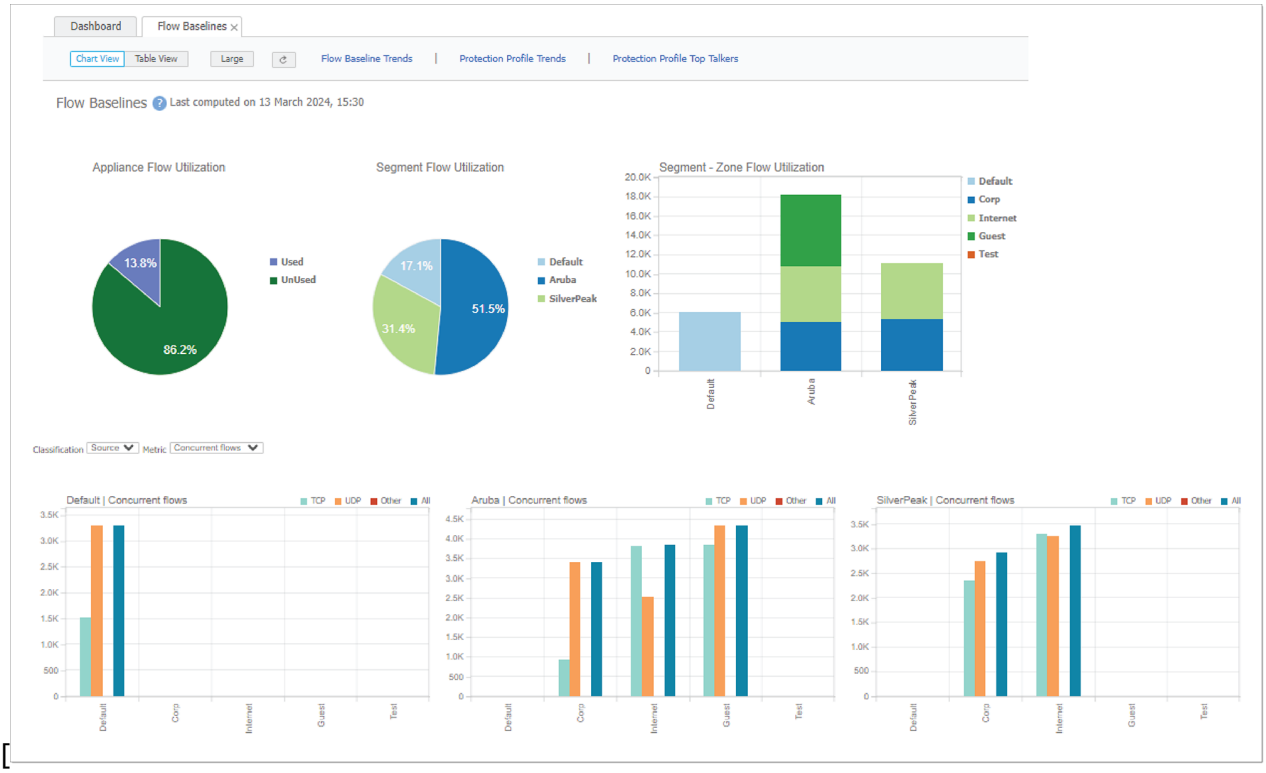
Option	Description
Time period	Click **Real Time** to enable live statistics for all available interfaces. Click a predefined time period (**1h**, **4h**, **1d**, **7d**) to display statistics over the last hour, four hours, day, or seven days. Click **Custom** and set your own custom time range to display statistics for that time period.

Option	Description
Large	Click this option to toggle the size of the charts between small (default) and large.
Source	Click this option to filter the list to view charts relevant for the thresholds configured at the source level.
Zone	Click this option to filter the list to view charts relevant for the thresholds configured at the zone level.
Refresh	Click the **Refresh** button to fetch data again for the selected time period.
Granularity	When a custom time period is defined, select the granularity level to be applied to the charts (**Minute**, **Hour**, or **Day**).

Flow Baselines

Monitoring > Security > Flow Baselines

The Flow Baselines tab shows the baseline data computed in the previous interval for the selected appliance. The data is refreshed after each baseline computation interval. These charts are only available if you have enabled baseline learning for the appliance.



[

The first set of charts shown includes the following:

- Appliance flow utilization – Indicates used and unused flow capacities. If Smart burst is configured, unused indicates the spare capacity that can be used by Smart burst.
- Segment flow utilization – Shows how flow capacity is disbursed across the segments.
- Segment-zone flow utilization – Shows flow capacity use for all zones in a segment.

Individual charts that show the baseline data for each configured segment are also shown. The segment charts show the most recent baselines computed for each protocol within each segment zone. By default, the charts show baseline data for concurrent flows at the source level. To see baseline data for other network statistics, select the following filters to customize the data on the individual segment charts.

- **Classification**: Filter to show data by either the source level (**Source**) or the zone level (**Zone**).
- **Metric**: Select the metric to use for stats (**Flows per second**, **Concurrent flows**, or **Embryonic flows**).

You can toggle the individual segment data between **Chart View** and **Table View**. When using Table View, the following columns appear in the table.

Column	Description
Segment	Name of the segment.
Zone	Name of the zone within the segment.
Classification	The level at which network statistics are measured. The baseline is computed at the Zone level or Source level.
Metric	The configured DoS threshold metric (Flows per second, Concurrent flows, or Embryonic flows).
TCP	The most recent baseline computed for TCP flows.
UDP	The most recent baseline computed for UDP flows.
ICMP	The most recent baseline computed for ICMP flows.
Other	The most recent baseline computed for other flows.
All	The most recent baseline computed for all flows.

You can customize settings for all the data using the controls at the top of the tab, as follows:

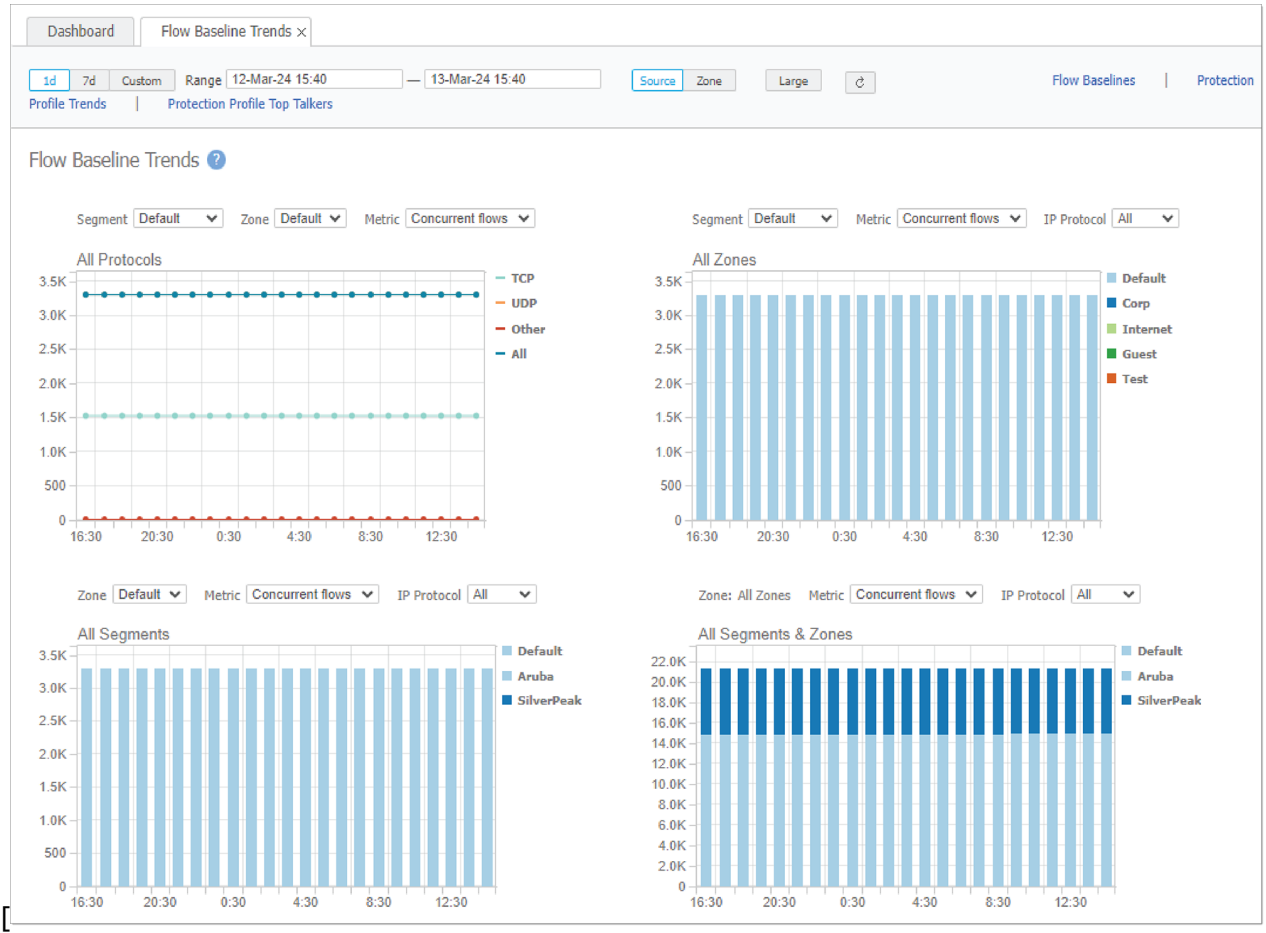
Option	Description
Chart View	Click to view the data in chart format.
Table View	Click to view the data for individual segments in table format.

Option	Description
Large	Click this option to toggle the size of the charts between smaller (default) and large.
Refresh 	Click the **Refresh** button to fetch data again for the selected time period.

Flow Baseline Trends

Monitoring > Security > Flow Baseline Trends

The Flow Baselines Trends tab shows historical statistics for flow baseline data for the selected appliance. You can use these charts to compare historical baseline data across the different zones, segments, and protocols in your network. These charts are only available if you have enabled baseline learning for the appliance.



All Protocols – This chart shows the historical baseline data that was computed at each interval for all protocols. By default, it uses the Concurrent flows metric and shows data for the default segment and zone. Use the filters above the chart to view this data for other segments or zones and using various metrics.

All Zones – This chart shows the baseline data for all zones within a segment. By default, it uses the Concurrent flows metric and shows data for the default segment and all protocols. Use the filters above the chart to view this data for a different segment and for various protocols.

All Segments – This chart shows the baseline data for a particular zone across all segments. By default, it uses the Concurrent flows metric and shows data for the default zone and all protocols. Use the filters above the chart to view this data for a different zone in the segment and for various protocols.

All Segments and Zones – This chart shows baseline data for all segments for all zones. By default, it uses the Concurrent flows metric and shows data for all protocols. Use the filters above the chart to view this data for various metrics and protocols.

Select the following filters to customize the data shown on the charts.

- **Segment**: Select the name of the segment from the drop-down list.
- **Zone**: Select the name of the zone from the drop-down list.
- **Metric**: Select the metric to use for stats (**Flows per second**, **Concurrent flows**, **Embryonic flows**, or **Embryonic Two-Way flows**).
- **IP Protocol**: Select the protocol (**All**, **TCP**, **UDP**, **ICMP**, or **Others**).

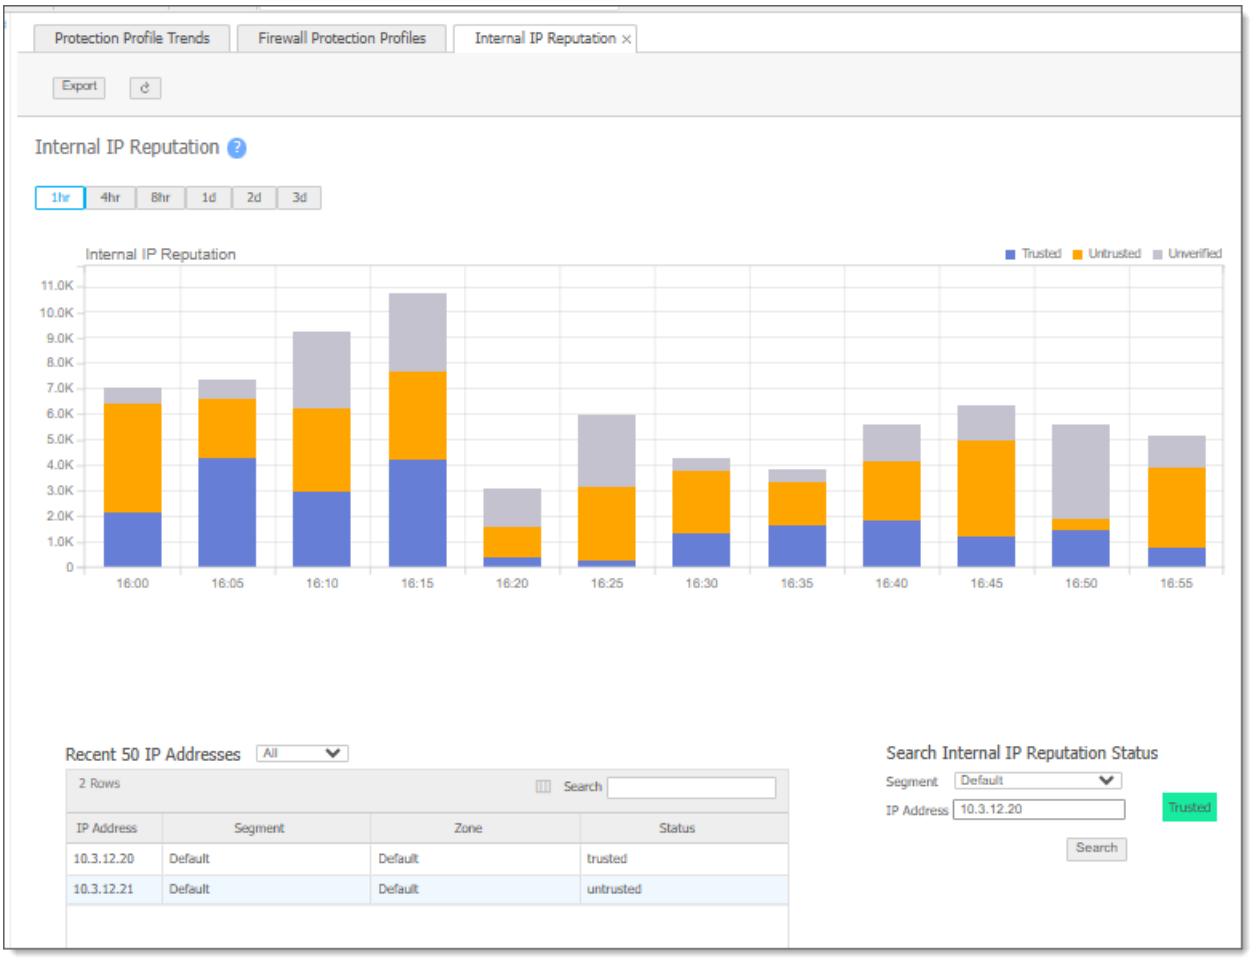
You can customize the chart settings using the controls at the top of the tab, as follows:

Option	Description
Time period	Click a predefined time period (**1d** or **7d**) to display statistics over the last day or seven days. Click **Custom** and set your own custom time range to display statistics for that time period.
Large	Click this option to toggle the size of the charts between smaller (default) and large.
Source	Click this option to filter the list to view baseline data computed at the source level.
Zone	Click this option to filter the list to view baseline data computed at the zone level.
Refresh 	Click the **Refresh** button to fetch data again for the selected time period.
Granularity	When a custom time period is used, select the granularity level to be applied to charts (**Hour** or **Day**).

Internal IP Reputation

Monitoring > Security > Internal IP Reputation

If Internal IP Reputation is enabled, this tab provides historical data about the reputation scores for IP addresses that have attempted TCP connections with an appliance. Internal IP Reputation is computed by ECOS. It is used in dynamic IP reputation management for the SYN Cookie and SYN Smart Cookie features for TCP connections.



- To view Internal IP Reputation data for an appliance, select the appliance in the appliance tree.
- The Internal IP Reputation chart shows the number of Trusted, Untrusted, or Unknown IP addresses that are present in the segment or zone where the Firewall Protection Profile is enabled. You can filter the data using **1hr**, **4hr**, **8hr**, **1d**, **2d**, or **3d**.
- To find the Internal IP Reputation status for an IP address, under Search Internal IP Reputation Status, select the segment to search within, enter the IP address, and click **Search**.
- To export the data from this tab to a CSV file, click **Export**.

The Recent 50 IP Addresses table provides the following information about the 50 most recent IP addresses in the segment or zone where the Firewall Protection Profile is applied with Enforce Strict 3-way TCP enabled. You can filter the data by selecting **All**, **Trusted**, or **Untrusted**.

Column	Description
IP Address	Numerical label for the IP address.
Segment	Firewall segment in which the IP address attempted a TCP connection with the appliance.
Zone	Firewall zone in which the IP address attempted a TCP connection with the appliance.
Status	Internal IP Reputation score for the address (Trusted or Untrusted).

Secure Web Services Top Talkers

Monitoring > Security > Secure Web Services Top Talkers

The Secure Web Services Top Talkers tab shows historical statistics related to Secure Web Services for selected appliances. Statistics are aggregated if more than one appliance is selected in the appliance tree.

NOTE: This feature provides statistics only for appliances that are licensed and configured for Secure Web Services. For more license information, see Secure Web Services.

Select one or more appliances in the appliance tree, and then click one of the following buttons to view the associated statistics:

Field	Description
Allowed Web Category	The table lists web categories representing sets of web domains and sites that users are allowed to access. It also provides hit counts related to the listed web categories. These categories are pre-defined by OpenText™ Threat Intelligence (BrightCloud) and are allowed using match criteria for web categories. To display a list of top URLs that users accessed for a listed web category, click the associated info icon in the Top URLs field. If the URL is implicitly allowed, or a web category to which it is assigned is not blocked, the URL is included in the Allowed URL table.
Denied Web Category	The table lists web categories representing sets of web domains and sites that users are not allowed to access. It also provides hit counts related to the listed web categories. These categories are pre-defined by OpenText Threat Intelligence (BrightCloud) and are denied using match criteria for web categories. To display a list of top URLs that users attempted to access but were denied, click the associated info icon in the Top URLs field. If the URL is implicitly denied, or a web category to which it is assigned is blocked, the URL is included in the Denied URL table.

Field	Description
Allowed URL	The table lists URLs that users were allowed to access. It also provides hit counts related to the listed URLs. These URLs were allowed because they belong to an allowed web category or are explicitly allowed using match criteria for the specific URLs. To display a list of top IP addresses that accessed a listed URL, click the associated info icon in the Top IPs field.
Denied URL	The table lists URLs that users attempted to access but were denied. It also provides hit counts related to the listed URLs. These URLs were not allowed because they belong to a denied web category, or the specific URLs were denied because match criteria for the specific URLs deny access. To display a list of top IP addresses that attempted to access a listed URL but were denied, click the associated info icon in the Top IPs field.
Top Src IP	The table lists top source IPs based on flow counts. To display a list of top URLs that were accessed by a listed source IP address, click the associated info icon in the Top URLs field.

To save the contents of the currently displayed table to a CSV file, click **Export**.

Security Log Viewer

Monitoring > *Security* > *Security Log Viewer*

On the Security Log Viewer tab, you can view real-time or stored (historical) security logs from EdgeConnect appliances. You can view logs from up to eight appliances at a time. Security logs are displayed in a tabular format; for detailed descriptions of the columns shown on this tab, see Security Log Attributes. You can also download security logs from this tab for review offline.

You can view the following types of security logs on the Security Log Viewer tab:

- IDS/IPS logs—Contain information about intrusion detection/prevention events.
- Firewall logs—Contain information about network traffic, connection states, and rule matches for zone-based firewall and Secure Web Services.
- Firewall protection profile logs—Contain information about threshold-based security events, such as DoS protection actions including Drop excess, Rapid aging, SYN Cookie, and Smart SYN Cookie.
- Stateful WAN drops logs—Contain information about inbound packet drops on WAN interfaces due to congestion, security policies, or resource limits.

You can also access the Global Firewall Policies and Firewall Drops tabs by clicking the links at the top of the page.

View Real-Time Security Logs

The default option for viewing logs on the Security Log Viewer tab is RealTime. Complete the following steps to view real-time security logs:

1. Navigate to **Monitoring > Security > Security Log Viewer**.
2. Select the log type.
3. Select up to eight appliances from the appliance tree.
4. Click **RealTime**.
5. Click **Run**.

   Orchestrator streams and displays real-time security logs for the selected appliances. Logs continuously stream in the background, and the table refreshes to display new logs every 5 seconds.

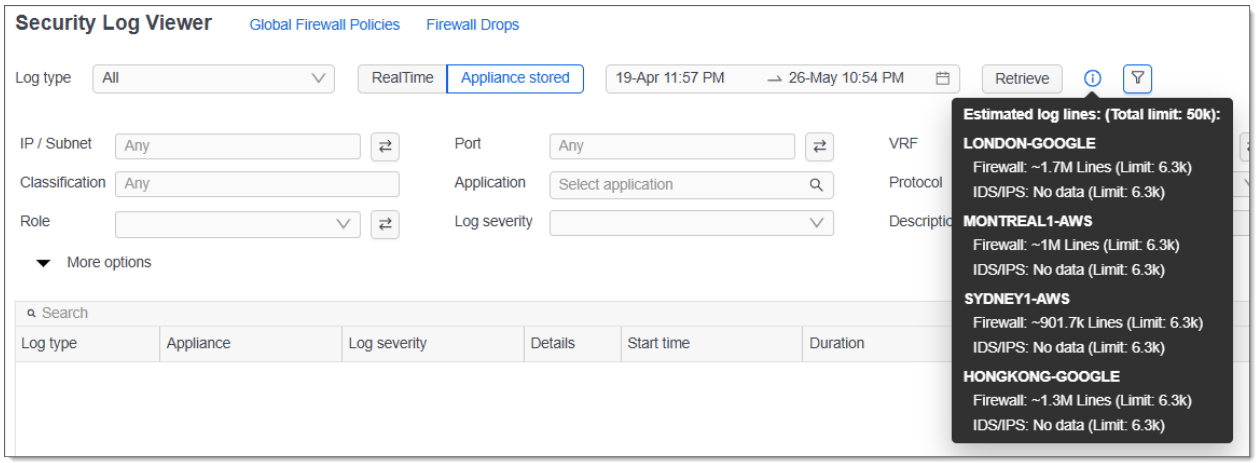
   - Change refresh interval—Click the refresh timer arrow and select an option from the menu (5 seconds, 10 seconds, or 20 seconds). The refresh timer is only active when real-time logs are streaming.
   - Pause log streaming—Move the toggle to enable Pause. Orchestrator continues to stream and accumulate logs in the queue for the selected appliances.
   - Resume log streaming—Move the toggle to disable Pause. When you resume streaming, the log table is marked with a thick line indicating the last log received before pausing. If you navigate to another Orchestrator tab and then return to the Security Log Viewer tab, the thick line will not persist.
   - Stop log streaming—Click **Stop**. The refresh timer becomes inactive.

   If you pause or stop streaming real-time logs, you can switch to another Orchestrator tab. When you return to the Security Log Viewer tab, the logs are still there and visible.

NOTE: Real-time log viewing is intended for active, short-duration troubleshooting while traffic or security events are being generated. It is not intended for unattended, long-running log collection. For historical analysis or logs generated over a longer period, use appliance stored logs.

View Stored Security Logs

When viewing stored security logs, the table has a default maximum line limit of 50,000 and one log is displayed per line. If you have multiple appliances selected, the 50,000 lines are distributed across the appliances, so an equal number of logs are displayed for each appliance. Before you retrieve logs in the selected date range, you can click the info icon (**i**) next to the Retrieve button to view a preview of how many logs will be shown for each appliance.



You can change the maximum line limit for retrieved logs in Orchestrator Advanced Properties. For more information, see Configure Maximum Line Limit for Security Logs.

If the selected date range contains more than 50,000 logs, the system notifies you that only the first 50,000 logs will be displayed. You can view all the logs in the range by clicking **Retrieve**, which loads logs in increments of 50,000. Continue clicking **Retrieve** to view the next 50,000 logs.

To view security logs that are stored (historical) on an appliance:

1. Navigate to **Monitoring > Security > Security Log Viewer**.
2. Select the log type.
3. Select up to eight appliances from the appliance tree.
4. Click **Appliance stored**.
5. Select a start date/time and click **OK**. Then select an end date/time and click **OK**.

   If you have multiple appliances selected, you can click the info icon (**i**) to see how many logs per appliance the system will display.

   NOTE: The available date range is based on the earliest and latest log dates for the selected appliance(s). If you selected multiple appliances, the range is limited to the period (earliest date and latest date) when security logs are available for all selected appliances. This ensures you cannot select a range with no data.
6. Click **Retrieve**.

   Orchestrator retrieves and displays the stored security logs for the selected appliances within the specified date range. When viewing appliance stored logs, you can navigate to another Orchestrator tab and the retrieved logs will continue to be displayed on the Security Log Viewer tab when you navigate back to it.

Configure Maximum Line Limit for Security Logs

To configure the maximum line limit, you must modify the Orchestrator Advanced Properties. The ability to modify Orchestrator Advanced Properties depends on your Orchestrator deployment.

- Self-hosted/on-prem Orchestrator users can edit the Orchestrator Advanced Properties using the following instructions.
- Orchestrator-as-a-Service (OaaS) users must open a TAC case to request changes to the Orchestrator Advanced Properties.

1. In Orchestrator, navigate to **Orchestrator > Software & Setup > Setup > Advanced Properties**.

   The Orchestrator Advanced Properties dialog box opens.

   WARNING: Use caution when modifying Orchestrator Advanced Properties. Verify your changes before applying them.

2. Set the following security log property values as necessary:

Security Log Property	Description	Range
securityRealtimeLogsLimit	Maximum number of real-time security logs that can be displayed (one line per log).	Allowable range: 10,000-200,000 lines. The default is 20,000.
securityStoredLogsLimit	Maximum number of stored (historical) security logs that can be displayed (one line per log).	Allowable range: 10,000-200,000 lines. The default is 50,000.

3. Click **Save**.

   You do not need to restart Orchestrator to activate the changes to the security log advanced properties.

Filter Security Logs

To filter displayed logs, click the filter icon to show the filter options. To view more filters, click **More options**. To apply selected filters, click **Apply**. To clear filters, click **Clear**.

The following table describes the filters that you can apply to the security log table.

Filter	Description
Log type	To view a specific type of security log, select an option from the Log type menu (Firewall, Stateful WAN drops, Firewall protection profile, or IDS/IPS). The default is All.
IP / Subnet	To view logs associated with a specific IP address or subnet, enter it in the IP / Subnet field; this filter shows logs that contain this IP address or subnet as the source or destination. To view logs associated with a specific source IP address or subnet, click the source/destination button and enter it in the Source field. To view logs associated with a specific destination IP address or subnet, click the source/destination button and enter it in the Destination field.

Filter	Description
Classification	To view IDS/IPS logs associated with a specific security or threat classification, enter it in the Classification field.
Role	To view logs associated with a specific network security/identity role, select it from the Role menu; this filter shows logs that contain this role for both source and destination traffic. To view logs associated with a specific network security/identity role for source traffic, click the source/destination button and select it from the Source menu. To view logs associated with a specific network security/identity role for destination traffic, click the source/destination button and select it from the Destination menu.
Username	To view logs associated with a specific username, enter it in the Username field.
Web category	To view logs associated with a specific type of web content, enter it in the Web category field.
Port	To view logs associated with a specific port, enter it in the Port field; this filter shows logs that contain this port for either source or destination traffic. To view logs associated with a specific source port, click the source/destination button and enter it in the Source field. To view logs associated with a specific destination port, click the source/destination button and enter it in the Destination field.
Application	To view logs associated with a specific application, select an option from the Application menu.
Log severity	To view logs by severity, select an option (**EMERGENCY**, **ALERT**, **CRITICAL**, **ERROR**, **WARNING**, **NOTICE**, **INFO**, or **DEBUG**). The default displays logs with all severity types.
Overlay	To view logs associated with a specific business intent overlay (BIO), enter it in the Overlay field.
Flow ID	To view logs associated with a specific flow ID, enter the ID in the Flow ID field.
VRF	To view logs associated with a specific routing segment (VRF), select it from the VRF menu; this filter shows logs associated with selected segment for both source and destination traffic. To view logs associated with a specific routing segment for source traffic, click the source/destination button and select it from the Source menu. To view logs associated with a specific routing segment for destination traffic, click the source/destination button and select it from the Destination menu.
Protocol	To view logs associated with a specific network protocol, select it from the Protocol menu.
Description	To view logs associated with certain keywords in the log description, enter the keywords in the Description field.
App protocol	To view logs associated with a specific application-layer protocol, enter it in the App protocol field.

Filter	Description
IP reputation	To view logs associated with a specific IP reputation score for source and destination traffic, enter it in the IP reputation field. To view logs associated with a specific IP reputation score for source traffic, click the source/destination button and enter it in the Source field. To view logs associated with a specific IP reputation for destination traffic, click the source/destination button and enter it in the Destination field.
Zone	To view logs associated with a specific network zone for source and destination traffic, select it from the Zone menu. To view logs associated with a specific network zone for source traffic, click the source/destination button and select it from the Source menu. To view logs associated with a specific network zone for destination traffic, click the source/destination button and select it from the Destination menu.
Event	To view logs associated with a specific Firewall protection profile event, enter keywords in the Event field.
Signature ID	For IDS/IPS events, enter the unique identifier for a signature profile to view logs associated with it.
HTTP method	To view logs associated with a specific HTTP request method, select an option from the HTTP method menu (**GET**, **POST**, **PUT**, **DELETE**, **PATCH**, **HEAD**, **OPTIONS**, **CONNECT**, or **TRACE**).

Security Log Attributes

The Security Log Viewer tab displays security logs in a table, with each column representing an attribute from the logs. The following table describes the information shown, including the security log type(s) that each attribute applies to.

Column (attribute)	Description	Applicable to Log Type
Log type	Type of security log (IDS/IPS, Firewall, Firewall protection profile, or Stateful WAN drops).	All
Appliance	Name of the appliance associated with the recorded log.	All
Log severity	Severity level of the logged event (EMERGENCY, ALERT, CRITICAL, ERROR, WARNING, NOTICE, INFO, or DEBUG).	All
Details	Click the info icon (**i**) to open the Log details dialog box, which contains additional attributes included in the log.	All
Start time	Timestamp (date and time) the event started.	Firewall
Duration	Total time the event lasted.	Firewall
End time	Timestamp (date and time) the event ended.	Firewall

Column (attribute)	Description	Applicable to Log Type
Classification	Security or threat classification assigned to the event.	IDS/IPS and Firewall protection profile
Zone	Network zone where the event occurred.	Firewall protection profile
Protocol	Network protocol in use when the event occurred.	All
Event	Short description of the logged event.	Firewall protection profile
Application	Application running when the event occurred.	Firewall
Zone ID	Identifier for the network zone associated with the logged event.	Firewall protection profile
Description	Short description of the logged event.	Stateful WAN drops
Src VRF	Routing segment where the traffic originated.	IDS/IPS
Threshold ID	Identifier for the threshold rule that was evaluated or triggered.	Firewall protection profile
Src domain	Source domain name of the originating host.	Firewall and IDS/IPS
Dst domain	Destination domain name of the target host.	Firewall and IDS/IPS
Dst VRF	Routing segment where the traffic was headed.	IDS/IPS
Metric	Metric setting for the Firewall Protection Profile threshold.	Firewall protection profile
Src port	Source port number used by the originating host.	Firewall, IDS/IPS, and Stateful WAN drops
Signature ID	Unique identifier for the signature profile in use during the logged IPS/IDS event.	IDS/IPS
Threshold value	Configured FPP threshold value during the logged Firewall protection profile event.	Firewall protection profile
Priority	Priority level assigned to the event.	IDS/IPS

Column (attribute)	Description	Applicable to Log Type
Threshold action	Action taken when FPP threshold was breached.	Firewall protection profile
Dst port	Destination port number targeted by the traffic.	Firewall, IDS/IPS, and Stateful WAN drops
Timestamp	Date and time when the security log was recorded.	IDS/IPS, Firewall protection profile, and Stateful WAN drops
Src zone	Network zone associated with the source of the traffic.	Firewall and IDS/IPS
Dst zone	Network zone associated with the destination of the traffic.	Firewall and IDS/IPS
Action	Action taken by the appliance for this traffic (for example, allow, deny, drop).	Firewall
VRF name	Name of the routing segment (VRF) where the event occurred.	Firewall and Firewall protection profile
Mode	Deployment mode the appliance was operating in during the event (IDS, IPS-Performant, or IPS-Inline).	IDS/IPS
Reason	Indicates why the action was taken by the appliance for the traffic involved in this event (flow create, flow end, implicitly deny).	Firewall
Flow ID	Unique identifier for the traffic flow associated with the event.	Firewall
Overlay	Business Intent Overlay (BIO) assigned to the traffic involved in the event.	Firewall
Dst address	Destination IP address involved in the event.	Firewall and Stateful WAN drops
Src address	Source IP address involved in the event.	Firewall, Firewall protection profile, and Stateful WAN drops

Column (attribute)	Description	Applicable to Log Type
Length	Length of the dropped packet in bytes.	Stateful WAN drops
Input interface	Network interface on which the traffic was received.	Stateful WAN drops and Firewall
Generator ID	Identifier of the sensor or rule generator that produced the alert.	IDS/IPS
Revision ID	Revision or version number of the signature or rule that matched.	IDS/IPS
Web reputation	Reputation score associated with the FQDN/URL involved in the event.	Firewall
DSCP	Differentiated Services Code Point from the IP packet header.	Firewall
TCP flags	TCP header flags set in the connection (for example, SYN, ACK, FIN).	Firewall
NAT src IP	Source IP address after Network Address Translation.	Firewall
NAT src port	Source port number after Network Address Translation.	Firewall
User device	Identifier or hostname of the user device that generated the traffic.	Firewall
HTTP response	HTTP response status code returned by the server (for example, 200, 404).	Firewall
Flow drop code (secure web services)	Code that indicates why a flow was dropped: 0x0—Flow was allowed. 0x01—Flow was denied because the source IP address has a bad reputation. 0x02—Flow was denied because the destination IP address has a bad reputation. 0x04—Flow was denied because the URL has a bad reputation. 0x08—Flow was denied because the URL web category is blocked. 0x10—Flow was denied because the URL is blocked.	Firewall
RX packets	Number of packets received for the logged flow or session.	Firewall
Username	Authenticated username associated with the traffic, if available.	Firewall
App protocol	Specific application-layer protocol detected (for example, HTTP/2, MQTT).	Firewall
Src role	Role assigned to the source host in role-based policies.	Firewall

Column (attribute)	Description	Applicable to Log Type
Dst role	Role assigned to the destination host in role-based policies.	Firewall
TX packets	Number of packets transmitted for the logged flow or session.	Firewall
RX octets	Number of bytes received for the flow or session.	Firewall
TX octets	Number of bytes transmitted for the flow or session.	Firewall
Ingress VRF ID	Identifier for the VRF on the ingress interface.	Firewall
Egress VRF ID	Identifier for the VRF on the egress interface.	Firewall
Host	Hostname or IP address of the device that generated or is associated with the log.	Firewall
Tag	User-defined or system-defined tag applied to the event for grouping or filtering.	Firewall
Direction	Traffic direction relative to the device (ingress or egress).	Firewall
Threat	Named threat or indicator of compromise associated with the event.	IDS/IPS
Src IP reputation	Reputation score associated with the source IP address.	Firewall
Dst IP reputation	Reputation score associated with the destination IP address.	Firewall
Output interface	Network interface on which the traffic was sent out.	Firewall
HTTP method	HTTP request method used (for example, GET, POST, PUT).	Firewall
Web category	Category assigned to the FQDN/URL (for example, Shopping, Social Networking, Business, and so on). For information about the web categories, visit https://brightcloud.com/tools/change-request.php, and then click the **Web Category Descriptions** link.	Firewall

Monitoring > Performance

The options under **Monitoring > Performance** focus on performance monitoring related to availability, IP SLA, internet breakout, and applications.

Availability

Monitoring > *Performance* > *Availability*

Use the Availability tab to view HPE Aruba Networking SD-WAN infrastructure availability data measured as a percentage where uptime (total time minus downtime) is divided by total time.

NOTE: To use the Availability feature, you must enable New Stats Collection. You should also discontinue legacy stats collection. For more information, see Stats Collector Configuration.

The Availability tab provides tables for five categories of availability data: Network Role Availability, Service Availability, Transport Availability, Interface Availability, and Reachability. For descriptions and details, see Availability Tables below.

By default, percent availability is shown as an aggregate of up to 50 reachable appliances you select in the appliance tree. To view percent availability by individual appliances, click **Show by Appliance**.

Reporting intervals are based on day, week, or month based on your selection. Data is not presented in the last 24-hour rolling format, but according to a calendar day, calendar week (Sunday through Saturday), or calendar month. Data will not yet be available on the tab if the entire day, week, or month selected has not elapsed.

Availability data provided on this tab is calculated based on the availability time zones and business hours configured for your appliances. If the business hours of operation are currently set to 24 hours, you might want to adjust them so that availability data better reflects your business operations. For example, perhaps you use a generator to power your network. At night, the generator shuts off. You would want to exclude those downtimes from availability calculations. Click **Availability Time Settings** in the tab header to configure each appliance's business hours and time zone.

NOTE: If you leave the business hours set for 24-hour availability, one minute of downtime is one minute over 24 hours. If you set business hours to 8 hours, then one minute of downtime is over those 8 hours. That one minute of downtime counts as a higher percentage because you are only calculating it over 8 hours of business hours versus 24 hours.

Currently, this feature does not monitor availability for unified fabric tunnels.

This feature does not calculate availability for clusters/sites or EdgeHA pairings. Appliances are treated as individual entities.

To export the data displayed on the tab to a CSV file, click **Export**.

Availability Tables

On the Availability tab, availability tables for Network Role, Service, Transport, Interface, and Reachability (that is, between Orchestrator and your appliances) present their respective uptime statistics. To display availability data for one or more of these tables, click the corresponding buttons at the top of the tab.

IMPORTANT: Renaming interface labels affects the calculation of availability statistics.

Instances of "No Data" in the tables indicate that no representative data exists. The following two examples assume that Show by Appliance is selected at the top of the tab.

- If “No Data” is displayed in the Local Breakout (Internet) column of the Service Availability table for an overlay, the overlay does not have local breakout as one of its transport options.
- If “No Data” is displayed in a service-related column (such as Zscaler Cloud) in the Service Availability table for an appliance, the appliance does not subscribe to the service.

Table	Description
Network Role Availability	Provides availability data by appliance network role (hub or spoke) for SD-WAN Fabric, Local Breakout (Internet), and Services. Appliance time zone and business hour information are displayed if Show by Appliance is selected at the top of the tab. Chart and info icons are displayed as well. Click a chart icon for SD-WAN Fabric, Local Breakout (Internet), or Services to view network role availability trends. Click the info icon for Services to open the Tunnels Availability dialog box, which shows the availability of tunnels related to services. Click a chart icon for a listed tunnel in this dialog box to view tunnel availability trends.
Service Availability	Provides service availability data by overlay for SD-WAN Fabric, Local Breakout (Internet), and individual services as defined in the BIO, such as HPE SSE Cloud and Zscaler Cloud. **NOTE:** The SD-WAN Fabric and Local Breakout (Internet) availability percentages are typically higher than underlay availability percentages. Chart and info icons are displayed in the Service Availability table if Show by Appliance is selected at the top of the tab. Click a chart icon for SD-WAN Fabric, Local Breakout (Internet), or a service to view service availability trends. Click the info icon for a service to open the Tunnels Availability dialog box, which shows the availability of tunnels related to the service. Click a chart icon for a listed tunnel in this dialog box to view tunnel availability trends.

Table	Description
Transport Availability	Provides transport availability data by appliance (and interface label if Show by Label is selected) for Underlay, Pass Through, and individual services as defined in the BIO, such as HPE SSE Cloud and Zscaler Cloud. To segregate data by interface label, select the **Show by Label** check box. **NOTE:** For interfaces with multiple IP addresses, statistics are provided for tunnels associated with physical interfaces, but not tunnels associated with alias interfaces. Chart and info icons are displayed in the Transport Availability chart if Show by Appliance is selected at the top of the tab. Click a chart icon for Underlay, Pass Through, or a service to view transport availability trends. Click the info icon for a service to open the Tunnels Availability dialog box, which shows the availability of tunnels related to the service. Click a chart icon for a listed tunnel in this dialog box to view tunnel availability trends.
Interface Availability	Provides availability data by appliance interface label. The appliance references the interface label to check interface status. (Policies are specified according to labels.) Interface availability does not distinguish between an administrator setting an interface down and the interface being down. It also does not reflect the uplink status for the radio component of an LTE modem. **NOTE:** For interfaces with multiple IP addresses, statistics are provided for the physical interface, but not individual IP aliases. Chart icons are displayed in the Interface Availability table if Show by Appliance is selected at the top of the tab. Click a listed interface's chart icon to view interface availability trends.

Table	Description
Reachability	Indicates appliance connectivity to the Orchestrator by way of Portal WebSocket and Orchestrator WebSocket. Appliance status is up if the appliance is connected to either the Orchestrator WebSocket or the Portal WebSocket. It also provides an Overall Reachability percentage if an Orchestrator OR Portal connection is available. Chart icons are displayed in the Reachability chart if Show by Appliance is selected at the top of the tab. Click a listed appliance's chart icon for Overall Reachability, Via Portal WebSocket, or Via Orchestrator WebSocket to view appliance reachability trends.

Configure Color Codes for Availability Thresholds

You can set thresholds for each availability category (Services, Transport, Interface, Network Role, and Reachability) based on your business requirements. Click the gear icon to open the Availability Thresholds dialog box. Thresholds determine the color coding of the availability statistics detail for any category. By default, green indicates 99.99% or greater, red indicates 99.00% or less, and yellow indicates any percentage in between. If desired, change these percentages, and then click **Save**. To return threshold percentages to their original default percentages, click **Restore Defaults**.

AppExpress Summary Tab

Monitoring > Performance > AppExpress Summary

From this tab you can view and monitor information about the AppExpress groups you have configured. There are buttons on this tab that allow quick access to Application Definitions, AppExpress Groups Tab, and Apply AppExpress Groups Tab. For detailed information about this feature including how to configure it, see About AppExpress.

The following table describes the information displayed on the AppExpress Summary tab. You can filter the view by selecting a specific appliance in the appliance tree.

Column	Description
Appliance	The name of the appliance on which the flow originates.
Application	The name of the application being monitored.
Group	The name of the AppExpress group that the appliance is part of. The group determines the available transport paths for the application and the frequency that the Ping QoE and User QoE are updated.

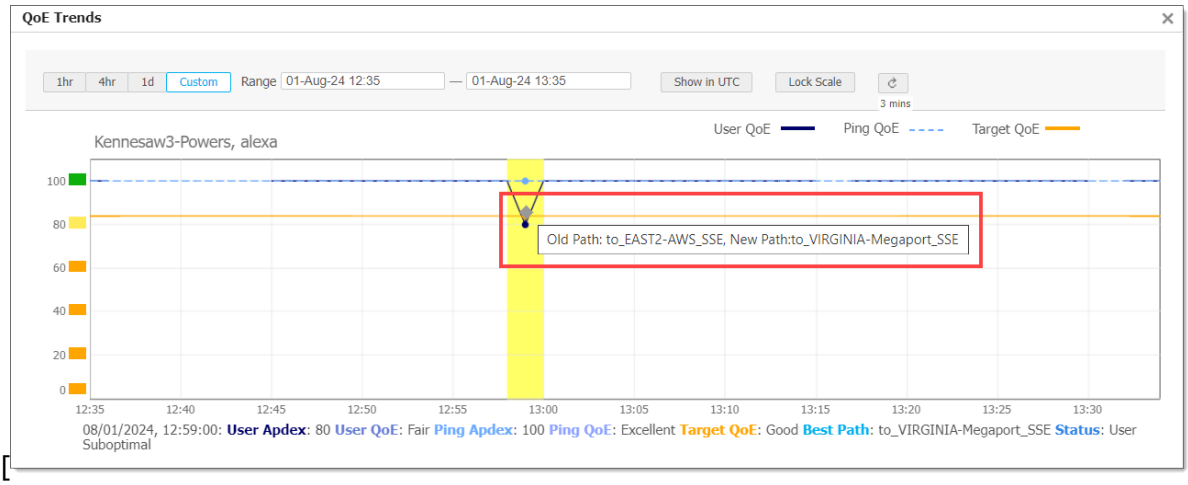
Column	Description
Target QoE	The quality of experience (QoE) target for the application. This target is what the system assesses the Ping QoE and User QoE against to determine the best path for each application.
User QoE	An Apdex score that is based on actual user flows. When users start using an application that is monitored, the EdgeConnect begins measuring the user flows to that application. This is a real-time QoE measure and is the basis for all the AppExpress decisions to redirect flows from one transport to another, for the application. To view trends for User QoE, click the chart icon in this column to open the QoE Trends chart. For information about the data shown on the chart, see QoE Trends.
Current Transports	The transport(s) that the application is currently using. Current transports can include third-party IPSec tunnels, local breakout labels, and SD-WAN peers.
Application Performance trend	The trend of how the traffic flow for this application has performed over time.

Column	Description
Status	Indicates if the traffic flow for this application on this appliance is meeting the optimal targets set for this application. Possible statuses: Initial – AppExpress is in provisioning and will soon enter the Learning state. Learning – AppExpress is in startup and needs time to accumulate Ping QoE scores. Ping Optimal – No User flows have been observed; however, AppExpress has found at least one Ping path that meets the Target QoE. When a user flow is observed, AppExpress will route the flow to the path shown in the Current Transports column. Ping Suboptimal – No User flows have been observed and there are no paths whose Ping QoE meets the Target QoE. User flows will be routed to the path with the highest Ping QoE. User Optimal – User flows have been observed and are meeting the Target QoE. New user flows will be routed to the path shown in the Current Transports column until the User QoE no longer meets the Target QoE. User Suboptimal – User flows have been observed; however, they are not meeting the Target QoE over the Current Transports. At the next User QoE interval, AppExpress will try the next best transport path sorted by Ping QoE score. Fallback – No user flows have been observed and pings to all paths are failing. AppExpress will revert to standard BIO handling for new flows.
Ping QoE	An Apdex score that is based on synthetic polling. Failed connections count toward the F boundary. Synthetic polling consists of pings that go out from an EdgeConnect appliance through the loopback interface for the appliance across each path. From these pings, the system determines the Ping QoE for each flow. To view trends for Ping QoE, click the chart icon in this column to open the Ping QoE Trends chart. For information about the data shown on the chart, see Ping QoE Trends.

QoE Trends

To view QoE trends (User QoE and Ping QoE) for traffic for an application that is monitored and steered using AppExpress, click the chart icon in the User QoE column for an appliance to open the QoE Trends chart. This chart displays trends for up to 20 applications.

- AppExpress must be enabled for the application for data to appear in the chart.
- The vertical axis of the chart is labeled using the EdgeConnect QoE ratings. These are fixed values of Excellent: 100-93, Good: 92-84, Fair: 83-64, and Best-Effort (<68).
- Displays the per-minute average of User QoE and Ping QoE for the current path.
- If the User QoE falls below the Target QoE, the flow is put on the next best path. When the path changes, the chart visually reflects the change with a yellow vertical band and a gray diamond. If you hover over the gray diamond, a text box indicates the Old Path and the New Path, as shown in the following figure.



Ping QoE Trends

To view Ping QoE Trends for all paths that are available on an appliance for a selected application, click the chart icon in the Ping QoE column for an appliance to open the Ping QoE Trends chart. This chart displays the trends associated with one application and one appliance only.

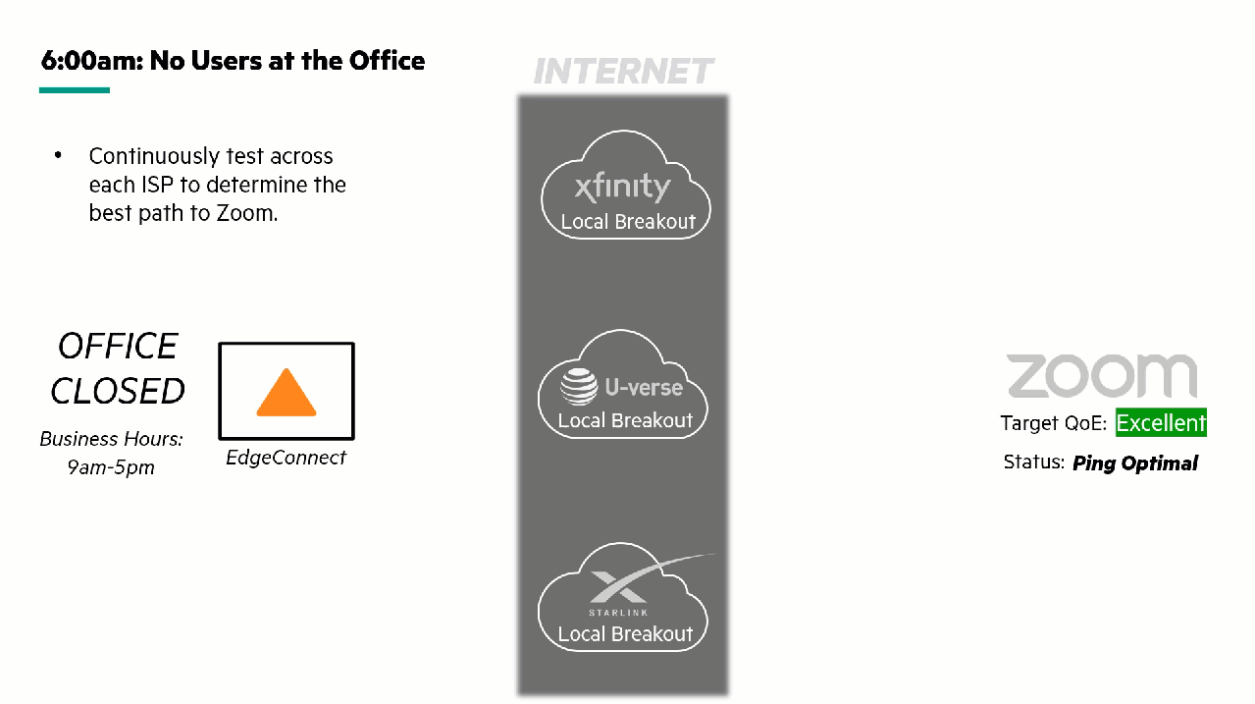
- AppExpress must be enabled for the application for data to appear in the chart.
- The vertical axis of the chart is labeled using the EdgeConnect QoE ratings. These are fixed values of Excellent: 100-93, Good: 92-84, Fair: 83-64, and Best-Effort (<68).
- Displays the per-minute average of Ping QoE for all possible paths on an appliance.
- The trends for each path are displayed as a separate chart within the window.
- If the Ping QoE falls below the Target QoE for a path, the chart is highlighted in yellow.

About AppExpress

For most enterprises, a handful of high-profile, high-impact applications drive the business. AppExpress allows you to optimize the user experience for high-impact applications. With AppExpress you can monitor the traffic flow for up to 50 applications and leverage synthetic polling and real-time user traffic observations to intelligently steer traffic. AppExpress automatically selects the best path for each of the 50 applications. See Determining the Best Flow for an Application Path and Transport Types for more information on how AppExpress does this. AppExpress works for internal and cloud-based applications.

AppExpress sends synthetic probes across all available paths—local breakout, backhaul, and third-party service tunnels—to applications and it determines which path appears to have the best latency and is most robust. AppExpress then places flows for applications on the best paths based on criteria that are set for each application.

The following example shows how AppExpress works for a common application, Zoom, throughout a typical business day.



Prerequisites

Before you begin using the AppExpress feature you must:

- Configure a loopback interface.
  - A loopback interface is required for AppExpress to function because it is the source of the synthetic probes that are sent to applications across backhaul and third-party tunnels. See Interface Labels and navigate to **Configuration > Overlays & Security > Interface Labels** in Orchestrator to create a LAN-side “LOOPBACK” label.

  - See Loopback Orchestration and navigate to **Configuration > Networking > Loopback Orchestration** in Orchestrator.

- Configure Stats Collector.
  - Distributed Stats Collector is required for AppExpress reporting and monitoring to fully function.
  - See Stats Collector Configuration.
- Enable AppExpress for each application that you want to monitor and steer.
  - See Application Definitions and navigate to **Configuration > Templates & Policies > Applications & SaaS > Application Definitions** in Orchestrator.
  - When enabling, select the **Monitor and Steer** option.
- Create AppExpress groups and add AppExpress applications to the groups.
  - If AppExpress is enabled for an application, but it isn't added to an AppExpress group then only monitoring of the application takes place. If you want AppExpress to also steer the application traffic, it must be part of an AppExpress group.
  - See AppExpress Groups Tab, and Apply AppExpress Groups Tab and navigate to **Configuration > Templates & Policies > Applications & SaaS > AppExpress Groups** in Orchestrator.

Determining the Best Path for an Application Flow

AppExpress uses three measures to determine the best path for an application flow.

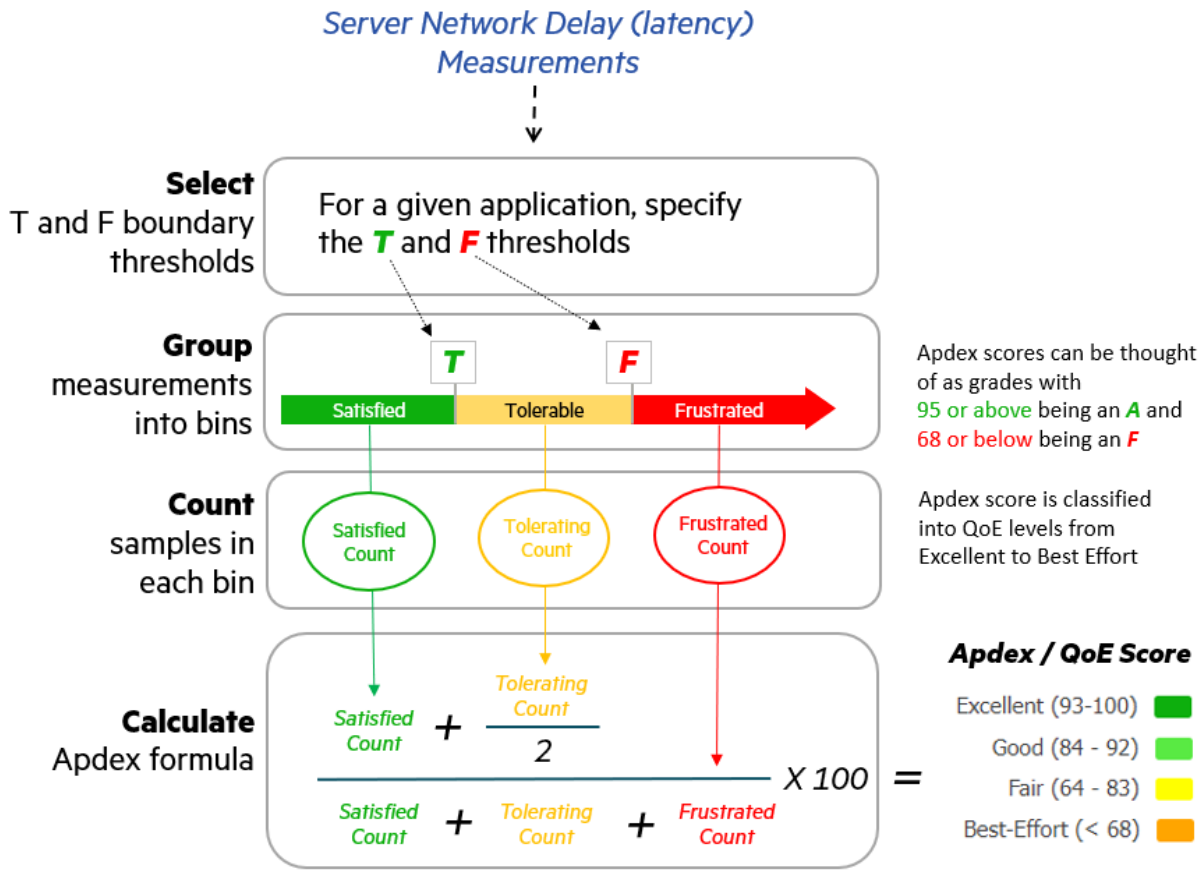
Target QoE

This is a user's target for an application, or how a user wants to experience an application. For example, if a goal is for users to have an excellent experience for Zoom calls, you would set the Target QoE to Excellent. For each application that is monitored through AppExpress, you must set a Target QoE (quality of experience) measure. The Target QoE is used to determine whether performance for a flow is optimal or not.

The Target QoE is calculated using Apdex. Apdex (Application Performance Index) is an open standard that defines a method to report, benchmark, and rate software application performance.

- Apdex normalizes a large sample set of measurements into a score between 0 and 100.
- The measurements are sorted into three bins:
  - Satisfied: Users are happy with their quality of experience.
  - Tolerable: Users aren't happy but aren't bothered enough to complain.
  - Frustrated: User experience is impacted; users are likely to complain.
- The point at which the performance of an application moves from Satisfied to Tolerable is the T boundary threshold.
- The point at which the performance of an application moves from Tolerable to Frustrated is the F boundary threshold.

- For Target QoE, failed User flows and failed synthetic pings contribute to the Frustrated bin.
- Apdex is calculated for flows once at the start of the flow.
- See https://www.apdex.org for more information about Apdex.



In AppExpress, the T and F boundary thresholds are the two User Experience Thresholds that you set when you add a new application definition. To do this, navigate to **Configuration > Templates & Policies > Applications & Saas > Application Definitions** and click **+Add New Application**. For more information about application definitions and setting the QoE measures for an AppExpress application, see Application Definitions.

Ping QoE

The Ping QoE is an Apdex score that is derived from synthetic probes (ICMP echo-request/response, TCP connect, HTTP, or HTTPS) that are sent across all available paths. For each AppExpress application, the system compares the Ping QoE against the Target QoE, to determine which path best meets the Target QoE, and then the flows for that application are put on that path.

User QoE

When users begin using an application, the system gathers data based on the real-time flows for the application. The User QoE is an Apdex score that is derived from these observed user flows. The system compares the User QoE against the Target QoE, and if the real-time flows for an application are no longer meeting the Target QoE, the system begins hunting for a different path that can meet the Target QoE. This process is continuous while an application is in use.

NOTE: AppExpress uses LAN to WAN flows to calculate User QoE and to gather application performance data. TCP and UDP flows are analyzed for User QoE determination. ICMP and other IP types are not supported by AppExpress. Also, EdgeHA flows are ignored by AppExpress.

Transport Types

There are three transport types that are handled by AppExpress.

Local Breakout

All Primary labels are considered. For information on Primary labels, see the Break Out Locally Using These Interfaces, Available Interfaces, and Link Selection section in Business Intent Overlays.

Third Party Service Tunnels

This includes SSE integrations such as Zscaler or Netskope. Services created using Service Orchestration are also included. All primary tunnels are considered. This means that AppExpress tries both the primary and secondary POPs (point of presence) for Zscaler, Netskope, etc. See the configuration information for Zscaler and Netskope.

Backhaul via BIOs to Hubs

AppExpress selects the backhaul peer based on the lowest peer priority configured, and it does not consider route metric or administrative distance. When no peer priority is configured, no paths are searched. AppExpress does not support passthrough paths. Up to 6 backhaul peers are supported.

IP SLA Summary

Monitoring > Performance > IP SLA Summary

The IP SLA Summary tab provides an at-a-glance view of IP SLA performance across all selected appliances. It summarizes latency, packet loss, availability, and up/down state for each monitored target, with results consolidated by appliance, interface, and target URL, so you can quickly assess the health of the destinations the IP SLA subsystem monitors.

To view or edit IP SLA rule configuration, click the **IP SLA** link at the top of the tab.

Time Range

Use the time-range selector at the top of the tab to control the aggregation window for the displayed statistics:

- 1hr - Aggregates statistics over the previous hour
- 4hr - Aggregates statistics over the previous 4 hours
- 1d - Aggregates statistics over the previous day

- 7d - Aggregates statistics over the previous 7 days
- Custom - Specify a start and end date/time

When you select **Custom**, a calendar opens so you can define an exact reporting window:

- Start and End fields - Set the beginning and end of the time range. Click a field to open the calendar picker.
- Date selection - Use the month and year navigation arrows to move between months, then click a day to set the date.
- Time selection - Set the hour and minute for the selected start and end dates so the window begins and ends at a precise time.
- Apply - Confirm your selection to refresh the table with data for the chosen range. The end date/time cannot be earlier than the start date/time, and the range cannot extend into the future.

Use the Granularity drop-down menu to set the reporting interval (for example, Hour).

Search

The Search field is located in the toolbar at the top of the tab, on the left side of the grid header directly above the table. Use it to filter the table by appliance name, interface, tunnel, or address/URL.

Table Columns

Column	Description
Appliance	The EdgeConnect appliance running the IP SLA rule. Expand the appliance row to view results per source interface or tunnel and per individual target. The number in parentheses indicates the total number of monitored targets.
Source Interface / Tunnel	The WAN interface or tunnel used to poll the target address. Shown at the interface grouping level.
Address / URL	The destination IP address or URL being monitored. At the appliance or interface level, a comma-separated list of all targets is shown. Individual targets are listed in expanded child rows.
Latency (ms)	Average round-trip latency, in milliseconds, over the selected time range.
Loss %	Average packet loss percentage over the selected time range.
Times Down	The number of down events recorded during the selected time range.

Column	Description
Availability %	Percentage of time the target was available during the selected time range, as derived from the Availability feature.
Charts	Click the chart icon to display the trends chart.
Active	Indicates whether the IP SLA rule is currently active (Yes/No).
State	Current state of the target or rule. At the appliance or interface level, a summary is shown (for example, “6/6 down” means all 6 targets are down). Individual rows show Up or Down. Down states are highlighted in red; partial-down states (some but not all targets down) are highlighted in orange.
Up time (last)	Duration of the most recent up period.
Down time (last)	Duration of the most recent down period.

Expandable Rows

The table uses a hierarchical, expandable layout:

1. Appliance level - Shows aggregated statistics across all interfaces and targets for the appliance.
2. Interface / Tunnel level - Expand an appliance row to see results grouped by the source interface or tunnel.
3. Target level - Expand an interface row to see individual results for each monitored address or URL.

Click the expand arrow (>) next to an appliance or interface name to drill down. Use the expand/collapse-all control to open or close all rows at once.

State Indicators

- Red - All targets for the rule or interface are down (for example, “6/6 down” or “Down”).
- Orange - Some but not all targets are down (for example, “1/6 down”).
- Green / No highlight - All targets are up.

Internet Breakout Trends

Monitoring > Performance > Internet Breakout

This tab displays trends for internet breakout traffic for each overlay. You can view trends for certain data about the internet breakout links including latency, loss, jitter, MOS (mean opinion score), and auto by selecting these options from the tab header. Each metric is displayed in a separate chart for each overlay. Orchestrator determines a value for a metric, such as latency, by measuring the traffic in all the tunnels within each link and averaging the results across the link for that metric.

Internet breakout is traffic that is sent from a local branch directly to the internet, rather than going through an IP Sec tunnel or out to a data center and through a firewall to the internet. Internet breakout traffic routing is determined during business intent overlay (BIO) configuration. To configure internet breakout traffic for an overlay, select the overlay and navigate to **Breakout Traffic to Internet & Cloud Services > Link Selection**.

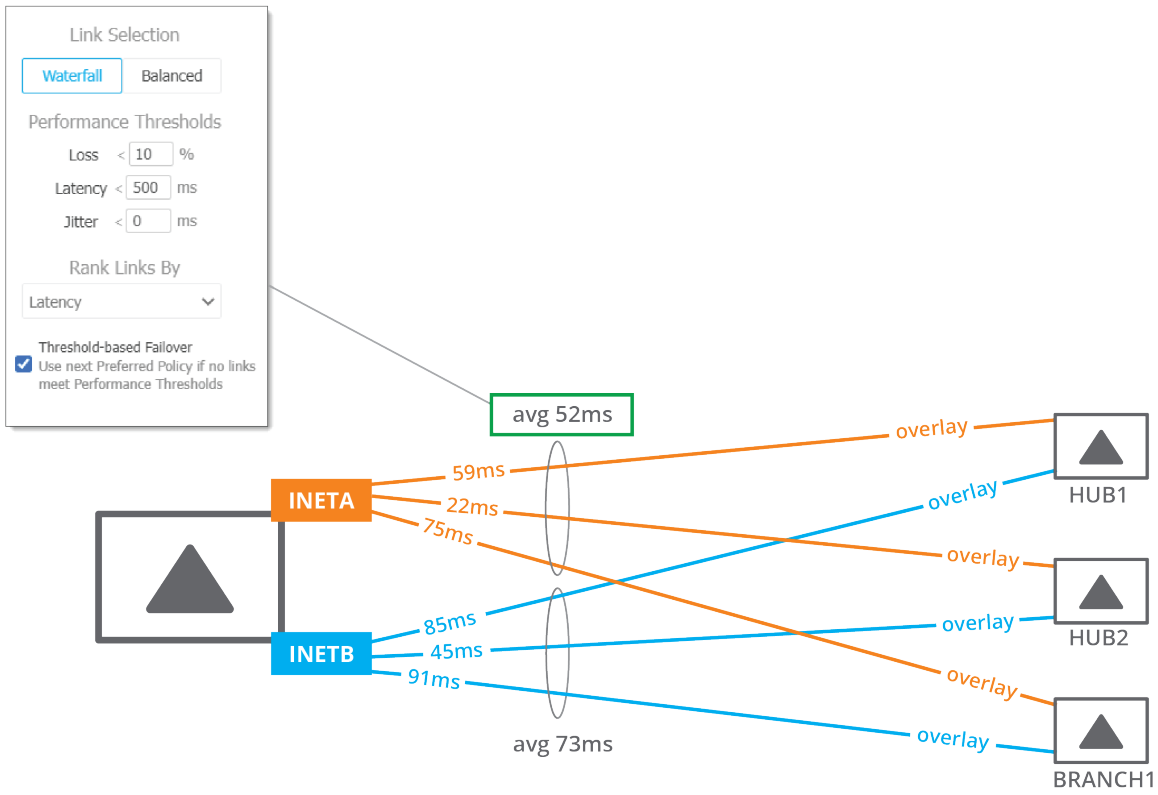
Internet Breakout Modes

There are two modes for managing internet breakout traffic, Waterfall and Balanced. For each BIO you choose which mode to use for internet breakout. Performance metrics for internet breakout traffic are captured and charted on the Internet Breakout tab when either mode is in place.

Waterfall

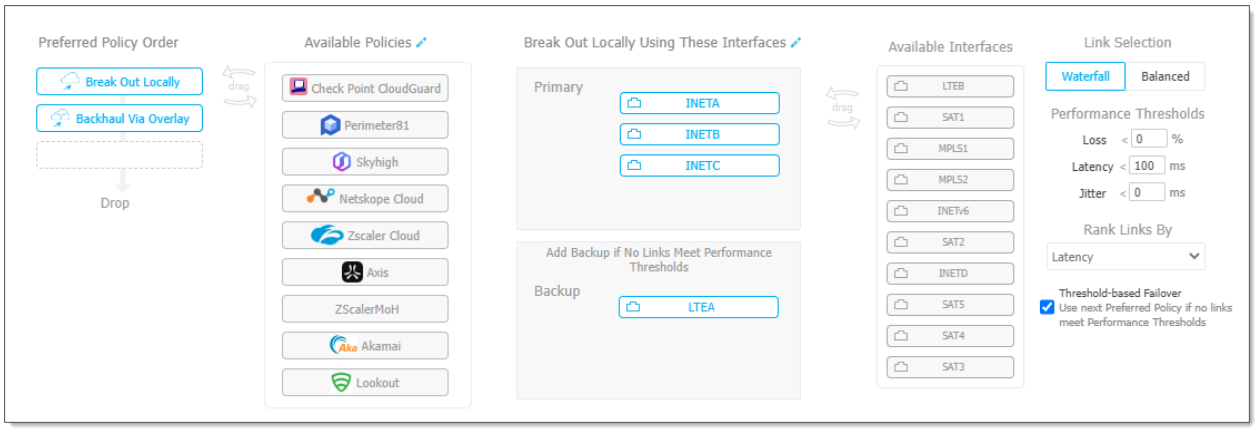
In waterfall mode, the system infers which internet breakout link is performing the best at that moment and it fills the link up with traffic until it reaches 80% then it routes the traffic to the next best link. Orchestrator uses the data gathered about the links through inference, the Rank Links By setting, and an algorithm to determine what is the best path to the internet through each of your labels. During link selection it also applies any performance thresholds you have set. If the primary link exceeds any of the thresholds, the system waterfalls the traffic to the next link.

The following figure shows an example of how Waterfall mode is applied and how it infers the best path to the internet.



In this example the following applies based on the Link Selection settings:

- INETA average inferred latency = 52ms
- INETB average inferred latency = 73ms
- INETA has the lowest inferred latency, so use INETA for internet breakout
- Route traffic to INETB when INETA reaches 80% or when INETA exceeds any performance thresholds
- Exclude any links that have an inferred loss greater than 10% or an inferred latency greater than 500ms
- If all links exceed the performance thresholds, use the next item in the Preferred Policy Order list. If there is no Preferred Policy entry beneath Local Breakout the traffic will be dropped.
  - In this example, if all labels exceed the performance threshold of 100ms of latency, the traffic will backhaul because “Backhaul Via Overlay” is listed beneath Local Breakout in the Preferred Policy Order.



Balanced

In balanced mode, Orchestrator uses weighted round robin and it distributes the traffic across the links evenly proportional to the amount of bandwidth on each interface. It determines the ratio based on the amount of bandwidth on each interface. You can set performance thresholds for loss, latency, and jitter, and if a link exceeds any threshold it is excluded from the available links.

Non-Default to Default Segment Internet Breakout

When all WAN interfaces are assigned to the Default segment, all internet-bound traffic exits through those interfaces. The system handles this behavior automatically, but the following considerations apply:

Key Points

- **Overlay control**: Overlays determine when to break out to the internet, regardless of the segment where traffic originates.
- **Firewall rules**: To allow traffic from non-default segments to use WAN interfaces in the Default segment, configure firewall rules that permit inter-segment traffic into Default.
- **Routing policy**: No additional inter-segment routing policy is required. Internet break-out handling is built in.
- **NAT behavior**: NAT on the WAN interface ensures return traffic flows back correctly to the originating segment.

Example Scenario

- Segment A sends traffic to the internet.
- All WAN interfaces are in the Default segment.
- Firewall rules allow Segment A Default segment.
- NAT ensures return traffic reaches Segment A.

Traffic Flow Diagram

[Segment A] —> [Default Segment WAN] —> Internet ˆ | | v Return Traffic via NAT

Configuration Steps

- Verify WAN segment assignment; ensure all WAN interfaces are assigned to the Default segment.
- Create firewall rules that allow traffic from non-default segments to the Default segment.
- Confirm NAT settings. NAT must be enabled on WAN interfaces for proper return traffic flow.
- Validate overlay policies to ensure overlays are configured for internet breakout as needed.

Best Practices

- Keep firewall rules minimal and specific to required traffic.
- Regularly audit NAT configuration on WAN interfaces.
- Monitor overlay breakout decisions for performance optimization.

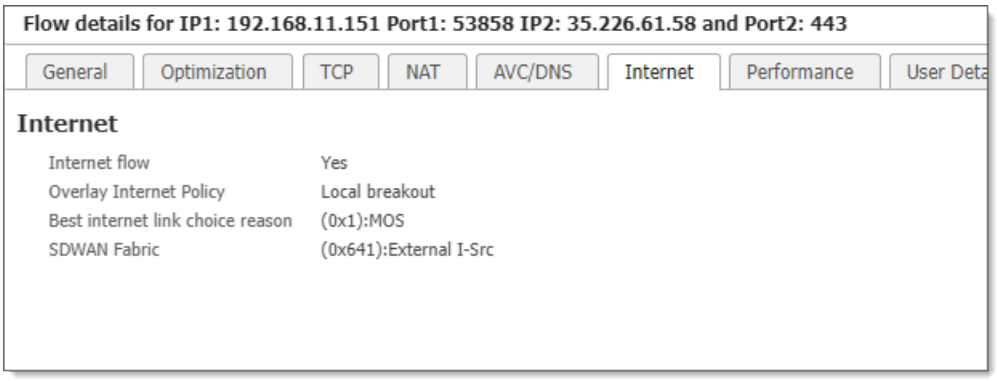
Troubleshooting

If traffic from non-default segments cannot reach the internet, check the following:

- Check firewall rules for inter-segment permissions.
- Verify NAT configuration on WAN interfaces.
- Confirm overlay breakout policies are active.

Additional Operational Notes

- Link Selection provides a session-affinity feature that pins all flows between an internal IP and an internet IP to the same Local Breakout label. Once an initial label is selected by the Link Selection mechanism, all flows between source and destination IPs stick to the same label until all flows between the two hosts are inactive for 2 minutes. Performance Thresholds override session-affinity.
- To see why a label was selected for a flow, refer to the Internet tab of the Flow Details. The “Best internet link choice reason” item shows why the Link Selection feature pinned a flow to a given label. In this example, “MOS” was configured in the Rank Links By section of the Link Selection settings.



The following table describes the possible Best internet link choice reasons.

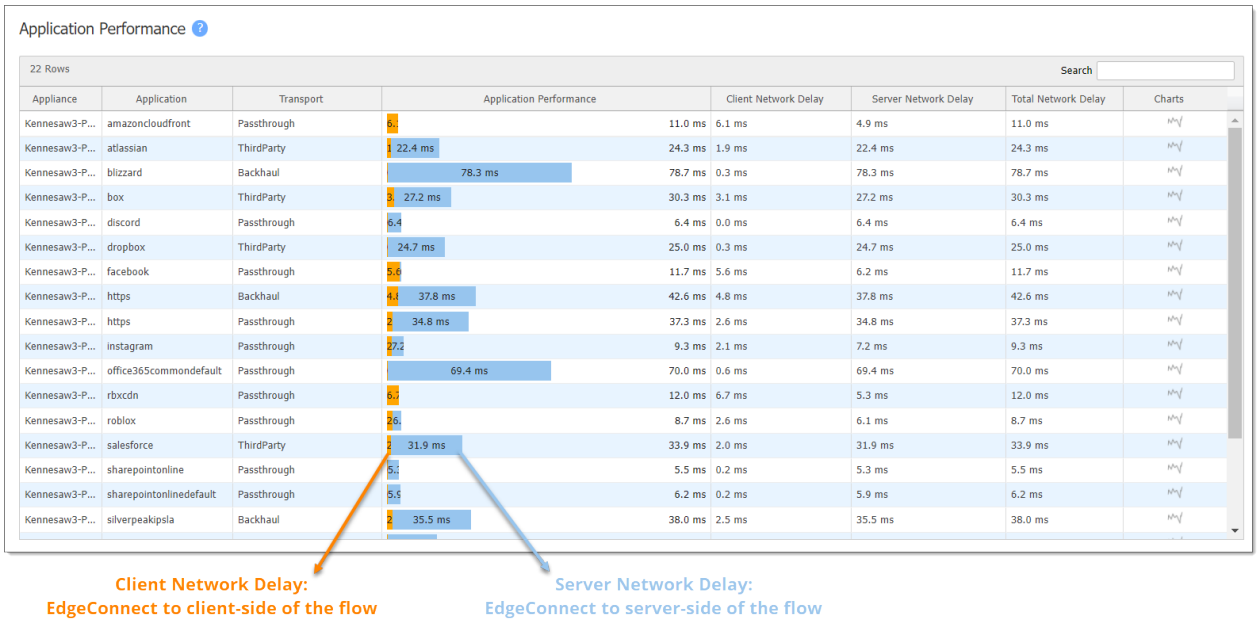
Reason	Description
MOS	Ranking of the links is done using the MOS (mean opinion score) quality attribute.
low-loss	Ranking of the links is done using loss.
Low-latency	Ranking of the links is done using the latency attribute.
Manual	No attribute is used to do the ranking, it follows the order configured by the user.
Session-affinity	The flow took the path of an earlier flow with the same src+Dst IP because that flow was started within the last 60s.
brownout-limit-crossed	One or more more links crossed the brownout threshold.
In-house-metric	Ranking of the links is done using an in-house (auto) metric
tunnel-down	The passthrough tunnel was momentarily down, so it was skipped.
No-primary-link	No primary links were found after ranking because either none were configured or all of the links crossed performance thresholds.
primary-bwfull	The bandwidth for all primary links crossed the performance threshold.
no-backup-link	No backup links were found after ranking because either none were configured or all of the links crossed performance thresholds.
bkup-bwfull	The bandwidth for all backup links crossed the performance threshold.
fallback-to-next-policy	All primary and backup links are browned out or blacked out, so it moves to the next preferred-policy.
all-links-blackout	All primary and backup links are blacked out.

Application Performance Summary

Monitoring > Performance > Application Performance

To display the summary view of the Application Performance tab, click the summary icon.

This view displays application performance data for up to 50 applications. It reflects a real-time measure of latency per flow. The Application Performance column displays latency measures for Client Network Delay (orange), which is from the EdgeConnect to the client-side of the flow, and for Server Network Delay (blue), which is from the EdgeConnect to the server-side of the flow.



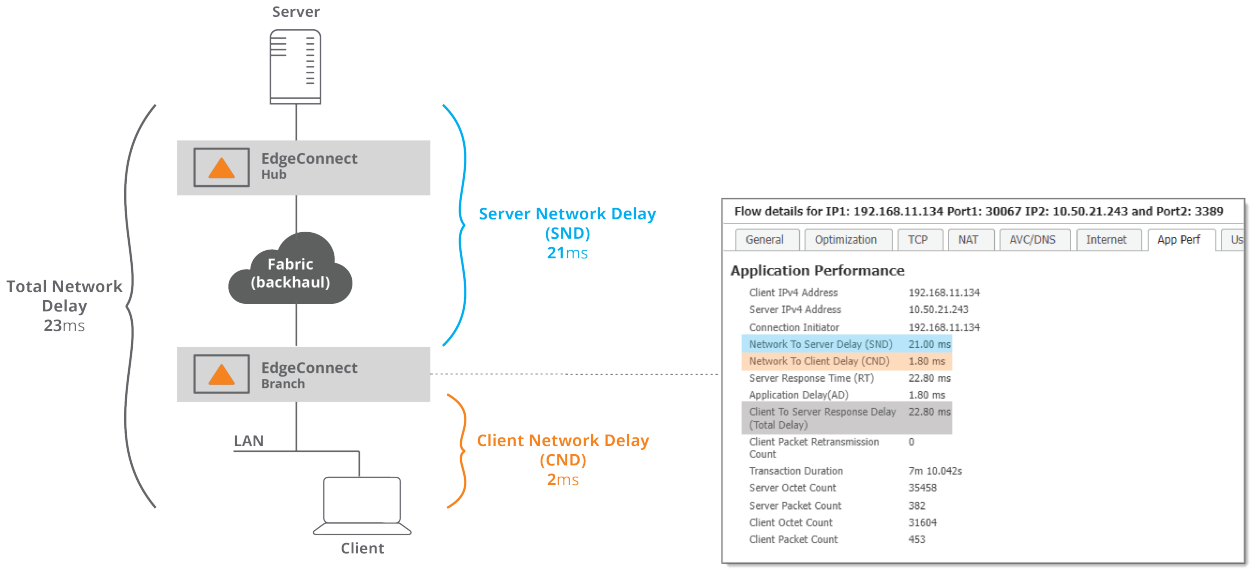
Latency is calculated using three different metrics. Each of these metrics are reflected in the summary view of the Application Performance tab:

- Client Network Delay (CND)
  - Represents the latency measured between the client and the EdgeConnect branch.
  - Corresponds to “Network to Client Delay (CND)” data on flow details.
- Server Network Delay (SND)
  - Represents the latency measured between the EdgeConnect hub and a server.
  - Corresponds to “Network to Server Delay (SND)” data on flow details.
- Total Network Delay
  - Represents the total latency measured.
    Client Network Delay + Server Network Delay = Total Network Delay
  - Corresponds to “Client to Server Response Delay (Total Delay)” data on flow details.

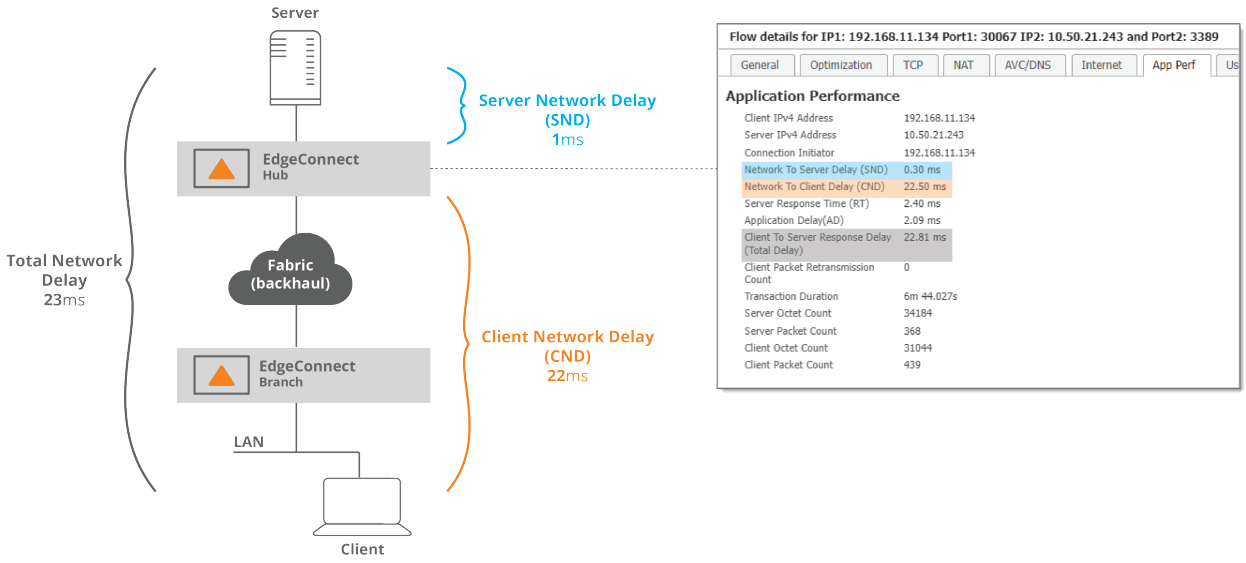
The latency metrics displayed on the summary view of the Application Performance tab are taken from the flow details for each flow. To view flow details, go to **Monitoring > Flows > Active & Recent Flows**, then click the info icon for any flow, and click **Performance** to view application performance data.

The Client Network Delay and Server Network Delay metrics differ when you are viewing data for an EdgeConnect at a branch versus data for an EdgeConnect at a hub that connects to a server.

When viewing data for an EdgeConnect at a branch, the Client Network Delay (orange) represents a smaller portion of the total delay. The graphic below depicts how each metric is measured for an EdgeConnect at a branch.



When viewing data for an EdgeConnect at a hub, the Client Network Delay (orange) represents a larger portion of the total delay. The graphic below depicts how each metric is measured for an EdgeConnect at a hub.



To view latency trends for an application over time, click the Historical Charts icon. The Application Trends tab opens and displays a chart for the application.

Application Performance Trends

Monitoring > Performance > Application Performance

To display the trends view of the Application Performance tab, click the trends icon.

This view displays application latency trends over time for up to 20 applications. Real-time latency measurements are charted for both Client Network Delay (orange), which is from the EdgeConnect to the client-side of the flow, and Server Network Delay (blue), which is from the EdgeConnect to the server-side of the flow.

Monitoring > Reporting

The options under **Monitoring > Reporting** focus on creating, managing, scheduling, and viewing Orchestrator reports.

Schedule and Run Reports

Monitoring > Reporting > Schedule & Run Reports

Use the Schedule & Run Reports tab to create, configure, run, schedule, and distribute reports.

You can specify what you want to include in your reports: appliances, data granularity (daily or hourly) for non-availability reports, traffic type, and types of charts to include. You can indicate whether the report should be regularly scheduled or be a single report with a custom time range. You can also specify email recipients for the report.

For availability reports, you can specify data granularity (daily or monthly). You can also indicate whether availability reports should include data for individual appliances or aggregated data (for up to 10 selected appliances). Availability reports are generated based on availability time settings you set up. For details, see Availability Time Settings.

Reports and statistics can help you isolate problems, investigate questions, or perform analyses. Orchestrator reports fall into two broad categories:

- Statistics related to network and application performance. These provide visibility into the network, enabling you to investigate problems, address trends, and evaluate your WAN utilization.
- Reports related to status of the network and appliances. For example, alarms; threshold crossing alerts; reachability between Orchestrator and the appliances; scheduled jobs; network role, service, transport, and interface availability; and appliance reachability.

By default, Orchestrator emails a preconfigured subset of charts every day in a report named Global Report.

- You can also view current data by clicking the icon associated with a chart listed on the Schedule & Run Reports tab. The icon associated with the Availability section opens the Availability tab, which shows availability data related to network role, service, transport, and interface, and appliance reachability data.
- To view previously generated reports residing on the Orchestrator server, click **View Reports** at the top of the tab.

Availability reports are generated separately from non-availability reports and are sent separately to email recipients as well.

The following table describes various elements on the tab. Use them to create, configure, and schedule your reports.

Field	Description
Name	Name of the report. To display a currently configured report, select the report from the drop-down list. To create a new report, click **New Report**, enter a unique report name in the Create New Report dialog box that opens, and then click **Save**. To remove a report, select the report from the drop-down list, and then click **Delete** on the Delete Report dialog box that opens. **NOTE:** Orchestrator provides a report named Global Report, which is preconfigured with a set of selected charts. This report is automatically sent daily to the email addresses listed in the Email Recipients field.
Email Recipients	Email addresses for one or more report recipients. Use commas or semicolons to separate multiple email recipients. To send a test email message or to configure a different SMTP server, navigate to **Orchestrator > Software & Setup > Setup > SMTP Server Settings**. If a test email message does not arrive within minutes, check your firewall.
Appliances in Report	Appliances to include in the report. By default, all appliances in your network are included. To include only specific appliances, select them in the appliance tree, and then click **Use Tree Selection**. **NOTE:** Whether you want to generate availability reports by individual appliance or as aggregated data, do not list more than ten appliances in this field.

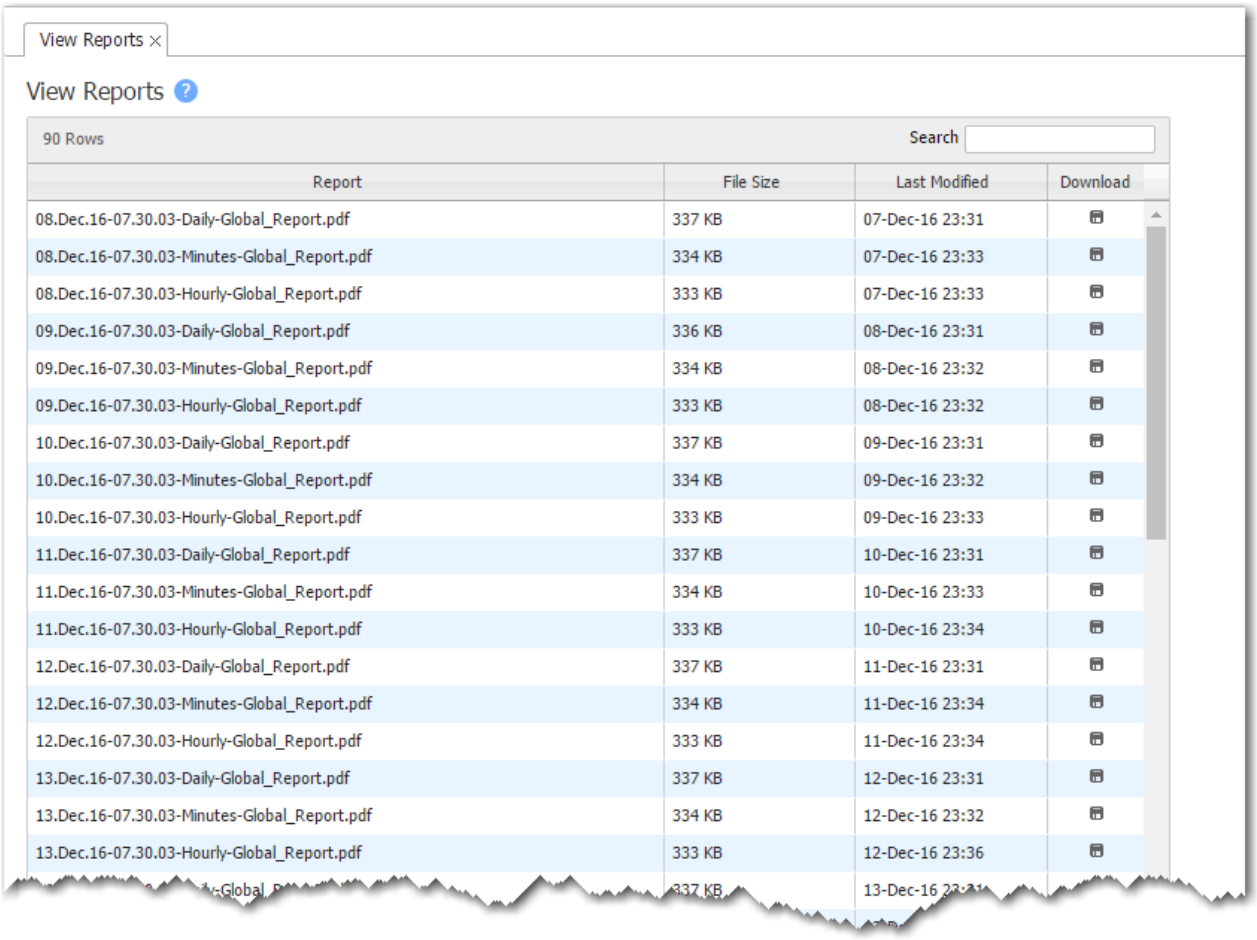
Field	Description
Data Granularity – Time Range	Granularity of the data to report. You can select one or both of the following: **Daily**: Specify the number of days. It is recommended that you specify no more than 14 days because increasing the scope uses additional memory. **Hourly**: Specify the number of hours (up to 24). **NOTE:** Daily and hourly ranges are not available for Health Map charts. Instead, these charts collect data for seven days for the day the report is run and for the six days preceding that date. For example, if you run the report on March 10, it includes data for March 4 through March 10. These daily granularity time ranges do not apply to availability reports.
Scheduled or Single Report	You can run your non-availability reports on a regularly scheduled basis or on demand as a single report. **NOTE:** Availability reports are generated based on availability time settings set up for your appliances. To set this up, navigate to **Orchestrator > Software & Setup > Setup > Availability Time Settings**. Select one of the following: **Run Scheduled Report**: Click the edit icon. The Schedule dialog box opens. Click **Daily**, **Weekly**, **Monthly**, or **Yearly**. Also specify the appropriate schedule criteria, and then click **OK**. **Run Single Report with Custom Time Range**: Specify the appropriate date and time range (for example, 2024-8-12 13:00 to 2024-08-19 13:00), and then click **Run Now**. To stop the report from running, click **Stop**. **TIP:** To specify the time zone for scheduled jobs and reports, navigate to **Orchestrator > Software & Setup > Setup > Timezone for Scheduled Jobs**.
Top	Maximum number of top reports to include in the report. Select **10**, **25**, **50**, **100**, or **1000** from the drop-down list.
Traffic Type	Type of traffic to include in the report. Select **Optimized Traffic**, **Pass-through Shaped**, **Pass-through Unshaped**, or **All Traffic** from the drop-down list.

Field	Description
Application Charts	Application charts to include in the report. Use the drop-down list at the top of this section to include data for **All Overlays**, **All Underlays**, or a listed overlay.
Tunnel Charts	Tunnel charts to include in the report.
Appliance Charts	Appliance charts to include in the report.
Availability	Availability charts to include in the report.
Availability Granularity	Granularity of availability data to include in the report (**Daily** or **Monthly**). Daily is the default setting. This granularity applies only to availability reports. Availability reports are generated based on the day or month indicated in the Scheduled or Single Report section of the tab. If Run Scheduled Report is used: For daily granularity, the report is generated for the previous day's availability data. For example, if the report is scheduled to run on June 17, the report shows data for June 16. For monthly granularity, the report is generated for the previous month's availability data. For example, if the report is scheduled to run on June 17, the report shows data for May. If Run Single Report with Custom Time Range is used: For daily granularity, the report shows availability data for all days within the start date and end date range. For monthly granularity, the report shows availability data for all days within the specified month. If the end of the month has not been reached, Orchestrator will not generate the report. The Orchestrator time zone is used to verify that the end of the month has been reached.
Reports	Select to produce availability reports **by Appliance** (for up to ten selected appliances) or as **Aggregated** data (for all selected appliances; this is the default setting).
Lock Scales for Trends	Indicates whether to automatically scale trend charts for specified scheduled reports or to lock scales. Toggle off to automatically scale; toggle on to lock scales.

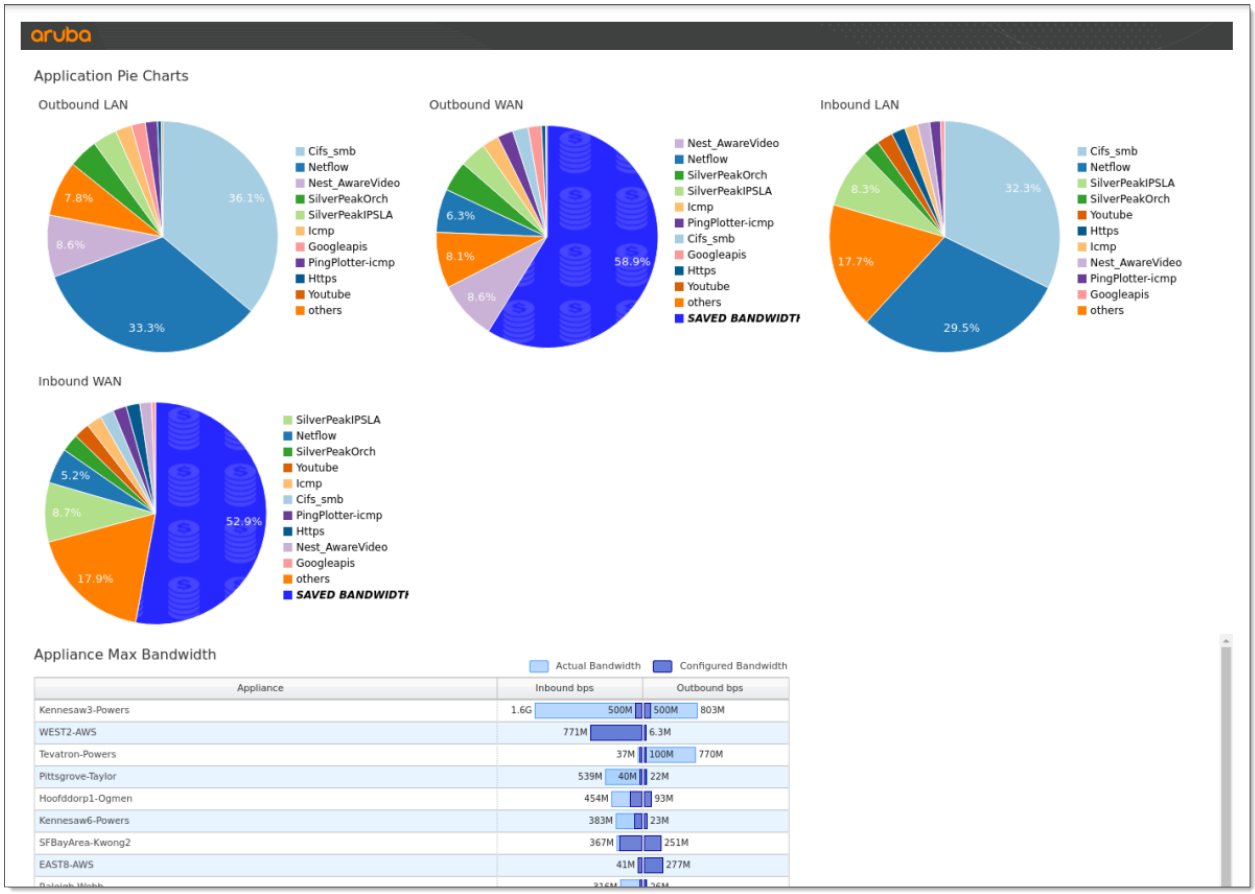
View Reports

Monitoring > Reporting > View Reports

Use this tab to **view** and **download** reports in PDF form. Reports can be filtered by keywords or sorted by **name**, **size**, or **date last modified**. These reports can also be emailed depending on the configuration set on the **Schedule & Run Reports** tab.



Sample Report

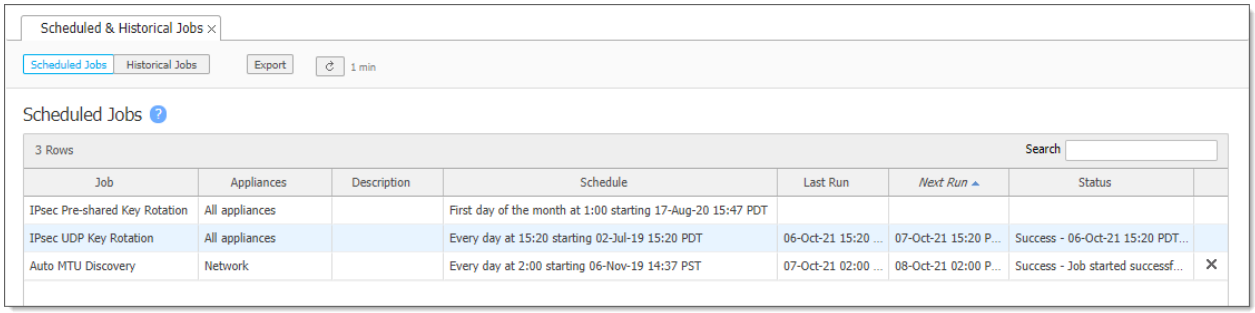


Scheduled and Historical Jobs

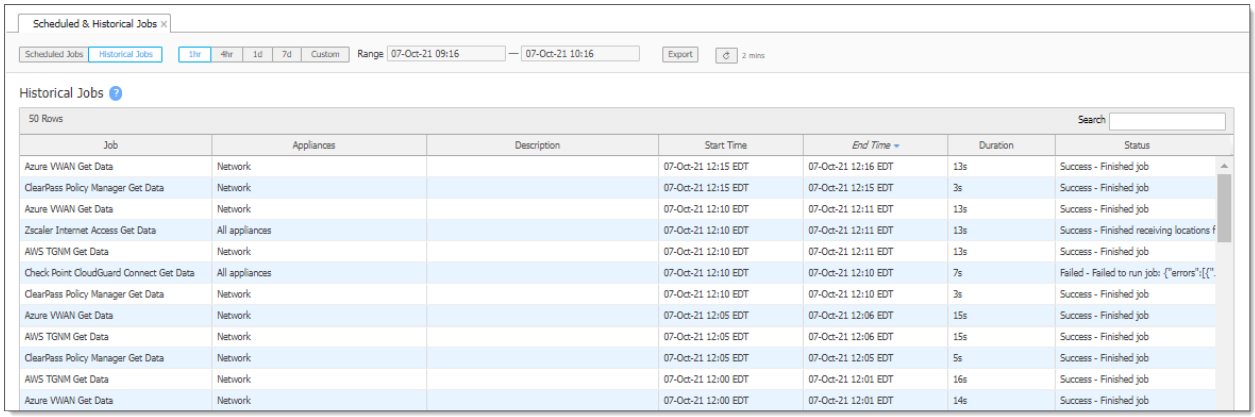
Monitoring > *Reporting* > *Scheduled & Historical Jobs*

This tab has two views:

- It provides a central location for viewing and deleting **scheduled jobs**, such as appliance backup and any custom reports configured for distribution.



- It provides a central location for viewing **historical jobs**.



Metered Usage Reports

Monitoring > Reporting > Metered Usage Reports

This topic is for future content.

The Metered Usage Reports menu option is visible on the Monitoring menu only if you have an active EC-SP-Metered license.

Monitoring > Bandwidth

The options under **Monitoring > Bandwidth** focus on reports related to performance, traffic, and appliance status. Additionally, Threshold Crossing Alerts are helpful in monitoring your network.

Overlay Interface Transport

Monitoring > Bandwidth > Overlay Interface Transport

The Overlay Interface Transport tab provides comprehensive overlay statistics, including the breakout split between third-party and local breakout, fabric overlays, and underlays.

The various charts that this tab provides typically illustrate the different transport methods such as overlay, underlay, and breakout used within the SD-WAN network. They can show the relationship between various interface types and the traffic they handle.

Key Information

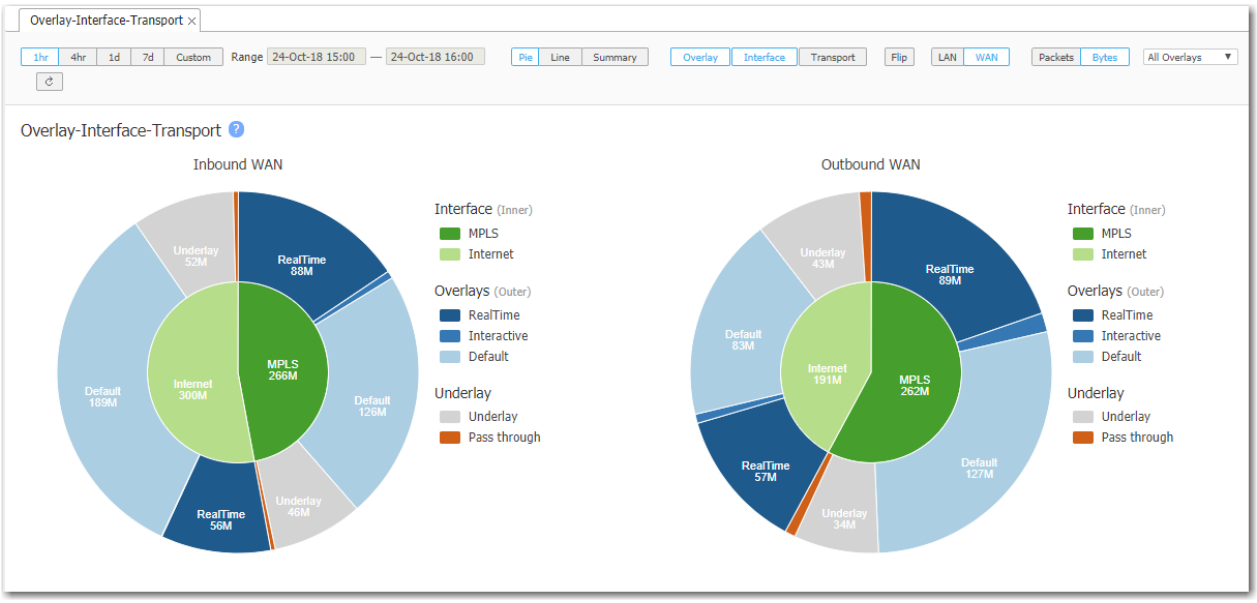
- Types of Transport: Details on the different transport mechanisms (such as MPLS, Internet, LTE, and so forth).
- Performance Metrics: Latency, jitter, packet loss, and throughput.

- Traffic Distribution: How traffic is distributed across different transport interfaces.
- Redundancy and Failover: Information on redundancy options and failover mechanisms.

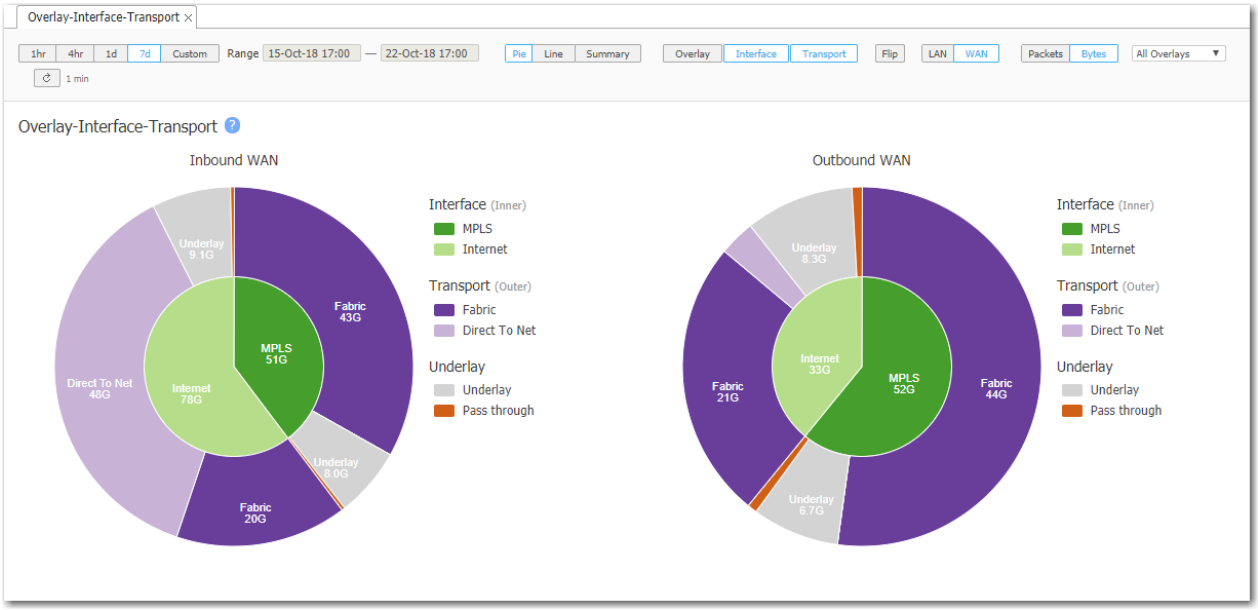
The Overlay Interface Transport tab provides the distribution of traffic across three dimensions (overlays, interfaces, and transport), which you can view individually or compared to another dimension by selecting two of the three dimensions. You can display the data in three ways: summary, trends, and pie charts.

NOTE: For IP Alias interfaces, the system does not provide statistics.

For example, in the pie charts view, select the Overlay and Interface options to see how the overlay traffic is distributed for your interfaces. You can use the Flip button as a toggle to switch (flip) the display of the inner- and outer-ring data.



You can also see how much traffic is transported from one EdgeConnect appliance to another on the SD WAN fabric (Overlays) compared to how much is broken out (local breakout or to the internet). The Underlay legend displays non-overlay traffic.



The types of transport traffic are:

- Underlay – Includes traffic that traverses directly over SD-WAN underlay tunnels (excluding BIO overlay traffic, which also uses these tunnels). For example, if you create a rule with match criterion *application x*, and then send it over to New_York_MPLS_MPLS (an underlay tunnel), the traffic is categorized as underlay traffic. In addition, underlay traffic includes data sent over actual SD-WAN underlay tunnels, generally including the following types of control messages:
  - Keep alive packets – Appliances use these packets to evaluate the reachability of remote peers and the health of connections.
  - Path characterization (pathchar) – Measures path characteristics (loss, latency, and jitter) for a specific path or tunnel.
- Services – Includes two types of traffic:
  - Any non-BIO traffic sent to non-SD-WAN locations over an encapsulated tunnel (that is, IPSec or GRE), such as Zscaler.
  - Any encapsulated traffic that matches a BIO match criterion sent to a cloud service.
- Passthrough – Includes any non-BIO traffic sent without any encapsulation to a destination.
- SP Overlay – Includes any traffic sent to SD-WAN peers over BIO-bonded tunnels.
- Breakout (local breakout or to the internet) – Includes any non-encapsulated traffic that matches a BIO match criterion sent to non-SD-WAN peers.

Interface Bandwidth Summary

Monitoring > Bandwidth > Interface Bandwidth

To display the summary view of the Interface Bandwidth tab, click the summary icon.

This view displays interface bandwidth data for one or more selected appliances. It provides a visual representation of various interface statistics, such as inbound and outbound traffic, firewall denies, average bandwidth utilization, and peak bandwidth utilization. You can export the data to a CSV file. From this view, you can also display the Interface Trends Charts, which show interface bandwidth statistics in chart form for easy analysis.

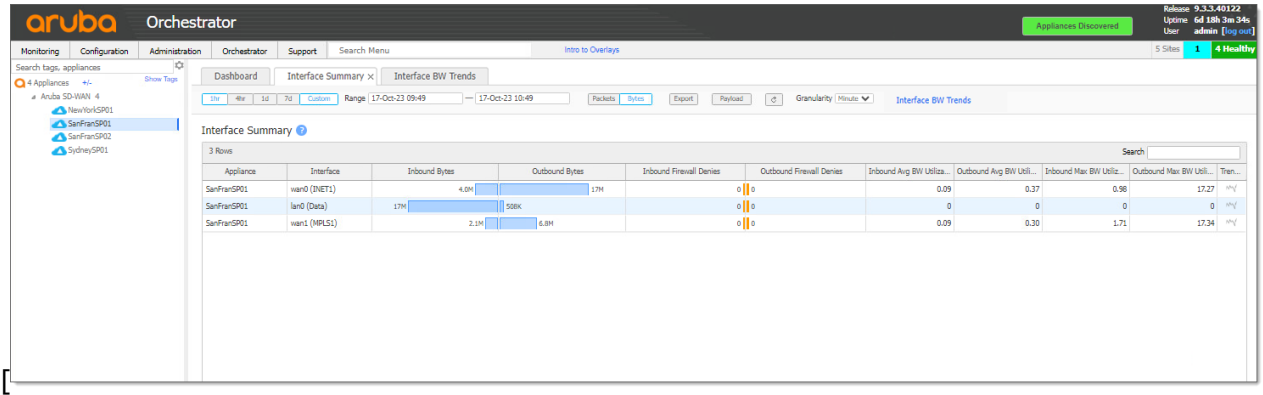
NOTE: For IP Alias interfaces, the system does not provide statistics.

The statistics in the summary view of the Interface Bandwidth tab, and also the charts, are provided for the selected time period.

- Interface bandwidth statistics are shown for one or more selected appliances, but the Interface Trends Charts display the statistics for only one appliance at a time. You can export the data for only one appliance at a time.
- The default time range is one hour. You can click a predefined time range (**1hr**, **4hr**, **1d**, or **7d**) or specify a custom time range by clicking **Custom** and entering start and end dates/times.
- Data can be expressed in packets or bytes. Packets displays statistics according to the number of packets sent and received. Bytes displays statistics for number of bytes sent and received.
- To view the Interface Trends Charts for an appliance without leaving the summary view of the Interface Bandwidth tab, click the chart icon in the Trends column for an appliance. The Interface Trends Charts dialog box opens, displaying the charts for the selected appliance. To view the data in the charts on a larger scale, click the trends icon to display the trends view of the Interface Bandwidth tab.

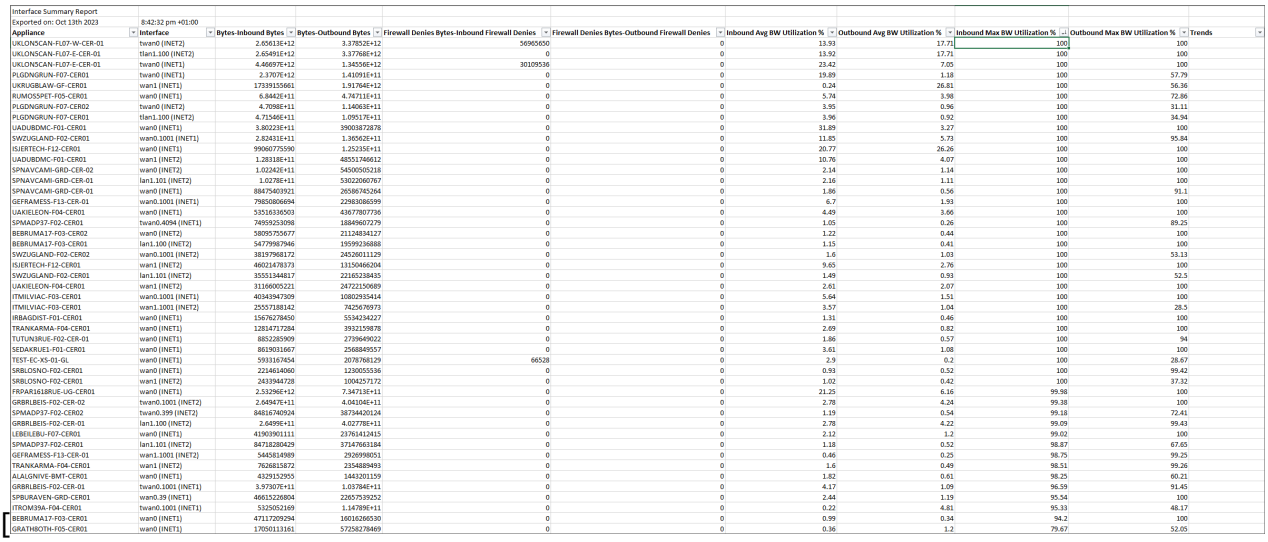
Interpreting the Interface Summary Report

The following example shows a custom Interface Summary Report displayed on the Interface Summary tab for a single appliance for one hour with minute granularity and data expressed in bytes.



[

The following example shows an example Interface Summary Report for 50 appliances exported on Oct 13, 2023 at 8:42 pm and data expressed in bytes.



Interface Summary Report										
Exported on: Oct 13th 2023	8:42:32 pm +01:00									
Appliance	Interface	Bytes-Inbound Bytes	Bytes-Outbound Bytes	Firewall Denies Bytes-Inbound Firewall Denies	Firewall Denies Bytes-Outbound Firewall Denies	Inbound Avg BW Utilization %	Outbound Avg BW Utilization %	Inbound Max BW Utilization %	Outbound Max BW Utilization %	Trends
UKLONSCAN-FL07-W-CER-01	tvan0 (INET2)	2.65613E+12	3.37852E+12	56965650	0	13.93	17.71	100	100	
UKLONSCAN-FL07-E-CER-01	tlan1.100 (INET2)	2.65491E+12	3.37768E+12	0	0	13.92	17.71	100	100	
UKLONSCAN-FL07-E-CER-01	tvan0 (INET1)	4.46697E+12	1.34556E+12	30109536	0	23.42	7.05	100	100	
PLGDNGRUN-F07-CER01	tvan0 (INET1)	2.3707E+12	1.41091E+11	0	0	19.89	1.18	100	57.79	
UKRUGBLAW-GF-CER01	wan1 (INET1)	17339153661	1.91764E+12	0	0	0.24	26.81	100	56.36	
RUMOSSPET-F05-CER01	wan0 (INET1)	6.8442E+11	4.76711E+11	0	0	5.74	3.98	100	72.86	
PLGDNGRUN-F07-CER02	tvan0 (INET2)	4.7098E+11	1.14063E+11	0	0	3.95	0.96	100	31.11	
PLGDNGRUN-F07-CER01	tlan1.100 (INET2)	4.71546E+11	1.09517E+11	0	0	3.96	0.92	100	34.94	
UADUBDMC-F01-CER01	wan0 (INET1)	3.80223E+11	3900387287B	0	0	31.89	3.27	100	100	
SWZUGLAND-F02-CER01	wan0.1001 (INET1)	2.82431E+11	1.36562E+11	0	0	11.85	5.73	100	95.84	
ISJERTECH-F12-CER01	wan0 (INET1)	9968775590	1.25233E+11	0	0	20.77	26.26	100	100	
UADUBDMC-F01-CER01	wan1 (INET2)	1.28318E+11	48551746612	0	0	10.76	4.07	100	100	
SPNAVCAMI-GRD-CER-02	wan0 (INET2)	1.02424E+11	54500505218	0	0	2.14	1.14	100	100	
SPNAVCAMI-GRD-CER-01	lan1.101 (INET2)	1.0278E+11	53022060767	0	0	2.16	1.11	100	100	
SPNAVCAMI-GRD-CER-01	wan0 (INET1)	88475403921	26598745264	0	0	1.86	0.56	100	91.1	
GEFRAMESS-F13-CER-01	wan0.1001 (INET1)	79850806694	22983086599	0	0	6.7	1.93	100	100	
UAKIELEON-F04-CER01	wan0 (INET1)	53516336503	43677807736	0	0	4.49	3.66	100	100	
SPMADP37-F02-CER01	tvan0.4094 (INET1)	74959253098	18849607279	0	0	1.05	0.26	100	89.25	
BEBRUMA17-F03-CER02	wan0 (INET2)	58095755677	21124834127	0	0	1.22	0.44	100	100	
BEBRUMA17-F03-CER01	lan1.100 (INET2)	54779987946	19599236888	0	0	1.15	0.41	100	100	
SWZUGLAND-F02-CER02	wan0.1001 (INET2)	38197968172	24526011129	0	0	1.6	1.03	100	53.13	
ISJERTECH-F12-CER01	wan1 (INET2)	46021478373	13150466204	0	0	9.65	2.76	100	100	
SWZUGLAND-F02-CER01	lan1.101 (INET2)	35551344817	22165238435	0	0	1.49	0.93	100	52.5	
UAKIELEON-F04-CER01	wan1 (INET2)	31166005221	24722150689	0	0	2.61	2.07	100	100	
ITMILVIAC-F03-CER01	wan0.1001 (INET1)	40343947309	10802935414	0	0	5.64	1.51	100	100	
ITMILVIAC-F03-CER01	wan1.1001 (INET2)	25557188142	7425676973	0	0	3.57	1.04	100	28.5	
IRBAGDST-F01-CER01	wan0 (INET1)	15676278450	5534234227	0	0	1.31	0.46	100	100	
TRANKARMA-F04-CER01	wan0 (INET1)	12814717284	393215987B	0	0	2.69	0.82	100	100	
TUTUN3RUE-F02-CER-01	wan0 (INET1)	8852285909	2739649022	0	0	1.86	0.57	100	94	
SEDAKRUE1-F01-CER01	wan0 (INET1)	8619031667	2568849557	0	0	3.61	1.08	100	100	
TEST-EC-X5-01-GL	wan0 (INET1)	5933167454	2078798129	66528	0	2.9	0.2	100	28.67	
SRBLOSNO-F02-CER01	wan0 (INET1)	2214614060	1230055536	0	0	0.93	0.52	100	99.42	
SRBLOSNO-F02-CER01	wan1 (INET2)	2433944728	1004257172	0	0	1.02	0.42	100	37.32	
FRPAR1618RUE-UG-CER01	wan0 (INET1)	2.53296E+12	7.34713E+11	0	0	21.25	6.16	99.98	100	
GRBRLBEIS-F02-CER-02	tvan0.1001 (INET2)	2.64947E+11	4.04104E+11	0	0	2.78	4.24	99.38	100	
SPMADP37-F02-CER02	tvan0.399 (INET2)	84816740924	38734420124	0	0	1.19	0.54	99.18	72.41	
GRBRLBEIS-F02-CER-01	lan1.100 (INET2)	2.6499E+11	4.0277BE+11	0	0	2.78	4.22	99.09	99.43	
LBEILEBU-F07-CER01	wan0 (INET1)	41903901111	23761412415	0	0	2.12	1.2	99.02	100	
SPMADP37-F02-CER01	lan1.101 (INET2)	84718280429	37147663184	0	0	1.18	0.52	98.87	67.65	
GEFRAMESS-F13-CER-01	wan1.1001 (INET2)	5445814989	2926998051	0	0	0.46	0.25	98.75	99.25	
TRANKARMA-F04-CER01	wan1 (INET2)	762681587B	2354889493	0	0	1.6	0.49	98.51	99.26	
ALALGNIVE-BMT-CER01	wan0 (INET1)	432915295S	144320115S	0	0	1.82	0.61	98.25	60.21	
GRBRLBEIS-F02-CER-01	tvan0.1001 (INET1)	3.97307E+11	1.03784E+11	0	0	4.17	1.09	96.59	91.45	
SPBURAVEN-GRD-CER01	wan0.39 (INET1)	46615226804	22657539252	0	0	2.44	1.19	95.54	100	
ITROM39A-F04-CER01	tvan0.1001 (INET1)	5325052169	1.14789E+11	0	0	0.22	4.81	95.33	48.17	
BEBRUMA17-F03-CER01	wan0 (INET1)	47117209294	16016266530	0	0	0.99	0.34	94.2	100	
GRATHBOTH-F05-CER01	wan0 (INET1)	17050113161	5725827846B	0	0	0.36	1.2	79.67	52.65	

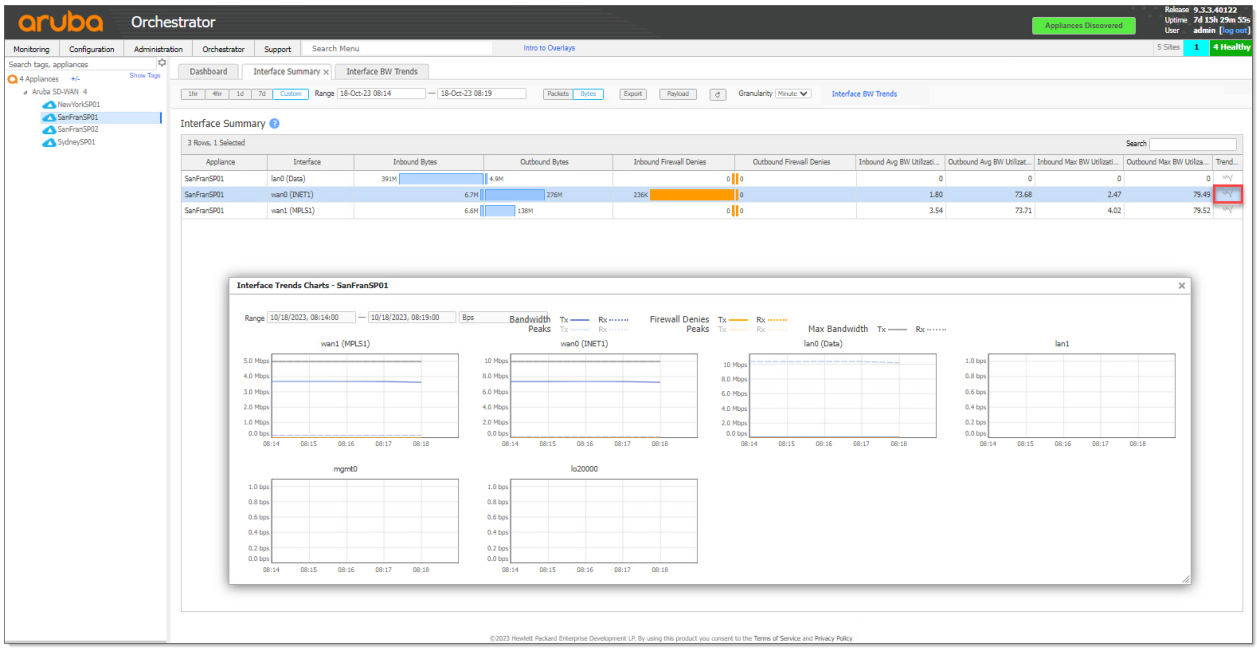
The previous examples provide two different output formats for the Interface Summary. Both report formats contain the following:

- Columns that show the total number of bytes sent and received on a particular interface within a specified time frame.
- Columns that track denied flows for inbound and outbound, which are counted for each flow that was denied during that period.
- Columns for average bandwidth utilization for inbound and outbound, which are calculated as follows:

  (Total number of bytes received or transferred) ÷ (Duration of the selected interval in seconds)

  The resulting numbers are expressed as percentages of the maximum configured bandwidth for that link. These statistics could be less useful if there are several idle periods within the selected time frame, as the idle periods bring down the averages.
- Columns for peak bandwidth utilization, which is measured as the maximum bandwidth utilization per second during the selected time frame, as reported by EdgeConnect. This statistic could be less meaningful if the traffic pattern for the link is mostly idle with occasional bursts up to link speed.

To accurately determine the bandwidth utilization for capacity planning, identify the “busy periods” during the workday and only use those time frames to calculate the average and maximum bandwidth utilization. Currently Orchestrator does not have an automated way to identify busy periods. However, you can use the Interface Trends Charts to get a sense of this information. To view the Interface Trends Charts, click the icon in the Trends column for an appliance, as shown in the following graphic, or click the trends icon to display the trends view of the Interface Bandwidth tab. For more information see, Interface Bandwidth Trends.



Interface Bandwidth Trends

Monitoring > Bandwidth > Interface Bandwidth

To display the trends view of the Interface Bandwidth tab, click the trends icon.

This view displays interface bandwidth data for a single selected appliance in real time or for a specific period. Real-time charts show the past five minutes of usage and refresh every second. By default, charts show transmit and receive statistics for bandwidth and firewall denies. You can toggle peak statistics or maximum bandwidth statistics on or off by clicking the sample indicator line next to each statistic name.

Peak bandwidth statistics display the highest (peak) bandwidth usage per interface that occurred during each interval. This helps to detect the highest points of bandwidth usage so that you can correlate them with possible performance issues or user complaints. If you look only at averages, brief periods of extremely high usage may be hidden, even if they affect network performance.

NOTE: For IP Alias interfaces, the system does not provide statistics.

To display the summary view of the Interface Bandwidth tab, click the summary icon.

You can customize the chart settings using the controls at the top of the tab, as follows:

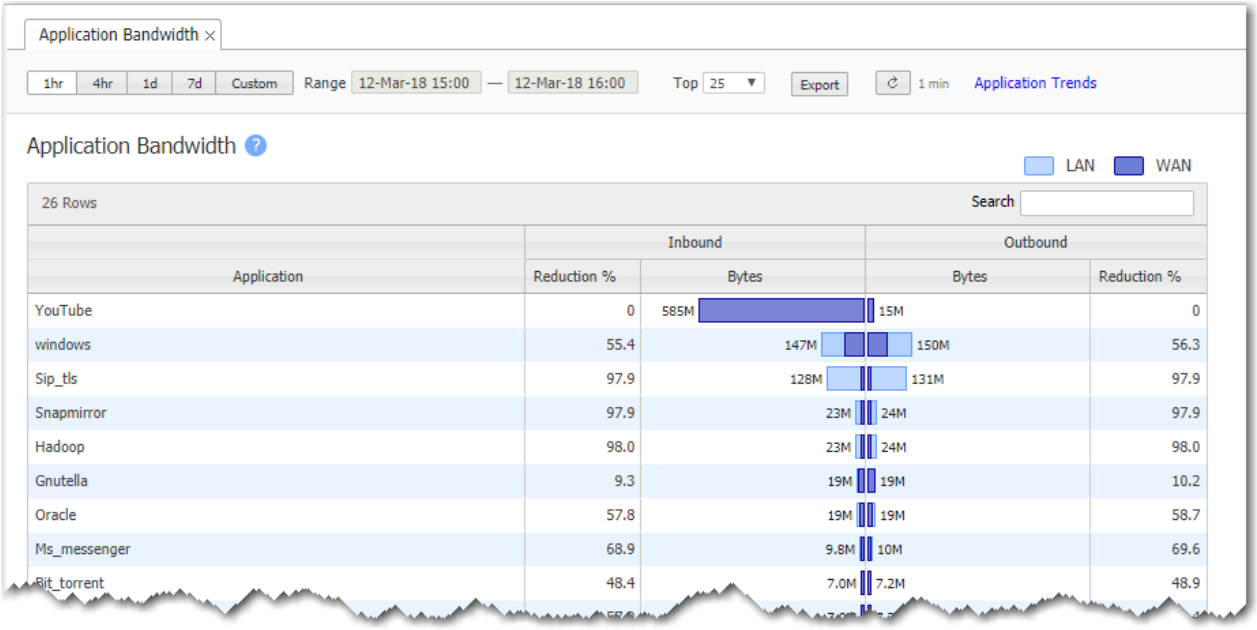
Option	Description
Time period	Click **Real Time** to enable live statistics for all available interfaces. Click a predefined time period (**1h**, **4h**, **1d**, **7d**) to display statistics over the last hour, four hours, day, or seven days. Click **Custom** and set your own custom time range to display statistics for that time period.
Packets/bps	Click **Packets** to display statistics according to the number of packets sent and received. Click **bps** to display statistics for bits per second sent and received.
Show in UTC	Click this option to toggle chart times between local appliance time or UTC.
Large	Click this option to toggle the size of the charts between smaller (default) and large.
Lock Scale	By default, each chart uses its own scale that is relative to the data displayed. Click this option to apply and lock the same scale to each chart.
Payload	By default, charts show complete bandwidth usage statistics—payload and also SD-WAN overhead. (SD-WAN overhead includes only tunnel control packet data.) To view bandwidth usage for payload only, click (enable) the **Payload** button.

Application Bandwidth Summary

Monitoring > *Bandwidth* > *Application Bandwidth*

To display the summary view of the Application Bandwidth tab, click the summary icon.

This view shows which applications have sent the most bytes.



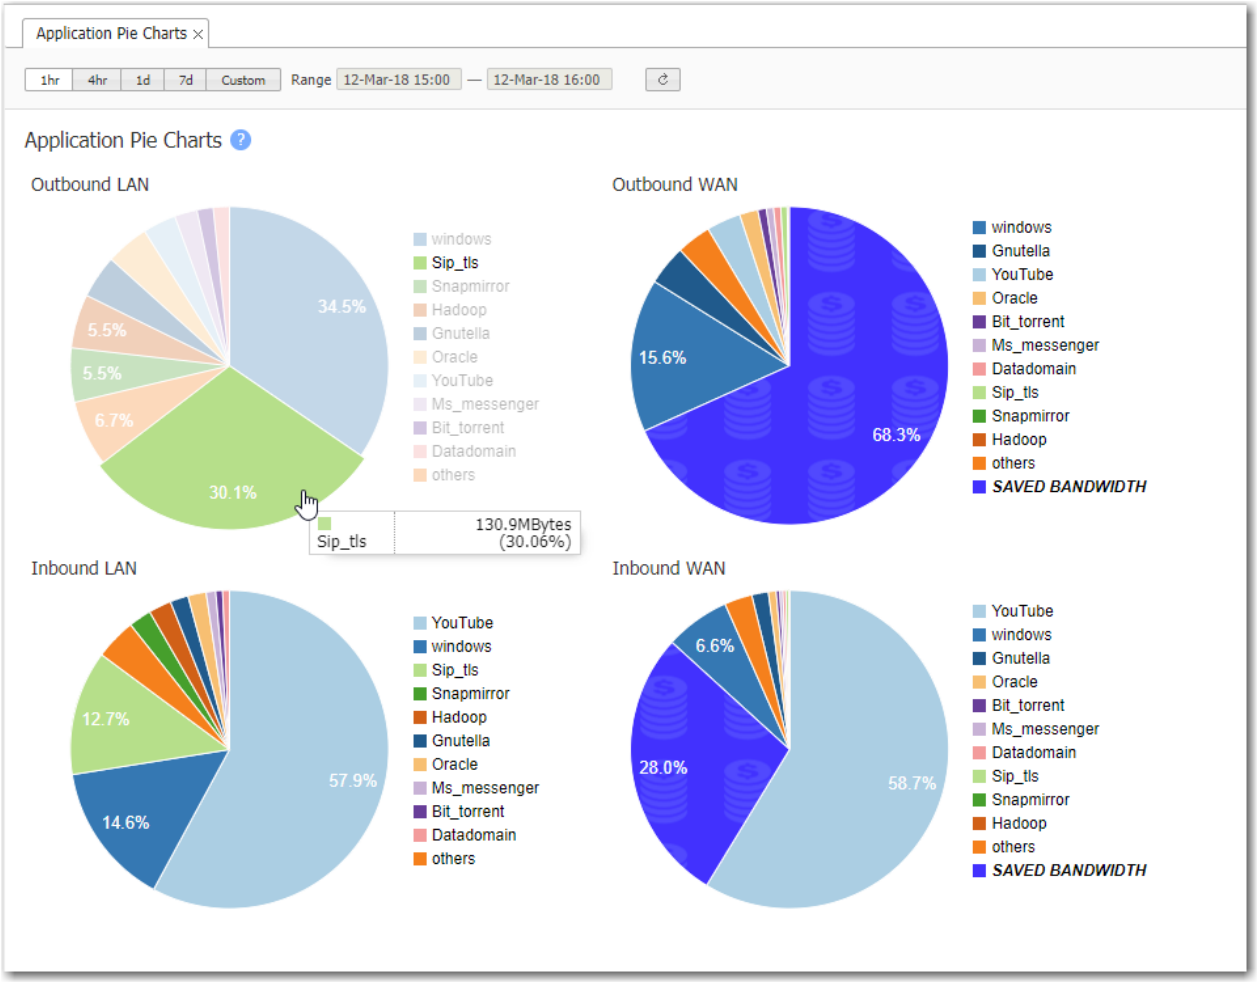
Application Bandwidth Pie Charts

Monitoring > *Bandwidth* > *Application Bandwidth*

To display the pie charts view of the Application Bandwidth tab, click the pie charts icon.

The charts in this view show what proportion of the bytes an application consumes on the LAN and on the WAN.

- Hovering over the charts and the legends reveals additional information.
- The WAN charts identify what percentage of the bandwidth the EdgeConnect appliance saved by optimizing the traffic.

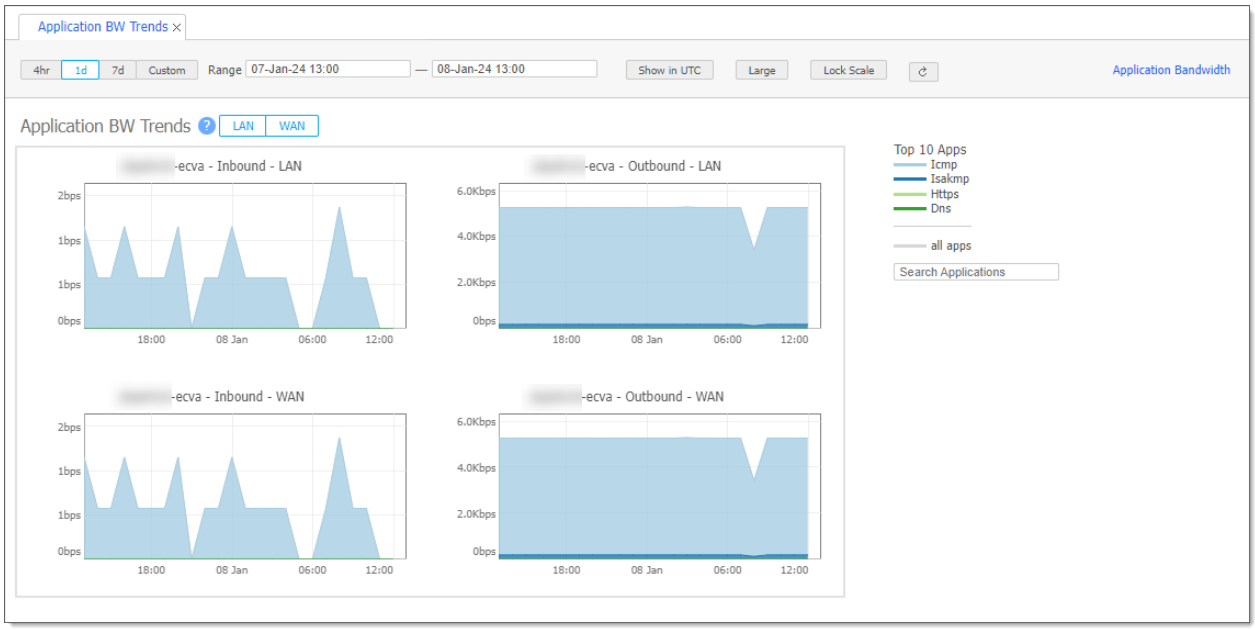


Application Bandwidth Trends

Monitoring > *Bandwidth* > *Application Bandwidth*

To display the trends view of the Application Bandwidth tab, click the trends icon.

This view displays application bandwidth usage over time.



User Bandwidth Summary

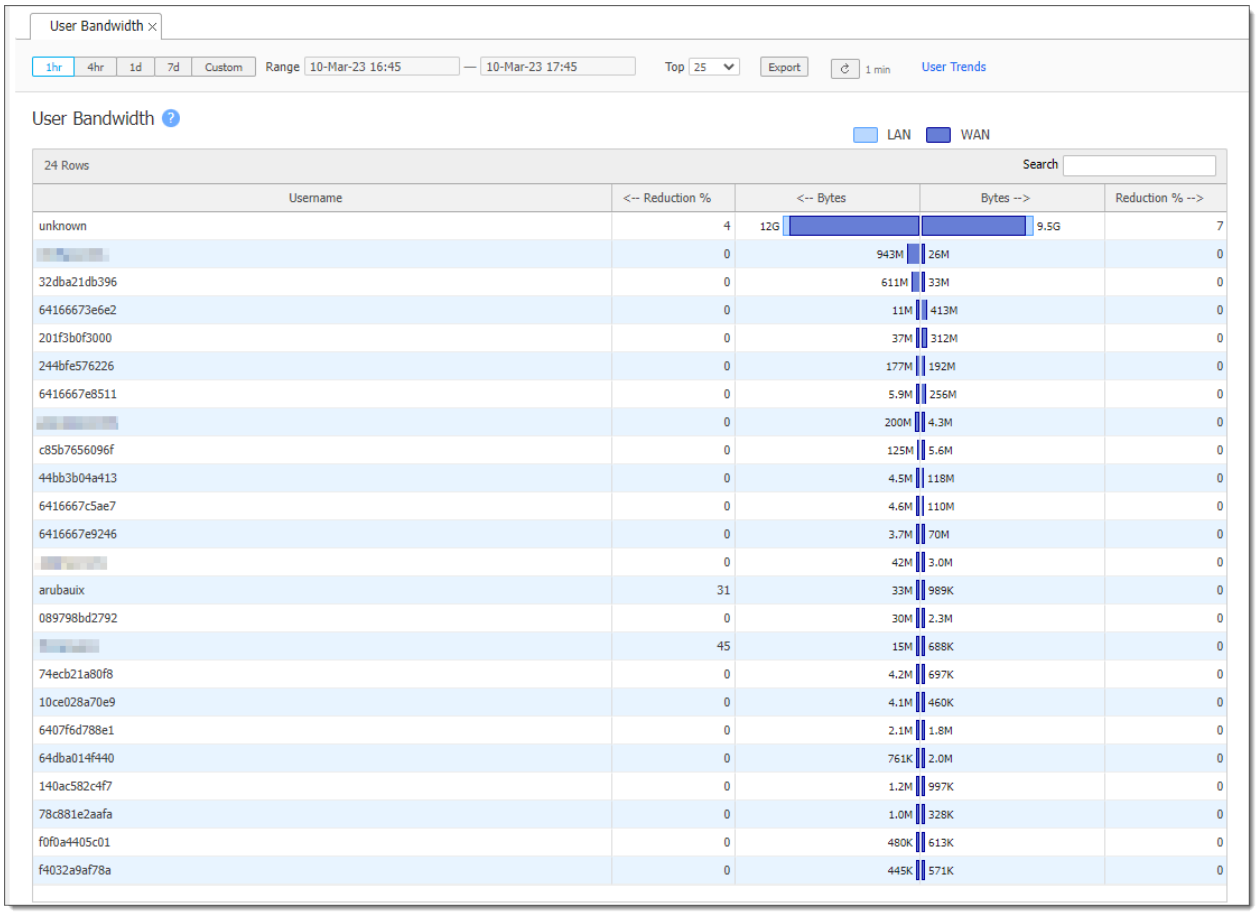
Monitoring > *Bandwidth* > *User Bandwidth*

To display the summary view of the User Bandwidth tab, click the summary icon.

This view displays a summary of user bandwidth data, which can help you identify users consuming the most bandwidth.

NOTE: User bandwidth is derived from RADIUS snooping.

Statistics are summarized for the selected time period. You can change the time period for which user bandwidth summary data is displayed by using the 1hr, 4hr, 1d, and 7d buttons at the top of the tab, or click **Custom** to specify a custom date range and granularity.



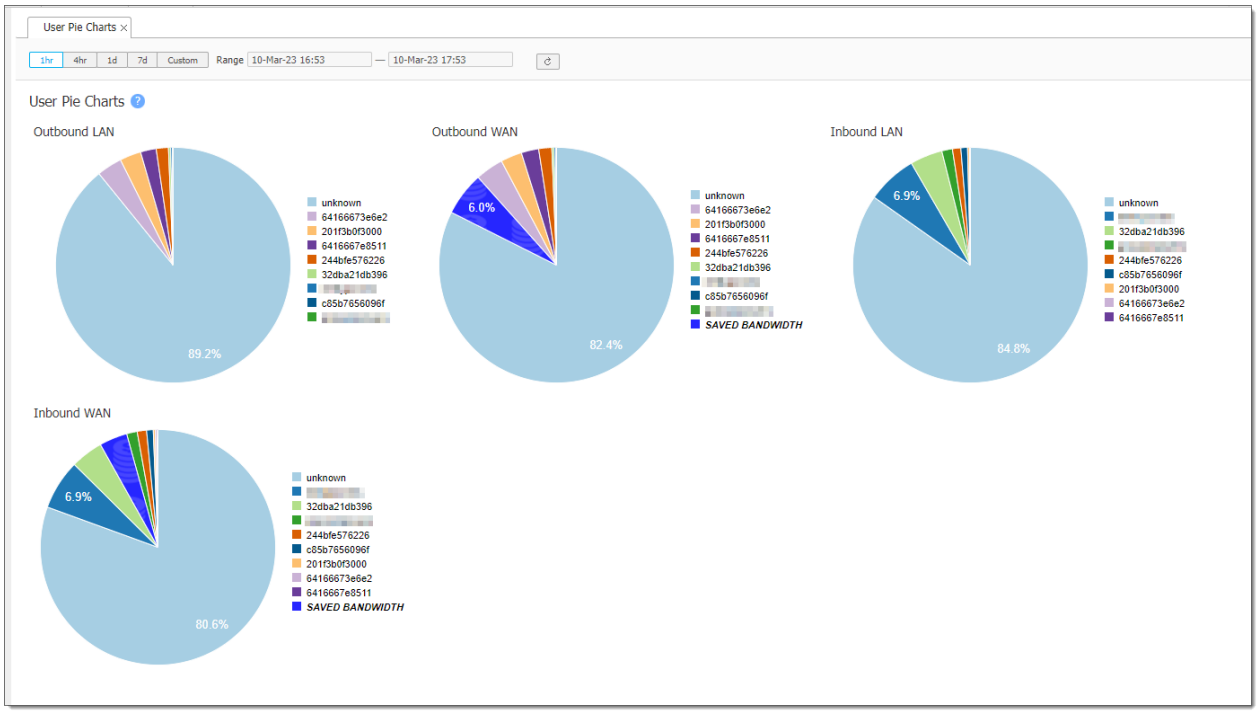
User Bandwidth Pie Charts

Monitoring > *Bandwidth* > *User Bandwidth*

To display the pie charts view of the User Bandwidth tab, click the pie charts icon.

The charts in this view show the proportion of bandwidth each user is consuming. Hovering over the charts and the legends reveals additional information. You can change the time period for which user bandwidth statistics are displayed by using the 1h, 4hr, 1d, and 7d buttons at the top of the tab, or click **Custom** to specify a custom date range and granularity.

NOTE: User bandwidth is derived from RADIUS snooping.



User Bandwidth Trends

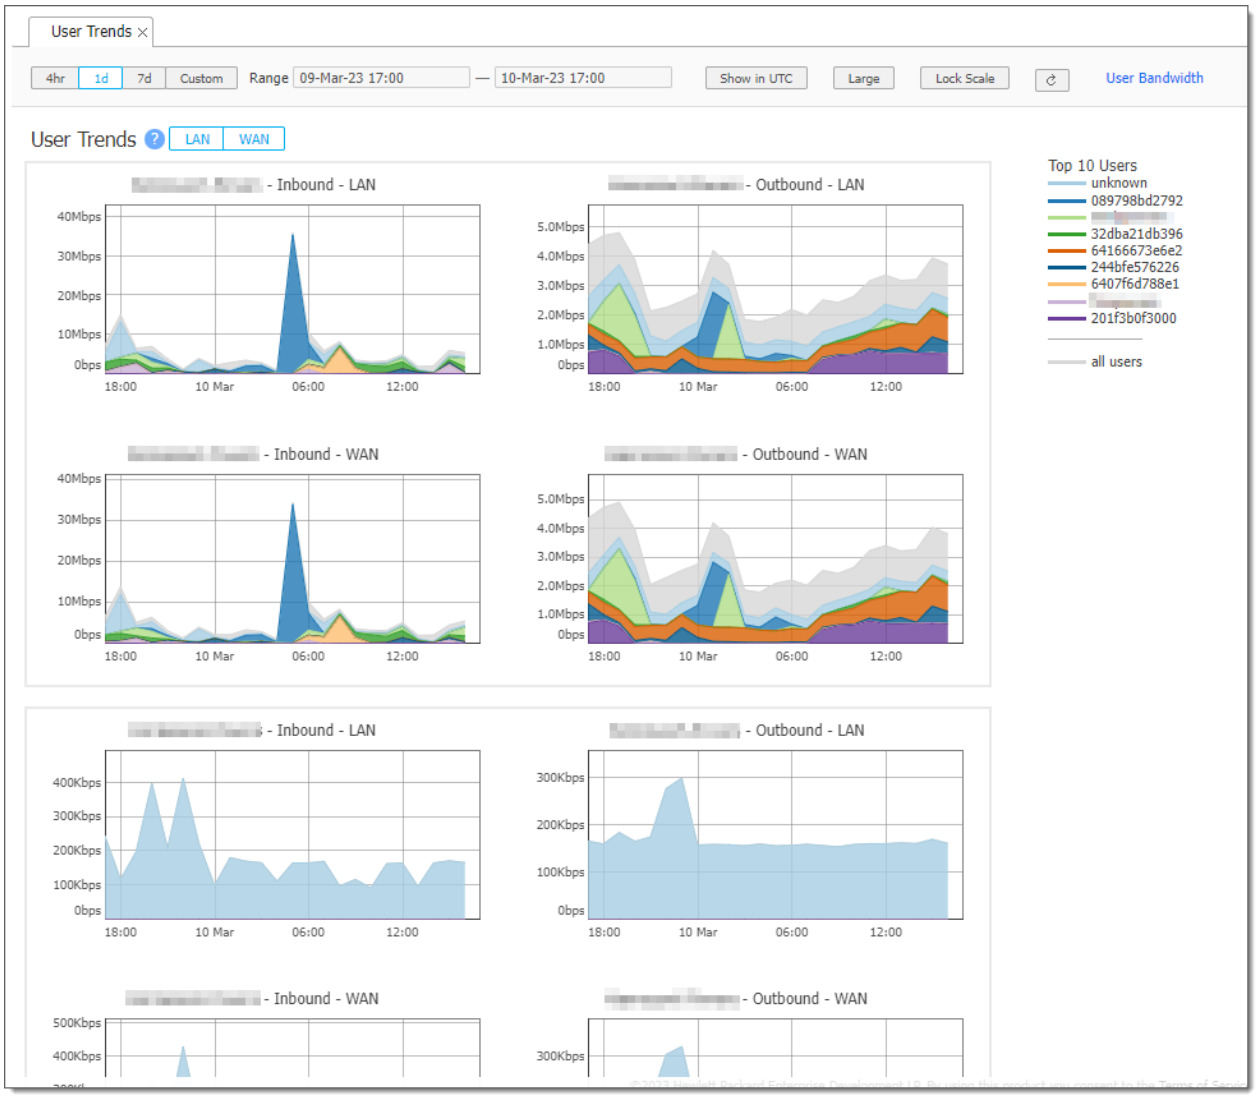
Monitoring > *Bandwidth* > *User Bandwidth*

To display the trends view of the User Bandwidth tab, click the trends icon.

This view displays user bandwidth trends over time for the selected appliances. In the appliance tree, select the appliances for which you want to view user bandwidth trends data. You can select up to 20 appliances.

NOTE: User bandwidth trends are derived from RADIUS snooping.

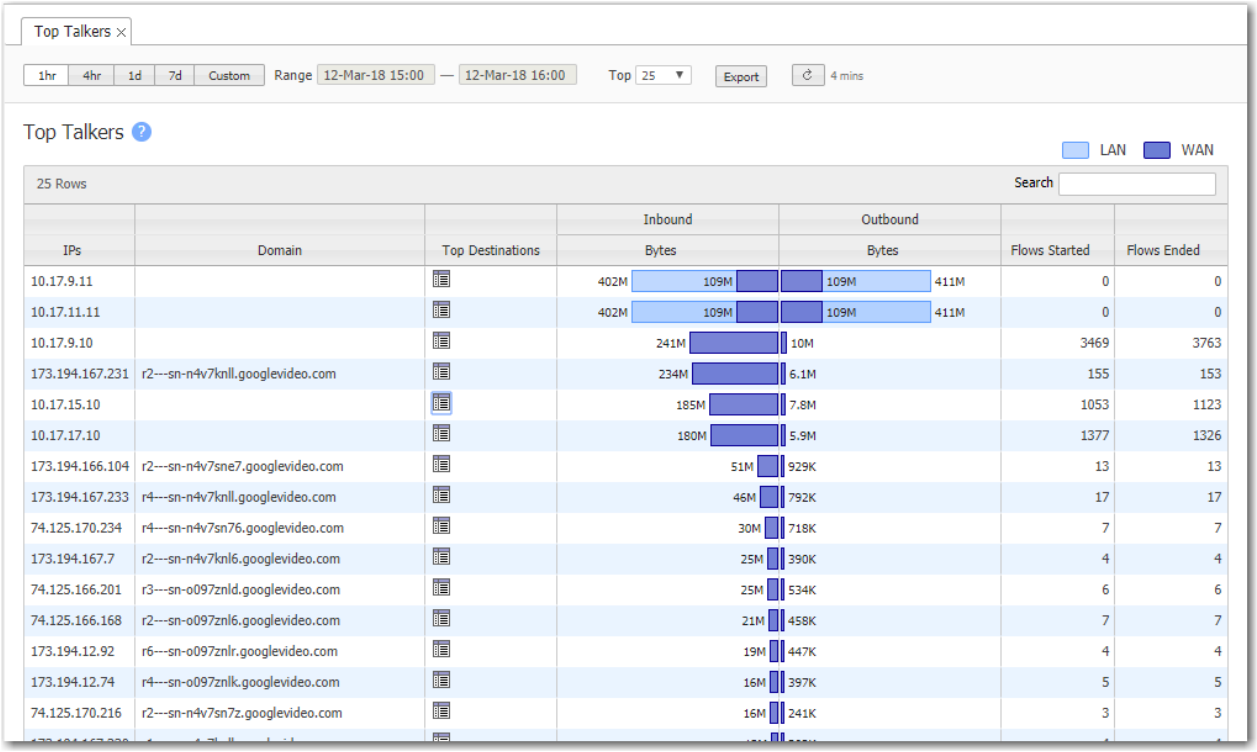
You can change the time period for which user bandwidth trends are displayed by using the 4hr, 1d, and 7d buttons at the top of the tab, or click **Custom** to specify a custom date range and granularity.



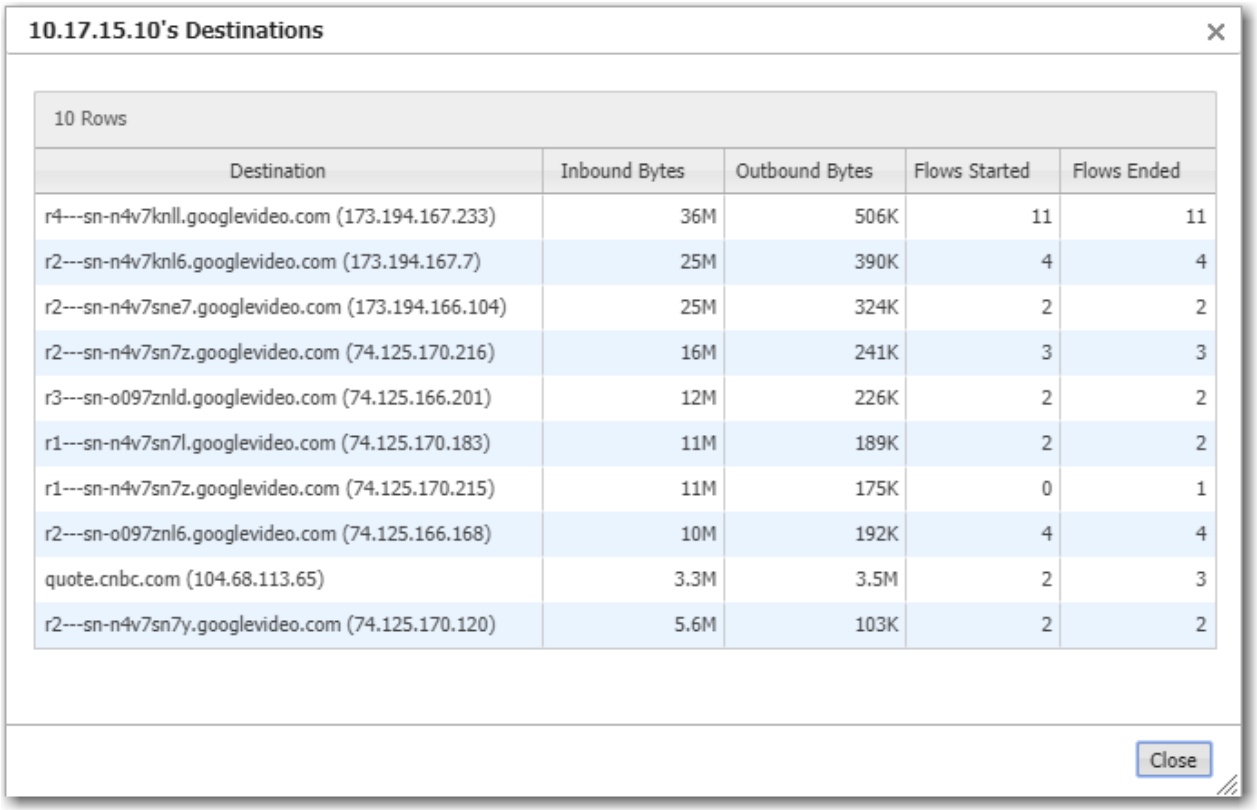
Top Talkers

Monitoring > Bandwidth > Top Talkers

The Top Talkers tab lists the IP addresses that use the most bandwidth.



You can also view each IP's destinations.

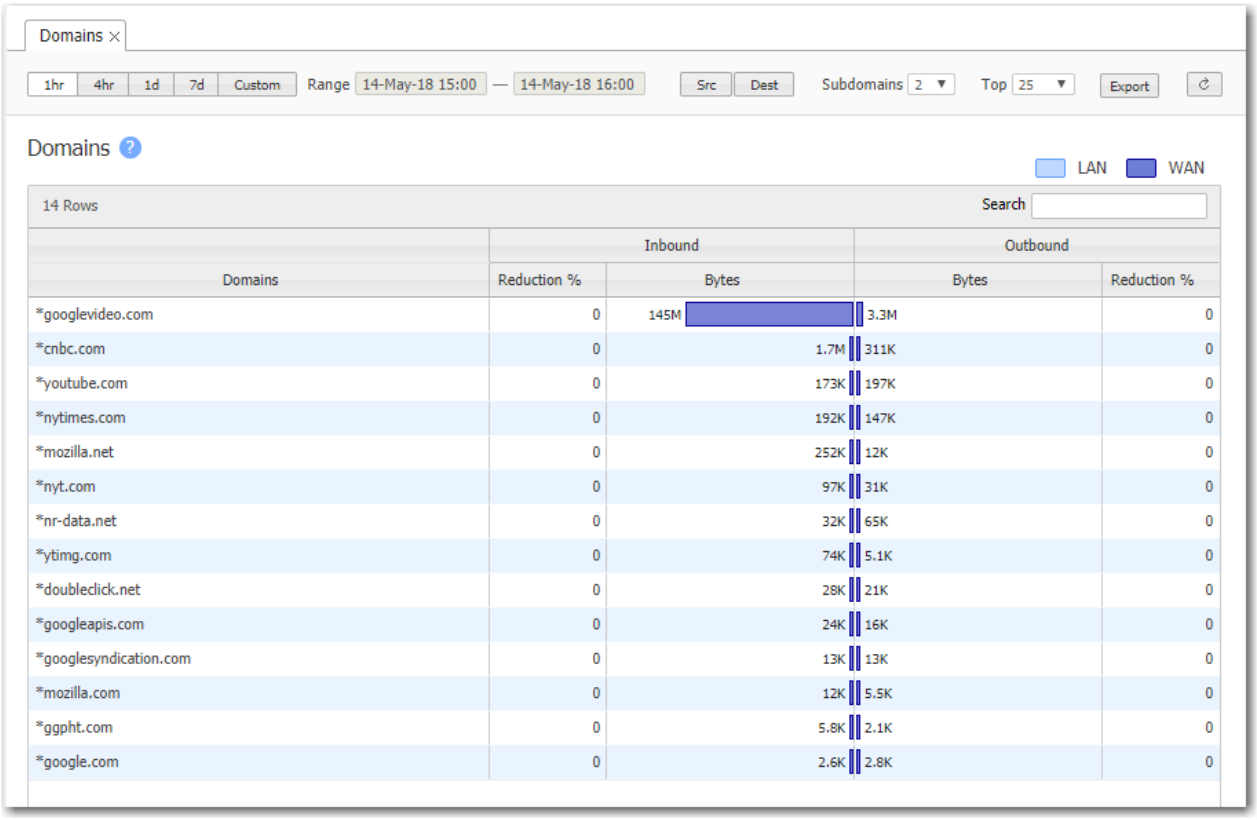


Top Domains

Monitoring > Bandwidth > Top Domains

The Top Domains tab lists the domains that use the most bandwidth.

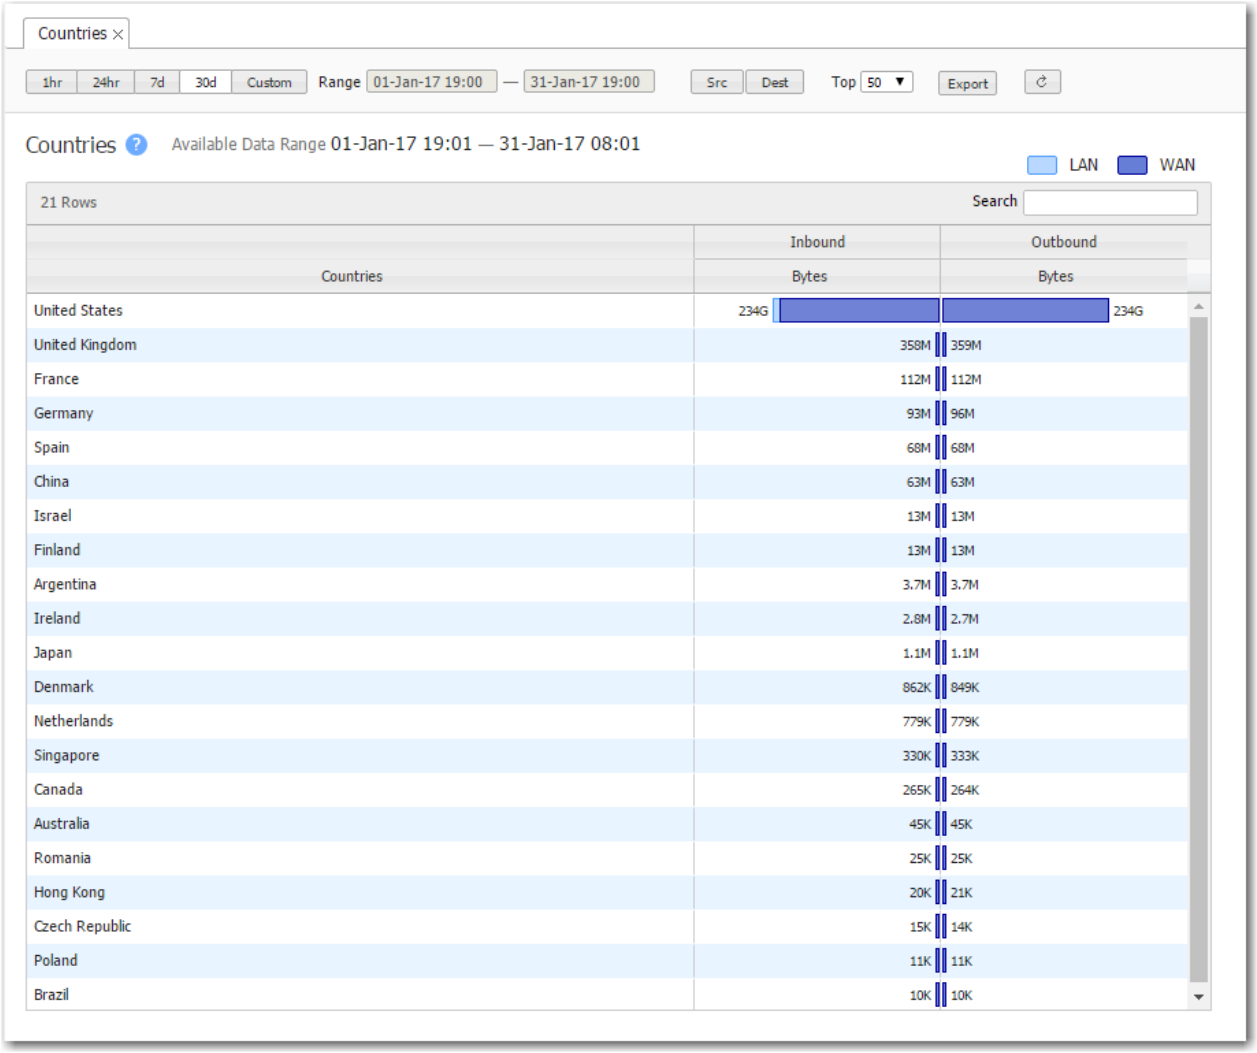
The number of **Subdomains** selected determines how the table aggregates subdomains for display. An asterisk (*) indicates that more subdomains would be displayed if a higher number were selected. This is not a filter, but rather a grouping convenience.



Top Countries

Monitoring > Bandwidth > Top Countries

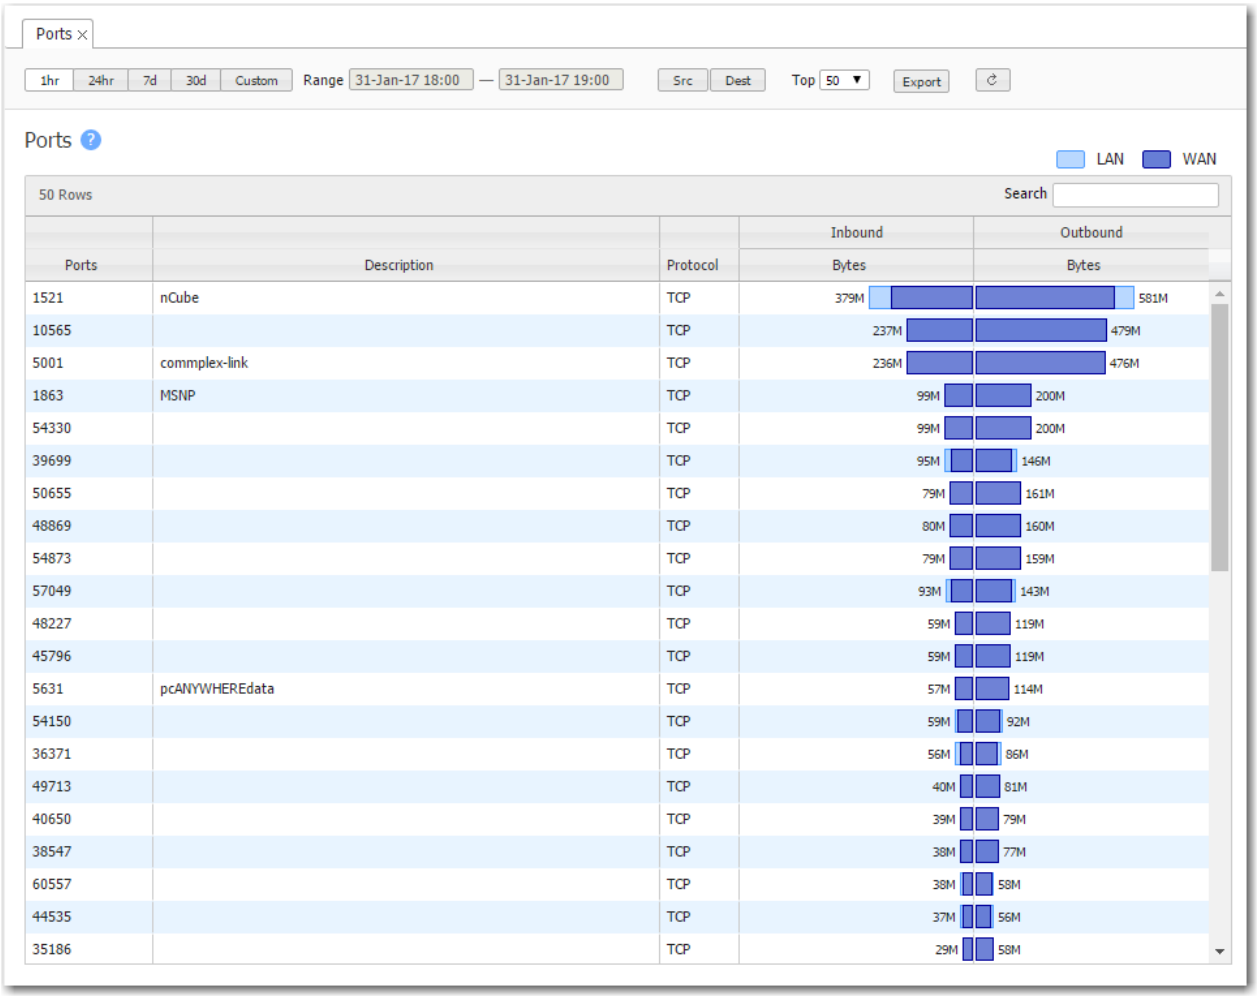
The Top Countries tab lists the countries that use the most bandwidth.



Top Ports

Monitoring > *Bandwidth* > *Top Ports*

The Top Ports tab lists the ports that use the most bandwidth.



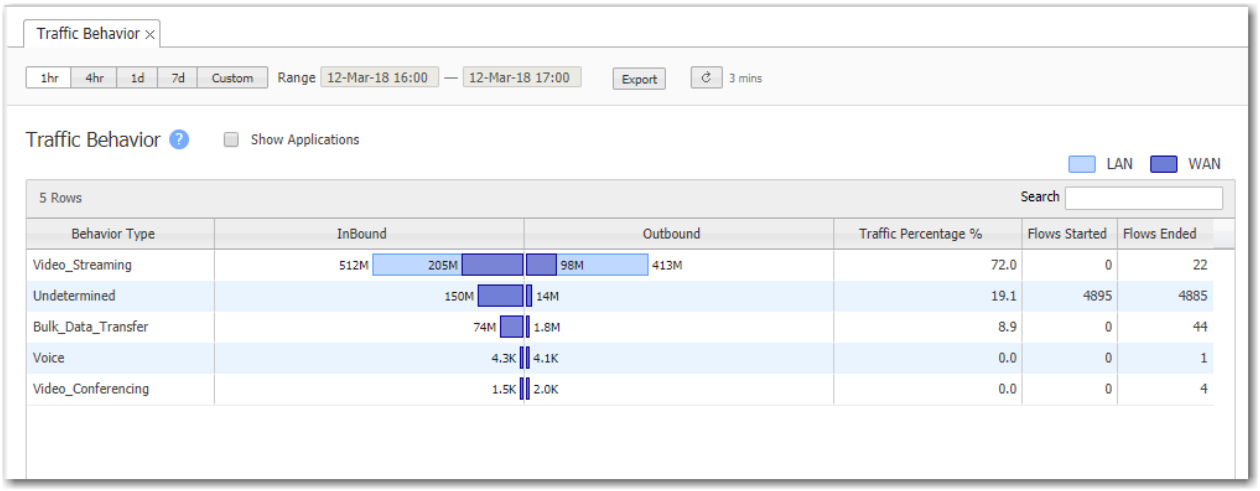
Traffic Behavior

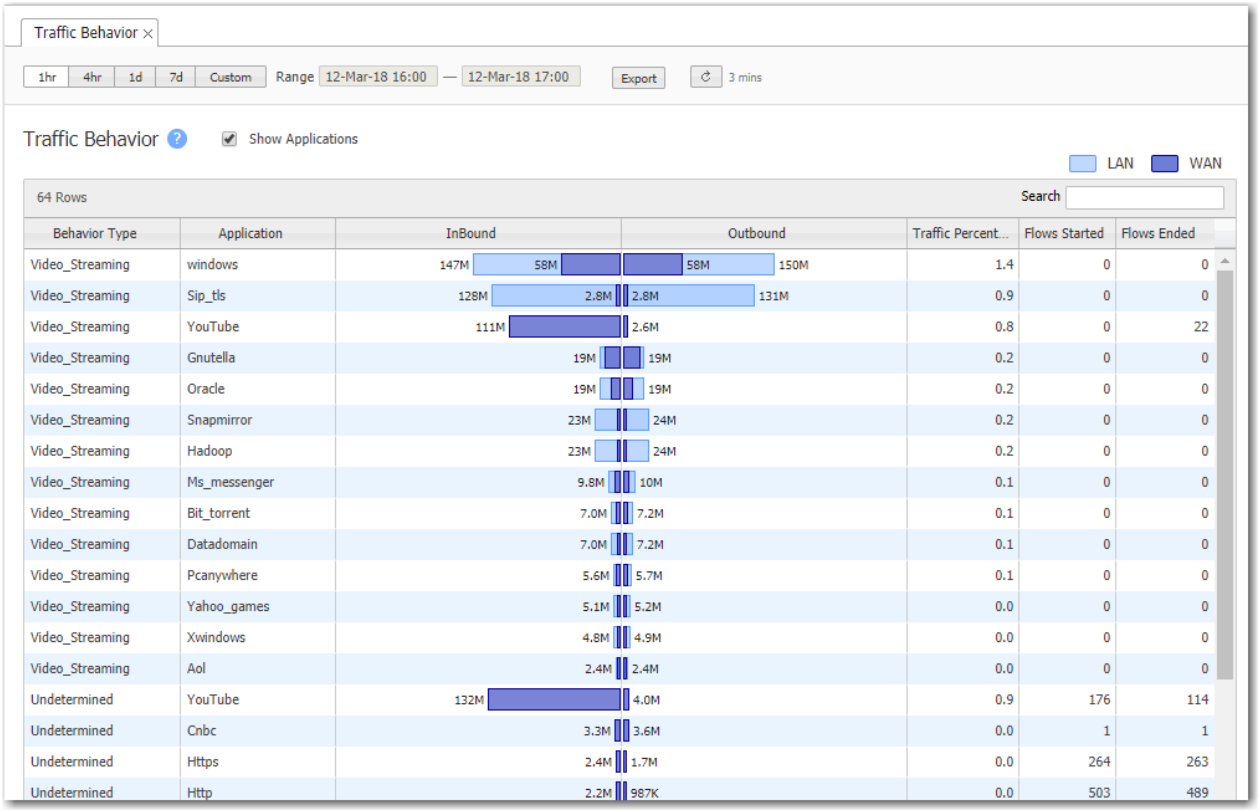
Monitoring > Bandwidth > Traffic Behavior

The Traffic Behavior tab identifies and categorizes traffic based on low-level characteristics of the data streams. The behavior types are:

- Voice
- Video Conferencing
- Video Streaming
- Bulk Data Transfer
- Interactive
- Undetermined

You can also specify these categories as match criteria when creating policies or ACLs (Access Control Lists).





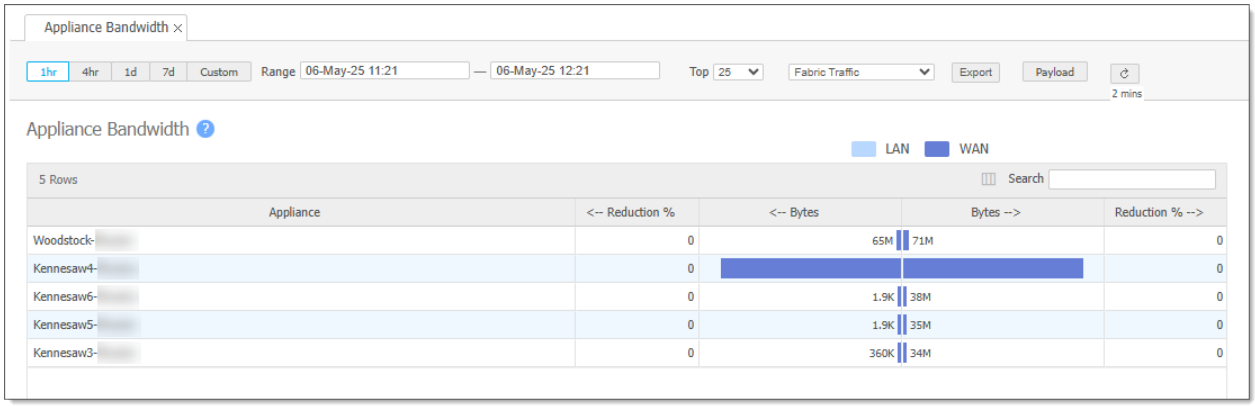
Appliance Bandwidth Summary

Monitoring > Bandwidth > Appliance Bandwidth

To display the summary view of the Appliance Bandwidth tab, click the summary icon.

This view lists the top appliances based on the total volume of inbound and outbound traffic before reduction. It shows how many bytes the EdgeConnect appliance saved when transferring data, aggregated over a selectable time period.

You can filter the table data by selecting a traffic type from the drop-down list: Fabric Traffic, Pass-through Shaped, Pass-through Unshaped, or All Traffic.



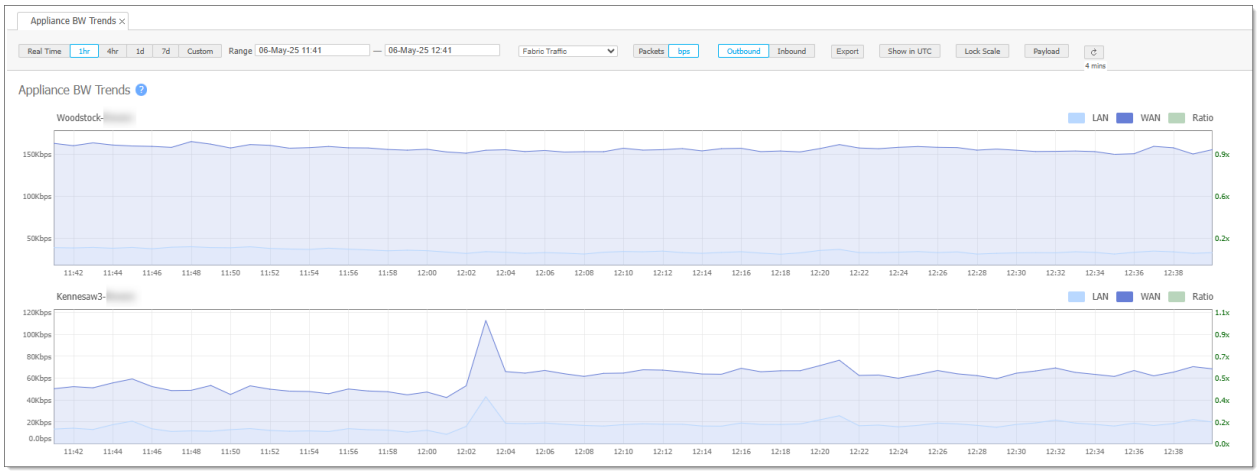
Appliance Bandwidth Trends

Monitoring > *Bandwidth* > *Appliance Bandwidth*

To display the trends view of the Appliance Bandwidth tab, click the trends icon.

This view displays appliance bandwidth usage over time.

You can filter the chart data by selecting a traffic type from the drop-down list: Fabric Traffic, Pass-through Shaped, Pass-through Unshaped, or All Traffic.



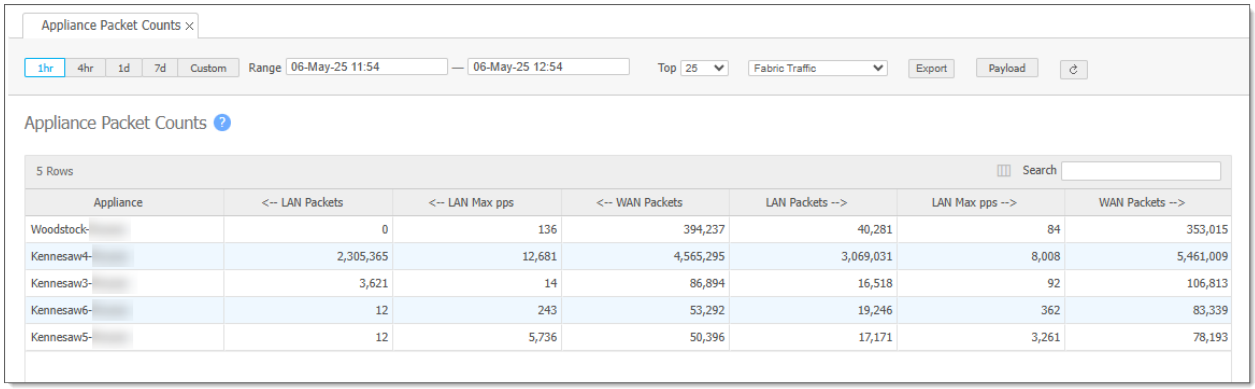
For each Business Intent Overlay, the Link Bonding Policy specified determines the bandwidth efficiency. To guarantee service quality levels, High Availability requires the most overhead, and High Efficiency requires the least. Charts show the total bandwidth used. The Payload option shows how much raw data is transmitted. At the same time, it exposes the Peaks option, which enables the viewing of peak transmissions.

Appliance Packet Counts

Monitoring > Bandwidth > Appliance Packet Counts

The Appliance Packet Counts tab lists the top appliances according to the sum of the inbound and outbound LAN packets, showing how much traffic was sent.

You can filter the table data by selecting a traffic type from the drop-down list: Fabric Traffic, Pass-through Shaped, Pass-through Unshaped, or All Traffic.



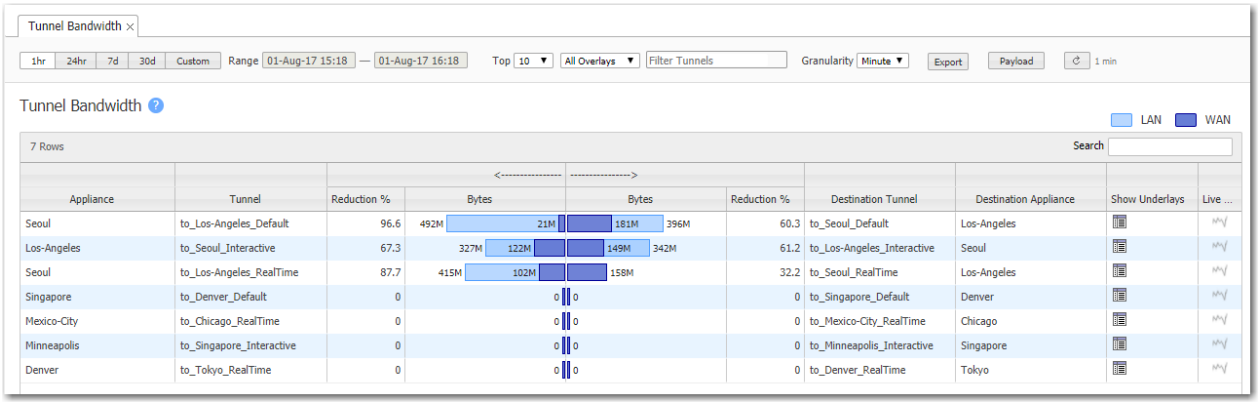
Tunnel Bandwidth Summary

Monitoring > Bandwidth > Tunnel Bandwidth

To display the summary view of the Tunnel Bandwidth tab, click the summary icon.

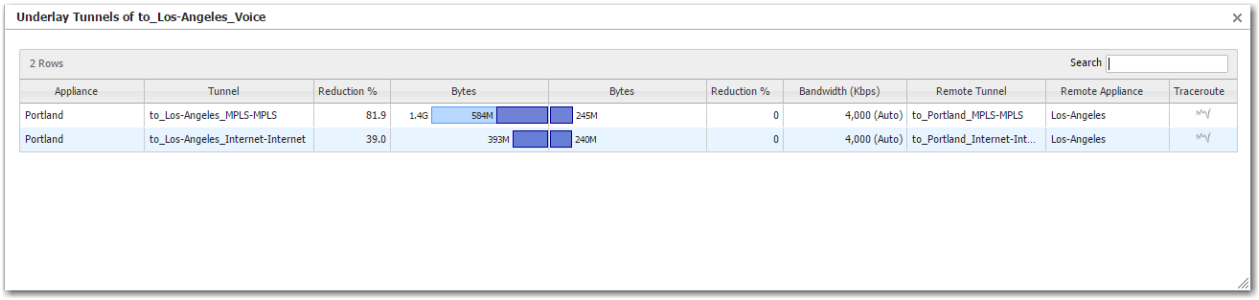
This view shows the tunnels that are sending the most bytes—that is, the most active tunnels.

NOTE: The tunnel bandwidth summary represents only the bonded tunnel statistics per overlay if the **All Overlays** drop-down option is selected. These charts do not render the overlay statistics corresponding to the breakout traffic.



Show Underlays

Underlays are actual IPSec tunnels and physical paths taken (such as MPLS). Overlays are logical tunnels created for different traffic types and policies (such as VoIP).



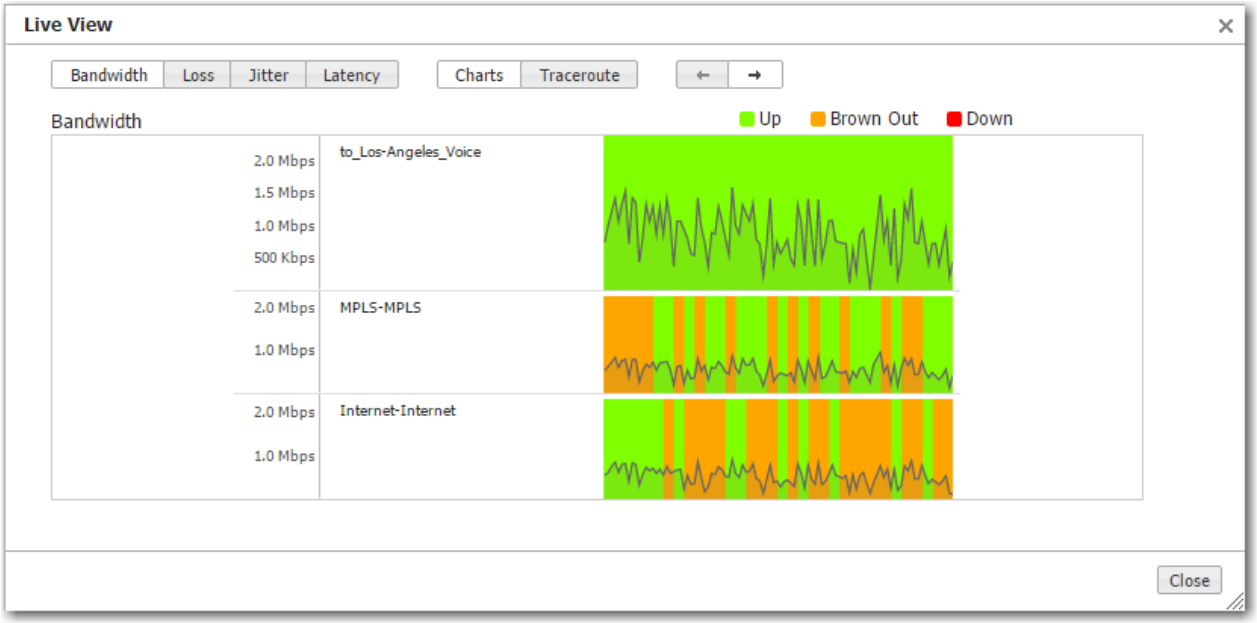
Traceroute

This shows trace route information between the tunnel source and destination IP addresses. It shows intermediate hops, their IP addresses, and the latency between each hop.



Live View

Live View shows the live bandwidth, loss, latency, and jitter on all the tunnels. For an overlay, it also shows live tunnel states—**Up**, **Browned Out**, or **Down**.



LiveView shows in real time how synergy is created to maintain coverage. The real-time chart shows the SD-WAN overlay at the top and the underlay networks at the bottom. The overlay is green and is delivering consistent application performance while both underlays are in persistent brown-out state.

Tunnel Bandwidth Pie Charts

Monitoring > *Bandwidth* > *Tunnel Bandwidth*

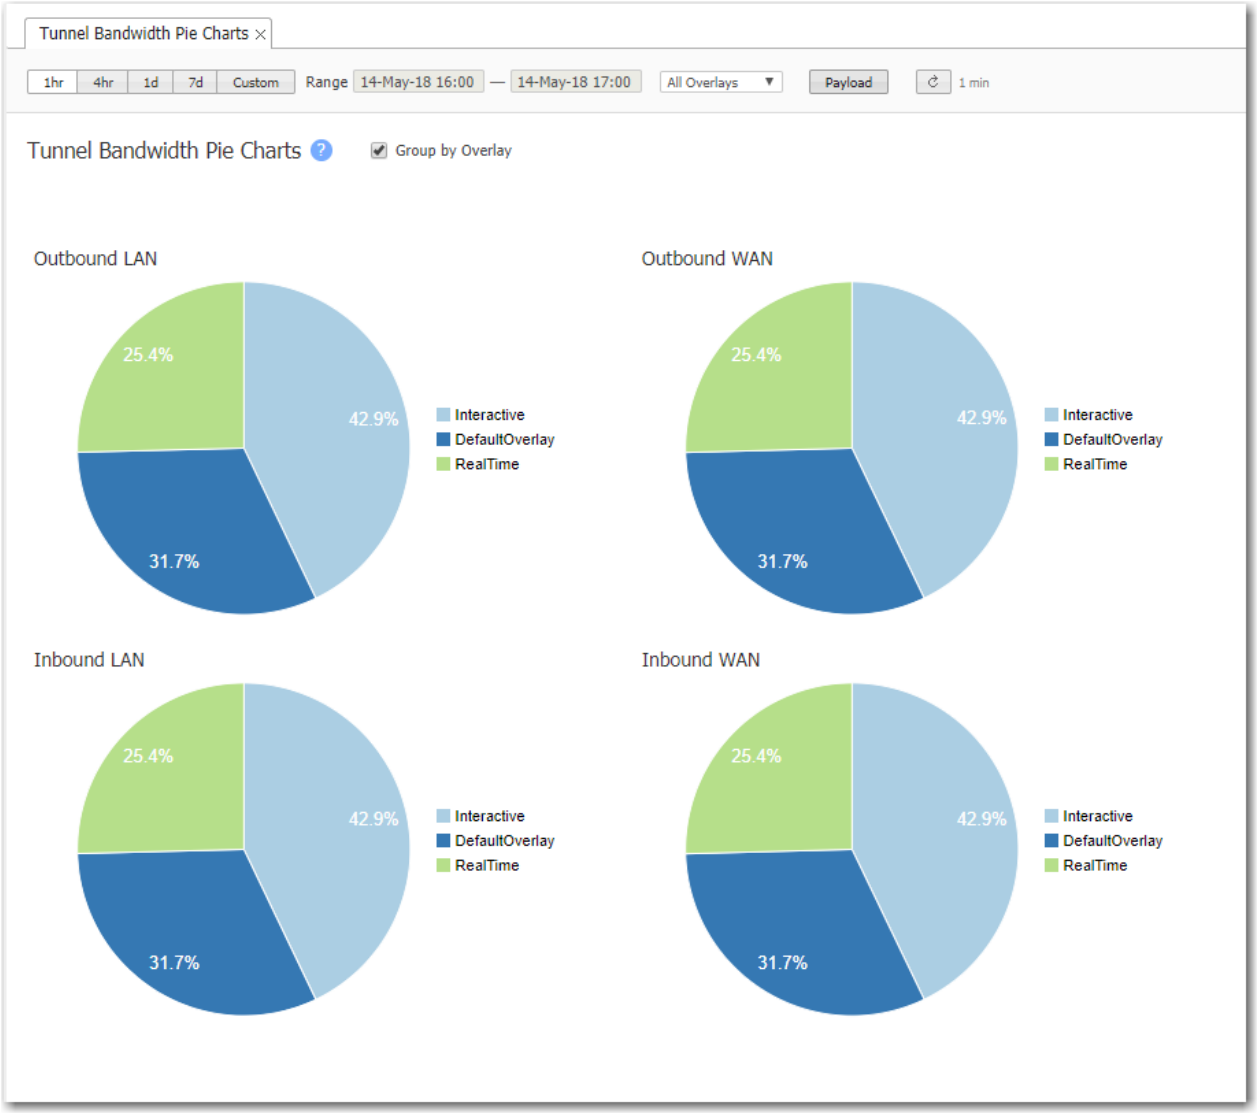
To display the pie charts view of the Tunnel Bandwidth tab, click the pie charts icon.

The charts in this view show the bandwidth usage distribution across different tunnels within the SD-WAN. They help in understanding which tunnels are consuming the most or least bandwidth. The WAN charts identify the percentage of bandwidth the appliance saved by optimizing the traffic.

NOTE: The tunnel bandwidth pie charts represent only the bonded tunnel statistics per overlay if the **All Overlays** drop-down option is selected. These charts do not render the overlay statistics corresponding to the breakout traffic.

Key Information

- Bandwidth Allocation: The percentage of total bandwidth used by each tunnel.
- Tunnel Identification: Specific tunnels and their corresponding bandwidth usage.
- Usage Patterns: Insight into which tunnels have high or low utilization.
- Efficiency: Helps identify if the bandwidth is being efficiently used or if there are any bottlenecks.



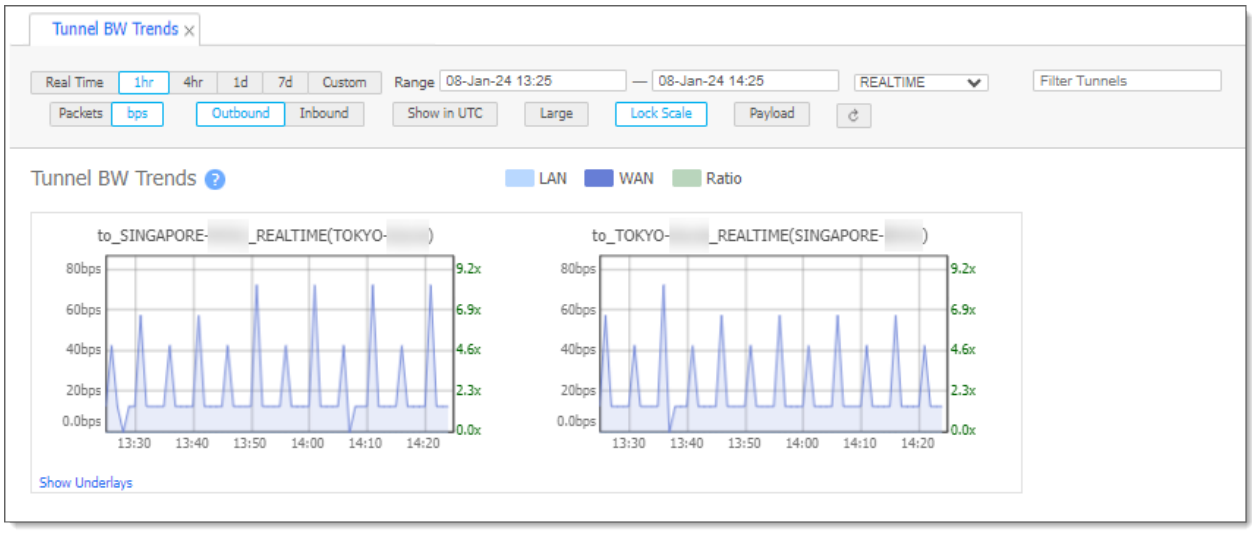
NOTE: For one-day statistics, Orchestrator saves statistics at midnight (UTC time) for the previous day. Therefore, clicking **1d** displays data for the entire day preceding midnight (UTC). For example, if you do this five hours after midnight (UTC), only recently generated statistics belonging to the day preceding midnight (UTC) will be displayed. The last five hours will not be reflected in the pie charts. To display statistics for the last 24 hours, click **Custom**, select **Hour** for Granularity, and then specify the exact 24-hour time range.

Tunnel Bandwidth Trends

Monitoring > *Bandwidth* > *Tunnel Bandwidth*

To display the trends view of the Tunnel Bandwidth tab, click the trends icon.

This view displays tunnel bandwidth usage over time.



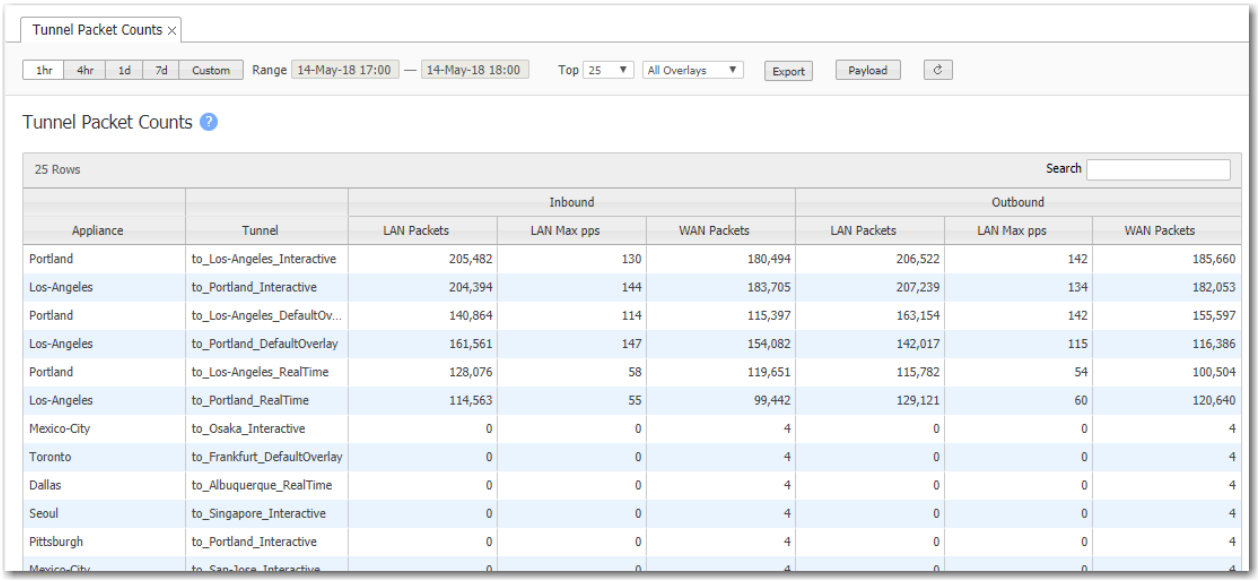
- For each Business Intent Overlay, the specified Link Bonding Policy determines the bandwidth efficiency.
- To guarantee service quality levels, High Availability requires the most overhead and High Efficiency requires the least.
- Charts display the total bandwidth used.
- The Payload option shows how much raw data is transmitted. At the same time, it exposes the Peaks option, which enables the viewing of peak transmissions.

NOTE: Underlay tunnels are a shared resource among overlays. Therefore, underlay charts display aggregated data.

Tunnel Packet Counts

Monitoring > Bandwidth > Tunnel Packet Counts

The Tunnel Packet Counts tab shows the tunnels that sent the most packets.



Appliance	Tunnel	Inbound			Outbound		
		LAN Packets	LAN Max pps	WAN Packets	LAN Packets	LAN Max pps	WAN Packets
Portland	to_Los-Angeles_Interactive	205,482	130	180,494	206,522	142	185,660
Los-Angeles	to_Portland_Interactive	204,394	144	183,705	207,239	134	182,053
Portland	to_Los-Angeles_DefaultOv...	140,864	114	115,397	163,154	142	155,597
Los-Angeles	to_Portland_DefaultOverlay	161,561	147	154,082	142,017	115	116,386
Portland	to_Los-Angeles_RealTime	128,076	58	119,651	115,782	54	100,504
Los-Angeles	to_Portland_RealTime	114,563	55	99,442	129,121	60	120,640
Mexico-City	to_Osaka_Interactive	0	0	4	0	0	4
Toronto	to_Frankfurt_DefaultOverlay	0	0	4	0	0	4
Dallas	to_Albuquerque_RealTime	0	0	4	0	0	4
Seoul	to_Singapore_Interactive	0	0	4	0	0	4
Pittsburgh	to_Portland_Interactive	0	0	4	0	0	4
Mexico-City	to_San-Jose_Interactive	0	0	4	0	0	4

Flows

Monitoring > *Bandwidth* > *Flows*

Use the Flows tab to view, filter, and manage flows for all your appliances or the appliances you select in the appliance tree.

The following table describes the filters displayed in the top portion of the Flows tab.

NOTE: For a detailed walkthrough of how to use the Flows tab, see HPE Aruba Networking's 5-part "Monitoring > Flows" video series here: How To and Technical Videos.

Filter	Description
Application	Includes built-in applications, custom applications, and user-created application groups. Select the text field and a list displays. Choose the application you want to apply to your flow or enter the exact application you want to apply.
App Group	Includes the application group created by the user. Select the text field and a list displays. Choose the application group you want to apply to your flow or enter the exact application group you want to apply.
Role	Specify the user role you want to apply.
User Name	Specify the name of the user you want to apply.
IP/Subnet	This shows the flows that match both SRC IP and DEST IP as the two endpoints if SRC:DEST is enabled. If not enabled, all sources will appear when the filter is applied. You can apply this filter by clicking **Enter** without selecting the Apply button if you want to do so.

Filter	Description
Port	This displays ports with SRC and DEST as the two endpoints if SRC: DEST is enabled. If not enabled, all ports will appear when the filter is applied.
Segment	Filters flows by the specified segment(s). **Any** will filter flows based on all segments as configured in routing segmentation. This filter is available only if routing segmentation is enabled. This filter works in conjunction with the Zone filter. For details, see the “Segment and Zone Filters” section below this table.
Zone	Filters flows by the specified firewall zone(s). **Any** will filter flows based on all firewall zones. This filter works in conjunction with the Segment filter. For details, see the “Segment and Zone Filters” section below this table.
VLAN	Identifies the Virtual Local Area Network of a packet. Enter the **VLAN ID** you want to apply to your flow in the text field.
DSCP	Select the desired DSCP from the list. You can choose any or a specified DSCP from the list.
Protocol	You can specify the protocol you want to apply to your filter. Select the text field and a list displays. You can select all or specify an individual protocol to apply.
Domain	Includes the domain you can specify to filter your flow. Use the format **.domain.__ or __.domain.[com, info, edu, org, net,* and so forth.*]* Select the text field and a list displays. Choose the domain you want to apply.
Overlay	Overlay to which the flow is applied. Overlays are defined on the **Business Intent Overlay** tab.
Transport	Select any of the three transport types: **SD-WAN**, **Breakout**, and **Underlay**. You can also apply a third-party service in this column if you have one configured.

Filter	Description
Flow Characteristics	You can apply any of the following flow characteristics to your flow: **WAN Optimization**, **Directly Attached**, **IPS Dropped**, **Pass-Through**, **Slow Devices**, **SYN Cookie Verified**, **Route Dropped**, **Firewall Dropped**, **Embryonic**, **Asymmetric**, and **AppExpress**. **NOTE:** You can select only one flow characteristic at a time. Slow Devices: For debugging. A slow device cannot receive data quickly enough from the EdgeConnect appliance. This causes the appliance to expend too many resources for this device at the expense of accelerating other devices. To counteract this, disable TCP acceleration for the slow devices in the Optimization Policy. Embryonic: For TCP, this is a flow that is in a state of formation (for example, three-way handshake is not complete). For UDP, ICMP, and other IP protocols, this is a flow for one-way traffic. ICMP Error packets without request are also considered embryonic. Dropped embryonic flows are highlighted in red. Other drops are also highlighted in red, including firewall policy drops, some system/routing drops, and IPS drops. For the Firewall Protection Profiles feature, some flows could drop because of security errors, such as not complying to strict three-way handshakes. These are highlighted in red as well.
Include EdgeHA	If not selected, EdgeHA flows are excluded (default). If selected, the flows between EdgeHA will be included.
Include Built-in	Includes the built-in policy flows. If not selected, they are excluded (default). If selected, they will be included.
Active/Ended	Select to apply an active or ended flow as a filter. If selected, the Started and/or Ended fields become available.
Started/Ended	Select the started or ended time of the flow from the drop-down list. If **Custom** is selected, use the provided fields to specify an exact date and time range. These fields are available only if Active or Ended are selected.
Duration	Shows flows that have lasted through a specific time frame. You can select **<** (less than) or **>** (greater than), and enter a specific duration (in minutes).
Bytes	You can specify whether you want to filter flows that have transferred their total bytes or within the last five minutes.
Filter	This list has all the saved filters. When selected, the filter configurations are loaded. See more information below about the **Filter** option.

Segment and Zone Filters

The Segment and Zone fields in the top portion of the Flows tab work together to filter flows that display in the Flows table based on your segment and firewall zone selections. The Segment filter is available only if routing segmentation is enabled.

The following table details the various Segment and Zone filter settings. Keep in mind that for both segment and zone, *either* means that filtered flows will be either sourced *from* or destined *to* the selected segment or zone. For a Segment *either* condition, only one Segment field is displayed; for a Zone *either* condition, only one Zone field is displayed.

- Click the Src <-> Dest icon (./Media/src-dest-icon.png)) associated with the Segment filter to toggle between displaying only one Segment field (the *either* option) and displaying two fields (*source* on the left and *destination* on the right).
- The From:To check box associated with the Zone filter functions similarly. Clear this check box to display only one Zone field (the *either* option). Select it to display the From and To fields.

Filter Settings	Result
For a Segment *either* condition, if the Segment field is set to Any	the Zone fields are disabled regardless of whether you select the From:To check box
For a Segment *either* condition, if the Segment field is set to a specific segment	only one Zone field is enabled (the Zone *either* condition) and the drop-down menu lists zones associated with only the selected segment; the From and To fields are not available
If the Segment *source* and *destination* fields are both set to Any	the Zone drop-down fields are disabled regardless of whether you select the From:To check box
If the Segment *source* field is set to Any, but the Segment *destination* field is set to a specific segment	the Zone *either* condition does not apply; only the To field is available and the drop-down menu lists zones associated with only the selected segment
If the Segment *source* field is set to a specific segment, but the Segment *destination* field is set to Any	the Zone *either* condition does not apply; only the From field is available and the drop-down menu lists zones associated with only the selected segment
If the Segment *source* and *destination* fields are both set to specific segments	the Zone *either* condition does not apply; both the From and To fields are available and the drop-down menus list zones associated with only the selected segments
If routing segmentation is disabled in your network	the Segment filter is not displayed; the Zone fields are enabled, but zones that are associated with only the default segment are listed in the drop-down menus

Filter Field

You can use the Filter field to configure specific filters. The drop-down menu displays a list of default filters you can apply to your flows. Click the edit icon to add, edit, or delete filters.

To add a filter:

1. Click the edit icon next to the Filter field.
2. Create a filter or select one from the list.
3. Click **+Add**.
4. Click **Save**.

You can also select the history tab with the two arrows next to the **Filter** field if you want to go back to a previously applied filter. A maximum of 20 previously applied filters can be saved.

Reset or Reclassify Flows

- You can **Reclassify** or **Reset** [Selected / All Returned / All] flows:
  - Resetting the flow kills it and restarts it. It is service-affecting.
  - Reclassifying the flow is not service-affecting. When policy changes occur, flow re-classification makes a best-effort attempt to conform the flow to the change. If the flow cannot be successfully "diverted" to this new policy, then an Alert asks if you want to reset.
  - **Selected** flows are individually selected; **All Returned** results from filtering (up to the max number of returnable flows); and **All** refers to all flows, visible or not.

Export

Click **Export** to save the contents of the Flows table to a CSV file.

Flow Detail

Click the info icon in the Detail column to display information about the flow. This information is primarily provided to assist Support with troubleshooting and debugging.

Inbound/Outbound Reduction %

The Inbound Reduction % and Outbound Reduction % columns in the Flows table refer to reduced WAN traffic relative to a specific appliance.

- Reduction % for *outbound* traffic = 100(Received from LAN – Transmitted to WAN)/Received from LAN
- Reduction % for *inbound* traffic = 100(Transmitted to LAN – Received from WAN)/Transmitted to LAN

NOTE: These columns are hidden in the Flows table by default. To display them, right-click any column header in the table to display a list of available columns, and then select the **Inbound Reduction %** and **Outbound Reduction %** check boxes.

Additional Information about Flows

This section provides general and version-specific information about flows.

Flows in Drop State

Appliances do not forward packets for flows in the "drop" state. These flow entries are kept in the Flows table for accounting purposes and to avoid reclassification storms. Dropped flows are aged out after reaching their configured timeouts. Timeouts are set using the UDP flow timeout and Non-accelerated TCP flow timeout fields on the System template or the System Settings subtab of the System Information dialog box. For details, see System Template or System Information.

Starting with Orchestrator 9.1, during a flow reclassify that results in a deny, the system actively ends the flow (the flow transitions to an end or reset behavior).

ICMP/UDP Flows

- For any non-TCP connection (such as icmp, UDP), a flow is deleted only from inactivity.
- The inactivity timeout is three minutes for this type of flow. For example, after a ping connection is stopped, the flow still appears in the "Current Flows" for three minutes. This setting can be modified by using the system template.

TCP Non-Accelerated Flows

- For a TCP connection, a flow is deleted under different timeouts. A half-open (single SYN) connection stays for two minutes if the connection does not establish correctly. A half-close (single FIN) or unclean-close (RST) deletes the connection after two minutes. A normal close (FIN-FIN) deletes the connection almost immediately.
- A TCP connection also has an inactivity timeout. If no activity is detected on an established TCP connection for 30 minutes (by default), the flow is deleted. This setting can be modified by using the system template.

TCP Accelerated Flows

For all ECOS versions, flow redirection applies to WAN Optimized (Boosted) TCP accelerated traffic only. Flow Redirection only applies to the TCP protocol; UDP traffic is not flow redirected.

NOTE: The WAN Optimization (formerly called Boost) TCP acceleration capability requires symmetric TCP traffic. When TCP traffic is asymmetric, the TCP acceleration capability requires Flow Redirection to be enabled. Flow Redirection only applies to Optimized TCP traffic. Non-Optimized traffic will not be flow redirected.

Starting from ECOS 9.4.1.0 and higher, the following limitations apply when you enable Flow Redirection:

- When segmentation is enabled, flow redirection works only for intra-segment traffic. It does not work for inter-segment traffic.
- Flow redirection does not work for IPv6 traffic.

- When Zone-Based Firewall is enabled, flow redirection will only work if the interface participating in flow redirection is in the Default zone.

Timeout is determined by the configured Keep Alive Timers.

- A heartbeat ACK is sent to idle endpoints after ten minutes.
- If the endpoints have closed, an RST is returned and the connection is deleted after two more minutes due to the unclean-close.

The timers can be modified per sequence number by using the Optimization Template.

- Idle Timeout: The period of time that a TCP connection has to be idle before a keep-alive is sent. (Default 600 seconds)
- Probe Interval: The time in seconds between each keep-alive probe. (Default 30 seconds)
- Probe Count: The number of times TCP probes the connection to determine whether it is alive after the keep-alive option has been activated. The connection is assumed to be lost after sending this number of keep-alive probes. (Default 8)
- **Auto Reset Flows** - Enables or disables the auto-reset of TCP flows. If a connection is seen by an appliance but after the handshake already completed, the connection would normally remain but without TCP Acceleration. If this feature is enabled, and a connection is reclassified in the Flows report, around 30 seconds later, it will be reset. When the endpoints re-establish the flow, it now will be subject to the optimization and route policies it matches. This feature is disabled by default. It can be enabled per sequence number by using the Optimization Template.

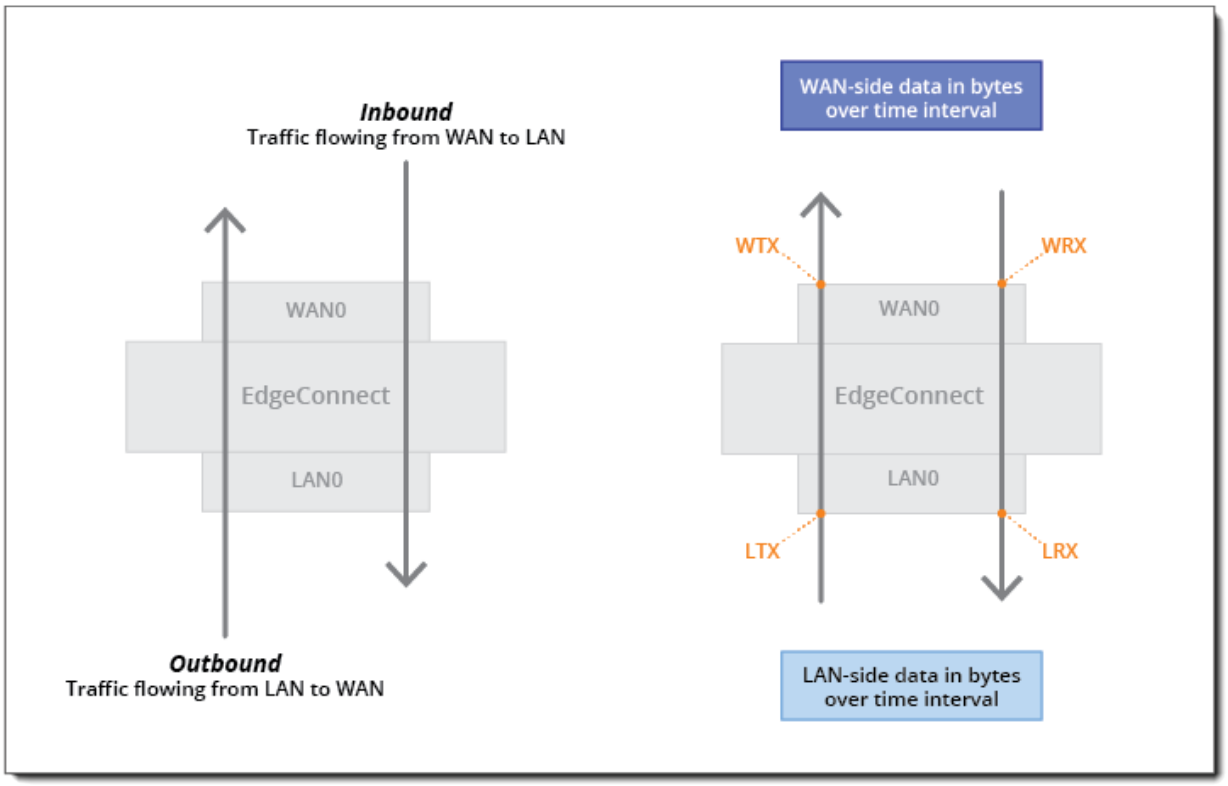
Outbound and Inbound

Outbound and *Inbound* in EdgeConnect refer to the direction of traffic as it flows from the LAN-side to the WAN-side of an appliance, or from the WAN-side to the LAN-side of an appliance. These are different from actual interface names, such as WAN0 or LAN0.

Description	Counter Type	Traffic Received On	Traffic Forwarded To
Inbound LAN	LAN TX	WAN-side interface	LAN-side interface
Outbound LAN	LAN RX	LAN-side interface	WAN-side interface
Inbound WAN	WAN RX	WAN-side interface	LAN-side interface
Outbound WAN	WAN TX	LAN-side interface	WAN-side interface

WAN optimization data reduction is calculated using the following formula:

Data Reduction % = (LAN Bytes - WAN Bytes) / LAN Bytes



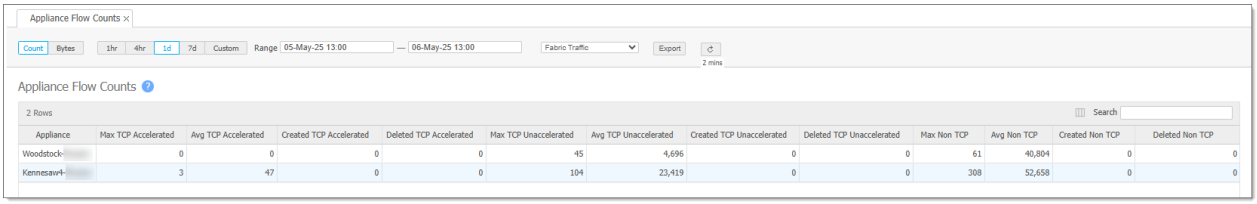
Appliance Flow Counts

Monitoring > Bandwidth > Appliance Flows

To display the counts view of the Appliance Flows tab, click the counts icon.

This view lists the top appliances based on the ones that had the most flows within a selected time period.

You can filter the table data by selecting a traffic type from the drop-down list: Fabric traffic, Pass-through shaped, Pass-through unshaped, or All traffic. When you filter on All traffic, the Created and Deleted columns display the number of new and ended flows for that same time period. Max column values are from a one-minute window within the time range.



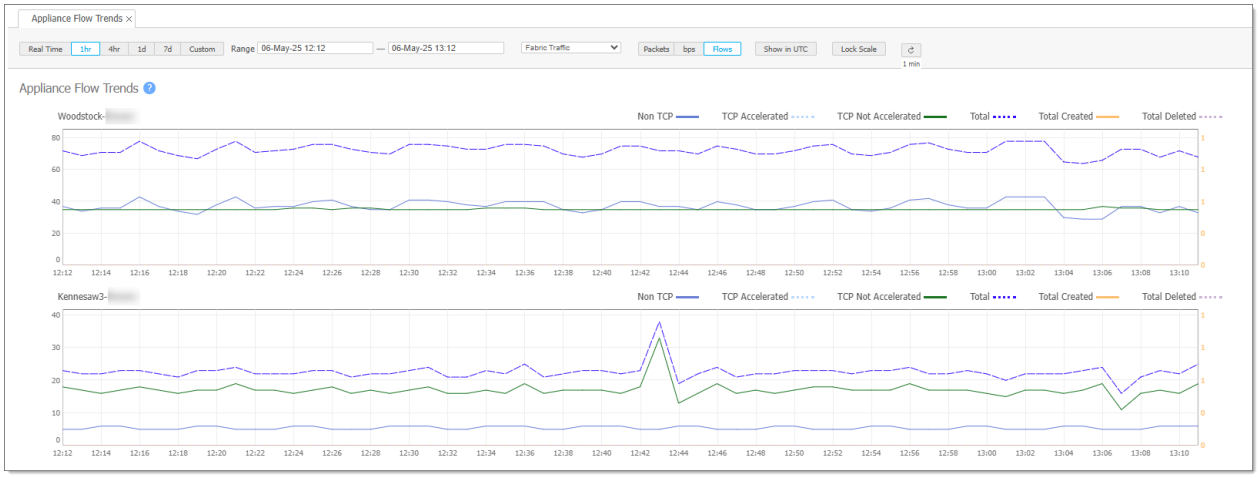
Appliance Flow Trends

Monitoring > Bandwidth > Appliance Flows

To display the trends view of the Appliance Flows tab, click the trends icon.

This view shows the number of flows, packets, and bits per second through the appliance over time. It also differentiates among TCP (accelerated and unaccelerated) flows and non-TCP flows.

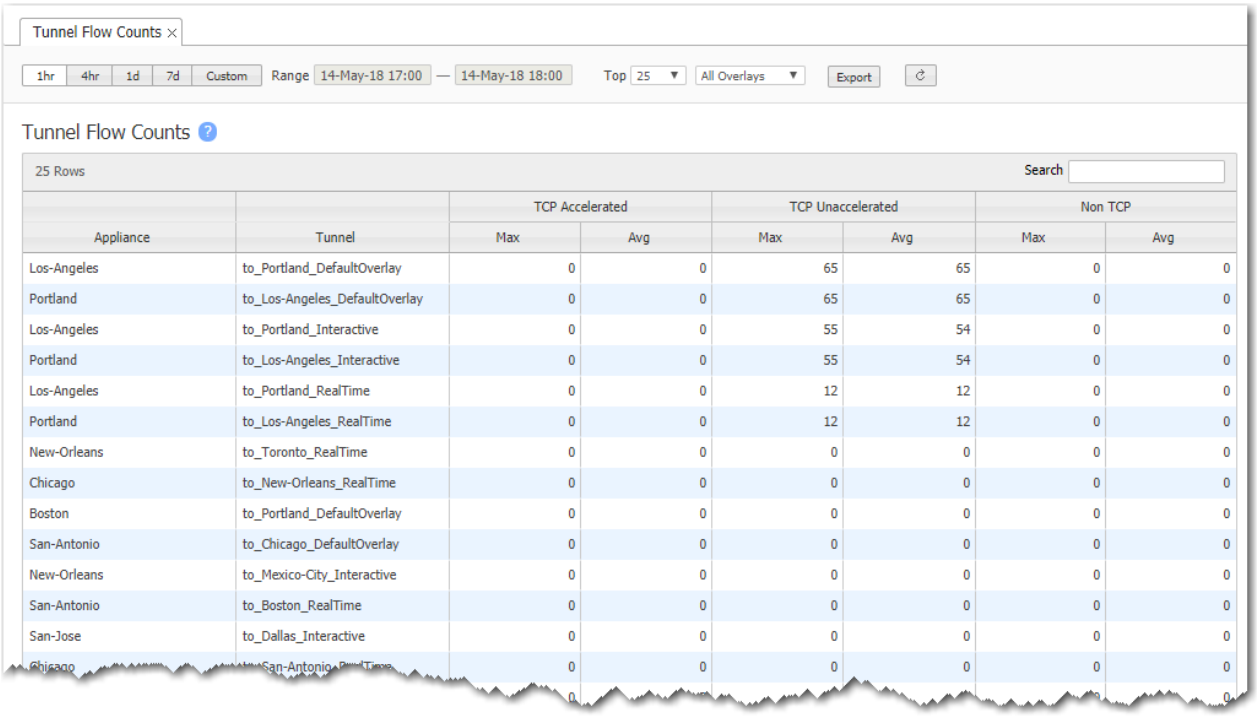
You can filter the chart data by selecting a traffic type from the drop-down list: Fabric traffic, Pass-through shaped, Pass-through unshaped, or All traffic.



Tunnel Flows

Monitoring > Bandwidth > Tunnel Flows

The Tunnel Flows tab lists tunnels with the most flows on average. It differentiates flows into TCP (accelerated and unaccelerated) and non-TCP, and also shows peak values.



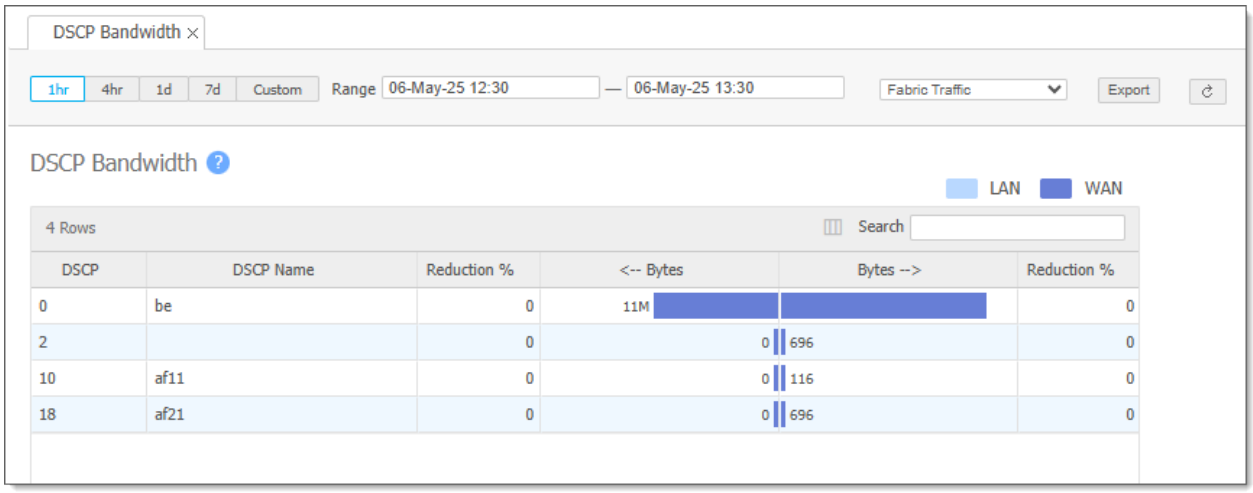
DSCP Bandwidth Summary

Monitoring > *Bandwidth* > *DSCP Bandwidth*

To display the summary view of the DSCP Bandwidth tab, click the summary icon.

This view shows the DSCP classes that are sending the most data.

You can filter the table data by selecting a traffic type from the drop-down list: Fabric traffic, Pass-through shaped, Pass-through unshaped, or All traffic.



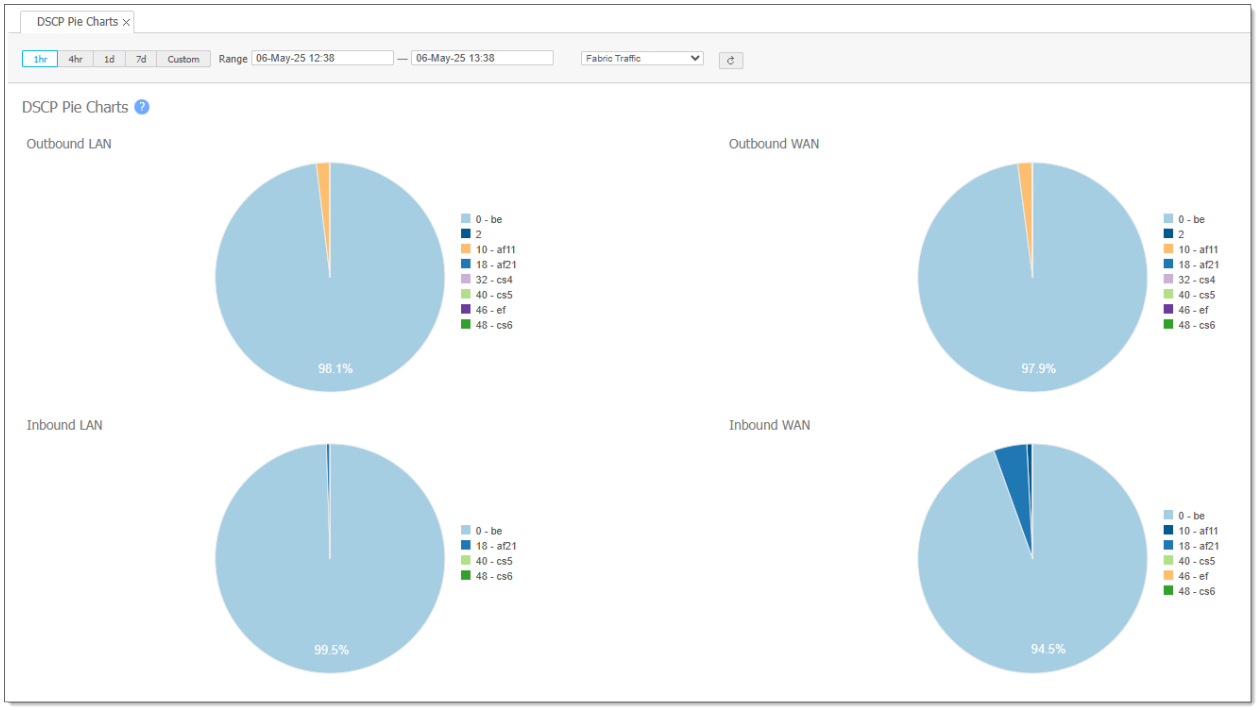
DSCP Bandwidth Pie Charts

Monitoring > Bandwidth > DSCP Bandwidth

To display the pie charts view of the DSCP Bandwidth tab, click the pie charts icon.

This view shows the proportion of traffic in each DSCP class. Hovering over the charts and legends reveals additional information.

You can filter the chart data by selecting a traffic type from the drop-down list: Fabric traffic, Pass-through shaped, Pass-through unshaped, or All traffic.



DSCP Bandwidth Trends

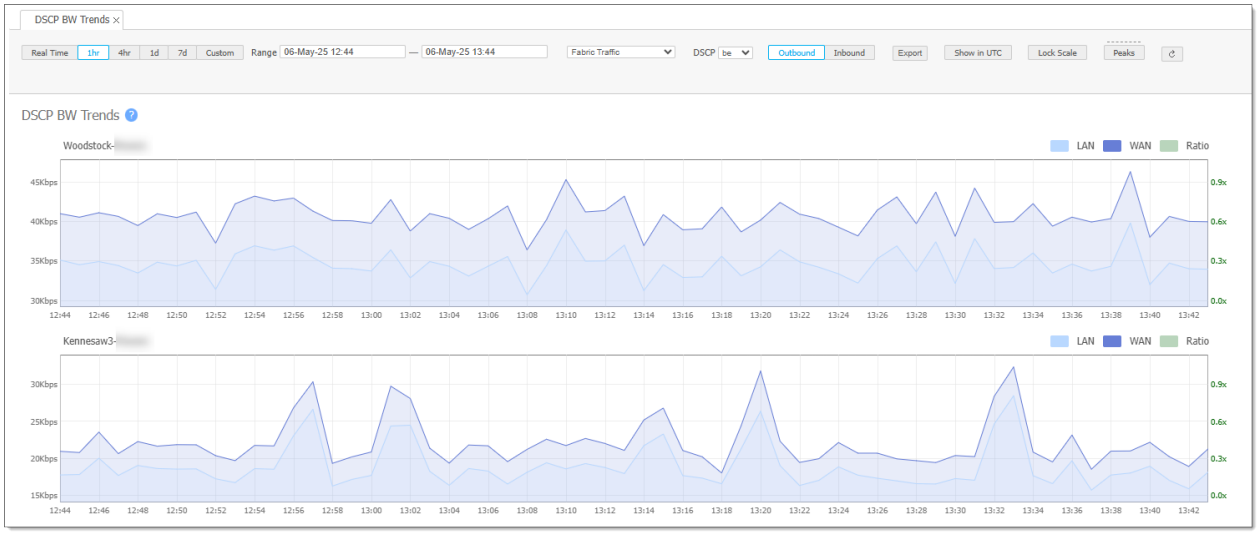
Monitoring > Bandwidth > DSCP Bandwidth

To display the trends view of the DSCP Bandwidth tab, click the trends icon.

Charts in this view of the DSCP Bandwidth tab show bandwidth usage for Differentiated Services Code Point (DSCP) classes in bits per second over time—for both LAN and WAN. By default, the charts show average bandwidth usage recorded within graphed intervals.

To display the highest (peak) bandwidth usage per DSCP class that occurred during each interval, click **Peaks**. This helps to detect and troubleshoot short traffic bursts of high bandwidth used by a DSCP class and ensure your Quality of Service (QoS) policies are effective and working as intended, which average bandwidth statistics might not reveal. Use peak bandwidth data to verify that your network can handle peak loads for critical traffic classes and identify if or when bandwidth exceeded expectations. In real-time mode, peaks are always shown.

You can filter the chart data by selecting a traffic type from the drop-down list: Fabric traffic, Pass-through shaped, Pass-through unshaped, or All traffic.



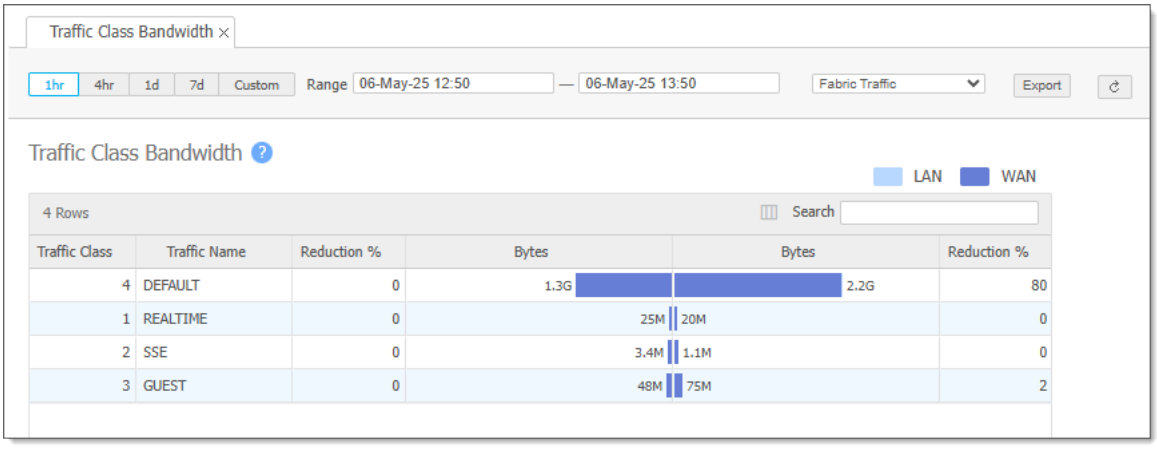
Traffic Class Bandwidth Summary

Monitoring > Bandwidth > Traffic Class Bandwidth

To display the summary view of the Traffic Class Bandwidth tab, click the summary icon.

This view shows the QoS traffic classes that are sending the most data.

You can filter the table data by selecting a traffic type from the drop-down list: Fabric traffic, Pass-through shaped, Pass-through unshaped, or All traffic.



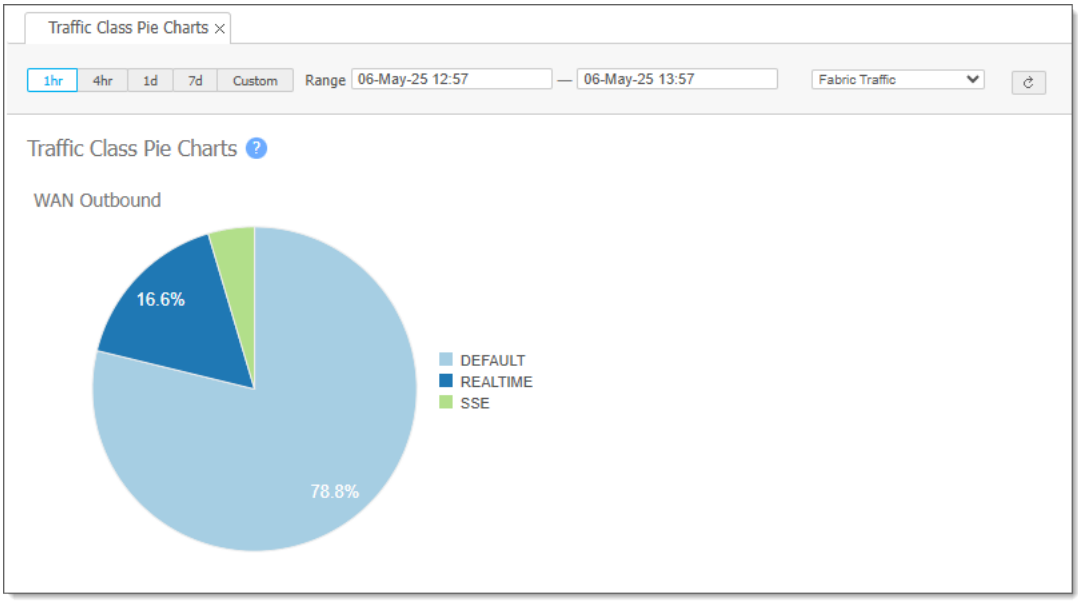
Traffic Class Bandwidth Pie Charts

Monitoring > Bandwidth > Traffic Class Bandwidth

To display the pie charts view of the Traffic Class Bandwidth tab, click the pie charts icon.

This view shows the proportion of traffic in each Traffic class. Hovering over the charts and legends reveals additional information.

You can filter the chart data by selecting a traffic type from the drop-down list: Fabric traffic, Pass-through shaped, Pass-through unshaped, or All traffic.



Shaper Bandwidth Summary

Monitoring > Bandwidth > Shaper Bandwidth

To display the summary view of the Shaper Bandwidth tab, click the summary icon.

This view displays a shaper bandwidth summary for all traffic classes on selected appliances. The shaper delays certain packet types to optimize overall network performance. For more information about shaping, see Shaper Tab and Shaper Template.

- Use the controls above the table to specify how much data—time and date range—you want to see in the summary.
- Use the **Top** filter to limit data according to top applications by total traffic bytes. You can include the top 10, 25, 50, 100, or 1000 applications.
- Click **Outbound** or **Inbound** to change the summary by traffic direction.

The following information is included in the shaper bandwidth summary:

Field	Description
Appliance	Name of the appliance that is shaping traffic to generate the shaper bandwidth summary.
Traffic class	Traffic classes defined by shaper parameters. The following four are pre-configured by Orchestrator: RealTime, CriticalApps, BulkApps, and Default. The user can configure the remaining six classes.
Total bytes	Total amount of bytes being shaped.
Shaped bytes	Amount of bytes used for shaping.
Shaped packets	Amount of packets used for shaping.
Average wait time (ms)	Specified amount of time Orchestrator waits until packets are dropped while shaping is in progress.
Drop packets	Amount of packets that have been reported as dropped due to expiration in the shaper queue.
Other drops	Refers to all other drops besides the expired drop packets.
Trends	Click the graph icon to see the Shaper Bandwidth Trends charts, which show Inbound and Outbound traffic trends in graphs.

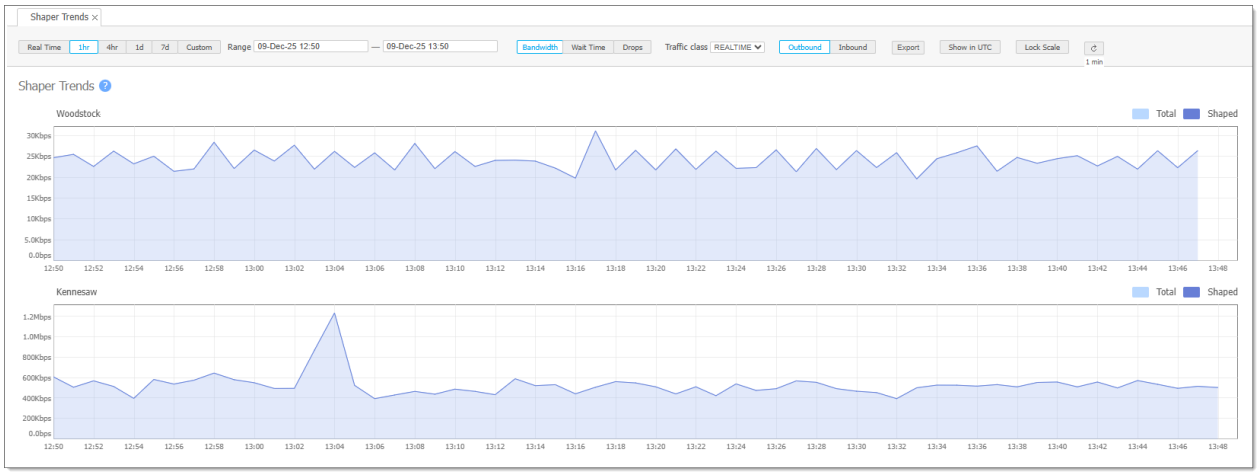
Shaper Bandwidth Trends

Monitoring > *Bandwidth* > *Shaper Bandwidth*

To display the trends view of the Shaper Bandwidth tab, click the trends icon.

This view provides a detailed overview of traffic management for selected appliances in your network. The charts display key metrics related to how network traffic is shaped and controlled over time for a selected traffic class. It shows the total volume of traffic arriving at the shaper, how much traffic is allowed to pass through, the packets dropped due to shaping constraints, and the average delay that packets experience while queued.

These statistics can help you monitor the effectiveness of traffic shaping policies, identify congestion or bottlenecks, and optimize network performance.



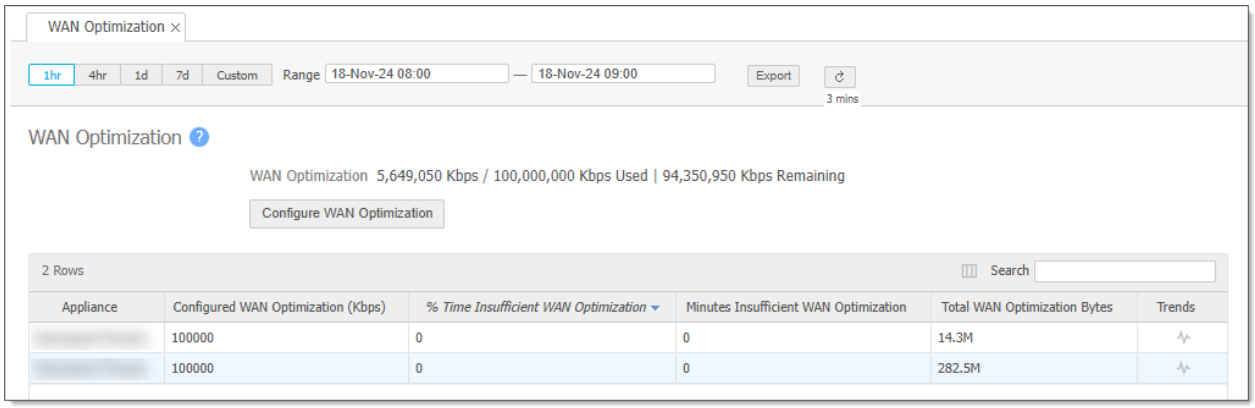
The following terms relate to the shaper feature.

Term	Description
Total bytes	Total number of bytes that arrived at the shaper during the selected period. This is the aggregate of shaped and dropped bandwidth.
Shaped bytes	Number of bytes allowed through the shaper during the selected period according to configured shaping policy. This is the bandwidth that passed through the shaper and was transmitted.
Shaped packets	Number of packets allowed through the shaper during the selected period according to configured shaping policy. These are the packets that passed through the shaper and were transmitted.
Average wait time (ms)	Average time, in milliseconds, that packets spent waiting in the shaper queue before being transmitted or dropped during the selected period.
Drop packets	Number of packets dropped by the shaper queue during the selected period. Packets are typically dropped because they expire after exceeding the allowed queue time.

WAN Optimization

Monitoring > Bandwidth > WAN Optimization

The WAN Optimization tab provides details about WAN Optimization configuration and usage for selected appliances.



You can change the time period for which to display WAN Optimization statistics. Click the **1hr**, **4hr**, **1d**, or **7d** button; or click **Custom** to specify a custom date and time period in the Range fields.

Field	Description
Appliance	Appliance name.
Configured WAN Optimization (Kbps)	WAN Optimization bandwidth (in kilobits per second) configured on the appliance.
% Time insufficient WAN Optimization	Percentage of time that WAN Optimization bandwidth was deemed insufficient by the WAN Optimization engine.
Minutes insufficient WAN Optimization	Amount of time (in minutes) that WAN Optimization bandwidth was deemed insufficient by the WAN Optimization engine.
Total WAN Optimization bytes	Total WAN Optimization bandwidth used over the specified time period.
Trends	Click the graph icon to display the WAN Optimization Trends chart, which shows Configured WAN Optimization, WAN Optimization, and Seconds Insufficient WAN Optimization trends for the appliance. Data represents activity during the specified time period.

When an application needs acceleration, the WAN Optimization engine increases throughput bandwidth to the licensed WAN Optimization amount. For example, if the licensed WAN Optimization limit is 300 Mbps, the engine can accelerate traffic to that amount per second. If the engine determines that it could accelerate traffic to higher speeds if more licensed bandwidth were available, it sends 300 megabits in the first second interval, 300 megabits in the next second interval, and marks that last interval as having insufficient WAN Optimization. Because WAN Optimization statistics are recorded on a per-minute basis, it marks the entire minute as having insufficient WAN Optimization.

Based on a licensed WAN Optimization limit of 300 Mbps:

- Transactional data that requires a transfer of, for example, just 75 MB (600 megabits) of data would result in the transfer of 300 megabits in the first second interval and 300 megabits in the next second interval with that last interval being marked as having insufficient WAN Optimization. WAN Optimization statistics would indicate one minute of insufficient WAN Optimization, even though only one second is actually marked as having insufficient WAN Optimization. In this case, you probably do not need to increase your WAN Optimization bandwidth license.
- Continuous transfers of data, such as occurs for backups and replications over a period of hours or days, require continuous WAN Optimization. In this case, WAN Optimization statistics, such as Minutes Insufficient WAN Optimization, are more precise than in the previous transactional data example. For example, if the transfer of 1 terabyte of data is needed, a licensed WAN Optimization limit of 300 Mbps would be insufficient because the WAN Optimization engine could speed this up to whatever the link speed allows.

Total WAN Optimization bandwidth available to your network is controlled by your license. You can purchase additional WAN Optimization bandwidth if needed. If you have licensed WAN Optimization bandwidth available, you can assign WAN Optimization to appliances on the Licenses tab or on an appliance's Deployment page. You can also configure WAN Optimization allocation by using Business Intent Overlays.

NOTE: Your network uses a single queue for WAN Optimization across all appliances. When that queue is completely utilized, appliances will have insufficient WAN Optimization for any additional demand.

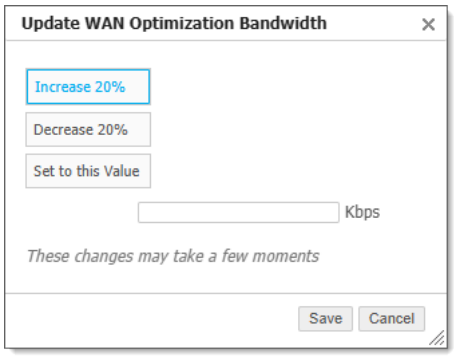
The WAN Optimization license functions as a shaper. When the WAN Optimization license is exceeded, TCP traffic is queued and shaped to optimize the bandwidth. Not having enough WAN optimization has an adverse effect on TCP applications. It is highly recommended that you obtain an amount of WAN Optimization equal to your entire WAN bandwidth.

For UDP traffic, not having enough WAN Optimization causes the traffic to be sent unoptimized. Because optimized UDP traffic involves compression of data, the traffic is sent uncompressed.

You can identify applications that use the WAN Optimization engine, and then use Access Control Lists (ACLs) and rules to permit or deny those applications from WAN Optimization. To identify your top application-related flows, use the Flows tab (Monitoring > Bandwidth > Flows > Active & Recent Flows) to list flows by total bytes sent.

Change WAN Optimization Configuration

To change WAN Optimization configuration, select one or more appliances in the table, and then click **Configure WAN Optimization**. The Update WAN Optimization Bandwidth dialog box opens.



Increase or decrease WAN Optimization bandwidth by 20%, or set it to a specific value in Kbps. Click **Save** to apply changes, or click **Cancel** to not apply changes and close the dialog box.

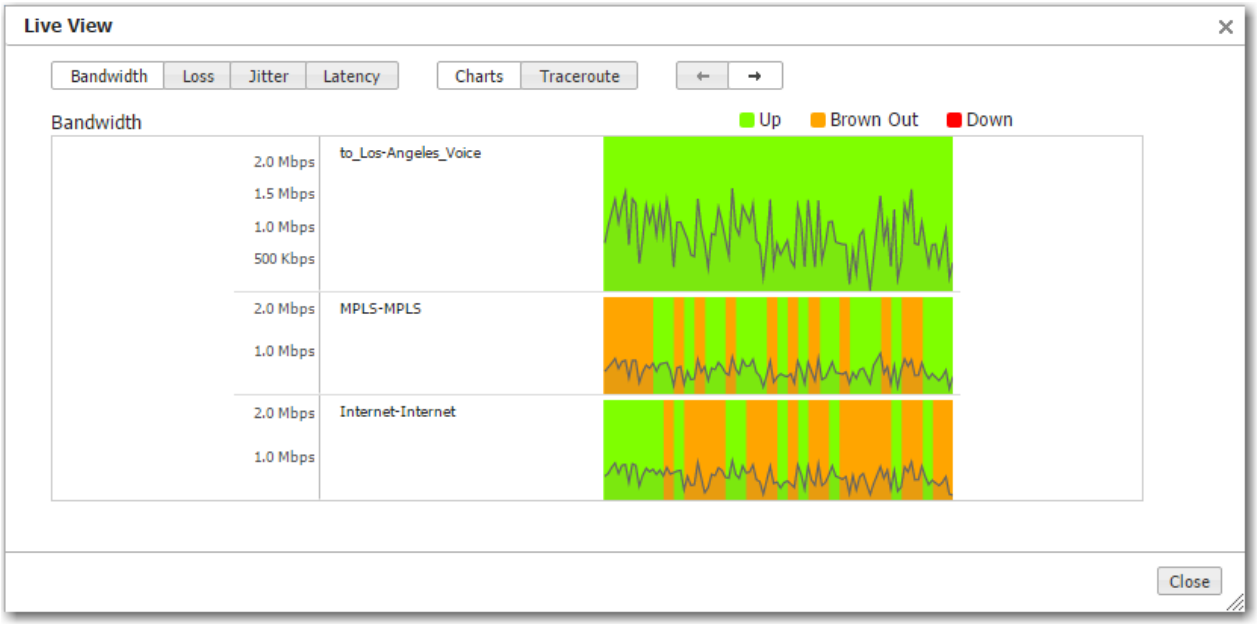
Monitoring > Tunnel Health

The options under **Monitoring > Tunnel Health** focus on reports related to tunnel health.

Live View

Monitoring > Tunnel Health > Live View

Live View shows the live bandwidth, loss, latency, and jitter on all tunnels. For an overlay, it also shows live tunnel states—**Up**, **Browned Out**, or **Down**.



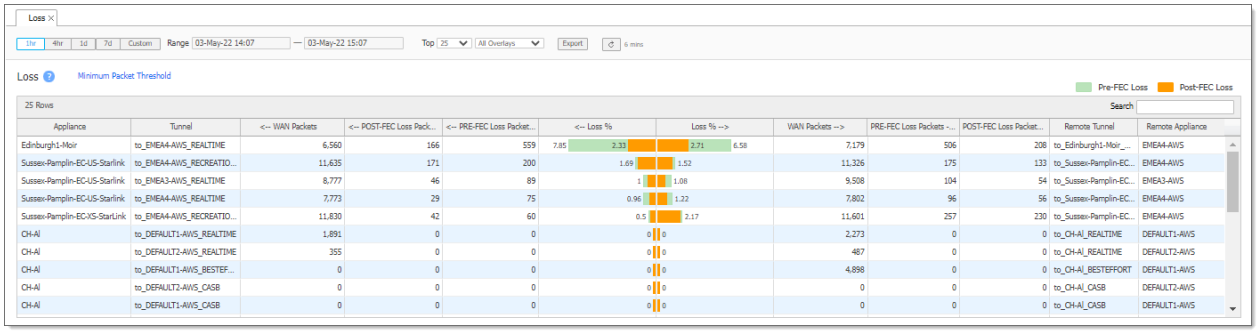
LiveView shows in real time how synergy is created to maintain coverage. The real-time chart shows the SD-WAN overlay at the top and the underlay networks at the bottom. The overlay is green and delivering consistent application performance while both underlays are in persistent brown-out state.

Loss Summary

Monitoring > Tunnel Health > Loss

To display the summary view of the Loss tab, click the summary icon.

This view shows tunnels that have the most dropped packets. Statistics are summarized for the selected time period.



Loss percentages, before and after Forward Error Correction (FEC), are determined by data that the local EdgeConnect observes. Two types of loss are measured:

- Pre-FEC Loss % – Percent of data packets lost before applying FEC / Total sent packets. This measure indicates what the packet loss would be if FEC were not applied.
- Post-FEC Loss % – Percent of data packets lost after applying FEC / Total sent packets. This measure indicates what the packet loss is after FEC is applied.

The total number of sent packets over the link is calculated based on three parameters:

- Total received packets (SUM_WRX_PKTS)
- Recovered packets from FEC (CORRECTED_PACKETS)
- Unrecovered packets after FEC (SUM_POST_LOSS)

Calculations are based on the following formulas:

- Total sent packets = SUM_WRX_PKTS + CORRECTED_PACKETS + SUM_POST_LOSS
- Packets lost in transmission (SUM_PRE_LOSS) = CORRECTED_PACKETS + SUM_POST_LOSS

Based on the above information, the Pre-FEC and Post-FEC Loss percentages are calculated as follows:

- Pre-FEC Loss (%) = SUM_PRE_LOSS * 100 / (SUM_WRX_PKTS + SUM_PRE_LOSS)
- Post-FEC Loss (%) = SUM_POST_LOSS * 100 / (SUM_WRX_PKTS + SUM_PRE_LOSS)

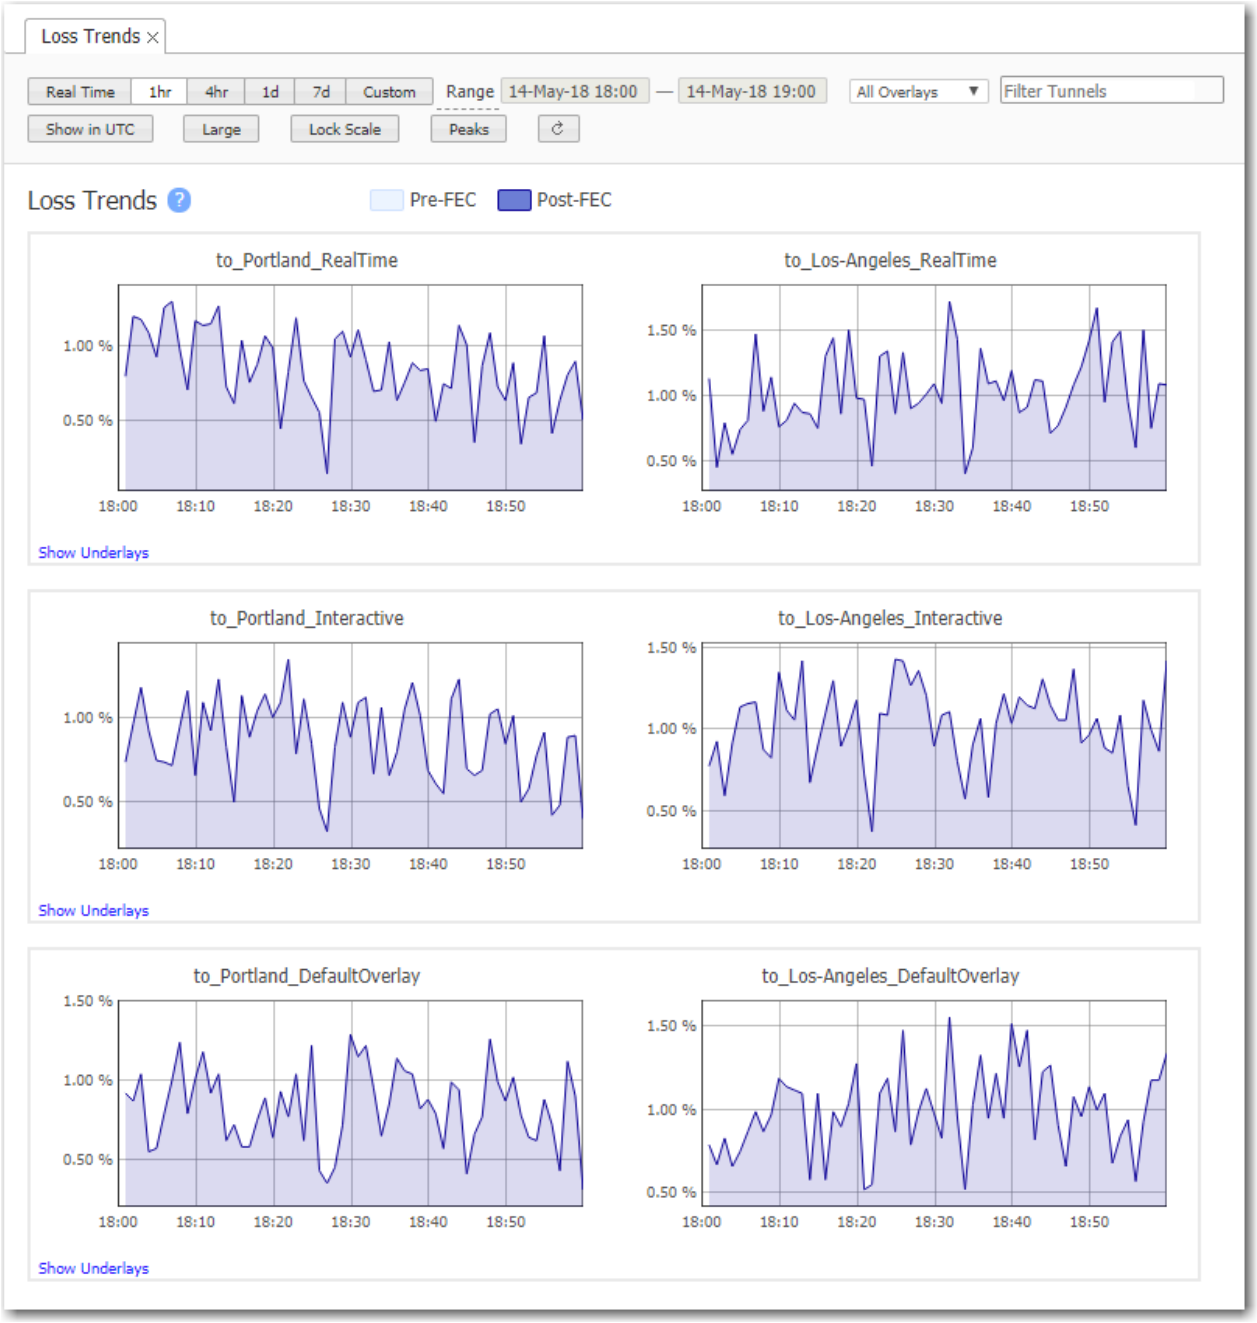
Loss Trends

Monitoring > *Tunnel Health* > *Loss*

To display the trends view of the Loss tab, click the trends icon.

Charts in this view of the Loss tab show tunnel packet loss as a percentage over time, before and after Forward Error Correction (FEC). By default, the charts show average packet loss recorded within graphed intervals.

To display the highest (peak) packet loss events that occurred during each interval, before and after FEC, click **Peaks**. This helps to quickly identify and investigate network performance issues that might be hidden by average statistics. Use peak loss data to spot significant spikes that could indicate transient network issues, congestion, or outages. In real-time mode, peaks are always shown.



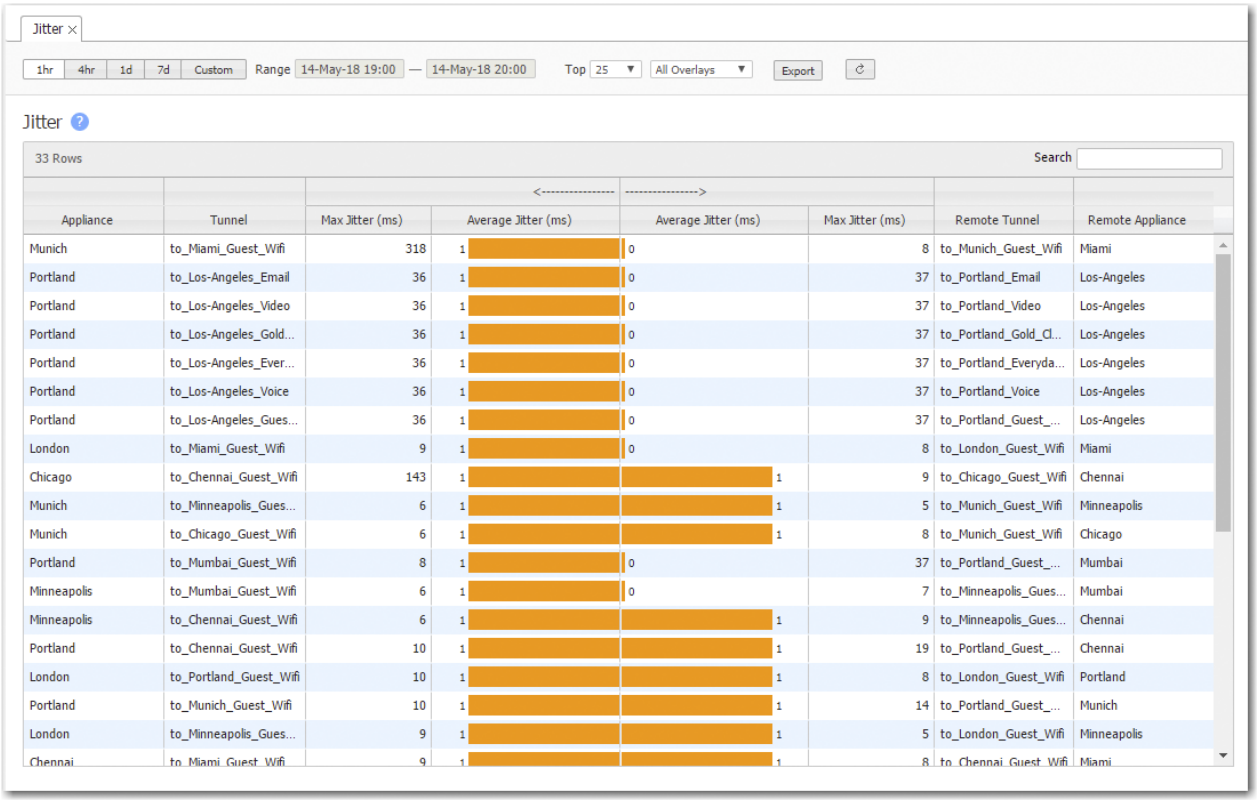
NOTE: Underlay tunnels are a shared resource among overlays. Therefore, underlay charts display aggregated data.

Jitter Summary

Monitoring > *Tunnel Health* > *Jitter*

To display the summary view of the Jitter tab, click the summary icon.

This view shows the tunnels that have the most jitter. Statistics are summarized for the selected time period. Jitter can be caused by congestion in the LAN, firewall routers, bottleneck access links, load sharing, route flapping, routing table updates, and timing drifts.



For jitter trend interpretation guidelines, see Jitter Trends

For more details about all tunnel monitoring statistics, see Tunnel Summary

Jitter Trends

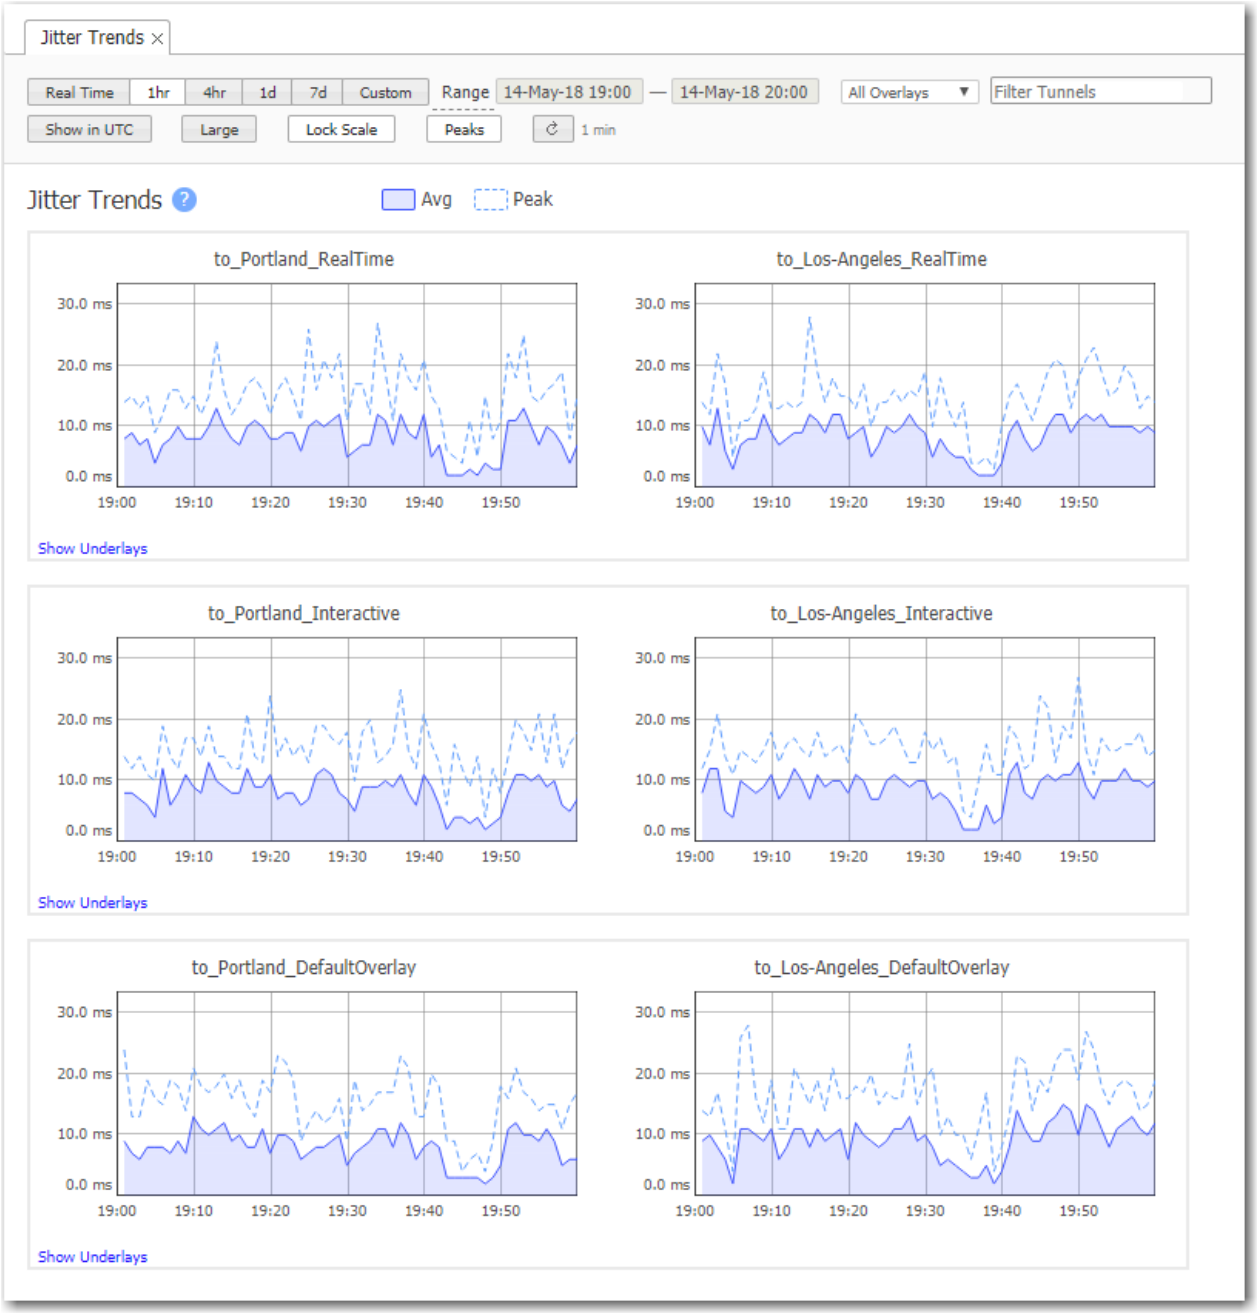
Monitoring > *Tunnel Health* > *Jitter*

To display the trends view of the Jitter tab, click the trends icon.

Charts in this view of the Jitter tab show tunnel jitter in milliseconds over time. By default, the charts show average jitter recorded within graphed intervals.

Jitter is the variation in the time between packets arriving, caused by network congestion, timing drift, or route changes. Achieving low jitter values is critical to maintaining the quality of real-time applications, especially voice, video, and other latency-sensitive services.

To display the highest (peak) jitter that occurred during each interval, click **Peaks**. This helps to quickly spot periods where network jitter spiked, which average jitter statistics might not reveal. If you look only at averages, short bursts of high jitter might go unnoticed, even though they can disrupt voice or video streams. In real-time mode, peaks are always shown.



NOTE: Underlay tunnels are a shared resource among overlays. Therefore, underlay charts display aggregated data.

Jitter Trend Interpretation Guidelines

Follow the guidelines below to interpret jitter trends in your network tunnel traffic and take appropriate actions to maintain optimal network and application performance.

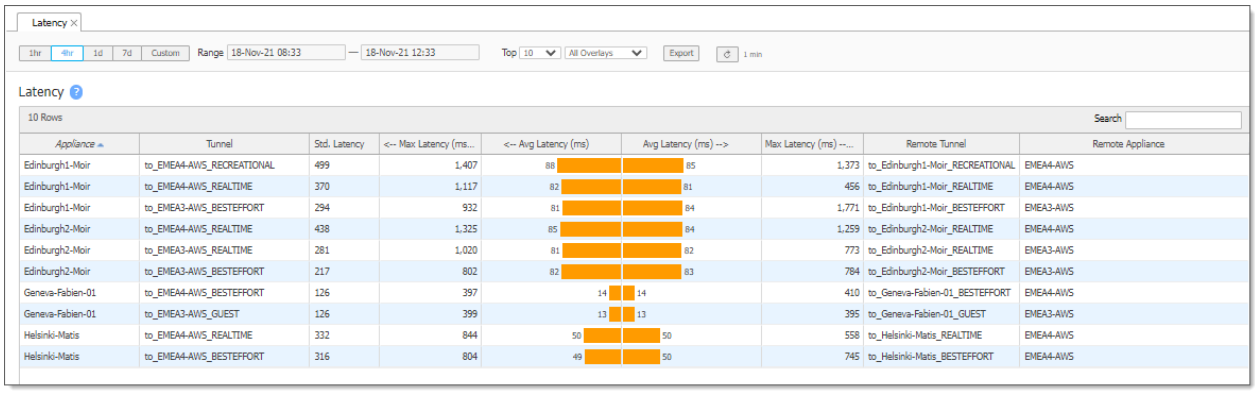
- Set Baselines: Establish a baseline for acceptable jitter levels based on the requirements of your applications. For example, VoIP typically requires jitter to be below 30 milliseconds for good call quality.
- Monitor Trends: Use monitoring tools to continuously track jitter values over time. Look for patterns or spikes that deviate from the baseline.
- Identify Patterns: Consistently high jitter indicates chronic network issues such as persistent congestion, poor routing, or inadequate bandwidth.
- Intermittent Jitter Spikes: May suggest temporary issues like traffic bursts, transient congestion, or route flapping.
- Gradual Increase in Jitter: Could indicate a growing network load or emerging congestion problems that need to be addressed.
- Correlate with Other Metrics: Check other performance metrics such as latency, packet loss, and throughput to gain a comprehensive understanding of network health. For example, high jitter combined with high packet loss may indicate severe congestion or faulty network equipment.
- Analyze Traffic Patterns: Determine whether jitter issues correlate with specific times of day, specific applications, or particular traffic patterns. This can help identify whether the problem is related to peak usage times or specific application demands.
- Investigate Network Links: Examine the performance of individual network links and paths. High jitter on a particular link may indicate a need for troubleshooting or reconfiguration. Consider performing a path trace to identify where jitter is introduced along the route.
- Quality of Service (QoS) Policies: Ensure that QoS policies are correctly configured and prioritize real-time traffic. Misconfigured QoS can lead to increased jitter and degraded application performance. Review and adjust QoS settings as necessary to ensure proper prioritization and handling of sensitive traffic.
- Capacity Planning: Assess whether your current network capacity meets the demands of your applications. If necessary, plan for additional bandwidth or infrastructure upgrades to alleviate congestion and reduce jitter.
- Implement Redundancy and Failover: Ensure that your configuration includes redundancy and failover mechanisms to maintain performance during link failures or congestion. Utilize multiple paths to distribute traffic and minimize the impact of jitter on critical applications.
- Use Advanced Features: Leverage dynamic path selection and application-aware routing to optimize traffic flows and reduce jitter.

Latency Summary

Monitoring > *Tunnel Health* > *Latency*

To display the summary view of the Latency tab, click the summary icon.

This view shows summary statistics for latency (transmission delay) on an in-band, end-to-end tunnel basis for the selected time/date range. Either overlay or underlay tunnels can be displayed, and anywhere between the top 10 to top 1000 tunnels are displayed by round-trip time (RTT).



Appliance	Tunnel	Std. Latency	<-- Max Latency (ms...	<-- Avg Latency (ms)	Avg Latency (ms) -->	Max Latency (ms) --...	Remote Tunnel	Remote Appliance
Edinburgh1-Moir	to_EMEA4-AWS_RECREATIONAL	499	1,407	88	85	1,373	to_Edinburgh1-Moir_RECREATIONAL	EMEA4-AWS
Edinburgh1-Moir	to_EMEA4-AWS_REALTIME	370	1,117	82	81	456	to_Edinburgh1-Moir_REALTIME	EMEA4-AWS
Edinburgh1-Moir	to_EMEA3-AWS_BESTEFFORT	294	932	81	84	1,771	to_Edinburgh1-Moir_BESTEFFORT	EMEA3-AWS
Edinburgh2-Moir	to_EMEA4-AWS_REALTIME	438	1,325	85	84	1,259	to_Edinburgh2-Moir_REALTIME	EMEA4-AWS
Edinburgh2-Moir	to_EMEA3-AWS_REALTIME	281	1,020	81	82	773	to_Edinburgh2-Moir_REALTIME	EMEA3-AWS
Edinburgh2-Moir	to_EMEA3-AWS_BESTEFFORT	217	802	82	83	784	to_Edinburgh2-Moir_BESTEFFORT	EMEA3-AWS
Geneva-Fabien-01	to_EMEA4-AWS_BESTEFFORT	126	397	14	14	410	to_Geneva-Fabien-01_BESTEFFORT	EMEA4-AWS
Geneva-Fabien-01	to_EMEA3-AWS_GUEST	126	399	13	13	395	to_Geneva-Fabien-01_GUEST	EMEA3-AWS
Helsinki-Matis	to_EMEA4-AWS_REALTIME	332	844	50	50	558	to_Helsinki-Matis_REALTIME	EMEA4-AWS
Helsinki-Matis	to_EMEA4-AWS_BESTEFFORT	316	804	49	50	745	to_Helsinki-Matis_BESTEFFORT	EMEA4-AWS

In this view, latency is a measure of the RTT within a tunnel in milliseconds. Values on the left display RTT as measured by the local appliance. Values on the right display RTT as measured by the appliance at the remote end of the tunnel.

Some column descriptions follow:

- Std. Dev. – Standard deviation (in milliseconds) of latency values for the tunnel within the specified period.

  Standard deviation is a measure of the amount of variation in a set of values. Low standard deviation indicates that the values tend to be close to the mean or expected value while a high standard deviation indicates that the values are spread over a wider range.
- Max latency (ms) – Maximum RTT value (in milliseconds) for the tunnel within the specified range.
- Avg latency (ms) – Average RTT value (in milliseconds) for the tunnel within the specified range.

High latency can negatively affect throughput in the network, most noticeably for TCP traffic. Physical distance has the most significant impact on latency. For example:

- If data is crossing the United States, you can expect delays from 60 to 120 milliseconds.
- International transmissions can normally experience delays up to 200 milliseconds.
- Satellite transmissions often have delays of about 1/2 second, and up to several seconds are possible.

High latency can also be caused by equipment (hop-by-hop delays), or by loss or congestion resulting from lost packets, lost acknowledgments, and necessary retransmissions.

TCP Acceleration (a function of WAN Optimization) can mitigate the impact of latency on throughput. In addition, path conditioning and packet re-ordering (a function of Business Intent Overlay link bonding) can mitigate the impact of loss and out-of-order packets on TCP throughput by reducing the number of retransmissions.

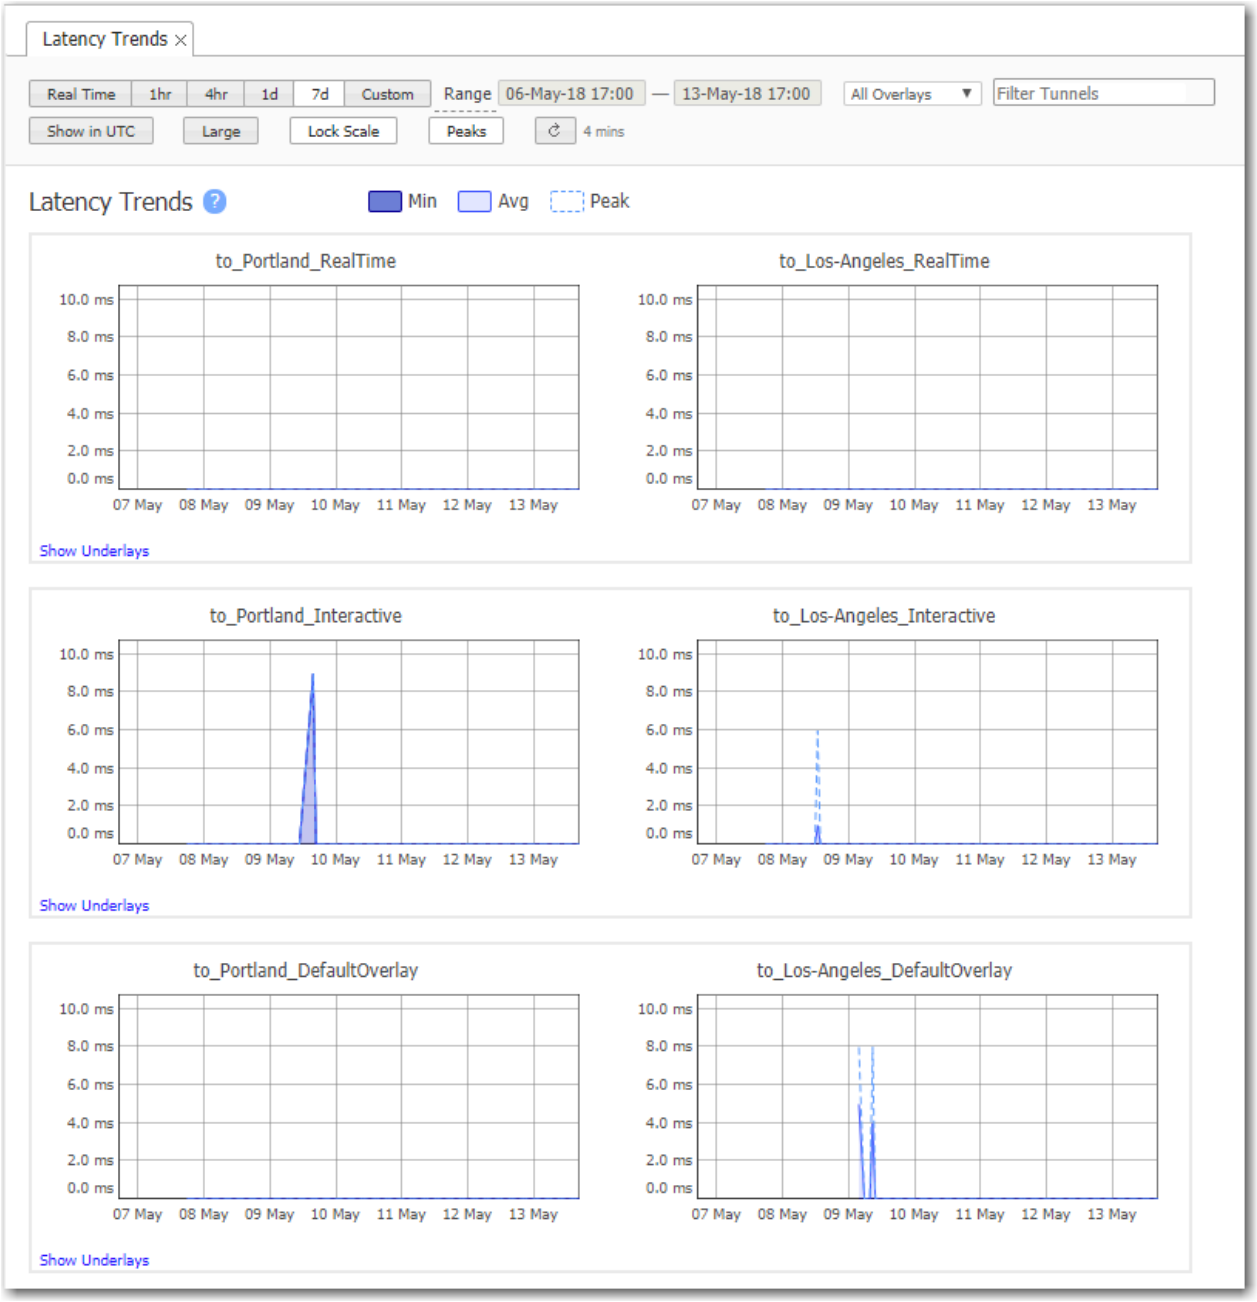
Latency Trends

Monitoring > *Tunnel Health* > *Latency*

To display the trends view of the Latency tab, click the trends icon.

Charts in this view of the Latency tab show tunnel latency in milliseconds over time. By default, the charts show average and minimum latency recorded within graphed intervals.

To display the highest (peak) latency incidents that occurred during each interval, click **Peaks**. This helps to quickly detect and address sporadic performance issues that average latency statistics might not reveal. Use peak latency data to spot significant spikes that could indicate transient network congestion, routing changes, or other issues. In real-time mode, peaks are always shown.



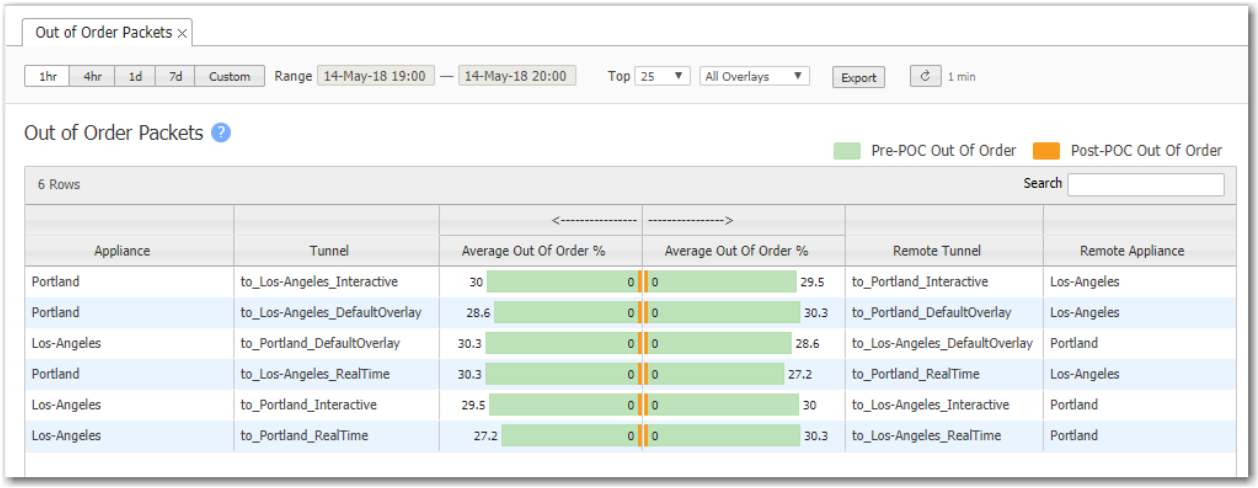
NOTE: Underlay tunnels are a shared resource among overlays. Therefore, underlay charts display aggregated data.

Out of Order Packets Summary

Monitoring > Tunnel Health > Out of Order Packets

To display the summary view of the Out of Order Packets tab, click the summary icon.

This view shows the tunnels that receive the most packets out of sequence relative to how they were sent.



Average out of order packets over the link is calculated based on three parameters:

- Out of order packets before packet order correction (SUM_PRE_POC)
- Out of order packets after packet order correction (SUM_POST_POC)
- Total received packets (SUM_WRX_PKTS)

Calculations are based on the following formulas. For these calculations, the (SUM_WRX_PKTS) must be greater than the threshold of 6000 minimum WRX packets. You can configure the threshold value from the Loss Summary tab.

- Inbound (WAN to LAN) average Out of Order Packets before Packet Order Correction (%) = SUM_PRE_POC * 100 / SUM_WRX_PKTS
- Outbound (LAN to WAN) average Out of Order Packets after Packet Order Correction (%) = SUM_POST_POC * 100 / SUM_WRX_PKTS

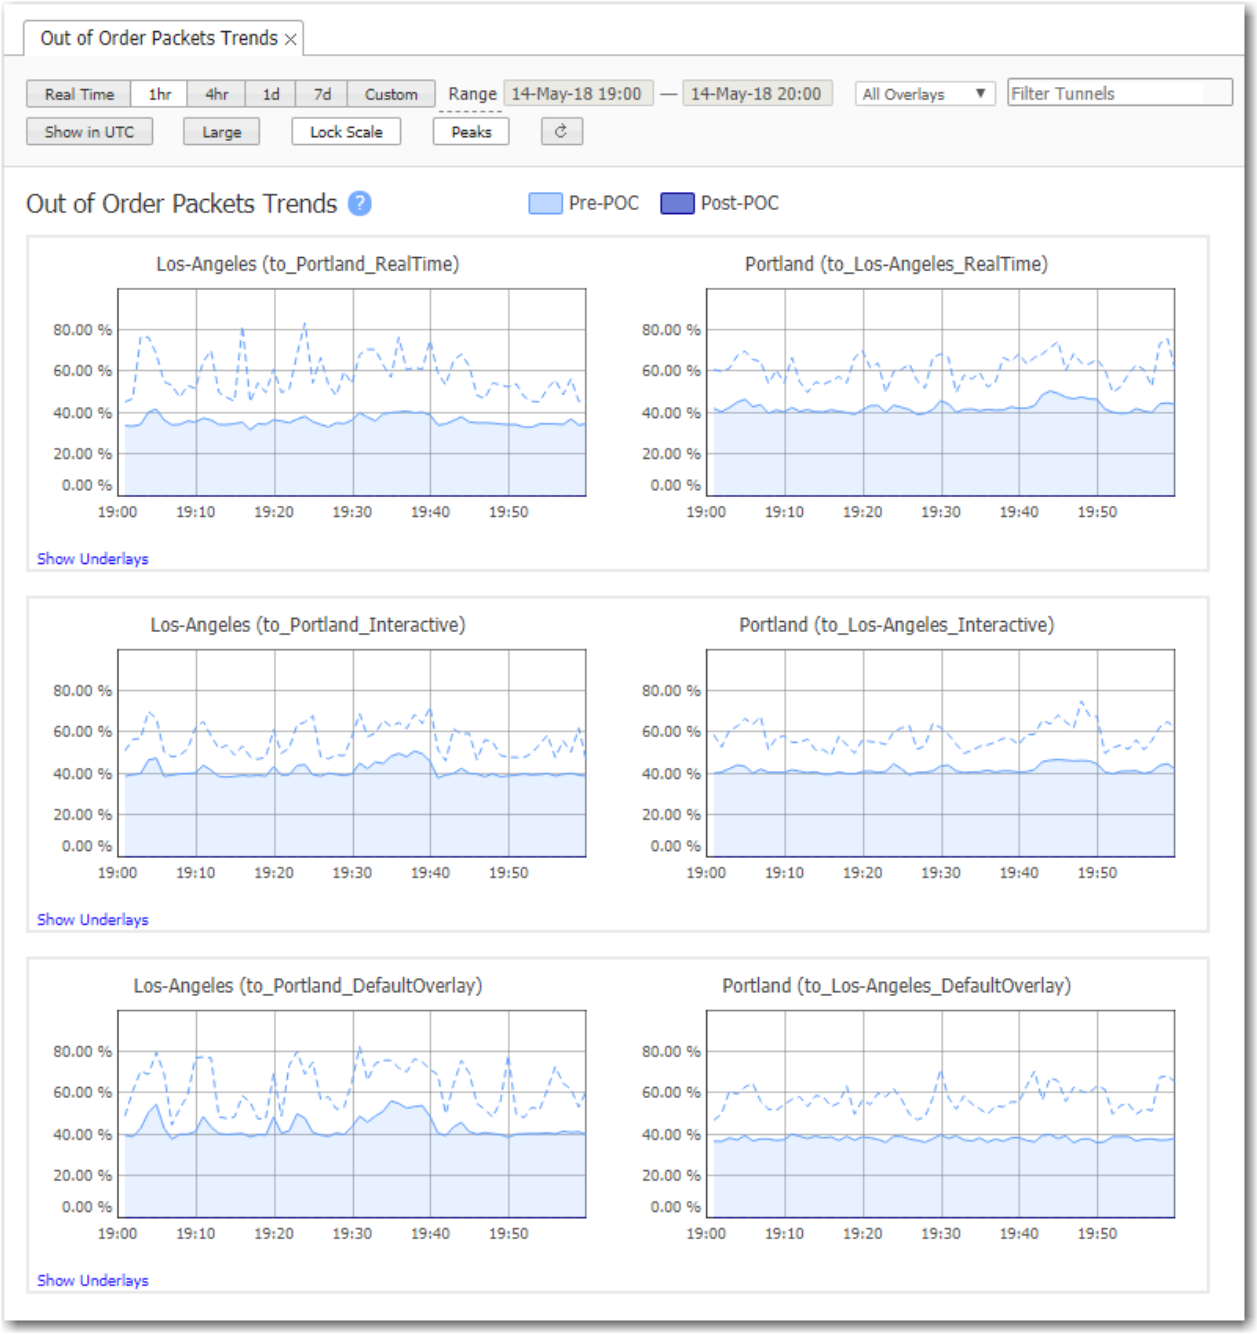
Out of Order Packets Trends

Monitoring > Tunnel Health > Out of Order Packets

To display the trends view of the Out of Order Packets tab, click the trends icon.

Charts in this view of the Out of Order Packets tab show tunnel packets that are out of order as a percentage over time, before and after Packet Order Correction (POC). By default, the charts show average out-of-order packet counts recorded within graphed intervals.

To display the highest (peak) out-of-order packet counts that occurred during each interval, click **Peaks**. This helps to quickly detect and address sporadic network problems that can affect application performance, which average out-of-order packet statistics might not reveal. Use peak data to easily identify intervals with significant spikes in out-of-order packets, indicating possible network instability or routing changes. In real-time mode, peaks are always shown.



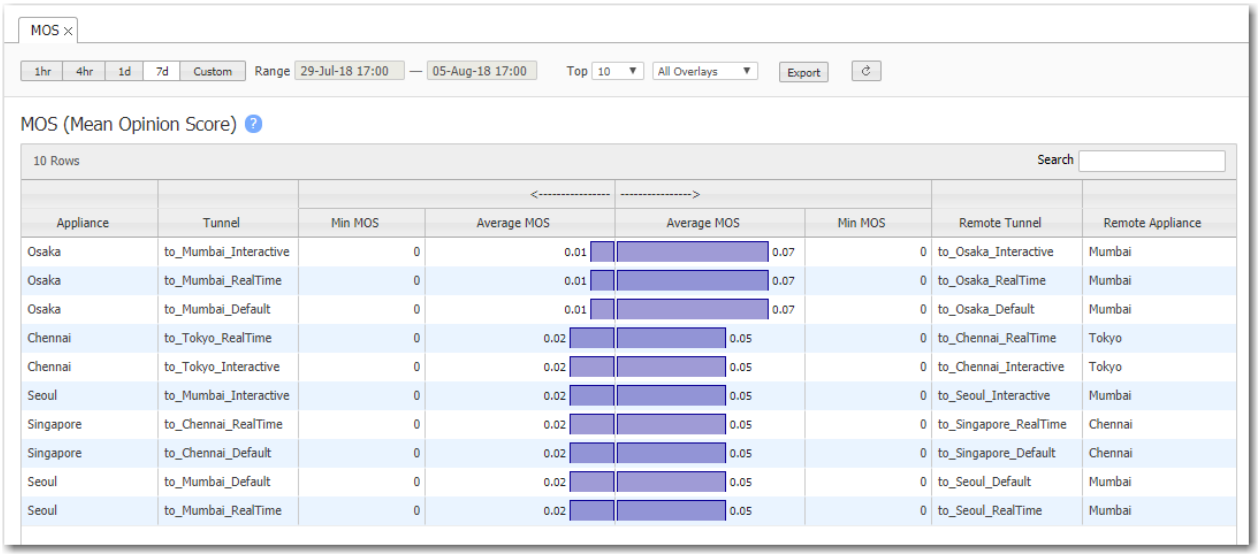
NOTE: Underlay tunnels are a shared resource among overlays. Therefore, underlay charts display aggregated data.

MOS Summary

Monitoring > *Tunnel Health* > *MOS*

To display the summary view of the MOS tab, click the summary icon.

The Mean Opinion Score (MOS) is a commonly used measure for video, audio, and audiovisual quality evaluation. Perceived quality is rated on a theoretical scale of 1 to 5. The higher the number, the better the quality.



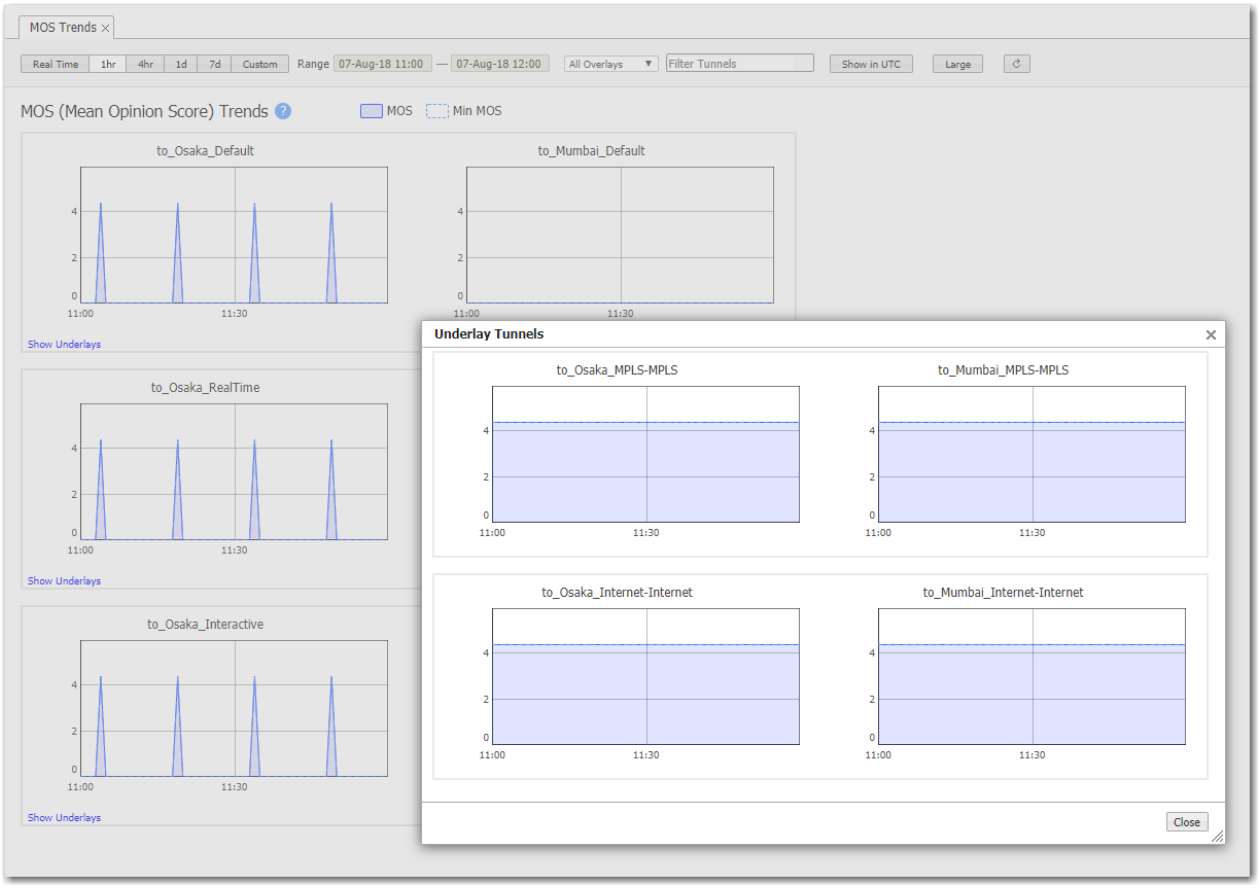
The value can be affected by loss, latency, and jitter. In practice, a value of 4.4 is considered an excellent quality target.

MOS Trends

Monitoring > *Tunnel Health* > *MOS*

To display the trends view of the MOS tab, click the trends icon.

The Mean Opinion Score (MOS) is a commonly used measure for video, audio, and audiovisual quality evaluation. Perceived quality is rated on a theoretical scale of 1 to 5. The higher the number, the better the quality.



- The value can be affected by loss, latency, and jitter. In practice, a value of 4.4 is considered an excellent quality target.
- The Min MOS value reports the worst score within a minute.

Tunnels Summary

Monitoring > Tunnel Health > Tunnels Summary

The Tunnels Summary tab summarizes tunnel statistics, including reduction, throughput, jitter, latency, and packet loss. For each Business Intent Overlay, the specified Link Bonding Policy determines the bandwidth efficiency. The data counts and transmission rates can be influenced by the policies you deploy. To guarantee service quality levels, High Availability requires the most overhead and High Efficiency requires the least. The table shows the total bandwidth used. The Payload filter removes overhead from the displayed values. The values reported are for the time period you specify.

The Tunnels Summary tab includes the columns listed in the table below.

Column	What is Reported	Comments
Tunnel	A logical connection established between two endpoints over a physical network infrastructure.	You can dynamically select the best path for data transmission based on real-time network conditions. Tunnels can be rerouted as needed to maintain optimal performance and reliability or can be configured to prioritize certain types of traffic.
Status	Indicates whether the tunnel is active (up) or inactive (down).	An active status means the tunnel is successfully established and can transmit data. An inactive status indicates a failure or disconnection.
<– LAN (inbound)	Amount of data transmitted to the LAN (post-WAN de-encapsulation).	Reports the amount of post-WAN de-encapsulated data transmitted through the tunnel from the WAN to the LAN.
<– WAN (inbound)	Amount of data received from the WAN (pre-WAN de-encapsulation).	Reports the amount of pre-WAN de-encapsulated data recieved through the tunnel from the WAN.
<– Reduction % (inbound)	Total traffic reduction for inbound traffic on this tunnel.	The calculation is as follows: Reduction % for inbound traffic = 100 (Received from LAN – Transmitted to WAN) / Received from LAN.
<– LAN Throughput (inbound)	Inbound WAN to LAN Traffic Rate Kbps (post-WAN de-encapsulation).	Reports the post-WAN de-encapsulated data transmission rate to the LAN.
<– WAN Throughput (inbound)	Inbound WAN Traffic Rate Kbps (pre-WAN de-encapsulation).	Reports the pre-WAN de-encapsulated data transmission rate through the tunnel.
LAN –> (outbound)	Amount of Data Transmitted from the LAN to the WAN (pre-LAN encapsulation).	Reports the amount of pre-LAN encapsulated data transmitted to the WAN tunnel from the LAN.

Column	What is Reported	Comments
WAN –> (outbound)	Amount of Data Transmitted to the WAN (post-WAN encapsulation).	Reports the amount of post-WAN encapsulated data transmitted through the WAN tunnel. The policies you deploy can effect this amount.
Reduction % –> (outbound)	Total traffic reduction for outbound traffic on this tunnel.	The calculation is as follows: Reduction % for outbound traffic = 100 (Received from LAN – Transmitted to WAN) / Received from LAN. The Reduction % shows the effect of the policies you deploy.
LAN Throughput –> (outbound)	Outbound LAN Traffic Rate Kbps (pre-LAN encapsulation).	Reports the amount of pre-LAN encapsulated data transmitted to the WAN tunnel from the LAN.
WAN Throughput –> (outbound)	Outbound WAN Traffic Rate Kbps (post-WAN encapsulation).	Reports the post-WAN encapsulated data transmission rate (kbits/second) through the tunnel.
Loss(packets)	The percentage of packets that are lost during transmission through the tunnel.	High packet loss can severely impact application performance and reliability.
Max Loss(packets)	The maximum number of packets that are lost during transmission through the tunnel.	Transient high packet loss can cause sporadic application performance and reliability issues.
Avg Jitter (ms)	The average variation in packet arrival times.	High jitter can affect the quality of real-time applications, leading to issues like voice call distortion or video lag.
Max Jitter (ms)	The maximum variation in packet arrival times.	Transient high jitter can cause sporadic degredation in the quality of real-time applications, leading to issues like sporadic voice call distortion or video lag.

Column	What is Reported	Comments
Latency (ms)	Reports the average time it took for a data packet to travel from one end of the tunnel to the other.	Low latency is crucial for real-time applications such as voice and video communications.
Max Latency (ms)	Reports the maximum time it took for a data packet to travel from one end of the tunnel to the other.	Transient high latency can cause sporadic degradation of real-time applications such as voice and video communications.

Configuration

The options under **Configuration** focus on how to configure Orchestrator.

Categories include the following:

- Overlays & Security
- Networking
- Templates & Policies
  - Policies
  - Templates
- Cloud Services

Configuration > Overlays & Security

The options under **Configuration > Overlays & Security** focus on configuring Business Intent Overlays (BIOs), interface labels, hubs, regions, deployment profiles, and internet traffic definitions. Other options are related to security, SSL certificates, appliance configuration and discovery, and licensing.

NOTE: Topics in this section relate to deploying a WAN optimization network or a software-defined Wide Area Network (SD-WAN). From a configuration standpoint, an SD-WAN uses Business Intent Overlays (BIOs), whereas a WANop network does not.

Business Intent Overlays

Configuration > Overlays & Security > Business Intent Overlays

Use the Business Intent Overlays (BIOs) tab to create separate, logical networks that are individually customized to your applications and requirements within your network. By default, there are several predefined overlays matching a range of traffic within your network.

The overlay summary table is used to easy compare values between your various configured overlays. You can select any link in the table and the Overlay Configuration dialog box launches. You can also temporarily save your changes before officially applying those changes to your overlay. An orange box around the edited item indicates the pending configuration updates. Click **Save and Apply Changes to Overlays** when you are ready to apply the changes and click **Cancel** if you want to delete the changes.

Overview

Orchestrator matches traffic to an ACL, progressing down the ordered priority list of overlays until it identifies the first one that matches. The matched traffic is then analyzed against the overlay's internet traffic configuration and forwarded within the fabric, or broken out to the internet based on the preferred policy order. If the software determines that the traffic is not destined for the internet, it refers to the WAN Link Bonding Policy configuration and forwards traffic accordingly within the overlay.

SD-WAN Traffic to Internal Subnets

You can access the Internal Subnets dialog box directly by navigating to **Configuration > Overlays & Security > Internal Subnets** as well as by navigating to **Configuration > Overlays & Security > Business Intent Overlays** and clicking the edit icon next to SD-WAN Traffic to Internal Subnets.

The Consider non-default routes as internal subnets option in the Internal Subnets configuration dialog box determines how the system identifies internal subnets for network orchestration and traffic routing purposes. When this option is selected, the platform treats all non-default routes (routes that are not the default route, typically 0.0.0.0/0) as internal subnets. This means that:

- Any subnet explicitly defined in the routing table, other than the default route, is considered part of the internal network of the organization.
- Traffic destined for these non-default routes is treated as internal traffic and managed accordingly (e.g., for traffic optimization, monitoring, or security policies).
- This can be useful in complex network environments where there are multiple internal subnets connected through specific routing rules, and you want the platform to include them in its orchestration logic.

If this option is unchecked, only the subnets explicitly defined as internal subnets in the configuration are treated as internal, and non-default routes are ignored.

Configuring segment IDs in the Internal Subnets table defines Internet traffic as any traffic that does not match the internal subnets, which can be IPv4 or IPv6 address spaces. Associating segment IDs with internal subnets logically separates traffic into different segments (such as corporate, guest, or IoT traffic). This segmentation enables the application of specific policies or routing behaviors to each segment. Traffic that matches the internal subnets can be tied to its corresponding segment ID, ensuring appropriate handling based on business or security

requirements. Traffic that does not match the internal subnets is automatically classified as Internet traffic. This distinction is critical for applying security and routing policies. Segment IDs allow for granular control over how traffic flows through the network. For example, traffic within internal subnets might be prioritized to use private MPLS links or other secure transport methods. Non-internal (Internet) traffic can take a direct Internet breakout path, reducing latency and offloading traffic from expensive private links. This ensures efficient utilization of network resources and improves application performance.

Overlay Configuration

You can begin to configure or modify a default overlay in the Overlay column. You can also select any icon on the Business Intent Overlay page, and the selected editor or dialog box opens.

Complete the following steps to configure your overlay.

1. Select the name of the overlay. The Overlay Configuration window opens. If you want to edit the default overlay or create a new overlay, enter the new name of the overlay in the **Name** field.
2. Select the **Match** field and choose the match criteria from the menu.
3. Click the edit icon next to the ACL field. To apply default ACLs or create your own, click **Add Rule** in the Associate ACL window.
4. Click **Save**.

Region

To view the associated region within your overlay, click **Regions** in the Region column in the overlay summary table. To modify, remove, or edit overlay settings for a selected region, select it from the **Region** drop-down list at the top right of the Overlay Configuration window. For more information about Regions, refer to the help on the tab.

Topology

Select the type of topology you want to apply to your overlay and network. You can choose between the following types of topology:

- **Mesh:** Choose **Mesh** if you want to make a local network.
- **Hub & Spoke:** Hubs are used to build tunnels in Hub & Spoke networks and route traffic between regions. If you choose **Hub & Spoke**, any appliance set as a hub will serve as a hub in any overlay applied to it. Hubs in different regions mesh with each other to support regional routing. To configure hubs, click **Hubs** at the top of the page.
- **Regional Mesh and Regional Hub & Spoke:** To streamline the number of tunnels created between geographically dispersed groups of appliances, you can assign appliances to Regions and select **Regional Mesh** or **Regional Hub & Spoke**.

1. At the top of the page, click **Regions**.
2. You can add and remove a region or view the status of each overlay within a selected region.

Building SD-WAN Using These Interfaces

You can select which WAN interfaces you want to use for each device to connect to the SD-WAN. First, you assign your traffic to go to the Primary interfaces. If the Primary interface is unavailable or not meeting the desired Service Level Objectives configured, either the Secondary or Backup interfaces are used depending on what you have configured. You can configure only Backup interfaces, only Secondary interfaces, or both Secondary and Backup interfaces. Move the desired interfaces between the **Primary**, **Secondary**, and **Backup** boxes. The interfaces are grayed out until they are moved into the boxes.

- **Cross Connect** – Allows you to define tunnels built between each interface label. By default, tunnels are formed between labels with the same name. For example, if you have INETA configured on two appliances that are both members of the same mesh overlay topology, those appliances will be connected via a tunnel over the INETA label. In most cases there will be more than one internet link at a given site, resulting in the use of INETA and INETB. In this case, it is necessary to cross-connect INETA to INETB by placing both labels into the same group. If both INETA and INETB are configured with "Group 1" then tunnels will be formed from INETA <> INETB and INETA <> INETA.
- **Show/Hide Secondary** – Click **Show Secondary** to display the Secondary box, so you can drag interfaces into the box to enable Secondary interfaces. If you do not enable Secondary interfaces, you can click **Hide Secondary** to close the Secondary box.
- **Add Secondary if Primary Are** – Specifies when the system should use the Secondary interfaces. Select either **Down** or **Not Meeting Service Levels**. If you have both configured, secondary interfaces will be used before Backup interfaces.
- **Add Backup if Above Are** – Specifies when the system should use the Backup interfaces. Select either **Down** or **Not Meeting Service Levels**. If you have Secondary interfaces configured, Backup interfaces will be used when both the Primary and Secondary interfaces are unavailable and not meeting the configured Service Level Objectives.

NOTE: The order that labels appear in the **Primary**, **Secondary**, and **Backup** boxes only matters when Custom bonding is used and Link Selection is set to "Waterfall" with Rank Links By set to "Fixed Order".

Service Level Objective (SLO)

Traffic is routed through the primary interfaces exclusively until the Service Level Objectives (SLOs) for Loss, Latency, or Jitter have been exceeded. If this occurs, backup interfaces are added to the overlay to help meet the specified SLO.

You should configure SLOs based on the tolerance of the application to network performance. You should not configure SLOs based on the type of network or expected performance of the network itself. SLOs are about the application, not the network. For example, for voice SLOs most customers find 250ms Latency, 50ms Jitter, and 10% Loss to be acceptable parameters.

For High Availability and High Quality "waterfall" overlay modes, when an underlay violates the Loss SLO, the underlay is not removed from the overlay until the overlay itself violates the SLO. For High Throughput and and High Efficiency "balanced" modes, when an underlay violates the Loss SLO, it is immediately removed from the overlay. The Exclude Links BIO setting controls this behavior and can be modified using the Custom link bonding policy.

The Exclude Links setting does not apply to Latency or Jitter SLOs. Those SLOs always operate with Exclude Links set to “on Underlay Brownout”.

NOTE: If all links are in violation of SLOs, the system acts as if no SLO is configured and all links are configured as primary.

Link Bonding Policy

You can select the following Link Bonding Policies when you need to specify the criteria for selecting the best route possible when data is sent between multiple tunnels and appliances. You can also select custom bonding, which enables you to customize link prioritization and traffic steering policies based on multiple criteria.

Field	Description
High Availability	High availability chooses the best performing path, uses the path until it is near full, then waterfalls traffic onto the next best performing path. All traffic receives 1:1 FEC (forward error correction) when a copy of the packet is placed on another transport. High availability link bonding policy type should be used only for real-time traffic, since it renders the effective bandwidth to 50%.
High Quality	High quality policy chooses the best performing path, uses the path until it is near full, then waterfalls traffic onto the next best performing path. Adaptive FEC is used to provide parity packets only if there is degradation of the circuit. High quality link bonding policy should be used as the default selection for all non-real-time traffic types.
High Throughput	High throughput policy load-balances packets across all transports performing below the SLO defined in the BIO. Adaptive FEC is used to provide parity packets only if there is degradation of the circuit. This link bonding policy is used only in unique circumstances.
High Efficiency	High efficiency policy load-balances packets across all transports performing below the SLO defined in the BIO. No FEC is used in this bonding policy. This link bonding policy is used only in unique circumstances.
Custom	If the current fixed overlay bonding modes are not flexible enough, the Custom link bonding policy allows you to fine-tune your network performance. Custom link bonding preserves existing bonding modes (HA, HQ, HT, HE) while allowing customization of link bonding characteristics on a per-overlay basis. This should only be used when absolutely necessary. If you select Custom, see the following table for information about the settings.

If you select Custom link bonding, enter the appropriate information for the following fields.

Field	Description
FEC Wait Time	Measured in milliseconds (ms). This controls how long to wait to fill a packet before sending. A lower number indicates more FEC overhead.
Exclude Links	This controls when an underlay is removed from an overlay during brownout conditions. **On overlay brownout** – Wait for the overlay to see a loss before removing any underlays from the overlay. This allows bandwidth to be used, but there could be increased latency due to path conditioning. **On underlay brownout** – Remove the underlay from the overlay as soon as it violates the brownout threshold. **NOTE:** The Exclude Links setting only applies to the Loss SLO. When underlays violate a Latency or Jitter SLO they are immediately removed from the Overlay regardless of the Exclude Links setting.
Link Reorder Frequency	This controls how aggressively underlays are evaluated and determines when to switch traffic from one link to another. It also controls the ranking and eligibility of links, which impacts Link Selection and brownout behavior. **Aggressive** – Changes in underlay performance are detected within a few seconds. This setting is best for high-speed networks, such as dual high-speed internet links. **Moderate** – Changes in underlay performance are detected within about a minute. **Conservative** – Changes in underlay performance take several minutes to detect. This is useful for situations where you want to be certain that the primary link is not performing as expected before switching to another link.
Path Conditioning	Measured as a percentage. This controls the amount of FEC employed for the overlay. More FEC means more overhead but a higher chance to recover lost or delayed packets. **NOTE:** For HQ and HA modes that use Waterfall, when two or more links are present FEC is transmitted on the second-best link. FEC is spread over all eligible underlays for HT and HE modes that use Balanced link selection.

Field	Description
Packet Reorder Wait Time	Measured in milliseconds (ms). This determines how long to wait for packet order correction (POC) to occur. When the wait time expires, all missed packets are declared as lost packets. If the EdgeConnect sees packets not arriving or arriving out of order, it dynamically increases this timer.
Link Selection	**Waterfall** – Cascades packets across eligible underlays based on one of five quality measures. Select one of the quality measures from the drop-down menu: **Overall Quality**, **Latency**, **Loss**, **Jitter**, or **Link Order**. Waterfall is used for both High Availability and High Quality overlay modes. **Balanced** – Per-packet load balancing across all eligible underlays based on one of three modes. Select one of the modes from the drop-down menu: **Link Capacity (Local)** – Fills the "most open" link first. The EdgeConnect fills whichever link has the most bandwidth available until all links have the same amount of absolute available bandwidth. This balancing mode is used by the High Throughput overlay mode. **Link Utilization (Local)** – Fills links proportionally to total link capacity. **Link Utilization (Local & Remote)** – Fills links based on tunnel capacity. This balancing mode is used by the High Efficiency overlay mode. For Link Utilization mode, the EdgeConnect tries to keep all links at the same percent (%) utilization. For example, if one link is at 50% utilization, the EdgeConnect fills the other links until all links are at 50% utilization.

QoS and Optimization

To further customize your overlay configuration, enter the appropriate information for the following fields.

Field	Description
FW Zone	Select the firewall zone you want to restrict traffic to from an overlay. **NOTE:** This field is disabled when an end-to-end zone-based firewall is enabled.
WAN Optimization	Select **Enabled** if you want to apply any purchased WAN Optimization to your overlay. Select **Disabled** if you do not want to apply WAN Optimization to your overlay.

Field	Description
Peer Unavailable Option	Select what the appliance should do when no peer is reachable via a tunnel. You can select a specific label, **Use Best Route**, or **Drop**. If you select a specific label, the appliance routes the traffic to that link. **Best Route**: When selected, the appliance searches for the next best route that is available. **Drop**: When selected, the appliance drops the connection.
Traffic Class	Channels traffic to the desired queue based on the applied service.
LAN DSCP	Select the DSCP you want to apply as a filter to the LAN interface.
WAN DSCP	Select the DSCP you want to apply as a filter to the WAN interface.
Flow Affinity	When enabled at a selected percentage, loads packets onto the best-performing path until that path is full, and then begins load-balancing packets onto the next-best underlay. Waterfall can be applied in High Quality, High Availability, and Custom bonding modes.

Breakout Traffic to Internet and Cloud Services

You can use the **Breakout Traffic to Internet & Cloud Services** to monitor and manage traffic coming to or from the internet.

Hub Versus Branch Breakout Settings

You can create different breakout policies for hubs. Any hub you select in the **Topology** section also displays at the top of the **Internet Traffic to Web, Cloud Services** tab. When you select an individual hub, the **Use Branch Settings** displays, selected, to the right of the screen. Complete the following steps to create a custom breakout policy for that hub:

1. Clear the **Use Branch Settings** check box.
2. Configure the now accessible parameters.
3. Click **OK**.

Preferred Policy Order and Available Policies

- You can move policies back and forth between the Preferred Policy Order and the Available Policies columns. You can also change their order within a column. The defaults provided are **Backhaul via Overlay**, **Break Out Locally**, and **Drop**.
- When you select **Break Out Locally**, confirm that any selected interface that is directly connected to the internet has **Stateful Firewall** specified in the deployment profile.

- You can add services (such as Zscaler, Fortigate, or Palo Alto). The service requires a corresponding internet-breakout (Passthrough) tunnel for each appliance traffic to that service. To add a service, select the edit icon next to **Available Policies**.
- The **Default** policy you configure for internet breakout is pushed to all appliances that use the selected Overlay. However, you might want to push different breakout rules to your hubs.

Break Out Locally Using These Interfaces, Available Interfaces, and Link Selection

- You can select the best internet breakout links by specifying the type of Link Selection; either **Waterfall** or **Balanced**.
  - If **Waterfall** is selected, links are ranked on the selected threshold, from best to worst, using an inference system that averages performance of all SDWAN fabric tunnels associated with a given label. In Waterfall mode, flows are routed across the best label until bandwidth utilization is above 80%. Once 80% utilization is reached flows will “waterfall” to the next-best label. For more information about Waterfall mode, see Internet Breakout Trends.
  - If **Balanced** is selected, flows are subjected to a weighted load-balancing algorithm. The weighting is proportional to the available bandwidth of the link.
  - For both **Waterfall** and **Balanced**, if a threshold is configured for **Loss**, **Latency**, or **Jitter**, the system removes the link from Local Breakout eligibility when it exceeds the threshold.
- You can choose to set IP SLA Rule destinations.
- If you select the **Threshold-based Failover** check box, the Preferred Policy Order is applied when all links violate the configured threshold. This setting is useful when you want the system to backhaul traffic during Local Breakout threshold violation. When this check box is cleared, all Local Breakout labels must be down for flows to fall to the next policy.
- For Local Breakout flows, the system uses session affinity to attempt to keep flows with the same source and destination IPs on the same link. You can change the session affinity timeout.

Complete the following steps.

1. In the Break Out Locally Using These Interfaces section, drag and drop available interfaces into the **Primary** or **Backup** boxes.
2. Under Link Selection, select **Waterfall** or **Balanced**, and enter the amount for the Performance Thresholds: Loss, Latency, Jitter, and Utilization.
3. If you selected **Waterfall**, select one of the following thresholds to rank links.

Field	Description
Auto	Default threshold if you do not specify the threshold for your links. The Auto metric uses combined loss and latency to derive the best link. This is the same metric used for determining the best underlay in HQ overlay bonding mode and is referred to as “Overall Quality” in the Link Selection section of the Custom Bonding configuration.
MOS	Inferred average MOS score for a given underlay.
Loss	Inferred average loss percentage as derived from the all Up-Active tunnels for a given underlay.
Latency	Inferred average latency for a given underlay.
Fixed Order	Links are sorted in the order specified under Break Out Locally Using These Interfaces. The link at the top of the Primary list is used first.

NOTE: Backup links are used when the passthrough tunnel is down due to IP SLA (Internet Protocol Service Level Agreement) settings, the primary link is down, or the inferred SLO (Service Level Objective) is exceeded for all primary links.

4. Click the edit icon next to Break Out Locally Using These Interfaces to change the default Local Breakout IPSLA endpoints.

   The IP SLA Rule Destination dialog box appears.

5. Click the **Enable IP SLA rule orchestration** toggle. Then enter information in the following fields to change the default Local Breakout IP SLA endpoints and create IP SLA rules.

Field	Description
Enable IP SLA rule orchestration	When enabled, Orchestrator automatically sets up IP SLA rules on all appliances where this overlay is applied.
Monitor	Select one of the three types of probes used to monitor IP SLA endpoints: Ping, HTTP, or HTTPS. **NOTE:** Using HTTPS causes additional CPU load and increased packets due to the overhead of SSL handshaking. Ping or HTTP are recommended.

Field	Description
Address	A comma separated list of hostnames or IP addresses to probe. A response from any of the destinations allows the system to validate the path. **NOTE:** For 9.6.0 or later deployments, your IP SLA Rule Destination addresses must be compatible with IPv6 if you want to use IPv6 local breakout addressing. Replace the contents of this field with the following IP SLA internet endpoints as one continuous string: *192.151.28.254,2600:9000:a60f:6c50:ca04:acb1:b183:9524,2600:9000:a70d:506e:85a8:bd96:7482:16,sp-ipsla.silverpeak.cloud*
Proxy Address	(optional for HTTP/HTTPS)
User Agent	(optional for HTTP/HTTPS)
HTTP Request Timeout	After an HTTP probe is sent, this is the length of time the system waits to hear back from the destination server.
Ping Interval/Polling Frequency	How frequently the ping or HTTP/HTTPS probe is sent. This value can be set to “1” for ICMP, however, this should be set to “2” or greater for HTTP/HTTPS.
Rolling average window for Loss and Latency	The rolling average for loss and latency for each destination. For a 1 second Keep Alive Interval, this would be a 5 minute rolling average. Reducing the sampling window could cause overly aggressive behavior.
Reachability	The system uses these values to determine if the probe can reach the destinations or not. A good value for these is 5, based on a 1 second Keep Alive Interval. Setting these any lower could cause false positives. **Mark Up after X Pings/X Sequential Successes** – Enter a numeric value. The system makes this many attempts to reach the destinations, and it marks the status of the tunnels as “up” after it receives this many successful responses from any of the destinations. **Mark Down after X Failed Pings/X Sequential Failures** – Enter a numeric value. The system marks the status of the tunnels as “down” after this many consecutive failed responses from ALL destinations.

Field	Description
Loss	The system uses these values as thresholds when calculating percentage loss to determine if the passthrough tunnels are “up” or “down”. **Mark Up after loss below X%** – Enter a percentage. The system marks the status of the tunnels as “up” if the percentage loss calculated from the best performing destination is below this threshold. **Mark down after loss above X%** – Enter a percentage. The system marks the status of the tunnels as “down” if the percentage loss calculated from the best performing destination exceeds this threshold. This means that all destinations in the Address field must have crossed this threshold for the system to invoke the “down” status.
Latency	The system uses these values as thresholds when calculating average latency to determine if the passthrough tunnels are “up” or “down”. These values are measured in milliseconds. **Mark Up after average latency below X** – Enter a numeric value. The system marks the status of the tunnels as “up” if the average latency calculated from the best performing destination is below this threshold. **Mark Down after average latency above X** – Enter a numeric value. The system marks the status of the tunnels as “down” if the average latency calculated from the best performing destination exceeds this threshold. This means that all destinations in the Address field must have crossed this threshold for the system to invoke the “down” status.
Loss OR Latency, Loss AND Latency	Select one of two options for combining the Loss and Latency metrics. **OR** – The system marks the status of the tunnels as “down” if either the Loss or the Latency thresholds are crossed. **AND** – The system marks the status of the tunnels as “down” if both the Loss and Latency thresholds are crossed.
Check IP SLA status every	How frequently EdgeConnect checks to see if the thresholds have been crossed. This is also how frequently a decision is made to move a tunnel in or out of service or to raise an IP SLA Down alarm. 30 seconds is the default. Setting this value much lower could cause false positives or tunnel flapping.

Field	Description
Route Label	Select a route label to apply to this orchestration. Orchestrator will automatically revoke routes with this route label. You can use this feature to add a static default route that will be revoked without needing to create an IP SLA rule. Depending on the IP SLA Condition set for the route label, the route will be marked DOWN if *any* of the IP SLA rules are DOWN or only when *all* IP SLA rules associated with the label are DOWN. See Route Labels.

6. (optional) Click the **Threshold-based Failover** check box.
7. Click the edit icon next to Link Selection and in the Session Affinity Settings dialog box, enter a value in the **Session Affinity Timeout** field. You can enter any value between 0 and 10,000 minutes.

   NOTE: Setting Session Affinity Timeout to "0" disables the Session Affinity feature.

Apply Overlays

Configuration > Overlays & Security > Apply Overlays

Use this page to **add or remove overlays** from appliances. If you select **Edit Overlays**, you will be redirected to the **Business Intent Overlay** tab for further customization. You can also view the status of the overlays if you select **View Status**.

Interface Labels

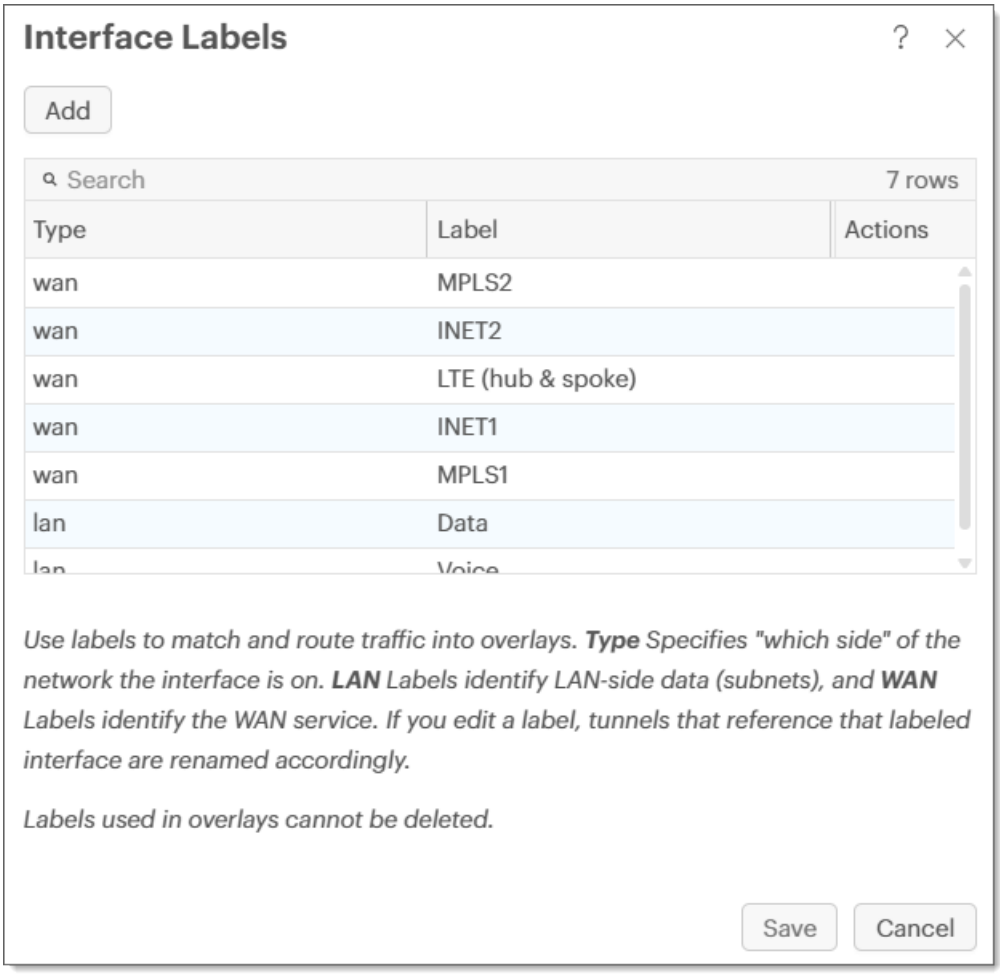
Configuration > Overlays & Security > Interface Labels

To make it easier to identify connections, you can create descriptive interface labels for each link type in your environment. Use labels to match and route traffic into overlays. The label type specifies "which side" of the network the interface is on. LAN labels identify LAN-side data (subnets), and WAN labels identify the WAN service, such as MPLS, Internet, or LTE. If you edit a label, tunnels that reference that labeled interface are renamed accordingly.

- LAN labels can be selected for a traffic access policy in a Business Intent Overlay (BIO), which in turn is applied to an appliance with those LAN labels. All traffic matching those interfaces is automatically processed by that BIO. If you use an ACL for a traffic access policy, the LAN label is ignored for that BIO.
- WAN labels are used by Orchestrator and BIOs to determine which interfaces on different appliances should be connected by tunnels built by Orchestrator. Orchestrator automatically pushes interface labels to appliances it manages.

Manage Labels

Use the Interface Labels dialog box to manage labels in Orchestrator, available under **Configuration > Overlays & Security > Interface Labels**.

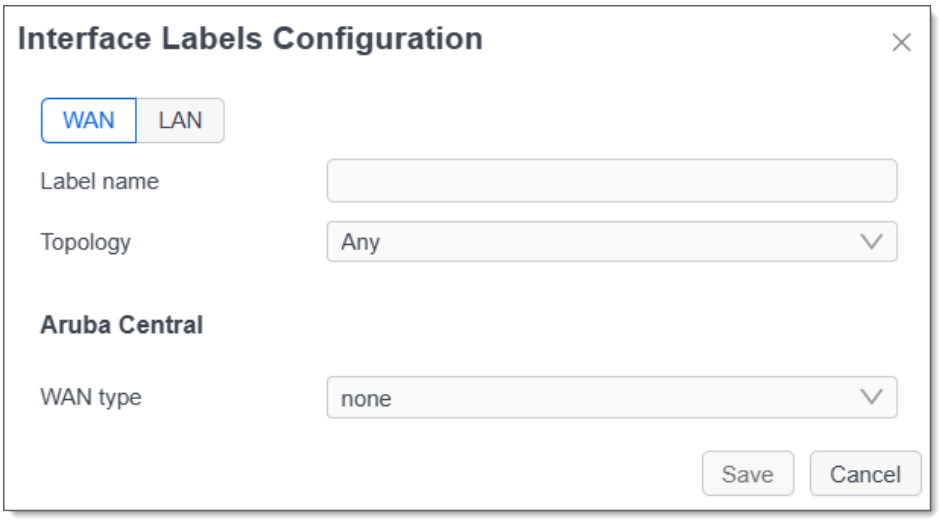


From this dialog box, you can create, edit, or delete labels.

Create a Label

1. Click **Add**.

   The Interface Labels Configuration dialog box opens.



2. Select **WAN** or **LAN** for the label type.
3. Enter a descriptive name in the **Label name** field.

   NOTE: For WAN labels, if you want to allow Orchestrator to build tunnels using this label in any topology, leave the Topology selection set to **Any**. If you want to override BIO settings and exclude this label in Full Mesh overlays, set Topology to **Hub & spoke**.
4. If you enabled HPE Aruba Networking Central, select the **WAN type**. Orchestrator appends the WAN type to uplink information for OTO WAN tunnels sent to HPE Aruba Networking Central.
5. Click **Save** to save your changes and close the dialog box. Otherwise, click **Cancel** to return to the list of interface labels.

Edit a Label

1. In the Interface Labels dialog box, click the edit icon for an existing label.
2. Select **WAN** or **LAN** for the label type—you cannot change the label type if the label is currently in use.
3. If you want to change the label name, modify it in the **Label name** field.

   NOTE: For WAN labels, if you want to allow Orchestrator to build tunnels using this label in any topology, leave the Topology selection set to **Any**. If you want to override BIO settings and exclude this label in Full Mesh overlays, set Topology to **Hub & spoke**.

   IMPORTANT: Renaming interface labels affects the calculation of availability statistics as provided on the Availability tab.
4. Click **Save** to save your changes and close the dialog box. Otherwise, click **Cancel** to return to the list of interface labels.

Delete a Label

1. In the Interface Labels dialog box, hover over the Actions column and click the **X** icon for the label you want to delete.

NOTE: Labels used in overlays cannot be deleted.

The label is deleted from the list but can be restored by closing the dialog box without saving.

2. To save your changes and permanently delete the label, click **Save**.

   WARNING: When deleting a label, a confirmation message warns you that deleted interface labels will be removed from all policies, interfaces, and deployment profiles that are currently using the label.

3. Click **Save** to confirm the removal. Otherwise, click **Cancel** to return to the Interface Labels dialog box.

Route Labels

Configuration > Overlays & Security > Route Labels

You can create route labels and apply them to static or announce-only routes. You then associate the route labels with one or more IP SLA rules. The IP SLA rules monitor specific destinations to dynamically control the status of static or announce-only routes. Using route labels with IP SLA rules improves network reliability by preventing traffic from being sent to unreachable destinations.

- Upgrades and new deployments are shipped with a default route label of "Internet."
- All route labels are available on all appliances.
- By default, if *all* IP SLA rules associated with a route label are in the DOWN state, the label will go DOWN, causing all routes assigned to it to also be marked DOWN and be removed from the routing table. If you set the IP SLA Condition to ANY, the route label and assigned routes will go DOWN if *any* linked IP SLA rule associated with the label is in the DOWN state.

Using route labels with IP SLA rules is a three-step process. You must:

1. Create a route label. See Route Labels Dialog Box. (If you need to create a route, see Routes Tab.)
2. Apply the route label to a static route. See Routes Tab.
3. Associate the route label with an IP SLA rule. See IP SLA Tab.

Route Labels Dialog Box

Use the Route Labels dialog box to add, edit, and delete route labels. You can create up to 127 route labels.

HINT: Click the Information icon to learn about viewing audit and server logs.

Add a Route Label

To add a route label:

1. Click **Add Label**. A new row is added to the table.
2. Click in the new cell in the Label column and then enter a name for the route label. Route Label names can be up to 15 characters and contain alphanumerical characters, dashes, and underscores.
3. Press **Tab** or click in the new cell in the IP SLA Condition column, and then select **ALL** or **ANY**.

   By default, if *all* IP SLA rules associated with a route label are in the DOWN state, the label will go DOWN, causing all routes assigned to it to also be marked DOWN and be removed from the routing table. If you set the IP SLA Condition to ANY, the route label and assigned routes go DOWN if *any* linked IP SLA rule associated with the label is in the DOWN state.
4. (Optional) Click in the Comment field and enter a comment.
5. Click **Save**.

Edit a Route Label

To edit a route label:

1. Click in the cell to edit the label, IP SLA Condition, or comment.
2. Click **Save**.

Delete a Route Label

To delete a route label:

1. Click the delete icon (**X**) in the last column of the route label you want to delete.
2. Click **Save**.
3. Click **OK**.

NOTICE: If you delete a route label, the route label is removed from IP SLA rules and static routes currently using that route label.

Hubs

Configuration > Overlays & Security > Hubs

On this tab, you can add, remove, and associate hubs to a specified region within the Regional Mesh or Regional Hub-and-Spoke topologies configured on the **Business Intent Overlay** tab.

You can specify whether a hub will re-advertise routes that were previously received from a spoke in the hub's region or a hub in another region.

NOTE: This feature requires appliance software version 9.1.0 or later.

You can also access the Regions tab and Business Intent Overlay tabs by clicking the links at the top of the page.

Complete the following steps to add a hub:

1. Start typing a name or select the appliance you want make a hub from the list.
2. Select one of the following:
   - **Re-Advertise Routes** – This hub will re-advertise its routes so that other appliances can learn them. This hub will also re-advertise routes learned from other EdgeConnect appliances within its region.
   - **Do Not Re-Advertise Routes (Stub Hub)** – This hub will not re-advertise routes learned from other regions or spokes. All local routes (static, directly connected, BGP, and OSPF) will still be advertised. Hubs that do not re-advertise their routes are stub hubs.
3. Click **Add Hub**.

To delete a hub, select the **X** icon next to the hub you want to delete.

NOTE: You must remove all overlays before you can revert a hub back to a spoke.

Deployment Profiles

Configuration > Overlays & Security > Deployment Profiles

Instead of configuring each appliance separately, you can create various **Deployment Profiles** and provision a device by applying the profile you want. For example, you can create a standard format for your branch.

TIP: For a smoother workflow, complete the DHCP Server Defaults tab (**Configuration > Networking > DHCP Server Defaults**) before creating Deployment Profiles.

You can use Deployment Profiles to simplify provisioning, regardless of whether you choose to create and use **Business Intent Overlays**.

NOTE: You cannot edit **IP/Mask** fields because they are appliance-specific.

Map Labels to Interfaces

- On the **LAN** side, labels are optional. They can be used as match criteria for Business Intent Overlay ACLs, such as **data**, **VoIP**, or **replication**.
- On the **WAN** side, labels identify the link type, such as **MPLS** or **Internet**. These labels are mandatory. They are used by Orchestrator to build Business Intent Overlay policies.
- To create or manage a global pool of labels, either:
  - Navigate to **Configuration > Overlays & Security > Deployment Profiles**, click the **Edit** icon next to Label, and make the appropriate changes, or

  - Navigate to **Configuration > Overlays & Security > Interface Labels**) and make the appropriate changes.

- The change you make to a label propagates automatically. For example, it renames tunnels that use that labeled interface.

LAN-side Configuration: Segments and Firewall Zones

EdgeConnect Segmentation (VRF) provides orchestrated layer-3 segmentation, Zone Based Firewall, and IDS—end-to-end across the SD-WAN fabric. Segment and zone policies are global in scope. They are managed on the **Configuration > Networking > Routing > Routing Segmentation (VRF)** tab.

Segments and zones are then assigned to LAN-side interfaces for each appliance by using the Deployment dialog box. By default, the Segment and FW Zone fields on LAN interfaces are set to the system-generated Default segment. You can select a different segment and firewall zone from the drop-down lists. These lists reflect the segments and zones that are set up on the Routing Segmentation (VRF) tab.

NOTE: The segment for WAN interfaces cannot be changed.

LAN–side Configuration: DHCP and Router Advertisements

- By default, the LAN IP does not act as a DHCP Server. Based on your configuration, you can set the interface to act as a DHCP relay server when the appliance is in Router mode.
- The global defaults are set in **Configuration > Networking > DHCP Server Defaults** and pre-populate this page. The other choices are **No DHCP/No RA** and having the appliance act as a **DHCP/BOOTP Relay**.
- Enter the LAN interface from the drop-down. Click **+IP** to add a specific IP address.
- Enter the IP address of the specific LAN interface above the **NO DHCP** link.
- To customize an individual interface on the Deployment Profiles tab, click the DHCP-related link under the IP/Mask field. The DHCP Settings / Router Advertisements dialog box opens.
- Before you can configure DHCP, you must navigate to Management Services and select an interface for DHCP Relay. See Management Services for more information.

If the LAN interface has an IPv4 IP address, click **V4** to display the DHCP configuration settings. See V4.

If the LAN interface has an IPv6 IP address, click **V6** to display the Router Advertisement settings. See V6.

V4

The following tables describe the various DHCP settings you can configure for LAN interfaces that have IPv4 IP addresses.

DHCP Server

Field	Description
Subnet Mask	Mask that specifies the default number of IP addresses reserved for any subnet. For example, entering **24** reserves 256 IP addresses.
Exclude first N addresses	Specifies how many IP addresses are not available at the beginning of the subnet's range.
Exclude last N addresses	Specifies how many IP addresses are not available at the end of the subnet's range.
Default lease, Maximum lease	Specify, in seconds, how long an interface can keep a DHCP–assigned IP address.
Default gateway	Indicates whether the default gateway is being used.
DNS server(s)	Specifies the associated Domain Name System servers.
NTP server(s)	Specifies the associated Network Time Protocol servers.
NetBIOS name server(s)	Used for Windows (SMB) type sharing and messaging. It resolves the names when you are mapping a drive or connecting to a printer.
NetBIOS node type	NetBIOS node type of a networked computer relates to how it resolves NetBIOS names to IP addresses. There are four node types: **B**-node – 0x01 Broadcast **P**-node – 0x02 Peer (WINS only) **M**-node – 0x04 Mixed (broadcast, then WINS) **H**-node – 0x08 Hybrid (WINS, then broadcast)
DHCP failover	Enables DHCP failover. To set it up, click the **Failover Settings** link.

DHCP/BOOTP Relay

Field	Description
Destination DHCP/BOOTP Server	IP address of the DHCP server assigning the IP addresses. This setting applies to the local interface only.

Field	Description
Common DHCP server for all segments	Select this check box to set the default values for all segments. **HINT:** You can reset the defaults in Management Services by setting the DHCP Relay interface to “any” and then selecting an interface label again. However, this might impact service. Or, you can manually reset the defaults by selecting the following values: Option 82 = enabled, Option 82 Policy = append, and select the following sub options: 1, 5, 10, 11, 151, and 152.
Distinct DHCP server per segment	Select this option to override the DHCP relay configuration set in the Manages Services tab with the settings you select in this dialog box.
Enable Option 82	When selected, inserts additional information into the packet header to identify the client’s point of attachment. This setting applies to all LAN-side interfaces on this appliance. **IMPORTANT:** Changing this setting will modify Option 82 settings on all LAN-side interfaces that are enabled as DHCP Relay.
Option 82 Policy	Tells the relay what to do with the hex string it receives. The choices are **append**, **replace**, **forward**, and **discard**. This setting applies to all LAN-side interfaces on this appliance. **IMPORTANT:** Changing this setting will modify Option 82 settings on all LAN-side interfaces that are enabled as DHCP Relay.

Field	Description
Sub Options	Select one or more of the following: **1** - Agent Circuit ID: Provides information about the interface or circuit through which the DHCP request was received. **5** - Link selection: Specifies the IP address used by the DHCP server to determine the appropriate subnet for addressing the DHCP client. **10** - Client Unicast/Broadcast Indication flag: Indicates whether the DHCP relay received the client packet as a unicast or broadcast packet. **11** - Server ID Override: Allows the DHCP relay agent to act as a proxy for the DHCP server to process unicast lease renewals. **150** - Link selection (Cisco proprietary): Provides information about a segment or VPN that is necessary to allocate an address to a DHCP client on that segment. **151** - VRF name/VPN ID **152** - VRF name/VPN ID Control Sub-Option OR Server ID Override (Cisco proprietary): Indicates whether the DHCP server supports sub option 151 (VRF Name/VPN ID). If this option is present in the reply from the server, the server does not support option 151.

V6

The following table describe the various router advertisement settings you can configure for LAN interfaces that have IPv6 IP addresses. The LAN clients can use these options to autoconfigure IPv6 addresses and to learn default gateway addresses.

NOTE: DHCP for IPv6 is not supported.

Setting	Description
Enable Router Advertisements	Specifies whether the router should send RA messages.
Managed Flag	Select this option to instruct IPv6 hosts to use DHCPv6 to obtain their IPv6 addresses in addition to any other configuration information.

Setting	Description
Other Flag	Select this option to instruct IPv6 hosts to use DHCPv6 to obtain additional configuration information, such as DNS server addresses and other network parameters.
Link MTU	Set the maximum transmission unit (MTU) size that can be transmitted without fragmentation. This helps ensure that all hosts on the network use the same MTU, avoiding issues related to packet fragmentation and reassembly.
Max Interval	Specify the maximum interval in seconds between unsolicited RA messages. This helps to control the frequency of RA messages.
Min Interval	Specify the minimum interval in seconds between unsolicited RA messages. This helps to control the frequency of RA messages.
Current Hop Limit	Set the default hop limit for IPv6 packets sent by hosts on the network. Hosts use this value to configure their own hop limit for outgoing packets.
Default Router Preference	Select **High**, **Medium**, or **Low** to set the preference level of the router for use as a default router. Hosts use this value to prioritize multiple routers on the same link.
Default Router Lifetime	Specify the lifetime in seconds of the default route that is advertised by the router. The hosts use this value to determine how long the router should be used as the default gateway.
Reachable Time	Specify the time in milliseconds that an IPv6 host considers a neighbor reachable after receiving a confirmation. This value maintains accurate and timely reachability information in the neighbor cache.
Retrans Timer	Specify the time in milliseconds between retransmissions of neighbor solicitation messages. This value reduces the frequency of retries when attempting to discover or confirm the reachability of neighbors on the network.

Add a Router Advertisement Prefix

Click **Add** and complete the following fields.

Considerations

- RA can be configured only on LAN side interfaces.
- Users can configure RA only on IPv6 configured interfaces.
- DHCPv4 server and RA cannot be configured on the same interface at the same time.
- DHCPv4 relay and RA cannot be configured on the same interface at the same time.

- RA, DHCPv4 Server, and DHCPv4 Relay cannot be enabled if there is an alias interface configured for the main/primary interface.
- A maximum of 10 prefixes can be configured in the RA configurations per interface.

Setting	Description
Prefix-id	The ID assigned to the prefix.
Prefix	The IPv6 prefix to advertise to hosts on the network. Hosts use this prefix to configure their IPv6 addresses and determine the network portion of the IP addresses.
Autonomous flag	Select whether the prefix can be used by hosts for SLAAC. When set to true, hosts can use the prefix to generate their own IPv6 addresses.
Onlink flag	Specifies whether the prefix is on-link, which affects how hosts handle routing for addresses within the prefix. If set to true, hosts assume that addresses within the prefix can be reached directly on the local network segment.
Valid Lifetime	Specify the duration in seconds for which the advertised prefix is valid.
Preferred Lifetime	Specify the duration in seconds (relative to the time the packet is sent) that addresses generated from the prefix via stateless address auto-configuration remain preferred.

WAN–side Configuration

Interface mode: Orchestrator release 9.5.2 and later supports IPv4 and IPv6 to meet the increasing demand for IP addresses. This feature also allows you to deploy appliances that support both IPv4 and IPv6 for a dual stack solution.

NOTE: The WAN interfaces of factory deployed appliances with ECOS version 9.5.2 and later support both IPv4 and IPv6 addresses for Zero Touch Provisioning (ZTP). Appliances with previous versions of ECOS only support IPv4.

IMPORTANT: ZTP for IPv6 is not supported for EC-Vs.

Select one of the following options for each WAN interface:

- **Static** – Configure the IPv4 or IPv6 address manually on the interface.
- **DHCPv4** – The IPv4 address is configured dynamically by DHCPv4 on the interface.
- **DHCPv6** – The IPv6 address is configured dynamically by DHCPv6 on the interface.
- **SLAAC** – The IPv6 address is configured dynamically by SLAAC (State Less Address Auto Configuration) on the interface.

- **DHCPv4 + DHCPv6** – This is a dual-stack mode that supports the simultaneous use of IPv4 and IPv6 addresses. The IPv6 address is configured dynamically by DHCP on the primary interface and the IPv4 address is configured dynamically by DHCP on the alias interface. All properties except Label are inherited from the primary interface and cannot be edited.
- **DHCPv4 + SLAAC** – This is a dual-stack mode that supports the simultaneous use of IPv4 and IPv6 addresses. The IPv6 address is configured dynamically by SLAAC on the primary interface and the IPv4 address is configured dynamically by DHCP on the alias interface. All properties except Label are inherited from the primary interface and cannot be edited.

NOTE: Only one DHCP configuration option is allowed for each physical interface.

Firewall Zone: Zone-based firewall policies are configured globally on the Orchestrator. A zone is applied to an **Interface**. By default, traffic is allowed between interfaces labeled with the same zone. Any traffic between interfaces with different zones is dropped. You can create exception rules (Security Policies) to allow traffic between interfaces with different zones. The firewall zones you have already configured will be in the list under **FW Zone**. Select the Firewall Zone you want to apply to the WAN you are deploying.

Firewall Mode: Four options are available at each WAN interface:

- **Allow All** permits unrestricted communication. Use this option with extreme caution and only if the interface is behind a WAN edge firewall.
- **Stateful** __*only__* allows communication from the LAN-side to the WAN-side.

  Use this if the interface is behind a WAN edge router.
- **Stateful with SNAT** applies Source NAT to outgoing traffic.

  Use this if the interface is directly connected to the Internet and you want to enable local internet breakout.
- **Harden**
  - For traffic inbound from the WAN, the appliance accepts __*only__* IPSec tunnel packets that terminate on an EdgeConnect appliance.
  - For traffic outbound to the WAN, the appliance __*only__* allows IPSec tunnel packets and management traffic that terminate on an EdgeConnect appliance.

VLAN Settings: With Orchestrator release 9.5.2 and later, you can assign multiple IP aliases to the same VLAN interface ID. Prior to Orchestrator release 9.5.2, Orchestrator only allowed you to assign one unique IP address to a VLAN ID. This new feature also allows you to assign VLAN IDs to interfaces in a dual stack solution with SLAAC.

Sub-interfaces behave the same as physical interfaces.

NOTE: If you modify the IP address or subnet mask of a sub-interface, all sub-interfaces and IP aliases with the same ID will be deleted and added back. The label, segment, and zone will remain unchanged. There will be a brief outage of all IPs while the interfaces are deleted and added back.

IP aliases and sub-interfaces will appear in all tabs where interfaces are listed or selectable (for example, the Dynamic table on the Interfaces tab). The first sub-interface will be the main interface. Additional IP aliases will be assigned with incremented interface numbers. For example, the main sub-interface will be wan0.120 and additional IP aliases will be assigned wan0.120:1, wan0.120:2, and so on.

NAT Settings: To change the NAT setting, click the NAT-related link under the Next Hop field on the WAN side. The NAT Settings dialog box opens.

Select one of the following options:

- If the appliance is behind a NAT-ed interface, select **NAT**.
- If the appliance is not behind a NAT-ed interface, select **Not behind NAT**.
- Enter an **IP address** to assign a destination IP for tunnels being built from the network to this WAN interface.

Shaping: You can limit bandwidth selectively on each WAN interface.

- **Total Outbound** bandwidth is always enabled.
- **Inbound Shaping** is disabled by default and can be enabled by clicking the Edit icon for a specific interface on the Shaper tab.

EdgeConnect Licensing: Only visible on EdgeConnect appliances.

- For additional bandwidth, you can purchase Plus, and then select it here for this profile.
- If you have purchased a pool of WAN Optimization for your network, you can allocate a portion of it in the **WAN Opt** field in a Deployment Profile. You can also direct allocations to specific types of traffic in the Business Intent Overlays.
- To view how you have distributed Plus and WAN Optimization, navigate to the **Configuration > Overlays & Security > Licensing > Licenses** tab.
- Select the appropriate licensing you have applied to your EdgeConnect appliance from the menu. The licenses will only display depending on the licenses you have for that particular account. You can select the following licensing options:
  - Mini
  - Base
  - Base + Plus
  - 50 Mbps
  - 200 Mbps
  - 500 Mbps
  - 1 Gbps
  - 2 Gbps
  - Unlimited

  NOTE: You must have the correct hardware to support the license selected.

BONDING

- EdgeConnect supports etherchannel bonding of multiple physical interfaces of the same media type into a single virtual interface. For example, **wan0** plus **wan1** bond to form **bwan0**. This increases throughput on a very high-end appliance and/or provides interface-level redundancy.
- For bonding on a virtual appliance, you would need to configure the host instead of the appliance. For example, on a VMware ESXi host, you would configure NIC teaming to get the equivalent of etherchannel bonding.
- Whether you use a physical or a virtual appliance, etherchannel must also be configured on the directly connected switch/router. Refer to the switch or router user documentation for configuring interface bonding.

A More Comprehensive Guide to Basic Deployments

This section discusses the basics of three deployment modes: **Bridge**, **Router**, and **Server** modes.

It describes common scenarios, considerations when selecting a deployment, redirection concerns, and some adaptations.

For detailed deployment examples, refer to the HPE Aruba Networking EdgeConnect SD-WAN documentation site for various deployment guides.

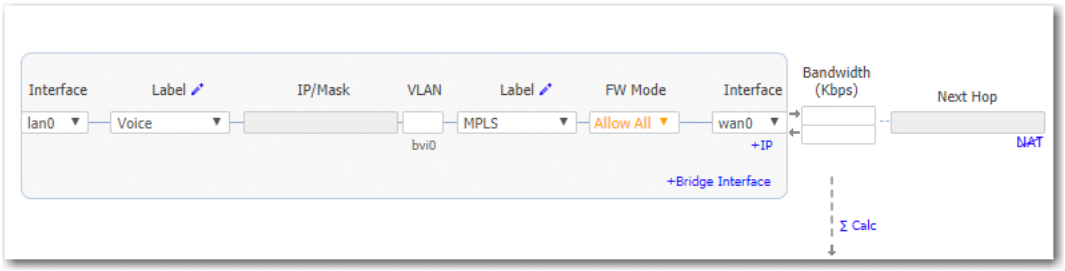
In Bridge Mode and in Router Mode, you can provide security on any WAN-side interface by **hardening the interface**. This means:

- For traffic inbound from the WAN, the appliance accepts __*only__* IPSec tunnel packets.
- For traffic outbound to the WAN, the appliance __*only__* allows IPSec tunnel packets and management traffic.

Bridge Mode

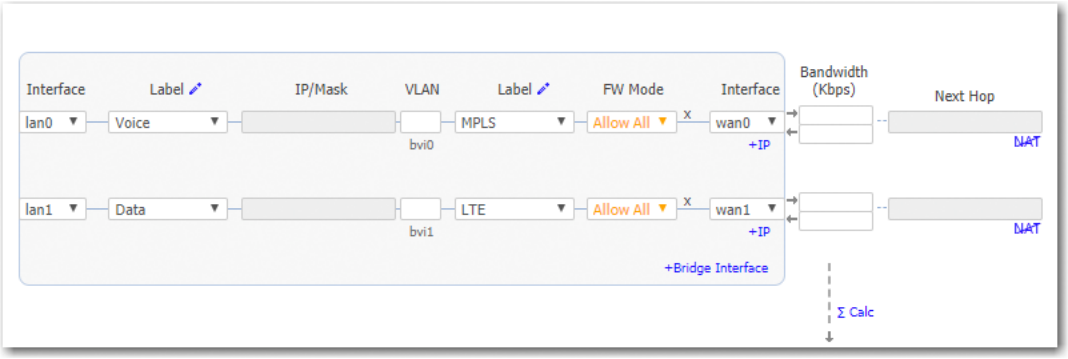
Single WAN-side Router

In this deployment, the appliance is in-line between a single WAN router and a single LAN-side switch.



Dual WAN-side Routers

This is the most common 4-port bridge configuration.



- 2 WAN egress routers / 1 or 2 subnets / 1 appliance
- 2 separate service providers or WAN services (MPLS, IPSec VPN, MetroEthernet, and so forth)

Considerations for Bridge Mode Deployments

- Do you have a physical appliance or a virtual appliance?
- A virtual appliance has no fail-to-wire, so you will need a redundant network path to maintain connectivity if the appliance fails.
- If your LAN destination is behind a router or L3 switch, you need to add a LAN-side route (a LAN next hop).
- If the appliance is on a VLAN trunk, you need to configure VLANs on the EdgeConnect appliance so that the appliance can tag traffic with the appropriate VLAN tag.

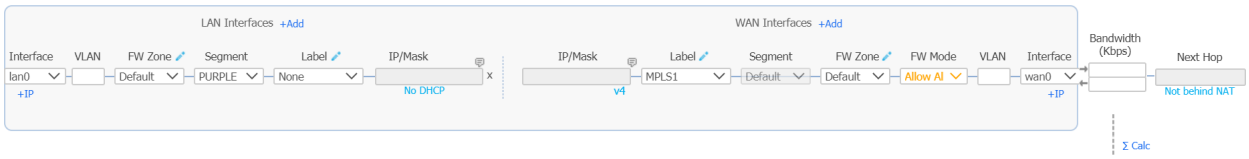
Router Mode

There are four options to consider:

1. Single LAN interface & single WAN interface
2. Dual LAN interfaces & dual WAN interfaces
3. Single WAN interface sharing LAN and WAN traffic
4. Dual WAN interfaces sharing LAN and WAN traffic

__*For best performance, visibility, and control, Options #1 and #2 are recommended because they use separate LAN and WAN interfaces.__* And when using NAT, use Options #1 or #2 to ensure that addressing works properly.

#1 - Single LAN Interface & Single WAN Interface

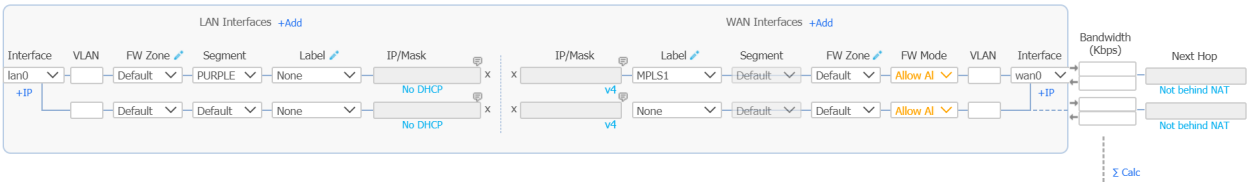


For this deployment, you have two options:

1. You can put EdgeConnect __*in-path__*. In this case, if there is a failure, you need other redundant paths for high availability.
2. You can put EdgeConnect __*out-of-path__*. You can redirect LAN-side traffic and WAN-side traffic from a router or L3 switch to the corresponding interface using WCCP or PBR (Policy-Based Routing).

To use this deployment with a single router that has only one interface, you could use multiple VLANs.

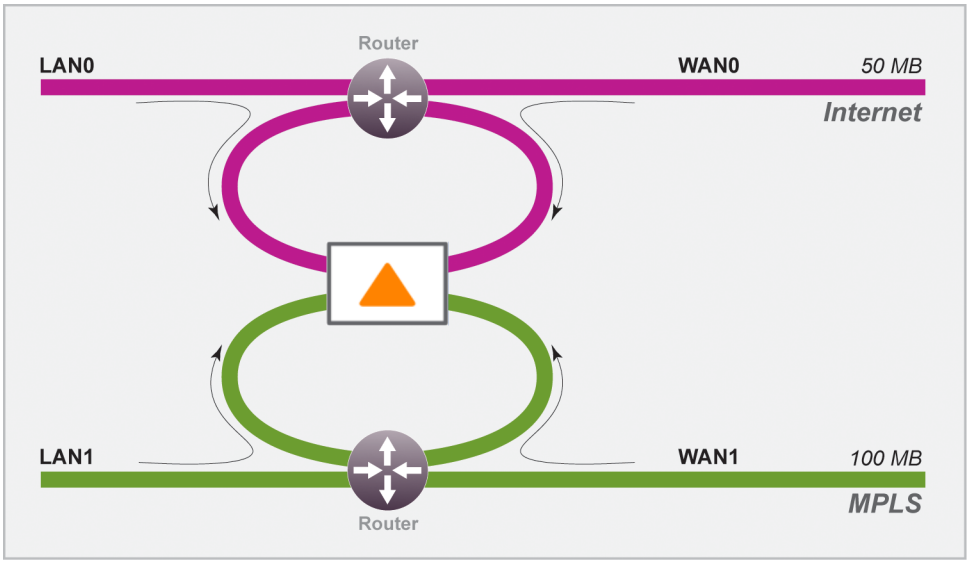
#2 - Dual LAN Interfaces & Dual WAN Interfaces



This deployment redirects traffic from two LAN interfaces to two WAN interfaces on a single EdgeConnect appliance.

- 2 WAN next-hops / 2 subnets / 1 appliance
- 2 separate service providers or WAN services (MPLS, IPSec VPN, MetroEthernet, and so forth)

Out-of-path dual LAN and dual WAN interfaces

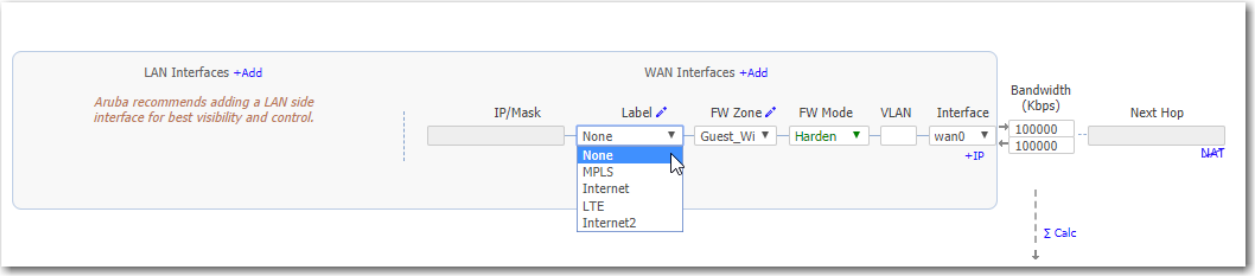


For this deployment, you have two options:

1. You can put EdgeConnect __*in-path__*. In this case, if there is a failure, you need other redundant paths for high availability.

2. You can put EdgeConnect __*out-of-path__*. You can redirect LAN-side traffic and WAN-side traffic from a router or L3 switch to the corresponding interface using WCCP or PBR (Policy-Based Routing).

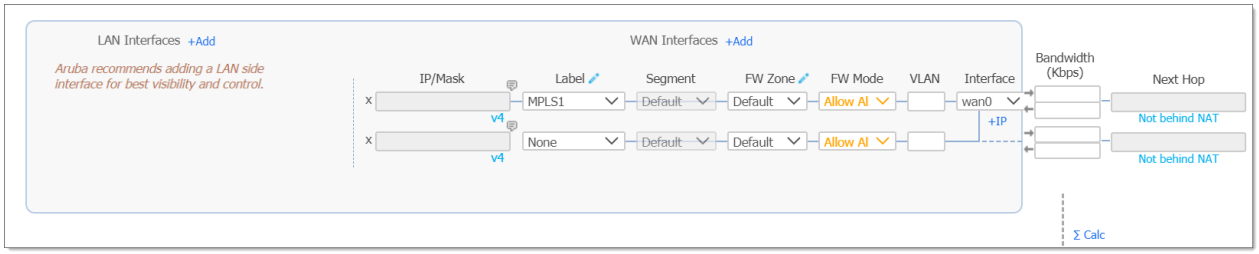
#3 - Single WAN Interface Sharing LAN and WAN traffic



This deployment redirects traffic from a single router (or L3 switch) to a single subnet on the EdgeConnect appliance.

- This mode only supports __*out-of-path__*.
- When using two EdgeConnects at the same site, this is also the most common deployment for high availability (redundancy) and load balancing.
- For better performance, control, and visibility, Router mode **Option #1** is recommended instead of this option.

#4 - Dual WAN Interfaces Sharing LAN and WAN traffic



This deployment redirects traffic from two routers to two interfaces on a single EdgeConnect appliance.

This is also known as **Dual-Homed Router Mode**.

- 2 WAN next-hops / 2 subnets / 1 appliance.
- 2 separate service providers or WAN services (MPLS, IPSec VPN, MetroEthernet, and so forth).
- This mode only supports __*out-of-path__*.
- For better performance, control, and visibility, Router mode **Option #2** is recommended instead of this option.

Considerations for Router Mode Deployments

- Do you want your traffic to be **in-path** or **out-of-path**? This mode supports both deployments. In-path deployment offers much simpler configuration.
- Does your router support VRRP, WCCP, or PBR? If so, you might want to consider out-of-path Router mode deployment. You can set up more complex configurations, which offer load balancing and high availability.
- Are you planning to use host routes on the server/end station?
- In the rare case when you need to send inbound WAN traffic to a router other than the WAN next hop router, use LAN-side routes.

Examine the Need for Traffic Redirection

Whenever you place an appliance out-of-path, you must redirect traffic from the client to the appliance.

There are three methods for __*redirecting outbound packets from the client to the appliance__* (known as **LAN-side redirection**, or **outbound redirection**):

- **PBR** (Policy-Based Routing) – Configured on the router. No other special configuration required on the appliance. This is also known as **FBR** (Filter-Based Forwarding).

  If you want to deploy two EdgeConnects at the site for redundancy or load balancing, you also need to use VRRP (Virtual Router Redundancy Protocol).
- **WCCP** (Web Cache Communication Protocol) – Configured on both the router and the EdgeConnect appliance. You can also use WCCP for redundancy and load balancing.
- **Host routing** – The server/end station has a default or subnet-based static route that points to the EdgeConnect appliance as its next hop. Host routing is the preferred method when a virtual appliance is using a single interface, **mgmt0**, for datapath traffic (also known as Server Mode).

  To ensure end-to-end connectivity in case of appliance failure, consider using VRRP between the appliance and a router, or the appliance and another redundant EdgeConnect.

How you plan to optimize traffic also affects whether you also need __*inbound redirection from the WAN router__* (known as **WAN-side redirection**):

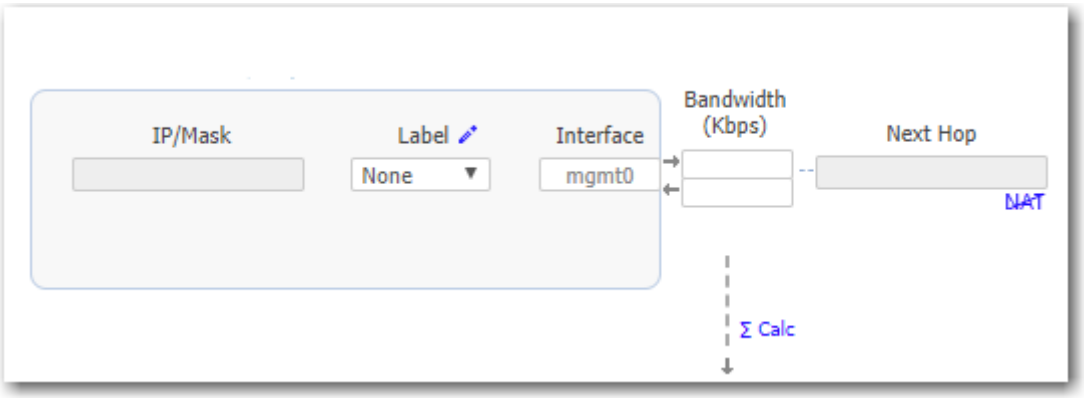
- If you use **subnet sharing** (which relies on advertising local subnets between EdgeConnect appliances) or **route policies** (which specify destination IP addresses), you only need LAN-side redirection.
- If, instead, you rely on **TCP-based** or **IP-based** auto-optimization (which relies on initial handshaking __*outside__* a tunnel), you must also set up inbound *and* outbound redirection on the WAN router.
- For TCP flows to be optimized, both directions must travel through the same client and server appliances. If the TCP flows are asymmetric, you need to configure flow redirection among local appliances.

A tunnel must exist before auto-optimization can proceed. There are three options for tunnel creation:

- If you enable **auto-tunnel**, the initial **TCP-based** or **IP-based** handshaking creates the tunnel. This means that the appropriate LAN-side and WAN-side redirection must be in place.
- You can allow the *Initial Configuration Wizard* to create the tunnel to the remote appliance.
- You can create a tunnel manually on the **Configuration > Networking > Tunnels > Tunnels** page.

Server Mode

This mode uses the **mgmt0** interface for management and datapath traffic.



ADD DATA INTERFACES

- You can create additional data-plane Layer 3 interfaces to use as tunnel endpoints.
- To add a new logical interface, click **+IP**.

Deployment - EdgeHA

EdgeHA mode is a high availability cluster configuration that provides appliance redundancy by pairing two EdgeConnect devices together.

When a deployment profile configures two EdgeConnect appliances in EdgeHA mode, the resilient cluster acts as a single logical system. It extends the robust SD-WAN multipathing capabilities such as Business Intent Overlays seamlessly across the two devices as if they were one entity.

With EdgeHA mode, a WAN uplink is physically plugged into a single one of the EdgeConnect appliances but is available to both in the cluster. For WAN connections that perform NAT (for example, a consumer-grade Broadband Internet connection), it means that only a single Public IP needs to be provisioned in order for both EdgeConnect devices in the EdgeHA cluster to be able to build Business Intent Overlays using that transport resource.

NOTE: EdgeHA does not support Active/Active deployments (that is, equal-cost multi-path [ECMP] routing). Active/Passive deployments are supported; the primary EdgeHA appliance carries all traffic to and from the LAN side of the EdgeHA cluster. You can configure Active/-Standby (Backup) by using Virtual Router Redundancy Protocol (VRRP), Border Gateway Protocol (BGP), Open Shortest Path First (OSPF) protocol, or Bidirectional Forwarding Detection (BFD) protocol.

Enable EdgeHA Mode

1. In the appliance tree, select the appliance, and then right-click to select **Deployment** from the contextual menu. The appliance's Deployment page appears.
2. Select the **EdgeHA** check box.
3. Configure the interfaces (LAN and WAN–side) on both EdgeConnect devices to reflect the WAN connections that are plugged into each one of the respective appliances.

   NOTE: Both EdgeConnect devices will be able to leverage all WAN connections regardless of which chassis they are physically plugged into. It is, however, important to match the deployment profile interface configuration to the actual chassis the WAN connection is physically, directly connected to.
4. Select the physical ports on the respective EdgeConnect appliances that you will connect to each other using an Ethernet cable (RJ-45 twisted pair or SR optical fiber).

   NOTE: You can choose any LAN or WAN port combination for this HA Link that is available on the respective EdgeConnect chassis. You must match the media type and speed for both ends of the HA link. (For example, 1 Gigabit-Ethernet RJ-45 to RJ-45 or 10 Gigabit-Ethernet multimode fiber LC-connector-to-LC-connector). Also, note that you cannot use MGMT ports for the HA Link; only LAN or WAN ports.

IPSec over UDP Tunnel Configuration

For both EdgeConnect appliances in a high availability cluster to be able to share a common transport connection, you must set the tunnel type to IPSec over UDP mode.

See Tunnel Settings in the Orchestrator (**Orchestrator > Orchestrator Server > Tools > Tunnel Settings**).

NOTE: *If you are deploying a network with EdgeConnect appliances running VXOA 8.1.6 or higher and Orchestrator 8.2 or higher, the tunnel type is already set to IPSec over UDP mode by default.*

VRRP Configuration

Typically, in a branch site deployment, you will choose to configure the cluster with a VRRP protocol and assign a VIP (virtual IP) address to the cluster.

- Set the VRRP priority of the preferred LAN-side Primary EdgeConnect to **128**.
- Set the other, Secondary appliance's VRRP priority to **127**.

LAN-side Monitoring

The IP SLA feature should be configured to monitor the LAN-side VRRP state in order to automatically disable subnet sharing from that appliance in the case of a LAN link failure.

For more information, refer to the IP SLA configuration guide.

Internal Subnets

Configuration > Overlays & Security > Internal Subnets

Use the Internal Subnets dialog box to define which subnets are internal to your SD-WAN fabric and which are external (internet-bound). This classification determines how appliances route traffic, as follows:

- Internal traffic: Destinations that match an entry in the Internal Subnets table are considered part of your internal network. This traffic is sent over configured overlay tunnels to other branches or hubs within the SD-WAN fabric. EdgeConnect appliances ship with default private addresses: RFC 1918 for IPv4 and ULA for IPv6. Traffic marked as Internal is not subject to local breakout, HPE SSE, Zscaler, Netskope, or Service Orchestration handling.
- External traffic: Destinations that do not match any entry in the Internal Subnets table are internet-bound. External traffic enters the Breakout Traffic to Internet & Cloud Services pipeline as depicted on the Business Intent Overlays tab. This traffic is routed based on the preferred policy order configured in the Business Intent Overlay (BIO).

The Internal Subnets configuration is part of the Business Intent Overlays tab and is common across all BIOs and all appliances within the fabric. For more information, see Business Intent Overlays.

Add Internal Subnets

You can configure internal subnets for the global segment or for specific VRF segments. The Internal Subnets table supports both IPv4 and IPv6 subnets (up to 512 for each).

NOTE: You can include non-default routes as internal subnets to significantly extend the number of destinations that can be classified as internal. See Consider Non-default Routes as Internal Subnets below.

To add an internal subnet:

1. Navigate to **Configuration > Overlays & Security > Internal Subnets**.

   The Internal Subnets dialog box opens.

2. Click **Add**.

   A second Internal Subnets dialog box opens.

3. From the **Segment** drop-down list, select **All** for the global segment or select a specific VRF segment. Listed VRF segments are defined on the Routing Segmentation (VRF) tab.

4. In the **Subnet / Mask** field, enter the subnet in CIDR notation (for example, 10.0.0.0/8 for IPv4 or 2001:db8::/32 for IPv6).
5. Click **Save**.
6. On the initial Internal Subnets dialog box, click **Save**.

NOTE: To edit an internal subnet, hover over the Actions column for the subnet and click the edit icon. To delete an internal subnet, hover over the Actions column for the subnet and click the delete icon.

Add or Replace Internal Subnets in Bulk

You can add internal subnets in bulk to the Internal Subnets table or replace the current list of internal subnets in the table with a new list.

1. On the Internal Subnets dialog box, click **Bulk Add/Replace**.

   The Bulk Add/Replace Internal Subnets dialog box opens.
2. From the **Segment** drop-down list, select **All** for the global segment or select a specific VRF segment.
3. In the unlabeled box below the Segment field, enter subnets in CIDR notation separated by commas.
4. To add the subnets to the Internal Subnets table, click **Save**, or click **Replace** to replace the subnets currently listed in the table with these new subnets.
5. On the Internal Subnets dialog box, click **Save**.

Consider Non-default Routes as Internal Subnets

In addition to manually configured subnets, you can include non-default routes as internal subnets. This provides a way to significantly extend the number of destinations you can classify as internal. It ensures that internal traffic is not inadvertently sent to the internet due to a missing entry in the Internal Subnets table.

Manually configured internal subnets are evaluated first. If no match is found, Orchestrator does a non-default route lookup.

To set this up, on the Internal Subnets dialog box, select the **Consider non-default routes as internal subnets** check box, and then click **Save**.

Global Firewall Policies

Configuration > Overlays & Security > Security > Global Firewall Policies

Use the Global Firewall Policies tab to define firewall policies for all appliances across your network. These firewall policies manage traffic between firewall zones.

- Zones are created on the Orchestrator. A zone is applied to an interface.

- By default, traffic is allowed between interfaces labeled with the same zone. Any traffic between interfaces with different zones is dropped. You can create exception rules (firewall policies) to allow traffic between interfaces with different zones or between their segments and firewall zones.
- When Routing Segmentation (VRF) is enabled:
  - By default, traffic is allowed between interfaces labeled with the same zone and the same segment. Any traffic between different zones or between different segments is dropped.
  - Define your firewall policies from the Routing Segmentation (VRF) tab.
  - Do not use templates. If a firewall policy template is applied while segmentation is enabled, it applies only within the default segment. It overrides the default security policy defined on the Routing Segmentation (VRF) tab. This prevents a traffic disruption when segmentation is enabled for the first time, and during a migration to segments. After the migration process finishes, remove the security policy template.
- If Routing Segmentation is disabled, define firewall policies by creating templates. You can then apply the template to interfaces and overlays.
- The Global Firewall Policies tab can be viewed in both Matrix View (default) and Table View. Click the icons at the top of the tab to toggle between views.
- To view real-time or stored (historical) security logs from EdgeConnect appliances, click **Security Log Viewer**.
- To view statistics on various flows, packets, and bytes dropped or allowed by a zone-based firewall for a given time range, click **Firewall Drops**. For information about troubleshooting flows that were denied by the firewall with the reason "outbound pkt new dst zone" or "Zone change detected on outbound packet," see this troubleshooting video.
- Clicking the edit icon in the Actions column of an appliance opens the Global Firewall Policy dialog box, which shows the applied firewall policy.

  NOTE: The edit icon is conditionally available. It is enabled for single-appliance selections and disabled for multi-appliance selections.

Global Firewall Policies Tab

Use the Global Firewall Policies tab to manage your firewall policies by adding or modifying rules.

1. Using the icons at the top of the tab, click either **Matrix View** or **Table View**.
2. Select a **Source segment** and a **Destination segment** from the drop-down menus.
3. Select a default logging level from the **Log 'deny all' events at level** drop-down menu.
4. In Matrix View, click the cell at the intersection of the source and destination zones to open the Zone <*name*> to Zone <*name*> dialog box, and then click the edit icon in the Actions column to open the rule editor. In Table View, click the edit icon in the Actions column of a rule to open the rule editor.
5. To create a new rule, click **Add rule**.

6. Modify the following fields for a new or existing rule:

Field	Description
Priority	Priority of the rule.
Match criteria (optional)	Click to open the Match Criteria dialog box and add or modify match criteria for the rule.
Action	Select the action to apply to traffic matching the rule: __Allow—__Matching traffic is allowed. __Deny—__Matching traffic is denied. __Inspect—__Matching traffic is inspected by the Intrusion Detection System (IDS).
Enabled	Select the check box to enable the rule. Clear the check box to disable the rule.
Logging	Select the logging level to apply when logging matches the specific rule. To disable logging for matching traffic, select **None**.
Tag (optional)	Enter a tag to include in logs for matching events.
Comment (optional)	Add comments about the rule.

7. Click **Done**.
8. Click **Done** again.

Wildcard-based Prefix Matching Rules

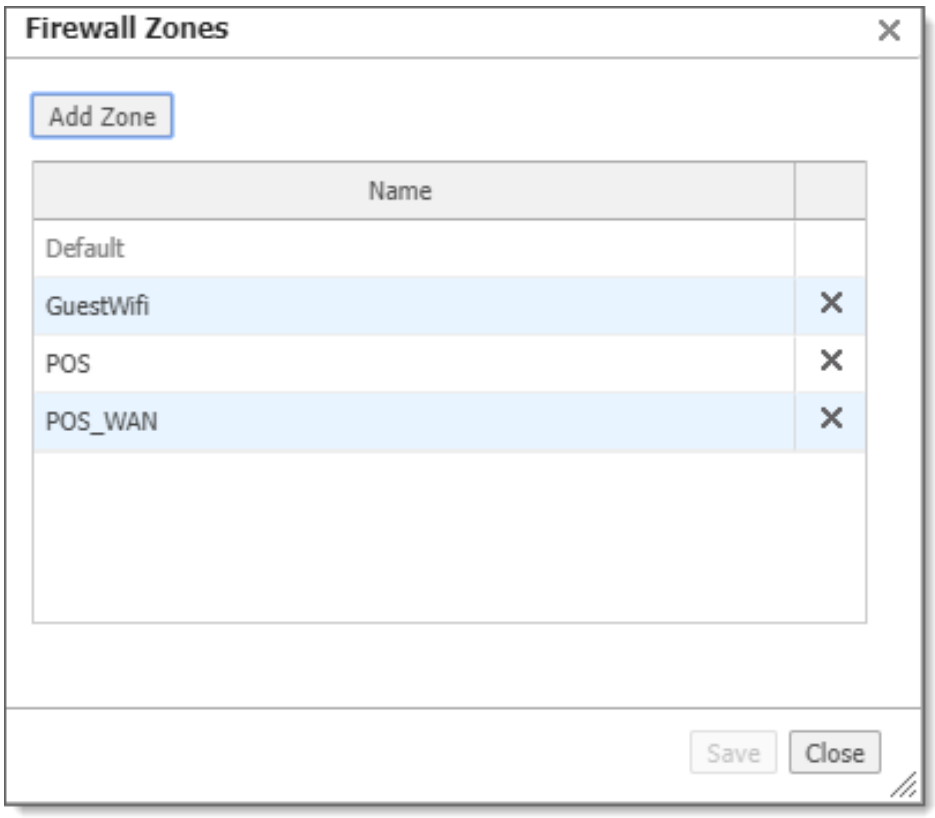
- When using a range or a wildcard, the IPv4 address must be specified in the 4-octet format, separated by the dot notation. For example, **A.B.C.D**.
- To specify a range, use a single dash (-). For example, **128-129**.
- To specify a wildcard, use an asterisk (*).
- You can use a range and a wildcard in the same address, but each octet can contain only one of them or the other. For example, **10.136-137.*.64-95**.
- Use a wildcard only to represent a complete octet. For example, **10.13*.*.64-95** is not supported. Use **10.130-139.*.64-95** to specify this range.
- These rules also apply to IPv6 addresses.
- CIDR notation and (range or wildcard) are mutually exclusive in the same address. For example, **192.168.0.1-127/24** is not supported. Use either **192.168.0.0/24** or **192.168.0.1-127**.
- Prefix-matching rules apply to the following policies only: Route, QoS, Optimization, NAT, Security, and ACLs.

Firewall Zone Definition

Configuration > Overlays & Security > Security > Firewall Zone Definition

Zone-based firewalls are created on the Orchestrator.

- A zone is applied to an **Interface**.
- By default, traffic is allowed between interfaces labeled with the same zone.
- Any traffic between interfaces with different zones is dropped.
- Users can create exception rules (Security Policies) to allow or deny traffic between interfaces within the same or different zones.



NOTE: “Default” will always be the initial default zone. You cannot have another zone named “Default”.

NOTE: The name of your firewall cannot exceed 16 characters and cannot contain any special characters. It can contain alphanumeric characters and underscores only.

Appliance Security Policies

Configuration > Overlays & Security > Security > Appliance Security Policies

The Appliance Security Policies tab lists the security policies for your appliances. These security policies manage traffic between firewall zones.

- Zones are created on the Orchestrator. A zone is applied to an interface.
- By default, traffic is allowed between interfaces labeled with the same zone. Any traffic between interfaces with different zones is dropped. You can create exception rules (security policies) to allow traffic between interfaces with different zones or between their segments and firewall zones.
- When Routing Segmentation (VRF) is enabled:
  - By default, traffic is allowed between interfaces labeled with the same zone and the same segment. Any traffic between different zones or between different segments is dropped.
  - Define your security policies from the Routing Segmentation (VRF) tab.
  - Do not use templates. If a security policy template is applied while segmentation is enabled, it applies only within the default segment. It overrides the default security policy defined on the Routing Segmentation (VRF) tab. This prevents a traffic disruption when segmentation is enabled for the first time, and during a migration to segments. After the migration process finishes, remove the security policy template.
- If Routing Segmentation (VFR) is disabled, define security policies by creating templates. You can then apply the template to interfaces and overlays.
- To toggle between views, click the **Matrix View** (default) or **Table View** filter at the top of the tab.
- To view statistics on various flows, packets, and bytes dropped or allowed by a zone-based firewall for a given time range, click **Firewall Drops**. For information about troubleshooting flows that were denied by the firewall with the reason "outbound pkt new dst zone" or "Zone change detected on outbound packet," see this troubleshooting video.
- To define policies on all appliances within your network, click **Global Firewall Policies**. Use the matrix and table views to further specify your policies.
- Clicking the edit icon in the Actions column of an appliance opens the Appliance Security Policy dialog box, which shows the applied security policy. Any changes you make are local to that appliance. Making changes from this dialog box is not recommended.

  NOTE: The edit icon is available when you select a single appliance. It is disabled when you select multiple appliances.

Appliance Security Policies Tab

Use the Appliance Security Policies tab to manage your security policies by adding or modifying rules.

1. Using the icons at the top of the tab, click either **Matrix View** or **Table View**.
2. Select a **Source segment** and a **Destination segment** from the drop-down menus.
3. Select a default logging level from the **Log 'deny all' events at level** drop-down menu.

4. In Matrix View, click the cell at the intersection of the source and destination zones to open the Zone <*name*> to Zone <*name*> dialog box, and then click the edit icon in the Actions column to open the rule editor. In Table View, click the edit icon in the Actions column of a rule to open the rule editor.
5. To create a new rule, click **Add rule**.
6. Modify the following fields for a new or existing rule:

Field	Description
Priority	Priority of the rule.
Match criteria (optional)	Click **Match everything** to open the Match Criteria dialog box and add or modify match criteria for the rule.
Action	Select the action to apply to traffic matching the rule: __Allow—__Matching traffic is allowed. __Deny—__Matching traffic is denied. __Inspect—__Matching traffic is inspected by the Intrusion Detection System (IDS).
Enabled	Select the check box to enable the rule. Clear the check box to disable the rule.
Logging	Select the logging level to apply when logging matches the specific rule. To disable logging for matching traffic, select **None**.
Tag (optional)	Enter a tag to include in logs for matching events.
Comment (optional)	Add comments about the rule.

7. Click **Done**.
8. Click **Done** again.

Wildcard-based Prefix Matching Rules

- When using a range or a wildcard, the IPv4 address must be specified in the 4-octet format, separated by the dot notation. For example, **A.B.C.D**.
- To specify a range, use a single dash (-). For example, **128-129**.
- To specify a wildcard, use an asterisk (*).
- You can use a range and a wildcard in the same address, but each octet can contain only one of them or the other. For example, **10.136-137.*.64-95**.
- Use a wildcard only to represent a complete octet. For example, **10.13*.*.64-95** is not supported. Use **10.130-139.*.64-95** to specify this range.
- These rules also apply to IPv6 addresses.

- CIDR notation and (range or wildcard) are mutually exclusive in the same address. For example, **192.168.0.1-127/24** is not supported. Use either **192.168.0.0/24** or **192.168.0.1-127**.
- Prefix-matching rules apply to the following policies only: Route, QoS, Optimization, NAT, Security, and ACLs.

Firewall Protection Profiles

Configuration > Overlays & Security > Security > Firewall Protection Profiles

Use the Firewall Protection Profiles tab to add or modify a protection profile on any appliance with a firewall, to enable baseline learning, and to manage denial of service (DoS) thresholds and corresponding response actions on designated appliances, segments, and zones.

Baseline Learning, Auto Rate Limit, and Smart Burst

When you enable baseline learning, the system establishes a baseline for network performance during normal operations. Flow baselines are established at regular intervals by analyzing network statistics and identifying new patterns based on observed data. The default interval for baseline learning is 14 days, but you can set the interval to 7-56 days. Baseline calculations are based on a snapshot of a combination of metrics and are computed in the background.

The baselines provide a way to continuously collect and aggregate data about your network. Use this data to assess network zone capacities and platform performance levels. The baselines are also used to build the graphs and charts found on Flow Baselines and Flow Baseline Trends.

You can enable baseline learning without adding any firewall protection profiles. To use the Auto rate limit or Smart burst DoS thresholds, you must enable baseline learning because baseline calculations are used to calculate targets for both. Baseline learning, Auto rate limit, and Smart burst all require either an AS (Advanced Security) license or an AAS-DTD (Dynamic Threat Defense) license.

NOTE: To disable baseline learning, you must first remove any firewall protection profiles that use Auto rate limit or Smart burst. If these are not removed before you disable baseline learning, a notification appears on the screen, and you cannot proceed with disabling the feature.

Auto Rate Limit

Auto rate limit is a DoS threshold setting that uses baseline learning to compute the minimum DoS threshold. The maximum DoS threshold is configured above the baseline minimum. Auto rate limit helps assess network zone capacities and platform performance levels; normal versus oversubscribed. It allows significant burst but limits a zone to a percentage of appliance capacity. The maximum value is not sensitive to zone trends, and flows that exceed the maximum value are “tail dropped” even if the zone or appliance has flow capacity.

Smart Burst

Smart burst is a DoS threshold setting that uses baseline learning to compute both minimum and maximum DoS thresholds and allocates extra flow capacity. It uses a triple token bucket zone-based policer schema for burst management. Smart burst does the following:

- Optimizes the use of appliance capacity by zones.
- Protects baseline flow capacity for all zones.
- Automatically calculates reserve/spare flow capacity.
- Supports two levels of bursts to mitigate tail drops and manages spare capacity of appliances to support bursts.
  - The spare capacity is distributed among all zones and is called committed burst. Committed burst is the first level of burst capacity.
  - On a per second basis, unused committed burst in zones is made available as a second level of burst capacity that is referred to as excess burst. Unused excess burst capacity goes back to the respective committed burst periodically.

Enable Baseline Learning

This section describes how to enable baseline learning for one EdgeConnect appliance. To enable baseline learning for multiple EdgeConnect appliances using a template, see Firewall Protection Profiles Template.

NOTE: Baseline learning, Auto rate limit, and Smart burst require either an AS (Advanced Security) license or an AAS-DTD (Dynamic Threat Defense) license.

1. Hover over the Actions column for the appliance you want to configure and click **Edit**.

   The Firewall Protection Profile - <appliance name> dialog box opens.
2. Select the **Baseline learning** check box.
3. To customize the baseline learning settings, click **Baseline settings**.

   The Baseline Settings dialog box opens.
4. Enter the following information based on your network or click **Cancel** to use the default settings.

Field	Description
Data aggregation method	Technique used for data aggregation. The default is percentile. No other options exist at this time.
Data aggregation limit	Indicates what percentage of the sample data is used to determine baseline values. The default setting is 95%, which means the top 5% of the sample is discarded and the other 95% is considered when computing the baselines. You can enter a value between 75-100%.

Field	Description
Computation interval	Time that passes before the system computes new baselines. The default is 8 hours. For example, when using the default, the baselines are computed every 8 hours using the latest sample data collected during the Model training interval. This can be configured in 4-hour units (e.g., 4, 8, 12, and so on) up to 240 hours.
Model training interval	During this period, data is collected for various metrics every five minutes and is aggregated into a data file. This data is used to compute the baselines. The default is 14 days, the minimum is 7 days, and the maximum is 56 days. **NOTE:** This period should include a diverse set of data that covers various types of legitimate traffic and captures the characteristics that distinguish normal traffic from malicious traffic during an attack.
Baseline upper limit	Upper limit for the minimum baseline. An alarm is raised when this value is reached. This setting is useful if Auto rate limit is configured without Smart burst. The setting is a percentage of the maximum baseline value, which is set manually. The default is 90%. You can enter a value between 50-100%.
TCP inactivity timeout	Inactivity timeout used for TCP flows created using burst support levels. Inactive flow gets deleted after this timeout. The default is 300 seconds. You can enter a value between 30-1800 seconds.
Headroom for baseline plus	Percentage of headroom that is added to the baseline. The default is 20%. You can enter a value between 5-100%.
Per-source limit for committed burst	Committed burst for a zone is available to all sources in the zone. This determines the percentage of committed burst in a zone that one source can use. The default is 50%. You can enter a value between 1-50%.
Reserve flow capacity distribution	Spare flow capacity is distributed among all zones by Smart burst using different methods (**Proportional** or **Equal**). The default method is Proportional.
Excess burst credit interval	On a per second basis, the zone is supposed to use a portion of committed burst capacity. Unused committed burst capacity of zones is made available as excess burst capacity every second. After this interval of time, unused excess burst capacity goes back to the respective committed burst. The default is 30 seconds. Enter a value between 30-100 seconds.
Minimum reserve capacity limit	Minimum amount of reserve flow capacity that should be available before Smart burst redistributes new reserve capacity after a baseline computation interval. Smart burst continues with previously distributed capacities if the minimum reserve capacity limit is not available. The default is 20%. You can enter a value between 10-50%.

5. Click **Done** to close the Baseline Settings dialog box.
6. Click **Save** to close the Firewall Protection Profile - <appliance name> dialog box and save your changes.

SYN Cookie, Smart SYN Cookie, and Dynamic IP Reputation Management

SYN cookie is a protection feature of the EdgeConnect firewall. It mitigates distributed denial of service (DDoS) attacks that exploit the TCP three-way handshake process—specifically the initial SYN (synchronize) packet—referred to as a SYN flood attack. SYN flood attacks can exhaust system resources, including connection limits and network bandwidth, preventing legitimate connections. SYN cookie dynamically analyzes traffic patterns. It automatically activates to ensure uninterrupted availability and integrity of network services.

SYN cookie is a Max action for DoS threshold that is applicable to the Embryonic Two-Way flows (half-open connections) metric. If the SYN cookie Max action is active when an Embryonic Two-Way flow exceeds the Max value, EdgeConnect generates a SYN cookie challenge (SYN-ACK packet) for the requesting client. If the requesting client responds with a valid ACK packet (SYN cookie challenge packet), the session is allowed, flow creation in EdgeConnect starts, and the server-side session starts. The SYN cookie feature allows genuine clients to establish TCP sessions and mitigates SYN flood attacks by preventing malicious clients from creating half-open connections with the server that use server resources.

The EdgeConnect firewall confers the following flow statuses on TCP connections.

- Embryonic Flow: The EdgeConnect has received the SYN packet from the requesting client.
- Embryonic Two-Way Flow: The EdgeConnect has received the SYN-ACK packet from the server. If the EdgeConnect never receives the SYN-ACK packet from the server, the flow status remains as Embryonic.
- Established Flow: The EdgeConnect has received the ACK packet from the requesting client. If the EdgeConnect never receives the ACK packet from the requesting client, the flow status remains as Embryonic Two-Way.

NOTE: Smart SYN cookie and dynamic IP reputation require either an AS (Advanced Security) license or an AAS-DTD (Dynamic Threat Defense) license.

Smart SYN Cookie

Smart SYN cookie uses smart functionality to increase the efficiency of the SYN cookie feature. You can select it as a Max action for a DoS threshold value for TCP embryonic two-way flows.

It uses dynamic IP reputation management to evaluate source IP addresses and assign a reputation score to them. The system uses this reputation score to determine how to respond to a TCP connection request. Depending on the score, it responds with a SYN cookie, generates a normal flow, or adds the source IP address to the blocklist.

Dynamic IP Reputation Management

Smart SYN cookie uses dynamic IP reputation management to assign a reputation score to source IP addresses. The system builds and continuously updates the reputation score based on the history of "good" (trusted) transactions for an IP address. The reputation score is built using the following data:

- Internal IP Reputation: Computed by ECOS for all IP addresses using both egress and ingress flows. You must enable Internal IP Reputation for the Smart SYN cookie feature to work properly. Internal IP Reputation is applied to both egress and ingress flows.
- External IP Reputation: Provided by BrightCloud/Webroot feed "IP Reputation".

Enable Internal IP Reputation

This section describes how to enable Internal IP Reputation for one EdgeConnect appliance. To enable it for multiple EdgeConnect appliances using a template, see Firewall Protection Profiles Template.

NOTE: Smart SYN cookie and dynamic IP reputation require either an AS (Advanced Security) license or an AAS-DTD (Dynamic Threat Defense) license.

1. Hover over the Actions column for the appliance you want to configure and click **Edit**.

   The Firewall Protection Profile - <appliance name> dialog box opens.
2. Select the **Internal IP reputation** check box.
3. To customize the reputation settings, click **Reputation settings**.

   The Internal IP Reputation dialog box opens.
4. Enter the following information based on your network or click **Cancel** to use the default settings.

Field	Description
Successive Good Transactions (Trusted)	Number of consecutive good (trusted) transactions that an IP address must incur to be considered trusted. The default is 20 transactions. You can enter a value between 5-1024.
Successive Bad Transactions (Untrusted)	Number of consecutive bad (untrusted) transactions that an IP address can incur before it is considered untrusted. The default is 20 transactions. You can enter a value between 5-1024.

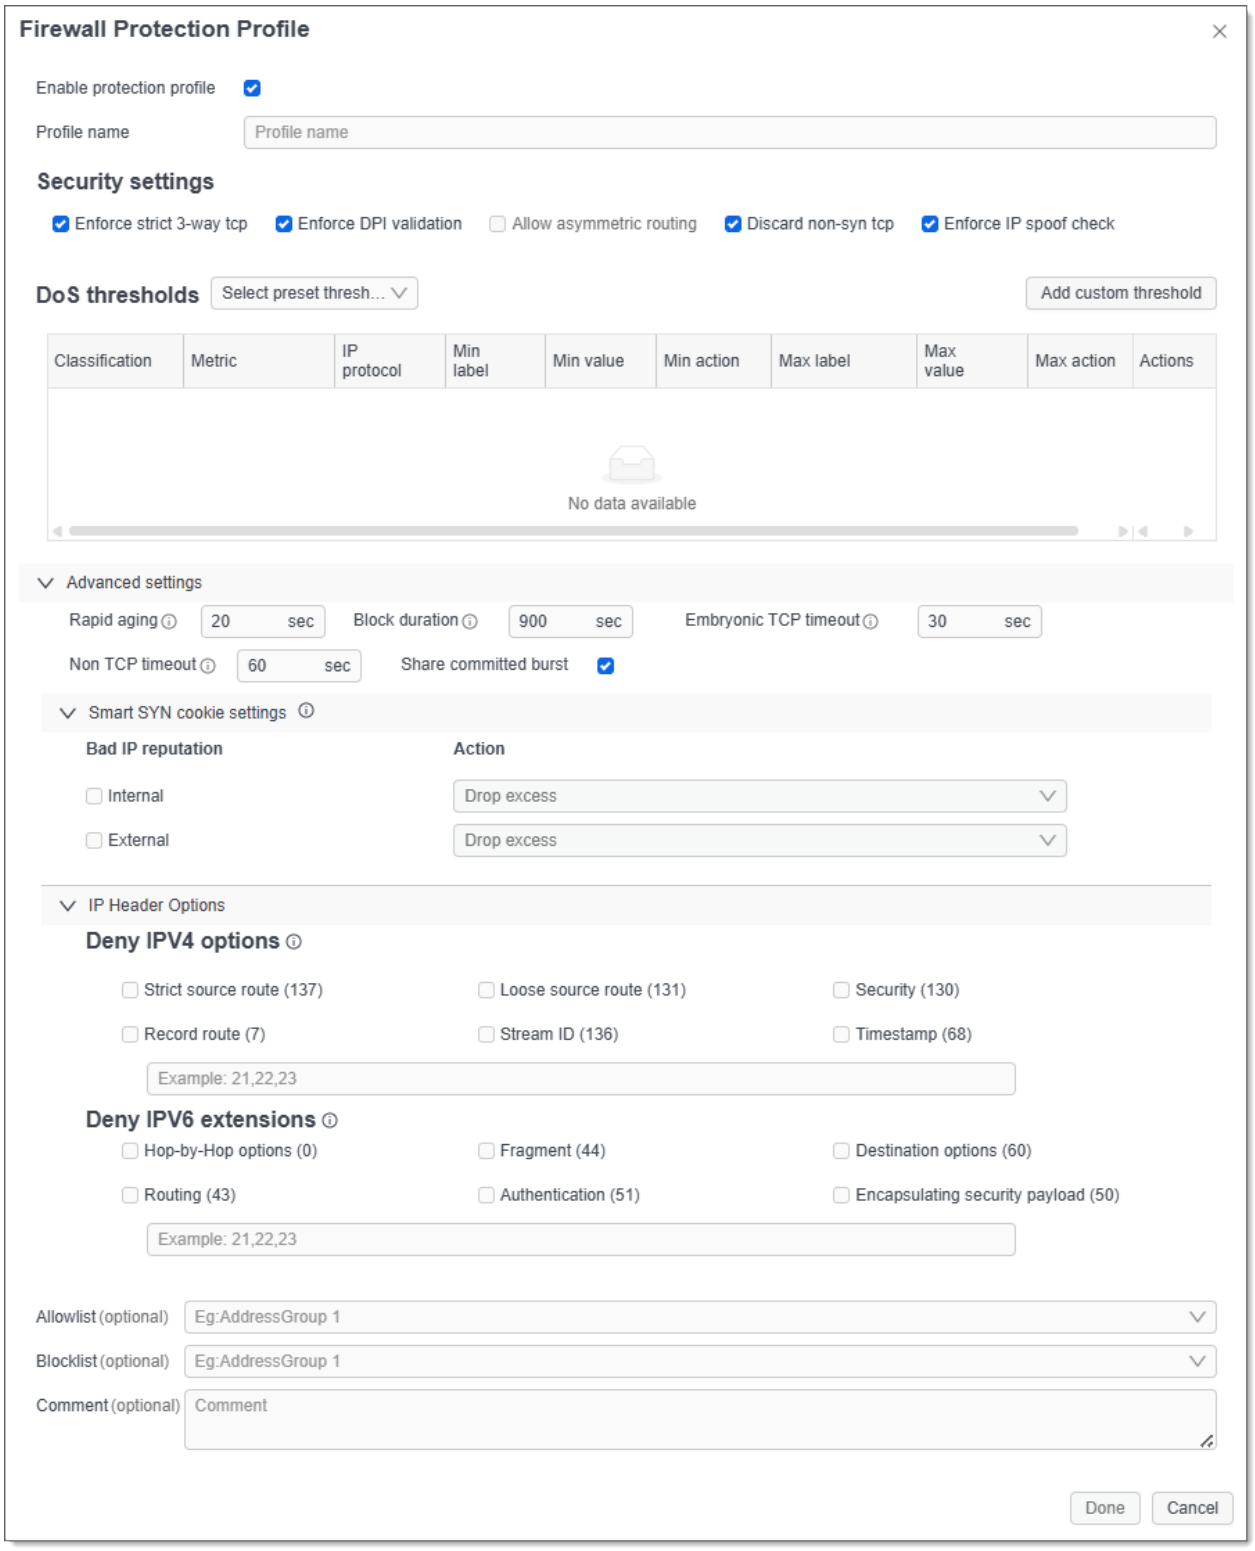
5. Click **Save** to close the Internal IP Reputation dialog box.
6. Click **Save** to close the Firewall Protection Profile - <appliance name> dialog box and save your changes.

Create a Firewall Protection Profile

1. Navigate to **Configuration > Overlays & Security > Security > Firewall Protection Profiles**.
2. Hover over the Actions column for the appliance you want to configure and click **Edit**.

   The Firewall Protection Profile - <appliance name> dialog box opens.
3. In the Profiles section, click **Add**.

   The Firewall Protection Profile dialog box opens.



4. Enter a name for the profile.
5. Select or clear any of the Security settings check boxes.

NOTE: When asymmetric routing is configured, strict three-way TCP enforcement and deep packet inspection (DPI) validation cannot be performed. To enable these settings, turn off asymmetric routing.

6. From the DoS thresholds drop-down list, select a preset threshold (**Lenient**, **Moderate**, **Strict**, **Auto rate limit**, or **Smart burst**).

   NOTE: To clear your selection and reset the DoS thresholds table, choose **Select preset threshold** from the drop-down list.

   NOTE: To use Auto rate limit or Smart burst, you must enable baseline learning first. These options only appear in the drop-down list after baseline learning is enabled.

7. *(Optional)* Add exceptions to the following fields:

Field	Description
Allowlist	Enter an existing Address Group. Any IP address contained within the Address Group will be exempt from DoS threshold analysis. The Allowlist does not exempt flows from the options shown in the Security Settings section.
Blocklist	Enter an existing Address Group to explicitly block any IP address contained within the configured Address Group.

8. *(Optional)* Click **Advanced settings** and set the following fields:

Field	Description
Rapid aging	Set a threshold value (in seconds) to enforce the tearing down of TCP connections when the period of inactivity matches the configured value (for example, 30s).
Block duration	Enforce dynamic blocking of flows originating from a source for a specified duration (for example, 300s).
Embryonic TCP timeout	Set this value so that the firewall can tear down half-open TCP connections when the timeout value is reached (for example, 30s). While a TCP connection goes through the three-way handshake (SYN, ACK, SYN-ACK), an embryonic connection is a half-open connection that produces, for example, a SYN without the other two parts of the handshake. This is a popular form of denial of service (DoS) attack.
Non TCP timeout	Set this value so that the firewall can tear down unidirectional/half-open non-TCP connections when the timeout value is reached. Default is 60 seconds.
Share committed burst	Select this check box to enable unused committed burst to be shared with other zones. This check box is enabled by default. For critical zones, you can disable this option, which retains the committed burst capacity for the zone itself.

Field	Description
Smart SYN cookie settings	Bad IP Reputation **Internal**: Select the check box to enable, and then select the action to take if the internal IP reputation for an IP address that is requesting a TCP connection is Untrusted (Drop excess or Block source). Default is Drop excess. **External**: Select the check box to enable, and then select the action to take if the external IP reputation for an IP address that is requesting a TCP connection is Untrusted (Drop excess or Block source). Default is Drop excess. **NOTE:** Smart SYN cookie settings require AS or AAS-DTD license and internal IP reputation must be enabled.
Deny IPV4 options	For appliances running ECOS version 9.4.0.0 or later, select the check box for any of these items to deny their inclusion in the Options field of IPv4 packet headers: **Strict source route**, **Loose source route**, **Security**, **Record route**, **Stream ID**, and **Timestamp**. In the text box, enter any options using comma-separated format.
Deny IPV6 extensions	For appliances running ECOS version 9.4.0.0 or later, select the check box for any of these optional extension headers of IPv6 packets to deny their inclusion: **Hop-by-Hop options**, **Fragment**, **Destination options**, **Routing**, **Authentication**, and **Encapsulating security payload**. In the text box, enter any extension options using comma-separated format.

9. Click **Done** to close the Firewall Protection Profile dialog box.
10. Click **Save** to close the Firewall Protection Profile - <appliance name> dialog box and save your changes.

Set Firewall Protection Profile Thresholds

To view the threshold settings on an existing firewall protection profile, click the link in the Thresholds count, Min thresholds, or Max thresholds columns of the Firewall Protection Profiles table.

NOTE: You can use both preset and custom thresholds.

To set or edit the threshold settings:

1. Hover over the Actions column for the appliance you want to configure and click **Edit**.

   The Firewall Protection Profile - <appliance name> dialog box opens.
2. Hover over the Actions column for the profile you want to edit and click **Edit**.

   The Firewall Protection Profile dialog box opens.

3. Either select a preset threshold from the DoS thresholds drop-down list, or click **Add custom threshold**.

   The DoS Threshold dialog box opens.

   NOTE: To edit the settings for an existing preset or custom threshold, hover over the Actions column for the threshold and click **Edit**.

4. Set the following parameters.

   NOTE: To return preset threshold settings to their defaults, click **Reset to default**, and then click **Reset** to confirm the change. This does not impact custom thresholds.

Field	Description
Classification	Classify flows in two ways: **Zone-level:** Flows originating from multiple endpoints that are part of a single firewall zone. **Source-level:** All flows originating from a single endpoint or source device.
Metric	DoS thresholds can be configured with any or all of the four metrics available in a firewall protection profile: **Flows per second:** Rate of flow (fps). A single flow is a unidirectional set of packets containing common attributes (source and destination IP, ports, protocols). **Concurrent flows:** Number of flows that are active at a given point in time. **Embryonic flows:** An incomplete TCP connection where only the first step in the three-way handshake has been completed; the EdgeConnect has received the SYN packet from the requesting client. For other protocols, such as UDP or ICMP or other, this indicates uni-directional traffic and the flow is called Embryonic. **Embryonic Two-Way flows:** A half-open TCP connection occurs when only the SYN and SYN-ACK packets of the three-way handshake are completed. In this state, EdgeConnect has received the SYN-ACK from the server, but it has not received the final ACK from the Client.
IP Protocol	Select an IP protocol (**TCP**, **UDP**, **ICMP**, **Others**, or **All**) for use in threshold settings.

Field	Description
Min Label	Select the method used to determine the min value: **Baseline**: If selected, the min value is determined by the system using baseline learning, and the corresponding Value field shows “Dynamic”. This option is available only if Baseline Learning is enabled. **Custom**: If selected, you configure the min value by entering a percentage in the corresponding Value field.
Value	Minimum threshold value as a percentage of target appliance flow capacity. When this value is breached, the protection profile takes a corresponding minimum action. If Baseline is selected as the Min Label, the system determines this value, and it cannot be configured.
Action	Action to take when the min value is breached (**Log**, **Rapid aging**, **Drop excess**, **Block source**, or **SYN cookie**). Because this corresponds to the min value, a less intensive action can be configured.
Max Label	Select the method used to determine the max value: **Custom**: If selected, you configure the max value by entering a percentage in the corresponding Value field. **Baseline plus**: A buffer of 20% is added to the computed baseline when determining flow capacity. If selected, the max value is determined by the system using baseline learning and the corresponding Value field shows “Dynamic”. **Committed burst**: Reserve flow capacity is divided equally or proportionally among all zones configured for Smart burst. If selected, the max value is determined by the system using baseline learning and the corresponding Value field shows “Dynamic”. **Excess burst**: Continuously, on a per second basis, unused committed burst (distributed reserve flow capacity) is collected from all zones and shared as a second level of support for all zones. If selected, the max value is determined by the system using baseline learning and the corresponding Value field shows “Dynamic”.
Value	Maximum threshold value as a percentage of target appliance flow capacity. When this value is breached, the protection profile takes a corresponding maximum action. If Baseline plus, Committed burst, or Excess burst are selected as the Max Label, the system determines this value, and it cannot be configured.
Action	Action to take when the max value is breached (**Log**, **Rapid aging**, **Drop excess**, **Block source**, **SYN cookie**, or **Smart SYN cookie**). Because this corresponds to the max value, a more intensive action can be configured.

5. Click **Done** to close the Firewall Protection Profile dialog box.
6. Click **Save** to close the Firewall Protection Profile - <appliance name> dialog box and save your changes.

Add Profile Mappings

After you create a profile, you can map it to a segment and zone of your firewall to achieve the expected behavior.

To map a profile to a segment:

1. In the Profile Mappings section, click **Add**.
2. From the Segment drop-down list, select the segment you want to map to your profile. To search for a segment in the list, type the segment name.
3. From the Zone drop-down list, select the zone you want to assign to your profile. To search for a zone in the list, type the zone name.
4. From the Profile drop-down list, select the profile you created earlier. To search for the profile in the list, type the profile name.
5. Click **Done** to close the Profile Mapping dialog box.
6. Click **Save** to close the Firewall Protection Profile - <appliance name> dialog box and save your changes.

Add Firewall Protection Profile to a Template Group

1. Navigate to **Configuration > Templates**.
2. Select a template group to add the firewall protection profile to, and then click **Add/Edit**.

   Firewall Protection Profiles appears as a template under Active Templates > Policies.

View DoS Threshold Alarms

To view a list of alarms triggered when a DoS threshold is breached, navigate to **Monitoring > Summary > Alarms**, and then search for “DoS” in the search bar. For more information, see Alarms.

Schedule IPSec Key Rotation

Configuration > Overlays & Security > Security > IPSec Key Rotation

Use the Schedule IPSec Key Rotation dialog box to schedule the rotation of auto-generated IPSec pre-shared keys.

Failure Handling and Orchestrator Reachability

Orchestrator distributes key material to all EdgeConnect appliances in the network. Immediately before the end of a key rotation interval, Orchestrator activates new ephemeral key material for all of the EdgeConnect appliances in the SD-WAN network. For key activation, all the appliances should be reachable to Orchestrator. However, there are two cases of unreachability:

1. **Inactive appliances:** When appliances are inactive, they exist in the Orchestrator, but do not have tunnels configured to any active appliances.
2. **Temporary unreachability:** Temporary unreachability issues occur in cases where an EdgeConnect appliance reboots or if there is a link or communication failure. In this case, Orchestrator will not activate the new key material until all active appliances are reachable and have received the new key material or if the maximum activation wait time has been exceeded. If the appliance is unreachable for a period longer than the key rotation interval, it will be treated as an inactive appliance.

Re-authorization: Inactive appliances that become active at a later point in time will be authorized to receive the current key material. Only then will they be able to download configurations and build tunnels.

Schedule IPSec Key Rotation Dialog Box

The Schedule IPSec Key Rotation dialog box enables you to schedule your key rotation. The following tables provide details about the two sections in this dialog box.

SD-WAN IPSec UDP Key Material Rotation Section

Field	Description
Enable Key Rotation	Select this check box to enable key rotation.
Persist Key Material	If enabled, key material is stored on each appliance, ensuring data plane tunnels are built quickly after an appliance reboot (no dependency on Orchestrator). If disabled, new key material from Orchestrator is required after any reboot (Orchestrator reachability is critical).
Max Activation Wait	Maximum time (in hours) Orchestrator must wait before activating the new key material. This wait time applies only when unreachable appliances exist in the network *and* at least one tunnel is UP from a reachable appliance to an unreachable appliance. This gives you time to fix connectivity issues. After the wait time expires, Orchestrator activates the new key material on all reachable appliances. Generally, it is recommended to set this wait time to half of the rotation period.

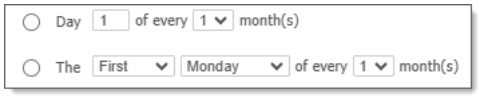
Field	Description
Rotation Period	Click the edit icon to open the Schedule dialog box, which enables you to set the rotation and the time for key material rotation to begin. Select **Hourly**, **Daily**, **Weekly**, **Monthly**, or **Yearly**, and then specify the appropriate schedule criteria. For monthly, see Monthly Schedule Rotation(./Mediately start a new key material rotation.
Key Material Lifetime	Amount of time a key material lasts. **CAUTION:** The lifetime must be at least three times the amount of the set Rotation Period.

SD-WAN IPSec Pre-Shared Key Rotation Section

Field	Description
Enable	Select this check box to enable.
Period	Click the edit icon to open the Schedule dialog box, which enables you to set the time for key rotation to begin. Select **Daily**, **Weekly**, **Monthly**, or **Yearly**, and then specify the appropriate schedule criteria. For monthly, see Monthly Schedule Rotation below for more details.

Monthly Schedule Rotation

When you schedule a *monthly* rotation on the Schedule dialog box using either of the last two options:



selecting every 1, 2, 3, 4, or 6 months determines the specific months during which rotation will occur, as indicated in the following table.

Rotation every …	JAN	FEB	MAR	APR	MAY	JUN	JUL	AUG	SEP	OCT	NOV	DEC
1 month	X	X	X	X	X	X	X	X	X	X	X	X
2 months	X		X		X		X		X		X	
3 months	X			X			X			X		

Rotation every …	JAN	FEB	MAR	APR	MAY	JUN	JUL	AUG	SEP	OCT	NOV	DEC
4 months	X				X				X			
6 months	X						X					

These months are fixed *regardless* of the day of the month you choose or when you save the monthly rotation configuration. For example, if you schedule monthly rotation on day 5 of every 3 months, and the current date is February 8th, the next rotation will occur on April 5th.

Inbound Port Forwarding

Configuration > Overlays & Security > Security > Inbound Port Forwarding

Inbound port forwarding allows traffic from the WAN to reach computers or services within a private LAN when you have a stateful firewall. It helps define and manage inbound traffic, remap a destination IP address and port number to an internal host, and create policies to manage branch devices from the WAN. Use this tab to define the desired inbound traffic.

Inbound Port forwarding is available in two modes when you add or edit a rule, depending on whether the Translate Mode is enabled or disabled. Inbound Port forwarding supports both IPv4 and IPv6 addresses.

When Translate Mode Is Disabled

The first operating mode for inbound port forwarding is when Translate Mode is disabled with inbound port forwarding. The LAN-side subnet with private IP addresses is allowed access through an inbound port forwarding rule (defined by you in the following steps) and exposes any external services. This requires LAN side private addresses to be routed on the WAN side. This represents the process of DMZ (Demilitarized Zone).

NOTES: 1. This mode is not common unless the port forwarding source is directly connected to the EdgeConnect or if the LAN-side device address is routed from the WAN side; 2. Inbound port forwarding does not support TFTP servers; 3. The DMZ feature is ineffective for local interface addresses. If you try to set up an inbound port forwarding rule to a local interface IP address on the same EdgeConnect appliance upon which the rule is created, the flow will not be successful. However, it will be successful if you enable translation with the same /32 destination IP.

To establish a DMZ connection when Translate Mode is disabled with inbound port forwarding, complete the following steps:

1. Go to the **Inbound Port Forwarding** tab.
2. Select the **Edit** icon next to **Appliance**.
3. Select **Add Rule**.

4. Complete each field with the appropriate information.

Field	Description
Source IP/Subnet	Source IP of the device accessing the internal host. This can be used to restrict access to specific IP addresses or ranges. IPv6 Source IP/Subnet addresses are supported.
Destination IP/Subnet	Address of the WAN interface IP. This is the IP address or subnet of the device within the private network to which the inbound traffic is directed. IPv6 Destination IP/Subnet addresses are supported.

When Translate Mode Is Enabled

When Translate Mode is enabled, the EdgeConnect WAN interface performs destination NAT to reach LAN-side device(s) from an external network.

Complete the following steps to enable the Translate Mode. This represents the process of DNAT (Destination Network Translation).

1. Go to the **Inbound Port Forwarding** tab.
2. Select the **Edit** icon.
3. Select **Add Rule**.
4. Select the **Translate** check box to enable Translate Mode.
5. Complete each field with the appropriate information.

Field	Description
Source Interface	Source interface name. The network interface from which the inbound traffic originates.
Source IP/Subnet	Source IP of the device accessing the internal host. This can be used to restrict access to specific IP addresses or ranges. IPv6 Source IP/Subnet addresses are supported.
Destination IP/Subnet	Address of the WAN interface IP. This is the IP address or subnet of the device within the private network to which the inbound traffic is directed. IPv6 Destination IP/Subnet addresses are supported.
Destination Port/Range	Port/range of the LAN device(s) that are managed remotely. This is the port number or range on the destination device that the inbound traffic targets.
Protocol	Select the protocol you want to apply: **UDP, TCP, ICMP, Any**. If you select **Any**, the Destination and Translated Ports have a default value that need to be between 0-100. If the value exceeds, 100 a warning appears.

Field	Description
Translated IP	IP address of the LAN device accessed inside your network.
Translated Port/Range	Port/range of the LAN device accessed inside your network.
Segment	Name of the segment being used.
Comment	Any additional details.

Additional Information

- **Interface Modes**

  Port forwarding is used only when you have 'stateful' or 'stateful+snat' configured on interfaces. It does not apply when you have 'Allow All' or 'Harden' configured.

- **Security Policies**

  *If 'security policies' are configured, make sure they allow the traffic specified in the port forwarding rules.

- You can also reorder the appliances associated with inbound port forwarding by selecting **Reorder** when adding a rule.

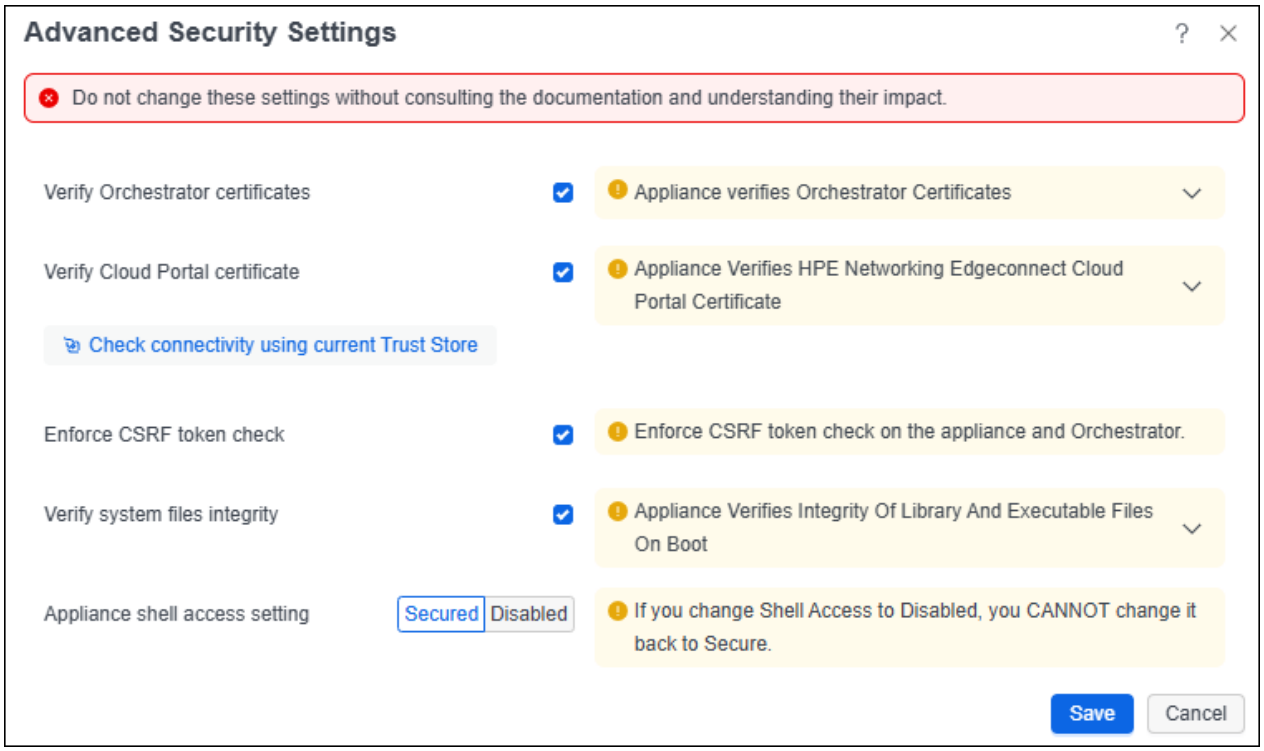
NOTE: 'Any' is a protocol option only on versions 8.1.9.4 and later.

Advanced Security Settings

Configuration > *Overlays & Security* > *Security* > *Advanced Security Settings*

Use the Advanced Security Settings dialog box to enable various security features for your network. Proceed with caution because setting these features can adversely affect your network. By default, all settings are automatically enabled for new Orchestrator installations starting with 9.3.1. Orchestrator upgrades retain previous advanced security settings except for *Perform Additional Identity Verification on Web Sockets*, which is always enabled and is no longer displayed as a setting on the Advanced Security Settings dialog box.

IMPORTANT: These settings can adversely affect your network. Understand the effects of changing them and proceed with caution. HPE Aruba Networking strongly recommends that these settings be always enabled, including Secure Shell Access.



The following security settings enable appliances to verify certificates. EdgeConnect appliances are pre-loaded with the Mozilla root store, which contains the root certificates of public certificate authorities (CAs). Appliances use this root store for cryptographic verification when opening Transport Layer Security (TLS) connections to Orchestrator and the Cloud Portal.

- **Verify Orchestrator and Stats Collector Certificates**

  Enables appliances to verify the Orchestrator and Stats Collector certificates. Disable this setting if any of the following statements are true:

  - Orchestrator or Stats Collector uses a self-signed certificate.
  - Orchestrator or Stats Collector is behind a proxy server.
  - Any appliances in your network are not configured with the Orchestrator or Stats Collector domain name.
  - Appliances are using Orchestrator as a proxy to reach the Cloud Portal, but Orchestrator does not have a valid certificate.
  - Orchestrator does not have a certificate signed by a public certificate authority (CA) or does not have the appropriate private CA root certificate.

- **Verify Cloud Portal Certificate**

  Enables appliances to verify the Silver Peak Cloud Portal certificate. Disable this setting if any of the following statements are true:

  - Any appliance in your network is behind a proxy server.
  - Orchestrator is not configured with the Cloud Portal domain name.

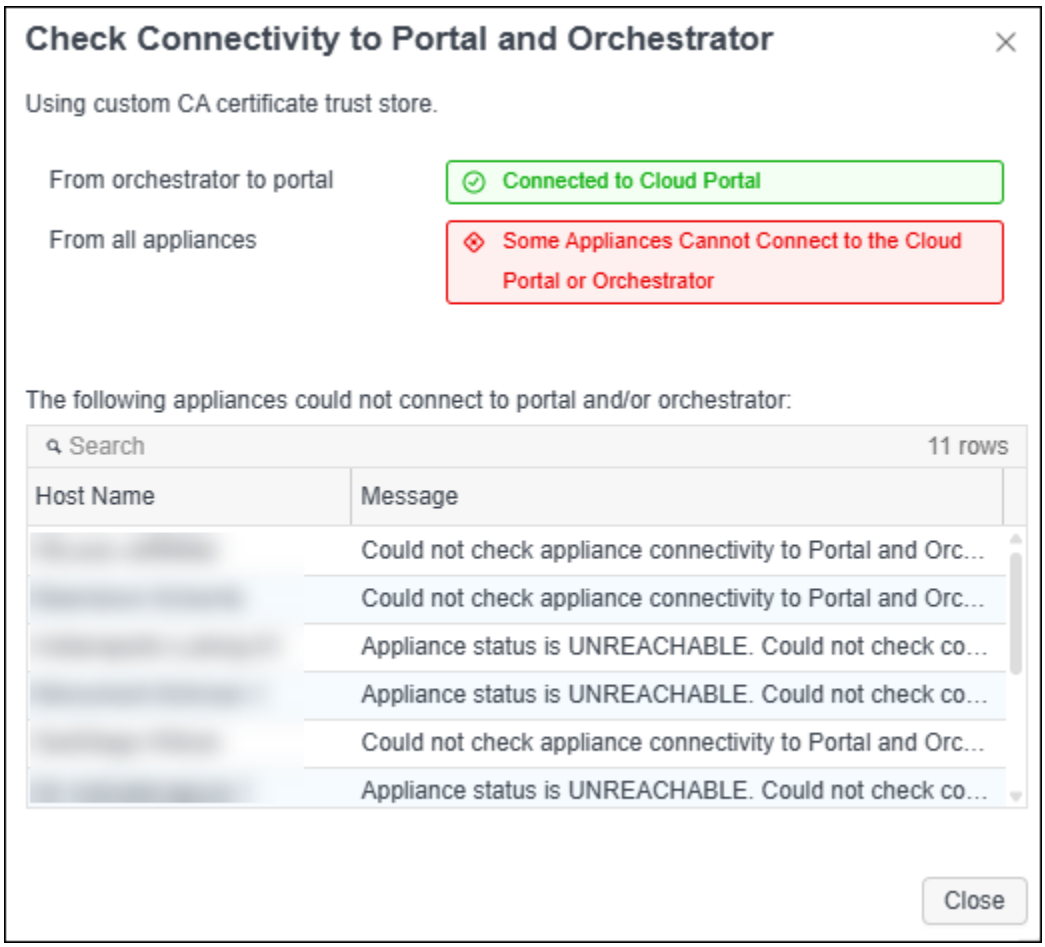
- Any appliance in your network is not configured with the Cloud Portal domain name.

To verify that your appliances can connect to the Cloud Portal and Orchestrator by using your current CA Certificate Trust Store, click **Check Connectivity Using Current Trust Store**.

NOTE: If you are not using a custom CA Certificate Trust Store, the verification process uses the default trust store. For details about custom trust stores, refer to Custom CA Certificate Trust Store.

IMPORTANT: If Common Criteria mode is enabled, this setting will be overridden and certificates will always be verified.

The Check Connectivity to Portal and Orchestrator dialog box opens.



The top portion indicates verification progress and results. The table explains any unsuccessful connections.

The Advanced Security Settings dialog box also displays the following security settings:

- **Enforce CSRF Token Check**

Enables Cross-Site Request Forgery (CSRF) token checking. Use this setting while using Orchestrator REST APIs and to avoid CSRF vulnerabilities. Before you enable this setting, be sure that the X-XSRF-TOKEN header in your script is set to the *orchCsrfToken* value returned by Orchestrator. This ensures that requests are legitimate and do not come from unauthorized sources, which helps prevent CSRF attacks and enhances security.

NOTE: It is highly recommended that you enable this setting. Any API scripts must be verified to ensure that X-XSRF-TOKEN is set appropriately.

- **Verify System Files Integrity**

  Enables verification of image signatures of binaries during the bootup process for appliances that are not FIPS certified. Appliances will verify the integrity of library and executable files. FIPS-certified appliances will ignore this setting and will always verify signatures of binaries at bootup. Be aware that enabling this setting increases bootup time by five or more minutes.

- **Appliance Shell Access Setting**

  Sets the level of Linux shell access to all appliances that Orchestrator manages. Orchestrator pushes this setting to all appliances.

  NOTE: When FIPS is enabled, shell access is disabled.

  - Open Shell Access – Allows users to fully use the appliance Linux shell. For new installations, this setting is no longer available.
  - Secure Shell Access – Restricts Linux shell use by requiring token access from Support. This is the default setting for new installations.
  - Disabled Shell Access – Completely locks down access to the appliance Linux shell.

    IMPORTANT: Once shell access is disabled, it cannot be reverted to secure shell access. You must redeploy to a new or remanufactured appliance.

Intrusion Detection/Prevention System

Configuration > Overlays & Security > Security > IDS/IPS

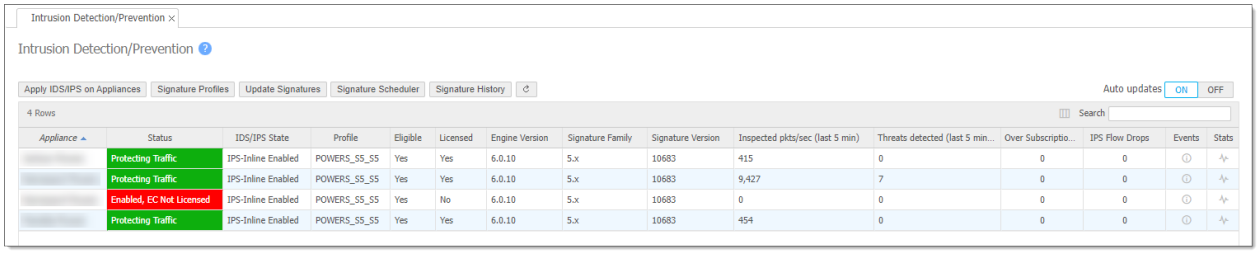
The Intrusion Detection/Prevention System (IDS/IPS) can monitor traffic for potential threats and malicious activity. It generates threat events based on preconfigured rules. Packets are copied and inspected against signatures downloaded to Orchestrator from the Cloud Portal. Orchestrator sends to appliances the signature file and any rules that have been added to the allow list.

- IDS designates traffic for inspection. To enable IDS:
  - For the zone-based firewall security policy, select **Inspect** as the action to apply to traffic that matches the rule. Matching traffic will be inspected by IDS/IPS. For details, see Appliance Security Policies.
  - On the Apply IDS/IPS dialog box, select the **IDS/IPS Mode** check box, and then select **IDS Only**. For details, see Apply IDS/IPS on Appliances below.
- IPS designates traffic for inspection as well as protection. To enable IPS:

- For the zone-based firewall security policy, select **Inspect** as the action to apply to traffic that matches the rule. Matching traffic will be inspected by IDS/IPS . For details, Appliance Security Policies.
- On the Apply IDS/IPS dialog box, select the **IDS/IPS Mode** check box, and then select **IPS-Performant** or **IPS-Inline**. IPS protects traffic by matching a signature, and then performing a configured action (Drop, Inspect, or Allow). For details, see Apply IDS/IPS on Appliances below.

Use the Intrusion Detection/Prevention tab to view IDS/IPS status, IPS/IDS state (IDS Enabled, IPS-Performant Enabled, IPS-Inline Enabled, and Disabled), signature profile information, appliance eligibility to run IDS/IPS, signature family and version, and IDS/IPS operational statistics. You can also modify the IDS/IPS configuration for appliances selected in the appliance tree.

The Auto updates ON and OFF buttons enable you to control whether signatures should be automatically updated. By setting this to OFF, you can make informed decisions before proceeding with signature updates. To make this evaluation, use the Signature History subtab to examine the differences between the signature rules in the latest active signature version on Cloud Portal and the production rules in your current signature version in Orchestrator.



Field	Description
Appliance	Name of the appliance.
Status	Status of IDS/IPS for the appliance, such as **Not Eligible**, **Protecting Traffic**, or **License applied but EC Not Eligible**.
IDS/IPS State	State of IDS/IPS on the appliance (**IDS Enabled**, **IPS-Performant Enabled**, **IPS-Inline Enabled**, or **Disabled**).
Profile	IDS/IPS signature profile applied to the appliance.
Eligible	Indicates whether the appliance is eligible to enable IDS/IPS. For more information, see Prerequisites for IDS/IPS below.
Licensed	Indicates whether the appliance is licensed to run IDS/IPS.
Engine Version	IDS/IPS engine version.
Signature Family	Proofpoint ETPro signature family, such as 4.x, 5.x, or 7.x.

Field	Description
Signature Version	IDS/IPS signature version, such as 10500 or 10729.
Inspected pkts/sec (last 5 min)	Number of packets inspected in the previous five minutes.
Threats detected (last 5 min)	Number of threats detected in the previous five minutes.
Over Subscription Drops	Number of dropped kernel flows (cumulative since IDS/IPS has been active) due to traffic oversubscription. This field displays the cumulative count of such dropped packets since IDS/IPS was enabled.
IPS Flow Drops	Number of dropped flows. These are drops that IPS does based on rule-based drops. This drop count is per minute.
Events	Click the info icon to view the most recent IDS/IPS events on the appliance. Click Export CSV to save the data to a CSV file.
Stats	Click the stats icon to view the following IDS/IPS statistics for the appliance: IPS Flow Drops per minute, Packets per second sent to the IDS/IPS, Threats Detected, Bits per second sent to the IDS/IPS, and Over Subscription Drops (Cumulative).

Prerequisites for IDS/IPS

- IDS/IPS can be enabled on the following ECOS releases and appliances.

  For IDS:

  - ECOS 9.1.x.x or later for all EdgeConnect appliances *except EC-XS (part numbers 200889 and 200900 only) and EC-US*
  - ECOS 9.1.x.x or later for EC-V deployments (minimum of 4 cores and 16 GB of RAM required)

  For IDS/IPS:

  - ECOS 9.2.x.x or later for all EdgeConnect appliances *except EC-XS (part numbers 200889 and 200900 only) and EC-US*
  - ECOS 9.2.x.x or later for EC-V deployments (minimum of 4 cores and 16 GB of RAM required)

- IDS/IPS can be enabled only on appliances running ECOS 9.1.0.0 or later. Appliances running an earlier version of ECOS will not be shown on the Intrusion Detection/Prevention tab.
- IDS/IPS is a licensed feature and can be enabled only on appliances that have been assigned the Advanced Security license. Refer to the help information for the Licenses tab (Configuration > Overlays & Security > Licensing > Licenses).

NOTE: IDS/IPS alarms are logged in standard syslog format. You can configure a logging facility for IDS/IPS and remote log receiver to send logs to a third party for additional review and analytics. See Advanced Reporting and Analytics below.

Apply IDS/IPS on Appliances

You can turn on or off IDS/IPS for appliances displayed in the table. You can also apply a different signature profile to the appliances. Orchestrator provides a Default signature profile.

1. Click **Apply IDS/IPS on Appliances**.

   The Apply IDS/IPS dialog box opens.
2. To apply or remove IDS/IPS, select the **IDS/IPS Mode** check box, and then select one of the following:

   - **Off** to turn off IDS and IPS for all appliances.
   - **IDS Only** to enable IDS on the appliances.
   - **IPS-Performant** to enable IPS-Performant on the appliances.
   - **IPS-Inline** to enable IPS-Inline on the appliances.

     NOTE: IPS-Inline mode can be applied only to appliances on ECOS version 9.4.1.0 or later. Appliances on earlier versions will ignore this mode.
3. To apply a different signature profile, select the **Profile** check box, and then select the appropriate signature family from the first drop-down list. The second drop-down list includes the Orchestrator-provided default profile that corresponds to the signature family you selected (*Default* for 4.x, *Default_S5* for 5.x, or *Default_S7* for 7.x) and any other signature profiles you have created. Select the signature profile you want to use. For more information about signature profiles, see Manage Signature Profiles below.
4. The Modification column displays the proposed changes, if any, for the appliances. To apply your changes, click **Save**. Or, to close the dialog box without making any changes, click **Cancel**.

Manage Signature Profiles

Signature profiles enable you to configure rules that are downloaded from the signature set on the Cloud Portal. Orchestrator provides the following default signature profiles:

- For the 4.x signature family, the *Default* signature profile
- For the 5.x signature family, the *Default_S5* signature profile

  NOTE: 5.x profiles can be applied only to appliances running IDS/IPS Engine Version 6.x.
- For the 7.x signature family, the *Default_S7* signature profile

  NOTE: 7.x profiles can be applied only to appliances running IDS/IPS Engine Version 7.x.

These default profiles include default settings for the signature rules. Default profiles are automatically used across all appliances. You can create additional signature profiles and override default rule settings by choosing different actions as needed. To open the Signature Profiles tab, click **Signature Profiles**. For information about creating signature profiles and modifying their rules, refer to the help information for the Signature Profiles tab.

Update Signatures on Appliances

You can immediately update signatures on your appliances or schedule the updates to occur when convenient for your organization—on a daily, weekly, monthly, or yearly basis.

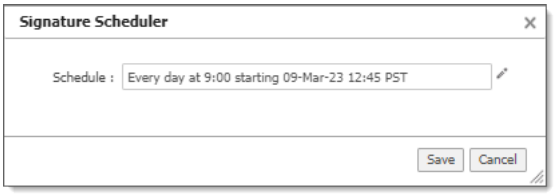
To immediately update your appliances with signature updates, ensure that **Auto updates** is set to ON, and then click **Update Signatures**. Orchestrator checks for any signature updates and pushes them to the appliances. This might take some time. You can check the audit log for status updates.

NOTE: If Auto updates is set to OFF when you click Update Signatures, the signatures will be downloaded but not activated on Orchestrator. This enables you to make informed decisions before proceeding with signature updates. Use the Signature History subtab to examine the differences between the signature rules in the latest active signature version on Cloud Portal and the production rules in your current signature version in Orchestrator. After evaluating, you can activate and push them to the appliances by setting **Auto updates** to ON and clicking **Update Signatures**.

To update based on a schedule:

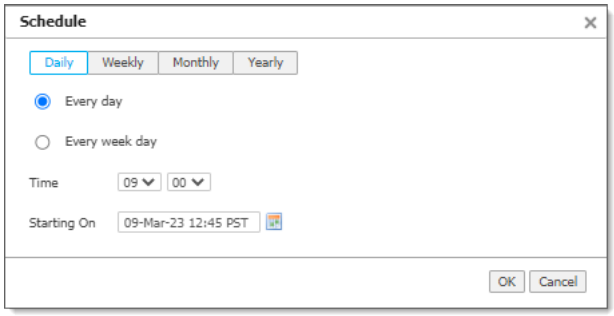
1. Click **Signature Scheduler**.

   The Signature Scheduler dialog box opens.



2. Click the edit icon.

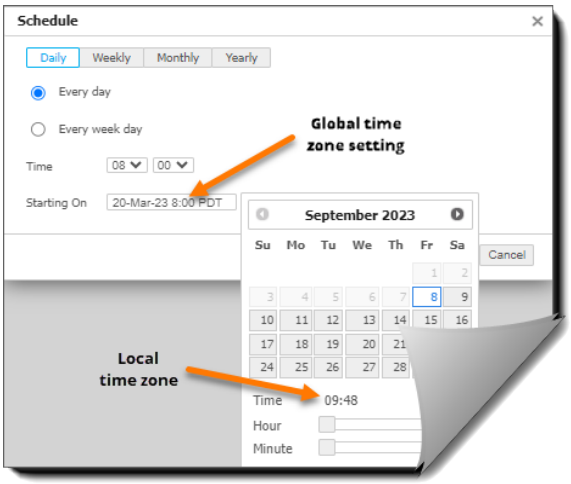
   The Schedule dialog box opens.



3. Click **Daily**, **Weekly**, **Monthly**, or **Yearly**.
4. Specify the appropriate schedule criteria, and then click **OK**.
5. On the Signature Scheduler dialog box, click **Save**.

   Orchestrator will automatically update your appliances according to your specified schedule.

NOTE: The time zone displayed in Signature Scheduler reflects the global time zone setting for scheduled jobs and reports (Orchestrator > Software & Setup > Setup > Timezone for Scheduled Jobs). The time zone displayed on the calendar when you click the calendar icon on the Schedule dialog box reflects your local time zone.



View Signature History

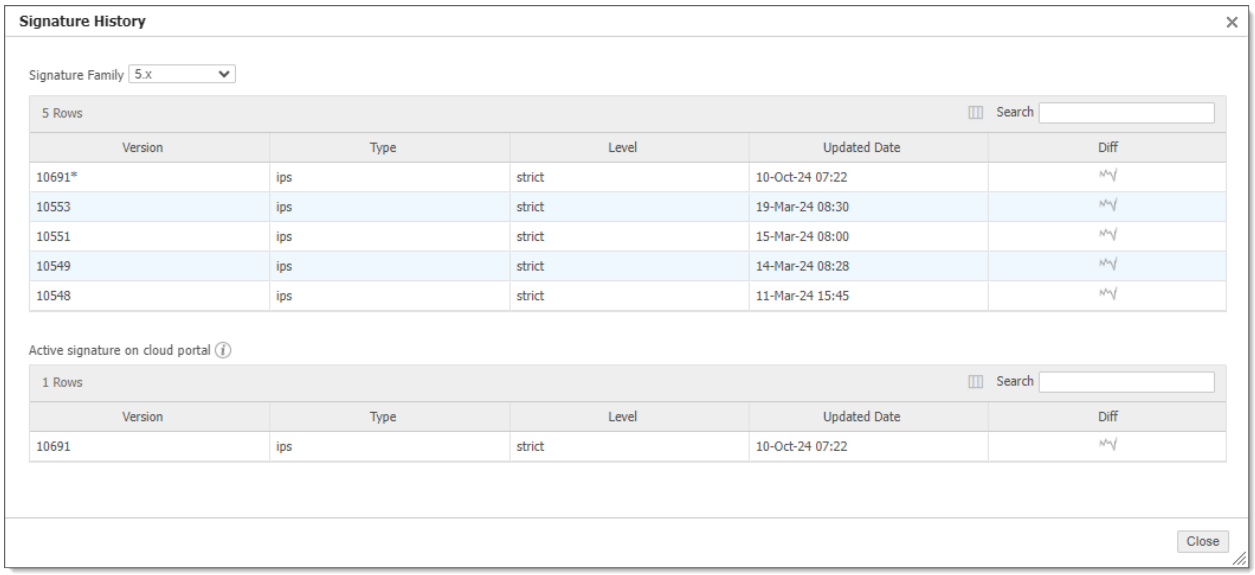
You can view a history of signature versions on your Orchestrator, including both the current and previous versions. Select the signature family from the drop-down menu at the top (such as 4.x, 5.x, or 7.x) to access the signature history. This provides the history for the past five signature versions for each active signature family on the network. You can also review the differences between:

- A signature version listed in the top table and the preceding listed version.
- The current version (listed in the top table) and the active signature version on cloud portal (listed in the bottom table). To view these differences, Auto updates must be set to OFF on the Intrusion Detection/Prevention tab.

Differences shown include modified, deleted, and added rules.

1. Click **Signature History**.

   The Signature History dialog box opens with a historical table of signature versions displayed based on signature family.



You can select a different signature family from the **Signature Family** drop-down list.

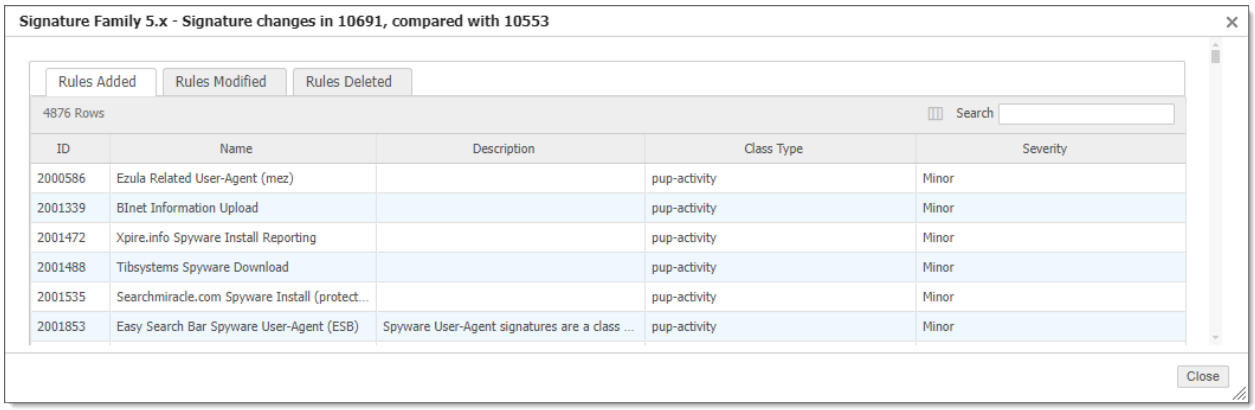
The Level column indicates the rules package applied to the IDS/IPS subsystem with varying levels of rules that control the strictness of inspection on the device. This setting does not indicate a restriction level on IPS usage. Rather, it specifies the intensity of IDS/IPS inspection (from lenient to strict) based on the selected rules package. Only the Strict rule package is currently supported. The Updated Date column indicates when the signature version was last updated.

The bottom table (Active signature on cloud portal) shows the active signature version on the Cloud Portal.

2. To view the differences, click the appropriate chart icon in the Diff column.

   The Signature Family dialog box opens.

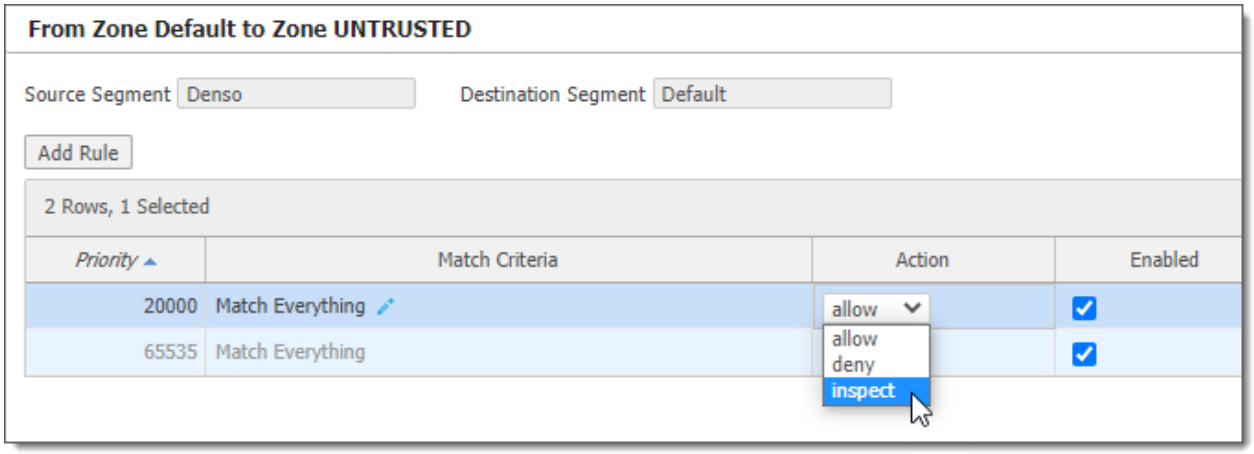
   NOTE: The title of this dialog box varies based on the signature family selected on the Signature History dialog box.



Use the tabs on this dialog box to view lists of new, modified, or deleted signatures in the signature version compared to the previous one.

Specify Traffic to Be Inspected

You can specify the traffic to be inspected according to source and destination zone, as well as specify detailed match criteria, using Appliance Security Policies (Configuration > Overlays & Security > Security > Appliance Security Policies).



With the addition of IDS/IPS, firewall actions have the following meanings:

- **allow:** Allow traffic and do not inspect.
- **deny:** Deny traffic and do not inspect.
- **inspect:** Allow traffic and inspect.

NOTE: No traffic will be inspected until rules with the inspect action are specified in the security policy.

For more information, see the following tabs in Orchestrator:

- **Templates (Security Policies)**: Configuration > Templates & Policies > Templates
- **Routing Segmentation**: Configuration > Networking > Routing > Routing Segmentation (VRF)

Advanced Reporting and Analytics

For users who are using or trying Splunk, you can install the HPE Aruba Networking EdgeConnect Security App for Splunk application to enable advanced reporting and analytics using the IDS/IPS alarms forwarded from EdgeConnect appliances. Search Splunkbase for “EdgeConnect” or click this link to search in your browser.



Follow the instructions provided to install and configure the application.

Signature Profiles

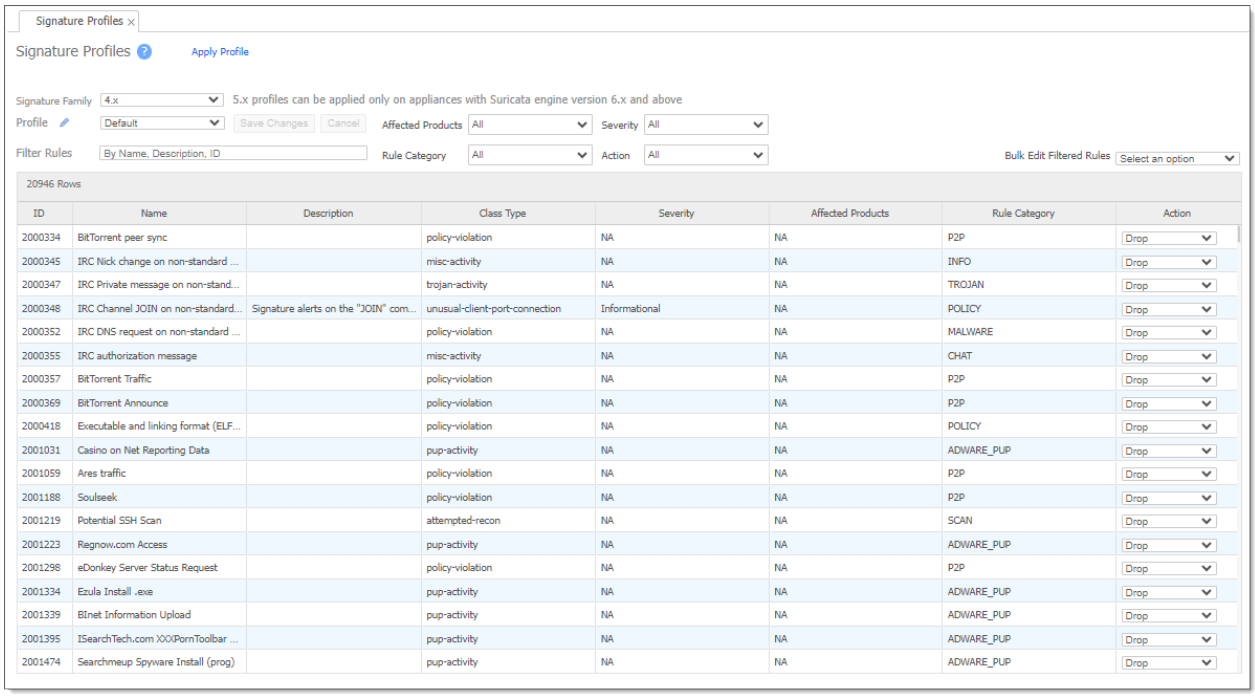
Signature profiles enable you to configure rules that are downloaded from the Cloud Portal. Orchestrator provides the following default signature profiles:

- For the 4.x signature family, the *Default* signature profile
- For the 5.x signature family, the *Default_S5* signature profile
- For the 7.x signature family, the *Default_S7* signature profile

These default profiles include default settings for the signature rules. Default profiles are automatically used across all appliances. You can create additional signature profiles and override default rule settings by choosing different actions as needed.

By default, all rules included in the signatures list are enabled on all appliances where IPS is enabled. The default action is to drop traffic when a rule is triggered. However, for certain traffic or in some other cases, you might want to specify different actions for IPS to take.

1. To open the Signature Profiles tab, click **Signature Profiles** on the Intrusion Detection/Prevention tab (Configuration > Overlays & Security > Security > IDS/IPS).



2. Select the appropriate signature family from the **Signature Family** drop-down list.

   NOTE: You can apply profiles for the 5.x signature family *only* to appliances running IDS/IPS Engine Version 6.x, and profiles for the 7.x signature family *only* to appliances running IDS/IPS Engine Version 7.x.

3. Initially, the Profile field indicates that rules for the default signature profile (*Default* for the 4.x signature family, *Default_S5* for the 5.x signature family, or *Default_S7* for the 7.x signature family) are displayed on this tab. To change the displayed signature profile, select the appropriate profile from the **Profile** drop-down list.

   To create signature profiles, see Create a Signature Profile below.

4. Use the **Filter Rules** field above the table to filter the list of rules. You can also use the filters to the right of the field to view rules by affected products, rule category, severity, and/or action.

5. To set the response for a specific rule, select one of the following actions from the drop-down list in the **Action** column. For multiple rules, select the appropriate rule rows in the table, and then select an action from the **Bulk Edit Filtered Rules** drop-down list.

   - **Drop:** Drop the traffic when a matching signature condition exists for the source, destination, or both.
   - **Inspect:** Continue the traffic flow to the destination after inspecting the traffic and raising an event for matching signature. This action detects the anomaly.
   - **Allow:** Excludes the rule from participating in IDS/IPS, rendering it no longer part of IDS/IPS processing.

NOTE: You can change signature rules for any custom signature profile, but you cannot change the Default signature profile. This ensures that you always have the original signature rules as provided by Orchestrator.

You can apply profiles to your appliances by clicking the **Apply Profile** link. For details, refer to the help information for the Intrusion Detection/Prevention tab.

Create a Signature Profile

When you create a signature profile, it will be selectable from the Profile drop-down list. Then you can change the rule actions for that profile as needed.

1. On the Signature Profiles tab, select the appropriate signature family from the **Signature Family** drop-down list
2. Click the edit icon associated with the Profile field.

   The Signature Profiles dialog box opens.
3. Click **+ Add**.

   The Add Signature dialog box opens.
4. Verify that the appropriate signature family is indicated.
5. In the **Profile Name** field, enter a signature profile name, and then click **Ok**.

   The new signature profile appears on the Signature Profiles dialog box.

   NOTE: If your newly created signature profile is based on signature family 5.x (or when previously existing signature profiles based on signature family 4.x are migrated during ECOS upgrade), Orchestrator appends the profile name you provided with *_S5. *For example, if the profile name is* BankCo*, Orchestrator changes it to* BankCo_S5*. Similarly, if your signature profile is based on signature family 7.x, Orchestrator appends the profile name with* _S7* (for example, BankCo_S7).
6. Click **Save**.

Roles

Configuration > Overlays and Security > Security > Roles

This dialog box allows you to define and map roles that are used throughout the EdgeConnect SD-WAN Fabric. For example, you can map a role to a Group Policy Identifier (GPID) from an HPE Aruba Networking CX Switch to facilitate identity (role) awareness between SD-WAN Orchestrator and HPE Aruba Networking CX Switches. After you map a role to a GPID, you can specify the role in match criteria when creating Access Lists and policies in SD-WAN Orchestrator.

You can also use this tab to import roles and GPIDs from a CSV file.

Prerequisites

- This feature is only available for Orchestrator version 9.4.1 and ECOS version 9.4.1.0 and later.
- GPID values and roles must match the GPID values configured on the HPE Aruba Networking CX Switch.

For more information on configuring VXLAN and a VTEP source loopback interface on the EdgeConnect appliance, see VXLAN Tab or VXLAN Template.

Add Roles

Add all the roles that are required for each VXLAN segment. If the VXLAN segment encounters a role that is not defined in Orchestrator, that role is labelled “unknown”.

To add roles:

1. Click **Add Role**. A new row is added to the table.
2. Click in the new cell in the Role column, and then enter a name for the role. Role names are case sensitive. You can create up to 254 roles. “0” is reserved for the default role and “65535” is reserved for the unknown default role.
3. Press **Tab** or click in the new cell in the GPID column, and then enter the GPID of the HPE Aruba Networking CX Switch you want to associate this role with.
4. Click **Save**.

You can filter flows and policies by role.

Import CSV

To import roles and GPIDs from a CSV file:

1. Click **Import CSV**.
2. Locate and select the CSV file to import, and then click **Open**.

   The Roles - Bulk Upload dialog box opens.

Color	Description
Green	Indicates a new GPID and role mapping. This role and GPID will be added to the table.
Yellow	Indicates a duplicate GPID. The existing role for the GPID will be replaced with the role in the file you are importing.
Red	Indicates a duplicate role. The duplicate role will not be added to the table even if the GPID does not match.

3. Review the roles and GPIDs to be imported.
4. Click **Save** to import the file or click **Cancel** to close the dialog box without making any changes.

Trust Store

Configuration > Overlays & Security > Security > Trust Store

Starting with Orchestrator 9.7.0, on the Trust Store tab you can manage a single, orchestrated trust store for Orchestrator and EdgeConnect functions that use Certificate Authority (CA) certificates, including IPSec tunnel peer authentications and TLS client connections. The trust store contains:

- Built-in CA certificates that Orchestrator receives from the Cloud Portal. These are well-known, globally trusted root CA certificates. The CA certificates required to establish trust with the Cloud Portal are included in the built-in certificates.
- Custom CA certificates that you add.

NOTE: Certificates that are used for WAN Optimization (previously called Boost) should not be added to the Trust Store. These should be added to either the SSL WAN Optimization Certificates tab or the SSL WAN Optimization CA Certificates tab (Configuration > Overlays & Security > SSL for WAN Optimization). For more information see, SSL WAN Optimization Certificates and SSL WAN Optimization CA Certificates.

Each certificate in the trust store can be deployed to or withdrawn from appliances. Only certificates that are deployed are synced to appliances. The trust store can automatically connect to the Cloud Portal daily to check for updates for built-in certificates. For more information, see Enable/Disable Auto Sync With Cloud Portal.

- You cannot edit or delete built-in certificates. To view information about a built-in certificate, hover over the Actions column and click the info (Details) icon (**i**).
- You can edit and delete custom certificates.
- To test connectivity between Orchestrator and the Cloud Portal and between Orchestrator and all appliances, click **Test Connectivity**.
- To add a custom CA certificate, click **Add Certificate**. For more information, see Add a Custom CA Certificate to the Trust Store.

The following table describes the information shown on the Trust Store tab for each certificate. Some columns are shown by default.

Column	Description
Alias	Unique identifier for the certificate. For custom certificates, the user enters the alias when the certificate is added to the trust store.
Issuer	Trusted CA that issued and signed the certificate.
Issued to	Entity that the certificate was issued to.
Expiration	Expiration date of the certificate. A red icon displays next to the expiration date for expired certificates before they are removed. A yellow icon displays next to the expiration date for certificates that are expiring in the next 60 days.

Column	Description
Source	Indicates how the certificate was added to the trust store. Built-in – The certificate was added to the trust store from the Cloud Portal. Custom – A user added the certificate to the trust store.
Sync to appliances	Status of a certificate that indicates whether the certificate has been deployed to or withdrawn from appliances. Deployed – The certificate is synced with all appliances. Deployment pending – The certificate is ready to sync with appliances but is waiting for Orchestrator to restart. Withdrawn – The certificate will not be synced with appliances. Withdrawal pending – The certificate is ready to be withdrawn from appliances but is waiting for Orchestrator to restart.
Date added	Date that the certificate was added to the trust store.
Certificate	Shows the certificate data in PEM-encoded format.
Certificate hash	A unique fixed-length string generated by applying a cryptographic hash function to the CA certificate data.
CA type	Role and trust level of the certificate (Root CA or Intermediary CA).
Key type	Cryptographic algorithm used to generate the public and private key pair for the certificate.
Key size	Length of the cryptographic key used for encryption and signing, measured in bits (for example, 2048 or 4096 bits).
Signature algorithm	Cryptographic hash function and algorithm used to authenticate the certificate.
Actions	Hover to view available actions for this row (Details, Edit, and Delete). Edit and Delete are not available for Built-in certificates.

Enable/Disable Auto Sync with Cloud Portal

The trust store can automatically connect to the Cloud Portal daily to check for updates for built-in certificates.

- Click **Auto sync on** to enable the trust store to automatically connect to the Cloud Portal. If Orchestrator receives updates to built-in certificates from the Cloud Portal, a minor alarm is raised, and you must restart Orchestrator to apply the built-in certificate updates.
- Click **Auto sync off** to disable the trust store from syncing with the Cloud Portal.

View Built-In Certificate Details

To view the details of a built-in certificate:

1. Hover over the Actions column for the certificate and click **Details**.

   The Certificate Details dialog box opens.
2. To expand the Certificate field and view the certificate data in PEM-encoded format, click **More**.
3. To copy the certificate, click **Copy**.

To view the details of a custom certificate, hover over the Actions column for the certificate and click **Edit** to open the Add/Edit Certificate dialog box.

Add a Custom CA Certificate to the Trust Store

NOTE: If you are managing CA certificates in an Orchestrator version earlier than 9.7.0, see Manage Certificates in Orchestrator Versions Earlier than 9.7.0.

To add a custom CA certificate to the trust store:

1. On the Trust Store tab, click **Add certificate**.

   The Add/Edit Certificate dialog box opens.
2. Enter an alias (unique identifier) for the certificate.
3. In the Certificate field, paste the certificate content.

   If the certificate is valid, the fields in the Certificate details section of the dialog box are populated with the appropriate information.
4. Click Done to close the Add/Edit Certificate dialog box.
5. On the Trust Store tab, click **Save**.

   The system prompts you to confirm the changes you are applying to the trust store.
6. Click **Confirm** to apply the changes.

   The certificate appears in the trust store with a status of *Deployment pending* in the Sync to Applications column.
7. To sync the certificate with all appliances, you must restart Orchestrator.

   After Orchestrator restarts, the status on the Trust Store tab in the Sync to appliances column should show *Deployed*.

Deploy/Withdraw a Certificate

Each certificate in the trust store can be deployed to or withdrawn from appliances. Only certificates that are deployed are synced to appliances. CA certificates that are added to the trust store are in *Deployment pending* state until you restart Orchestrator.

To deploy a certificate:

1. Find the certificate on the Trust Store tab and in the Sync to appliances column, move the toggle from **Withdrawn** to **Deployment pending**.

2. Restart Orchestrator.

   After Orchestrator restarts, on the Trust Store tab the status for the certificate shown in the Sync to appliances column changes to *Deployed*. The certificate is now synced to all appliances.

To withdraw a certificate:

1. Find the certificate on the Trust Store tab and in the Sync to appliances column, move the toggle from **Deployed** to **Withdrawal pending**.
2. Restart Orchestrator.

   After Orchestrator restarts, on the Trust Store tab the status for the certificate shown in the Sync to appliances column changes to *Withdrawn*. The certificate is now withdrawn from all appliances.

Manage Certificates in Orchestrator Versions Earlier than 9.7.0

Orchestrator versions earlier than 9.7.0 ship with the following:

- A default trust store with well-known, globally trusted root CA certificates. The default trust store contains Certificate Authority (CA) certificates required to establish trust with the Cloud Portal.
- A Custom CA Certificate Trust Store.

To manage certificates in Orchestrator versions earlier than 9.7.0:

- You must add the default root CA certificates from the default trust store to the Custom CA Certificate Trust Store (one-time action).
- Enterprises deploying self-hosted Orchestrators must install their end entity certificate (HTTPS server certificate) in the Orchestrator instance. If the end entity certificate is not in the collection of default root CA certificates from the default trust store, enterprises must add the root CA certificate associated with the issuer to the Custom CA Certificate Trust Store.
- If you want your Orchestrator and appliances to establish connectivity with any of the following services, you must add the certificates for these services to the Custom CA Certificate Trust Store:
  - Remote authentication servers, such as OAuth, JWT, or SAML
  - Remote log receivers
  - Netskope
  - Zscaler
  - Azure
  - HPE Aruba Networking ClearPass Policy Manager

NOTE: If you are using Orchestrator as a Service (OaaS), ensure that you have copied the root CA certificates from the default trust store to the Custom CA Certificate Trust Store, and verified communications from the appliances to both Orchestrator and Cloud Portal before you enable the Custom CA Certificate Trust Store.

In Orchestrator versions earlier than 9.7.0, follow these steps to add well-known, globally trusted certificates from the default trust store to the Custom CA Certificate Trust Store:

1. Click **Add Default Certificates**.

   The custom trust store is populated with multiple default certificates.

2. Click **Apply Changes**.

To enable the custom certificate trust store:

NOTE: You must either add the root CA certificates from the default trust store or upload at least one certificate before you can enable the custom trust store.

1. Click **Test Connectivity to Portal** to validate that appliances can successfully connect to Orchestrator and Cloud Portal using the custom CA.
2. Click the **Use Custom Certificate Trust Store** check box.
3. Click **Apply Changes**.

To add a CA certificate to the custom certificate trust store:

1. Click **Add Certificate to Custom Trust Store**.

   The Add/Edit Custom Certificates dialog box opens.

2. Enter a meaningful alias for the certificate in the **Alias** field. For example, for an Orchestrator web server certificate use “Orchestrator_HTTPS” or for a Syslog server use “Syslog_HTTPS”.
3. Paste the root certificate into the **Certificate** field.
4. Click **Save**.

NOTE: After adding root CA certificates to the Custom Trust Store, Orchestrator must be restarted.

End Entity Certificates Tab

Configuration > Overlays & Security > Security > End Entity Certificates

The EdgeConnect platform consists of Cloud Portal, Orchestrator, and EdgeConnect gateways running EdgeConnect OS (ECOS). Historically, EdgeConnect supported Certificates for a few limited applications such as verifying identity between the triad of Cloud Portal, Orchestrator, and associated EdgeConnect Gateways. Certificates could also be installed so the web server for the Orchestrator UI and ECOS UI can be trusted by browsers with their built-in trust store. Release 9.4 introduced the use of end entity certificates for IPSec tunnel peer authentication.

The following are use cases for end entity certificates in 9.4 and later:

- Orchestrator web server certificate
- EdgeConnect web server certificate
- Syslog client certificate for EdgeConnect

- IKE-based IPSec tunnels

Whenever an EdgeConnect appliance or Orchestrator is required to authenticate itself with a certificate, it must generate the private key associated with the certificate. EdgeConnect appliances and Orchestrator can use end entity certificates to do this during the creation of TLS and IP Sec connections between SD-WAN components, as well as with third party services. While the legacy feature of uploading private keys is still supported, it is best practice to use orchestrated appliance end entity profiles for EdgeConnect appliance certificates and to manually generate CSRs for Orchestrator certificates.

On the End Entity Certificates tab, you can view and manage all end entity certificates created for your EdgeConnect appliances and Orchestrator. From this tab you can do the following:

- Configure an EST (Enrollment over Secure Transport) server profile to enroll certificates for use with your EdgeConnect appliances.
- Create orchestrated appliance end entity profiles that allow for automated enrollment of certificates using an EST server.
- Manually create a certificate signing request (CSR) and add an end entity certificate for an appliance or Orchestrator, and manually create a labeled profile that enables EdgeConnect appliances to generate CSRs.
- Click **Appliance** to view end entity certificates for your appliances.
- Click **Orchestrator** to view end entity certificates for Orchestrator.

The information in the following table is displayed for each end entity certificate on this tab.

Column	Description
Hostname	The hostname of the appliance or the Orchestrator instance.
Label	The name assigned to the certificate.
Issuer	The issuer of the certificate.
Issued to	The entity to which the certificate is issued (common name on CSR).
Certificate	After the certificate is successfully enrolled, a link appears in this column that allows you to view and download the certificate.
Expiration date	The date when the certificate expires.

Adding End Entity Certificates

There are two methods for adding end entity certificates. The first method is fully orchestrated and is used for certificate applications on EdgeConnect appliances. The second method involves completing a CSR for one certificate at a time and can be used for Orchestrator or EdgeConnect appliances.

Method 1: This method automates certificate enrollment using an EST server and globally orchestrated end entity profiles. You can only use this method for EdgeConnect appliances to create certificate-based, orchestrated tunnels, which are used by Business Intent Overlays, ECOS web server certificate, or ECOS syslog client certificate. The workflow for this method is as follows:

- Prepare the Custom CA Certificate Trust Store
- Add an EST Server Profile
- Add an Appliance End Entity Profile
- Use an End Entity Certificate or Profile for a Service
  - To apply an orchestrated appliance end entity profile to IKE-based IPSec tunnels, which can be used in Business Intent Overlays, see Tunnel Settings Tab.
  - To use an orchestrated appliance end entity profile for the EdgeConnect HTTPS Certificate, see HTTPS Certificate Template.
  - To use an orchestrated appliance end entity profile for the syslog client certificate, see Logging Template.

Method 2: You can use this method to enroll certificates for use with both Orchestrator and EdgeConnect appliances. It is a manual method, and you must repeat the process for each EdgeConnect appliance.

- Prepare the Trust Store
- Manually Obtain a Signed End Entity Certificate
  - Create a Certificate Signing Request (CSR)
  - Send the CSR to the Certificate Authority (CA)
  - Obtain the Signed Certificate from the CA
  - Upload the Signed Certificate to the End Entity Certificate Tab
- Use an End Entity Certificate or Profile for a Service

End Entity Certificate Validation at the Time of Upload

At the time of manually uploading an end entity certificate, the following validation checks occur. If you are using globally orchestrated end entity profiles, these validation checks occur as part of the automated process of enrolling certificates. If any of these validations fail, upload of the end entity certificate fails.

- Revocation Status: Online Certificate Status Protocol (OCSP) is run by Orchestrator and EdgeConnect to verify that the intermediate CA certificates and end entity certificate are not revoked using the OCSP URLs present in each of these certificates. OCSP exception checking includes the following:
  - If communication cannot be established with the OCSP server, then the revocation check is ignored.
  - Nonce check:

    * If the OCSP server does not return a nonce, then the nonce test is ignored, and the revocation check continues.
    * If the OCSP server returns a nonce that does not match the nonce in the OCSP request, then the revocation check fails, and the end entity certificate is rejected.
    * If the OCSP server returns a nonce that matches the nonce in the OCSP request, then the revocation check continues.
  - If the OCSP response does not contain a status for the certificate that was requested, then the revocation check fails, and the end entity certificate is rejected.
  - If the OCSP response is not signed by the CA that issued the certificate or signed by an OCSP responder delegated by the CA (the delegated responder should have a valid certificate signed by the CA containing the OCSP signing purpose), then the revocation check fails, and the end entity certificate is rejected.
  - This update: If the OCSP response for "this update" is in the future, then the revocation check fails, and the end entity certificate is rejected.
  - Next update: If the OCSP response for "next update" is in the past, then the revocation check fails, and the end entity certificate is rejected.

  After completion of all the above checks, and if the OCSP response was valid, revocation status itself is determined and is assigned one of these values: good, revoked, or unknown. The certificate is accepted for "good" or "unknown" statuses. If the revocation status is "revoked" the certificate is rejected.
- Expiry Status: Each certificate in the chain is verified as not expired.
- Issuer Sequence Check: The signed end entity certificate that is being uploaded must contain the entire certificate chain and it must be contained in a single file. The system verifies that the end entity certificate chain comes first, followed by the intermediate CA certificate, and finally the root CA certificate.
- Digital Signature Validation: The digital signature for each certificate in the chain is validated.
- Check for CA Certificates in Custom CA Certificate Trust Store:
  - If the signed end entity certificate is for use with Orchestrator, Orchestrator checks that the root CA certificate in the end entity certificate chain is present in the Orchestrator Custom CA Certificate Trust Store.
  - For manually uploaded certificates, if the signed end entity certificate is for use with an EdgeConnect appliance, that specific EdgeConnect checks that the intermediate CA and root CA certificates in the end entity certificate chain are present in the EdgeConnect Custom CA Certificate Trust Store. For orchestrated certificates, the EdgeConnect only checks that the root CA certificate is present in the EdgeConnect Custom CA Certificate Trust Store.
- Common Name Comparison: The common name in the CSR is compared to and must match the common name in the signed certificate.
- Subject Alternative Name Comparison: Checks that the subject alternative names in the CSR are present in the signed certificate.

- Key Correspondence: At the time that the CSR is generated it contains only the public key; the private key is stored only on this specific EdgeConnect appliance. When the signed certificate is uploaded, it contains the public key from the CSR. The EdgeConnect appliance validates that the public key in the certificate mathematically corresponds to the private key it has stored (the EdgeConnect does not store the public key).
- Starting with the release that includes end entity certificate orchestration, if you select **TLS Server** or **TLS Client** in the Purpose field when creating an appliance end entity profile, an additional check is performed. The check validates that the Extended Key Usage field on the enrolled certificate contains the text *TLS WWW [Server|Client] Authentication*.

NOTE: If manual upload of an end entity certificate fails, navigate to **Orchestrator > Orchestrator Server > Tools > Audit Logs** and enter "end entity" in the Search field. In the search results, look for entries with "end entity upload action" in the Action field and find the recent failed upload, which will show "Failed" in the Success column. Hover over the Results column for additional information.

Certificate Expiry Checking

Once per day the expiry date is checked on all certificates (CA, intermediate, and end entity). If the expiration date is within 60 days an alarm occurs; this alarm persists if the user does not clear it. If the expiration date has passed, an alarm occurs indicating that the certificate is expired. Certificate expiry happens whether Common Criteria mode is enabled or not. For certificates that are enrolled using an EST server and an appliance end entity profile, during expiry checking if the Re-enrollment threshold percentage is met, the server automatically attempts to re-enroll the certificate.

Prepare the Trust Store

Before adding an EST server, creating orchestrated appliance end entity profiles, or manually creating a CSR, upload the following certificates to the trust store depending on whether you plan to use the end entity certificate for Orchestrator or EdgeConnect.

- If you are running an Orchestrator version earlier than 9.7.0, you must enable the trust store and add the default root CA certificates from the default trust store to the Custom CA Certificate Trust Store (one-time action). This is primarily required to ensure that the root CA certificate for Cloud Portal is in the trust store.
- Upload the root CA certificate for the CA that will sign the CSR for the Orchestrator HTTPS server certificate.
- Upload the root CA certificate for the CA that signed the end entity certificate for the syslog server (this is the syslog server to which Orchestrator sends its logs).
- Upload the root CA certificate for the CA that signed the end entity certificate for the EST server.
- If you are using the orchestrated EST-based method, upload the root CA certificate for any CA that will be signing all end entity certificates.

NOTE: The certificates for the CA must be in place before you create an appliance end entity profile otherwise validation of the certificate orchestration will fail.

- If you are using the manual method, upload both the root CA and intermediate CA certificates for the CA that will sign the CSR.

  NOTE: The certificates for the CA must be in place when you upload the signed certificate chain otherwise validation of the certificate will fail.

To upload the necessary certificates, navigate to **Configuration > Overlays & Security > Trust Store**. For instructions on how to enable the trust store and upload a certificate, see Trust Store.

EST Servers

You can configure profiles for EST servers that are used to enroll certificates for use with EdgeConnect appliances. Using an EST server with appliance end entity profiles provides an automated process for using globally orchestrated certificates for authentication. From this dialog box you can view the EST server profiles that are configured, delete EST server profiles, and add EST server profiles.

NOTE: An EST server must be reachable out-of-band without dependence on IPSec tunnels. EST management plane is configured in the Management Services template (**Configuration > Templates & Policies > Templates**) using the management service named “Other VRF mgmt Apps”.

Add an EST Server Profile

The following instructions describe how to configure an EST server profile to use for certificate enrollment.

1. Navigate to **Configuration > Overlays & Security > Security > End Entity Certificates**.
2. Click **EST Servers**.

   The EST Servers dialog box opens.
3. Click **Add EST Server**.

   The EST Server dialog box opens.
4. Complete the following fields.

Field	Description
Profile name	Enter the name of the EST server.
Host name	Enter the host name for the EST server.
Host	Enter the port to use for communicating with the EST server.
Username	Enter a username for the EST server. This is used during HTTPS basic or digest authentication.

Field	Description
Password	Enter the password for the EST server. This is used during HTTPS basic or digest authentication.
Challenge password	Enter a challenge password for the EST server. This is a field added in the CSR that is sent to the EST server. If this field is left empty, the tls-unique is used as a challenge password in the CSR.
Arbitrary label	When the system attempts to enroll a certificate, this label is added to the URL of the EST server. This label is used for csrattrs and cacerts, per RFC 7030. Requirements are EST server dependent. The default is blank. Example: https://est-service999.com/.well-known/est/rsa2048/cacerts Where /rsa2048 is the arbitrary label.
Arbitrary enrollment	When the system attempts to enroll a certificate, this label is added to the URL of the EST server. This label is used for simpleenroll, per RFC 7030. Requirements are EST server dependent. The default is blank. Example: https://est-service999.com/.well-known/est/rsa2048/simpleenroll Where /rsa2048 is the arbitrary label.
Arbitrary re-enrollment	When the system attempts to enroll a certificate, this label is added to the URL of the EST server. This label is used for simplereenroll, per RFC 7030. Requirements are EST server dependent. The default is blank. Example: https://est-service999.com/.well-known/est/rsa2048/simplereenroll Where /rsa2048 is the arbitrary label.
Retry interval	Enter a numeric value (in seconds). During initial enrollment or re-enrollment of a certificate, this is the amount of time the system waits before it contacts the EST server again. It keeps attempting at this interval until the enrollment or re-enrollment is successful.

5. Click **Save**.
6. Click **Save and Apply Changes**.
7. Click **Close**.

Appliance End Entity Profiles

An appliance end entity profile can be applied globally to multiple appliances. The profiles allow for scaled management of certificates using an EST server.

When you configure an appliance end entity profile, you specify the name of an EST server profile that you have already configured in Orchestrator. The EST server specified in the profile is used to enroll and reenroll certificates via EST protocol for any appliance that uses the profile. Each profile also has a designated purpose that you specify during configuration. That purpose determines the services for which the profile can be used, such as for IKE-based IPSec tunnel authentication or HTTPs authentication.

Add an Appliance End Entity Profile

1. Navigate to **Configuration > Overlays & Security > Security > End Entity Certificates**.
2. Click **Appliance End Entity Profiles**.

   The Appliance End Entity Profiles dialog box opens.
3. Click **Add Profile**.

   The Appliance End Entity Profile dialog box opens.
4. Complete the following fields.

Field	Description
Purpose	Select one of the following options: **SD-WAN** – When selected, certificates that are enrolled using this profile are only used for IPSec tunnel authentication. Profiles with this purpose appear on the Tunnel Settings tab for IKE-based IPsec tunnels and can be selected as the source for the end entity certificate for authentication. These certificate-based, orchestrated tunnels can be used in Business Intent Overlays. **TLS Server** – When selected, certificates that are enrolled using this profile are only used for HTTPS server authentication, for example EdgeConnect web UI. Profiles with this purpose appear on the HTTPS Certificate template and can be selected as the source for the end entity certificate for authentication. **TLS Client** – When selected, certificates that are enrolled using this profile are only used for Syslog client authentication (EdgeConnect as client). Profiles with this purpose appear on the Logging template and can be selected as the source for the end entity certificate for authentication in the remote log receiver configuration. **General** – When selected, certificates that are enrolled using this profile are used for IPSec tunnel, HTTPS server, or Syslog client authentication. Profiles with this purpose appear on the Logging template, HTTPS Certificate template, and the Tunnel Settings tab and can be selected as the source for the end entity certificate for authentication.
Signing Algorithm	Select the algorithm that will be used to authenticate the certificate. The default is rsa_2048_with_sha256
Re-enrollment threshold percentage	During expiry checking, the server automatically attempts to re-enroll a certificate after this percentage of the validity period for the certificate occurs. For example, if a certificate is valid for 12 months and the re-enrollment threshold percentage is set to 75, the server attempts to re-enroll the certificate after 9 months (75% of 12 months) have passed. Enter a percentage value. Minimum value is 50, maximum value is 95.

Field	Description
EST Server	Select an EST server profile. The EST server profile must already be configured. For more information on EST server profile configuration, see Add an EST Server Profile. You can select more than one EST server up to a maximum of four, by clicking **+EST**. If you have more than one EST server selected, the system attempts certificate enrollment using the EST servers in the order they are listed.
Certificate Information	These fields vary based on the Purpose you select. **Common Name** – Applies to profiles that have SD-WAN or General selected for Purpose. This information is auto generated using the appliance host name, such as “mysystem-ecva”.When a certificate is enrolled using this profile, this value is used as the Common Name. Example of how this appears on a certificate: CN=mysystem-ecva **User FQDN** – Applies to profiles that have SD-WAN or General selected for Purpose. The prefix is auto generated and consists of the appliance host name, such as “mysystem-ecva”. Enter an FQDN in the field, such “hpe.com”. When a certificate is enrolled using this profile, the hostname and the user-entered FQDN are included in the SAN (Subject Alternative Name) field for email. Example of how this appears on a certificate in the SAN field: email: mysystem-ecva@hpe.com **Host Name** - Applies to profiles that have TLS Server or TLS Client selected for Purpose. The prefix is auto generated and consists of the appliance host name, such as “mysystem-ecva”. Enter a domain name in the field, such as “hpe.com”. When a certificate is enrolled using this profile, the host name and the user-entered domain name are included in the SAN field for DNS. Example of how this appears on a certificate in the SAN field: DNS: mysystem-ecva@hpe.com **SAN - Domain Name** – Applies to profiles that have TLS Server or TLS Client selected for Purpose. This information is auto generated based on the domain name you entered in the Host Name field.

Field	Description
Additional Information (Optional)	These fields only appear if you have selected General for the Purpose. **Domain Component** – The prefix is auto generated and consists of the appliance host name, such as “mysystem-ecva”. Enter a system domain name in the field, such as “arubanetworks.com”. When a certificate is enrolled using this profile, this information is included in the SN (Subject Name) field, and it contains an entry for each dot separated part of the system domain name you entered. Example of how this appears on a certificate in the SN field: DC=com, DC=arubanetworks, DC=mysystem-ecva **Host Name** – The prefix is auto generated and consists of the appliance host name, such as “mysystem-ecva”. Enter a domain name in the field, such as “hpe.com”. When a certificate is enrolled using this profile, the host name and the user-entered domain name are included in the SAN field for DNS. Example of how this appears on a certificate in the SAN field: DNS: mysystem-ecva@hpe.com **IP Address** – Select a label from the drop-down menu. When a certificate is enrolled using this profile, the IP address associated with the label selected is included in the SAN field as an IP address. Example of how this appears on a certificate in the SAN field: IP Address=X.X.X.X **Organization Name** – Enter the name of the organization requesting the certificate, such as “Hewlett Packard Enterprise”. When a certificate is enrolled using this profile, this value is included in the SN. Example of how this appears on a certificate in the SN field: organizationName=Hewlett Packard Enterprise **Organizational Unit Name** – Enter the name of an internal department that handles the certificate within the organization, such as “HPE Aruba Networking”. When a certificate is enrolled using this profile, this value is included in the SN field. Example of how this appears on a certificate in the SN field: organizationalUnitName= HPE Aruba Networking **Add Serial Number to Subject** – When this check box is selected and a certificate is enrolled using this profile, the serial number for the appliance, as shown on the System Information tab, is included in the SN (Subject Name) field for serial number. Example of how this appears on a certificate in the SN field: serialNumber=XXXXXXXXXXX

5. Click **Save**.
6. Click **Close**.

After clicking Save, profile orchestration and certificate enrollment begins. During this time Orchestrator sends the profile to all EdgeConnect appliances, and then each appliance contacts the EST server to get a certificate and have it validated. When validation is successfully completed for a certificate, you can open the End Entity Certificates dialog box for an appliance and a View Certificate link appears in the CSR / Certificate column for the certificate.

Manually Obtain a Signed End Entity Certificate

The following instructions outline the process to manually add an end entity certificate. After you have prepared the Custom CA Certificate Trust Store, you create a CSR, obtain a signed end entity certificate, and manually upload the certificate for further use.

Create a Certificate Signing Request (CSR)

1. Navigate to **Configuration > Overlays & Security > Security > End Entity Certificates**.
2. Click **Appliance** to create a CSR for an appliance (EdgeConnect) or click **Orchestrator** to create a CSR for Orchestrator.
3. Click the edit icon next to the appliance or Orchestrator instance for which you want to create the end entity certificate and generate the CSR.

   The End Entity Certificates dialog box appears.
4. Click **Add Certificate**.

   The Add End Entity Metadata dialog box appears. Enter information in the following fields.

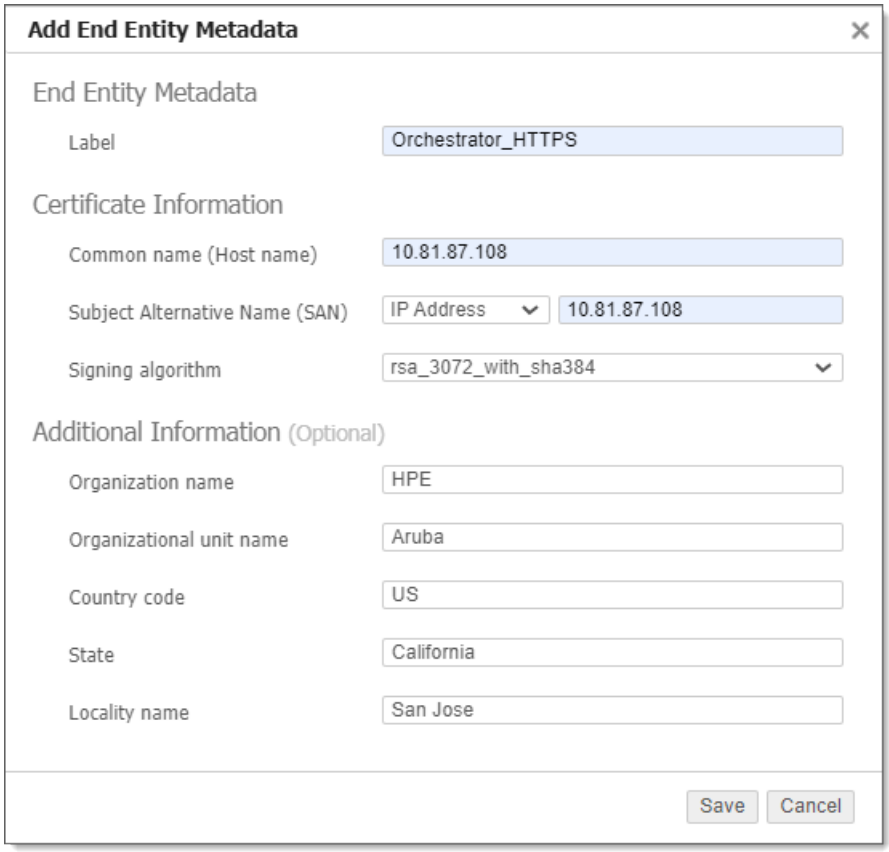
Field	Description
Label	Any string (for example, Orchestrator_HTTPS). The Label has significance to Orchestrator and EdgeConnect and does not need to be globally unique.
Common name (Host name)	Any string (for example, host name or IP address). The common name along with any information that is entered in the Additional Information (Optional) fields are what make up the Subject Name (SN) on the certificate.

Field	Description
Subject Alternative Name (SAN)	Select one of the following options from the drop-down menu and enter the required information. **FQDN** – Enter the fully qualified domain name that the certificate secures (for example, ecva.silverpeaksystems.net). **USER_FQDN** – Enter an ID in the format of an email address (local-part, @ symbol, domain suffix). Example: username@aruba.com This ID is used to identify the EdgeConnect appliance when establishing an IPSec tunnel; it does not send or receive emails. **IP Address** – Enter an IPV4 or IPV6 address that is secured by the certificate.
Signing algorithm	Select the algorithm that will be used to authenticate the certificate. If you are using Common Criteria mode, the recommended option is **ecdsa_secp384r1_with_sha384**. If you are not using Common Criteria mode, the recommended option is **ecdsa_secp256r1_with_sha256**.
Additional Information (Optional)	Entering the following information is optional. If entered, this information along with the common name are what make up the Subject Name (SN) on the certificate. **Organization name** – Enter the name of the organization requesting the certificate. Example: “HPE” **Organizational unit name** – Enter the name of an internal department that handles the certificate within the organization. Example: “Aruba” **Country code** – Enter the two-digit country code of the country where the organization is located. Example: “US” for USA. See allowable list following this table. **State** – Enter the state, province, region, or county where the organization is located. **Locality name** – Enter the village, town, or city where the organization is located.

Allowable Country Codes

'AF', 'AL', 'DZ', 'AS', 'AD', 'AO', 'AI', 'AQ', 'AG', 'AR', 'AM', 'AW', 'AU', 'AT', 'AZ', 'BS', 'BH', 'BD', 'BB', 'BY', 'BE', 'BZ', 'BJ', 'BM', 'BT', 'BO', 'BA', 'BW', 'BV', 'BR', 'IO', 'BN', 'BG', 'BF', 'BI', 'KH', 'CM', 'CA', 'CV', 'KY', 'CF', 'TD', 'CL', 'CN', 'CX', 'CC', 'CO', 'KM', 'CG', 'CD', 'CK', 'CR', 'CI', 'HR', 'CU', 'CY', 'CZ', 'DK', 'DJ', 'DM', 'DO', 'EC', 'EG', 'SV', 'GQ', 'ER', 'EE', 'ET', 'FK', 'FO', 'FJ', 'FI', 'FR', 'GF', 'PF', 'TF', 'GA', 'GM', 'GE', 'DE', 'GH', 'GI', 'GR', 'GL', 'GD', 'GP', 'GU', 'GT', 'GN', 'GW', 'GY', 'HT', 'HM', 'VA', 'HN', 'HK', 'HU', 'IS', 'IN', 'ID', 'IR', 'IQ', 'IE', 'IL', 'IT', 'JM', 'JP', 'JO', 'KZ', 'KE', 'KI', 'KP', 'KR', 'KW', 'KG', 'LA', 'LV', 'LB', 'LS', 'LR', 'LY', 'LI', 'LT', 'LU', 'MO', 'MK', 'MG', 'MW', 'MY', 'MV', 'ML', 'MT', 'MH', 'MQ', 'MR', 'MU', 'YT', 'MX', 'FM', 'MD', 'MC', 'MN', 'MS', 'MA', 'MZ', 'MM', 'NA', 'NR', 'NP', 'NL', 'AN', 'NC', 'NZ', 'NI', 'NE', 'NG', 'NU', 'NF', 'MP', 'NO', 'OM', 'PK', 'PW', 'PS', 'PA', 'PG', 'PY', 'PE', 'PH', 'PN', 'PL', 'PT', 'PR', 'QA', 'RE', 'RO', 'RU', 'RW', 'SH', 'KN', 'LC', 'PM', 'VC', 'WS', 'SM', 'ST', 'SA', 'SN', 'SC', 'SL', 'SG', 'SK', 'SI', 'SB', 'SO', 'ZA', 'GS', 'ES', 'LK', 'SD', 'SR', 'SJ', 'SZ', 'SE', 'CH', 'SY', 'TW', 'TJ', 'TZ', 'TH', 'TL', 'TG', 'TK', 'TO', 'TT', 'TN', 'TR', 'TM', 'TC', 'TV', 'UG', 'UA', 'AE', 'GB', 'US', 'UM', 'UY', 'UZ', 'VU', 'VE', 'VN', 'VG', 'VI', 'WF', 'EH', 'YE', 'ZM', 'ZW'

The following is an example completed CSR with label "Orchestrator_HTTPS".



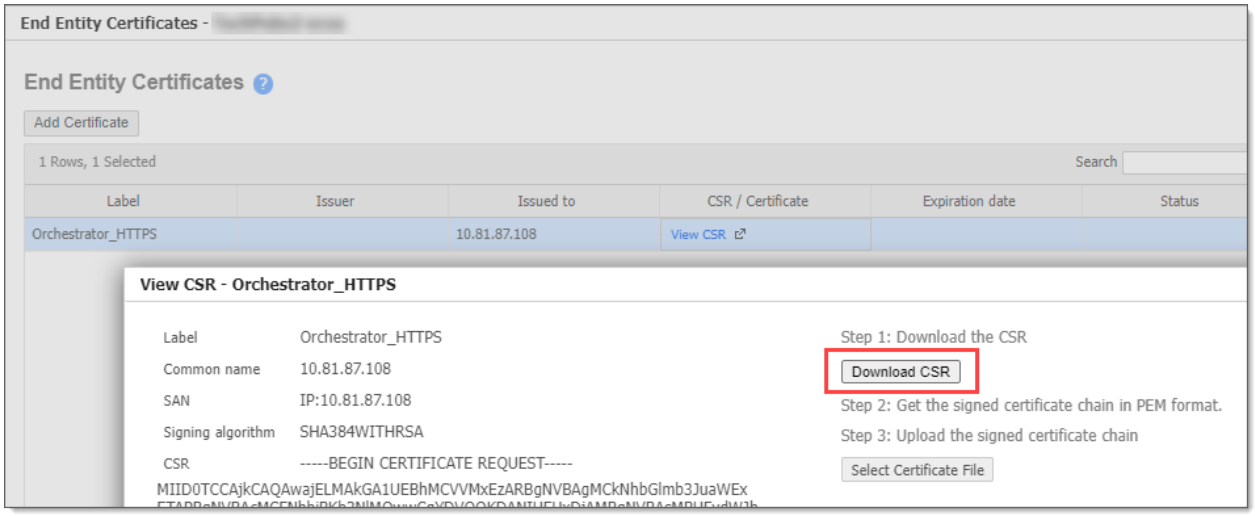
5. Click **Save**.

   The certificate appears in the list of certificates on the End Entity dialog box.

Send the CSR to the Certificate Authority (CA)

1. In the End Entity Certificates Dialog box find the certificate you created, and in the CSR / Certificate column click **View CSR**.

The View CSR dialog box opens.



2. To download the CSR as a .pem file, click **Download CSR**.

   The file is saved to your local directory.

3. Upload the CSR file to a trusted certificate authority (CA).

Obtain the Signed Certificate From the CA

When you receive the signed certificate from the CA, if there are multiple files you need to combine all the files into a single file, which includes the end entity certificate, all intermediate CA certificates, and the root CA certificates. This is necessary because you must upload the entire certificate chain in Orchestrator as a single file. The sequence of certificates in the single-file chain is important and should be as follows:

1. End entity certificate (top of file)
2. One or more certificates of the intermediate CA(s)
3. Self-signed root CA certificate

Upload the Signed Certificate to the End Entity Certificate Tab

NOTE: Any certificates in the chain that are expired will not be accepted by the EdgeConnect appliance.

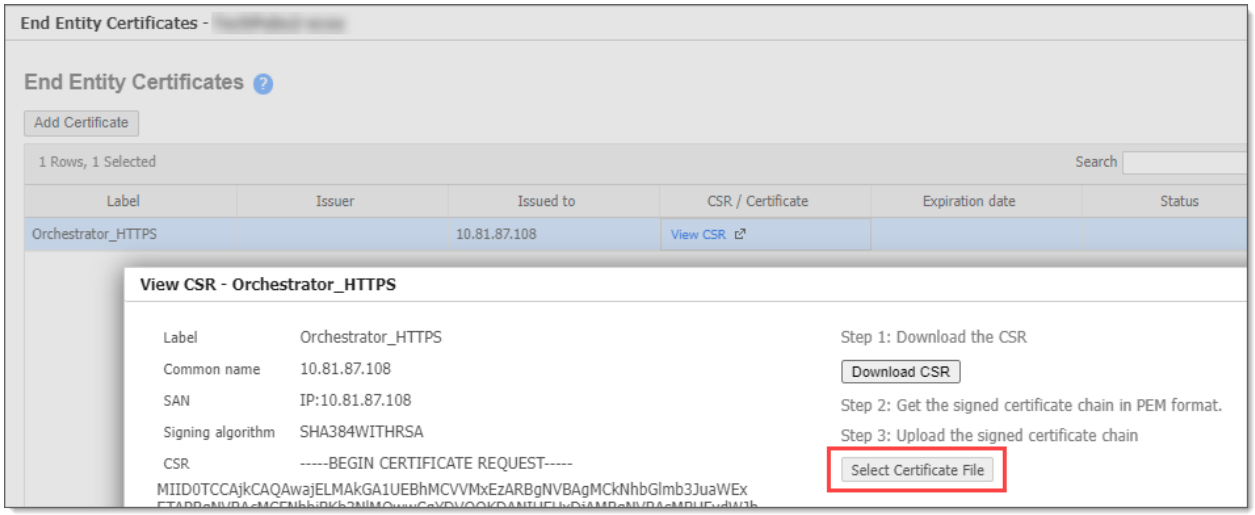
NOTE: The subject name on the certificate must match the subject name on the CSR. If these do not match, the EdgeConnect appliance will not accept the certificate.

1. Navigate to **Configuration > Overlays & Security > Security > End Entity Certificates**.
2. Click the edit icon next to the appliance or Orchestrator instance for which you created the CSR.

   The End Entity Certificates dialog box appears.

3. In the End Entity Certificates Dialog box find the certificate you created, and in the CSR / Certificate column click **View CSR**.

   The View CSR dialog box opens.

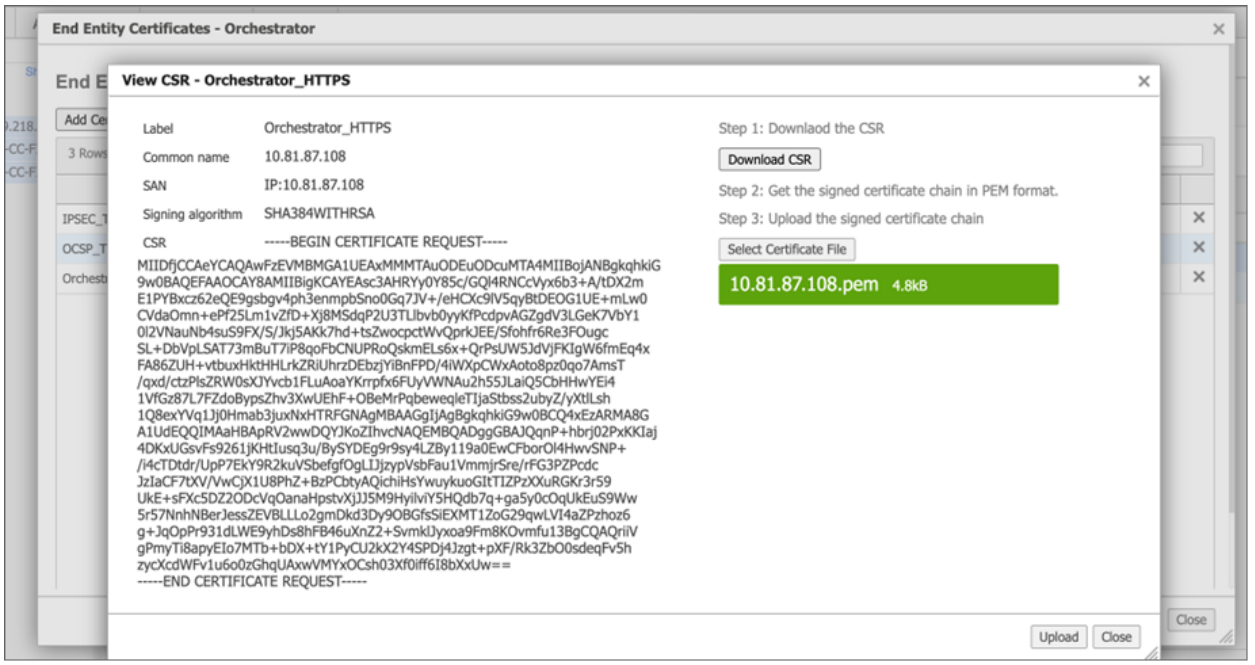


4. Click **Select Certificate File**, and then select the signed certificate chain.

   The file name of the end entity certificate appears in gray beneath the Select Certificate File button.

5. Click **Upload**.

   The system performs validation checks on the end entity certificate. If the certificate passes the validation checks and upload is successful, the file name turns green.



6. Click **Close**.

   After the certificate is successfully uploaded, on the End Entity Certificates dialog box the Issuer column contains the common name, the Expiration date column shows when the certificate expires, and the link in the CSR / Certificate column changes from **View CSR** to **View Certificate**. Click **View Certificate** to view the certificate.

Use an End Entity Certificate or Profile for a Service

You can use end entity certificates and appliance end entity profiles for the following services.

- HTTP server
  - To use a manually added end entity certificate for the Orchestrator HTTPS certificate, see Orchestrator HTTPS Certificate. The Orchestrator HTTPS certificate cannot be added using EST.
  - To use an orchestrated appliance end entity profile for the EdgeConnect HTTPS Certificate, see HTTPS Certificate Template.
  - To use a manually added end entity certificate for the EdgeConnect HTTPS Certificate, see HTTPS Certificate Tab.
- Syslog client server
  - To use an orchestrated appliance end entity profile for the syslog client certificate, see Logging Template.
  - To use a manually added end entity certificate for the syslog client certificate, see Remote Log Receivers.
- IKE-based IPSec tunnel
  - To apply an orchestrated appliance end entity profile to IKE-based IPSec tunnels, which can be used in Business Intent Overlays, see Tunnel Settings Tab.
  - To manually create IKE-based IPSec tunnels that use an end entity certificate, see Add or Modify a Manually Created Underlay Tunnel.

End Entity Certificates Dialog Box

From this dialog box you can view the end entity certificates for each appliance or Orchestrator, delete end entity certificates, and generate a new certificate signing request (CSR) to manually add a new end entity certificate.

The information in the following table is displayed for each end entity certificate on this dialog box. If you have created appliance end entity profiles, a separate row appears on the dialog box with a certificate for each appliance for each profile.

Column	Description
Label	The name assigned to the certificate.
Issuer	The issuer of the certificate.
Issued to	The entity to which the certificate is issued (common name on CSR).
Certificate	If you are in the process of manually adding a certificate for an appliance or Orchestrator and have created a CSR, a View CSR link appears in this column. After you have sent the CSR to a CA, obtained the signed certificate from the CA, and uploaded the signed certificate, a View Certificate link appears in this column that allows you to view and download the certificate. If you have created an appliance end entity profile, after enrollment with the EST server is completed, a View Certificate link appears in this column that allows you to view and download the certificate. A separate row will appear on the dialog box with a certificate for each appliance for each profile.
Expiration date	The date when the certificate expires.
Status	If certificate enrollment fails, a brief description of the reason for failure appears in this column, and the View Certificate link does not appear in the CSR / Certificate column.

Clients Table

Configuration > Overlays & Security > Security > Clients Table

The Clients Table tab provides at-a-glance details about LAN-side client devices in your network, which can assist in monitoring LAN-side hosts and troubleshooting issues. The Clients table on this tab lists client devices for all appliances in your network or those selected in the appliance tree. Client devices discovered by RADIUS snooping and/or Network Access Control (NAC) are listed in the table.

The maximum number of rows displayed in the Clients table is limited to 10,000. The maximum number of rows displayed for each appliance is 10,000 divided by the total number of appliances in your network or the number selected in the appliance tree. The following indicators are provided above the table:

- Displayed – Number of displayed clients (up to 10,000).
- Matched – Number of clients that match your query as defined by the filter selections.

The following filters, displayed at the top of the tab, work together to filter the Clients table:

- IP/Subnet – Filters on the specified user device IP address or subnet range (for example, 192.168.11.0/24).

- Segment – Filters on the selected segment. This filter is available only if routing segmentation is enabled.
- MAC Address – Filters on the specified user device MAC address (format XX:XX:XX:XX:XX:XX).
- Effective Role – Filters on the selected user device role.

After selecting the filtering criteria, click **Apply**. To clear filtering criteria, click **Clear**.

Descriptions of fields in the Clients table follow.

NOTE: The Clients table does not show real-time data.

Field	Description
Appliance	Name of the appliance. **NOTE:** If an appliance is not running on an ECOS version that corresponds with the current version of Orchestrator, relevant data for the appliance will not be collected nor displayed in this table.
MAC Address	MAC address of the user device.
User IP	IP address of the user device.
User Segment	Segment to which the client device is associated.
Interface	Interface through which the user device has been authenticated. Obtained from NACD.
User Name	Name of the user device. Obtained by RADIUS snooping.
User Group	User group assigned to the user device.
User Device	Operating system used for the user device.
User State	Current state of the user device. Active – Indicates active flows have occurred. Expired – Indicates no flows have occurred in the last hour. IP not acquired – Indicates that the user device IP address is not known, but the MAC address is known.
Authentication Method	Type of authentication method used (802.1x or MAC authorization). Obtained from NACD. “snooped” indicates that this was obtained by RADIUS snooping.
Effective Role	Role assigned to the user device. **NOTE:** Specific NADC and RADIUS snooping roles are in the NACD Role and Radius Snooping Role fields, which are not displayed in the table by default.
Session Start Time	Time when the user device was authenticated and the session was started.

Network Access Control (NAC)

Configuration > Overlays & Security > Security > Network Access Control (NAC)

When Network Access Control (NAC) is enabled on an appliance, the appliance authenticates traffic that accesses the network over untrusted interfaces. The appliance interprets the protocol packets and builds a RADIUS packet to get the station (client) authenticated with an external RADIUS server.

By default, authentication for all interfaces is set to *trusted*. When authentication is set to *trusted*, no authentication is required to access the network. When NAC security is enabled, the appliance authenticates stations that are trying to access the network using the policy you assign to the interface. This feature supports the EAP-TLS, EAP-TTLS, and EAP-PEAP methods for 802.1x authentication and the PAP method for MAC authentication.

Use the Network Access Control (NAC) tab to:

- Configure and modify NAC security for appliances using 802.1x and MAC authentication. There are two methods for enabling and managing NAC security: via templates or on a per-appliance basis.
  - Managing NAC security via templates is the recommended method. If you manage NAC authentication using templates, all settings are initially applied via the Network Access Control (NAC) configuration template for all appliances. Navigate to **Configuration > Templates & Policies > Templates** to display the Templates tab to add or edit a Network Access Control (NAC) template.
  - To enable or edit Network Access Control (NAC) on a per-appliance basis, see Enable or Edit Network Access Control (NAC).
- View configuration settings for NAC security for all appliances whether NAC is enabled via templates or on a per-appliance basis.
  - To refresh the NAC configuration table, click the **Refresh** button.
- Monitor and troubleshoot NAC authentication, including failures.
  - To view information about active stations, click **Active**. For more information, see Active Stations.
  - To view cached stations, click **Cache**. For more information, see Cached Stations.
  - To view stations that have been temporarily denied access, click **Deny list**. For more information, see Deny List Stations.
- View historical station data to track persistent authentication failures.
  - To view historical station data, click **History**. For more information, see Stations History.
- Download NAC security information from the NAC tab and subtabs by clicking **Export**.

Orchestrator supports the following maximum stations (MAC addresses), depending on the version you are running:

- 9.5.4 and earlier—Orchestrator supports authentication of up to 100 unique stations (MAC addresses) per interface, with 1024 unique MAC addresses allowed across all interfaces combined.
- 9.6.0 and later—Orchestrator supports 2048 unique stations (MAC addresses) per interface or across all interfaces combined.

The Configuration table on the Network Access Control (NAC) tab displays the following information.

Field	Description
Appliance	Name of the appliance for the Network Access Control (NAC) security settings.
LAN interface	LAN interface of the appliance to which the NAC policies are applied.
AAA profile	AAA profile applied to the appliance.
Auth type	Authentication type applied to the appliance.
Enabled	Indicates if NAC is enabled for the appliance or not.
Actions	Hover to view available actions for this row (Edit).

Active Stations

Use the NAC Active Stations subtab to review and monitor the authentication of all stations. To disconnect a station, select the entry for it in the table, and then click **Disconnect**. To select all entries in the table, click **Select all**. To refresh the table, click **Auto refresh**. To pause updates to the table, click **Pause**. By default, the table refreshes automatically.

Field	Description
Appliance	Name of the appliance.
MAC address	The MAC address of the station.
Identity	Identity on which the port is learned.
Interface	Port on which the identity of the station is learned.
Auth type	Authentication method (802.1x or MAC) used to verify the identity of the station.
Status	Indicates whether the station was authenticated or not.
Details	Click the information icon to display a complete list of NAC status details for the station.

Cached Stations

Use the NAC Cached Stations subtab to review and monitor MAC authenticated stations (clients) that are stored in the persistent local cache on an EdgeConnect appliance. Authentication survivability is enabled for the stations included in this table, which allows them to stay connected for a configured period of time in the event of a RADIUS server outage. After the survivability duration expires for a station it must reauthenticate. Authentication survivability is configured in the AAA profile. For more information, see AAA Profile.

NOTE: The Cached stations subtab only shows MAC authenticated stations. This feature is not available for 802.1x authentication.

The NAC Cached Stations table can have a maximum of 2048 entries. When the table reaches its capacity, new entries are discarded. To refresh the table, click **Auto refresh**. To pause updates to the table, click **Pause**. By default, the table refreshes automatically.

Column	Description
Appliance	Name of the appliance.
MAC address	MAC address of the cached station.
Interface	Port on which the identity of the station is learned.
Auth type	Authentication method used to verify the identity of the station; only MAC authentication is supported for caching.
Cache expiry	Configured duration after which cached data becomes invalid and the station must reauthenticate.

Deny List Stations

Use the NAC Deny List Stations subtab to review and monitor stations that are temporarily denied access due to excessive authentication failure. The NAC deny list is a temporary block list that prevents specific station MAC addresses from attempting authentication for a configurable period of time, and while on the deny list, any new authentication requests from the MAC address are silently dropped. You must enable the deny list in the AAA profile. For more information, see AAA Profile. A failure is defined as follows:

- A RADIUS authentication rejection for a MAC-based authentication request.
- A RADIUS response that indicates rejection due to policy, such as invalid credentials, an unknown device, or explicit deny rules.

The NAC Deny List Stations table can have a maximum of 2048 entries. When the table reaches its capacity, new entries are discarded. To unblock a station before the deny list duration expires for it, select the entry for it in the table, and then click **Unblock**. To select all entries in the table, click **Select all**. To refresh the table, click **Auto refresh**. To pause updates to the table, click **Pause**. By default, the table refreshes automatically.

Column	Description
Appliance	Name of the appliance.
MAC address	MAC address of the station that is temporarily denied access.
Interface	Port on which the identity of the station is learned.
Auth type	Authentication method (802.1x or MAC) used to verify the identity of the station.
Reason	Reason the station MAC address is temporarily denied access.
Deny expiry	Configured duration of time that the station MAC address is denied access.

Stations History

Use the NAC Stations History tab to track and monitor stations that fail authentication including repeat failures. If a station fails authentication repeatedly, the entry in the table is updated with each failure.

To refresh the table, click **Auto refresh**. To pause updates to the table, click **Pause**. By default, the table refreshes automatically.

Column	Description
Appliance	Name of the appliance.
MAC address	MAC address of the station.
Interface	Port on which the identity of the station is learned.
Reason	Reason the MAC address is temporarily denied access.
Total count	Number of times the MAC address has been denied access due to a failed authentication attempt. This number is updated in the table entry for a station each time it fails authentication.
Inactive time	Date and time when the MAC address was denied access due to a failed authentication attempt. This timestamp is updated in the table entry for a station each time it fails authentication.
Details	Click the information icon to display a complete list of NAC history details for the station.

Enable or Edit Network Access Control (NAC)

Setting Network Access Control (NAC) is a four-step process that is completed in the Network Access Control (NAC) dialog box.

1. Create an 802.1x/MAC authentication profile. See 802.1x/MAC Authentication Profiles.
2. Define the servers and optional server groups used for authenticating stations on the selected interface. See Server.

3. Add or edit the AAA profiles used for authentication. See AAA Profile.
4. Apply the Network Access Control (NAC) policies to the interface labels. See Apply Policies.

802.1x/MAC Authentication Profiles

Use the 802.1x/MAC tab to add or edit authentication profiles. You should create both 802.1x authentication and MAC authentication profiles. If the station is 802.1x compliant, the appliance will use the 802.1x profile to authenticate the station. If the station is not 802.1x compliant, the appliance will use the MAC profile to authenticate the station.

802.1x Authentication Profile Fields

1. On the Network Access Control (NAC) tab, click **Configuration**. Then select one or more appliances from the appliance tree, hover over the Actions column, and click the edit icon in the applicable table row.

   The Network Access Control (NAC) dialog box opens.
2. Click **Enable NAC**.
3. Click **Add** to add a new 802.1x authentication profile, or hover over the Actions column and click the edit icon to edit an existing 802.1x profile.

   The Add 802.1x Authentication Profile dialog box opens.

   NOTE: To modify an existing profile, edit the applicable fields and then click **Done**. To delete an 802.1x authentication profile, hover over the Actions column for the profile and click the **X** icon.
4. Complete the following fields.

Field	Description
Profile	Enter a name for the 802.1x profile.
Max auth failure	Maximum number of authentication failures allowed before the station is denied access. Default is 1, and the range is 1-5.
Max request	Maximum number of authentication requests that the appliance will send to the server. Default is 0, and the range is 0-3.
Identity requests interval	Interval (in seconds) between identity request retries. Default is 5 sec, and the range is 1-65535 sec.
Quiet period	Interval (in seconds) to wait before attempting to reauthenticate after a failed authentication. Default is 30 sec, and the range is 1-65535 sec.
Server retry count	Maximum number of retries that can be made on each server in a server group. If a server is not available, after the specified number of retries, Orchestrator attempts to access the next server in the server group. Default is 3, and the range is 1-5.

Field	Description
Server group retry period	Timeout duration (in seconds). If the appliance cannot reach the server within the specified period, the session times out. Default is 60 sec, and the range is 60-65535 sec.
Reauthentication	Select this check box to force the appliance to perform a reauthentication within the configured reauthentication interval.
Max reauthentication	Maximum number of reauthentication attempts. Default is 3, and the range is 1-10.
Reauthentication interval	Interval (in seconds) between reauthentication attempts. The configured interval is overridden if the RADIUS server provided the reauthentication period. Default is 86400 sec, and the range is 60-864000 sec.
Ignore EAPOL-START	Select this check box if the appliance should ignore the EAPOL-START messages after authentication.
Handle EAPOL-Logoff	Select this check box to handle the EAPOL-LOGOFF messages sent by the stations.

5. Click **Done**.

MAC Authentication Profile Fields

1. Click **Add** to add a new MAC authentication profile, or hover over the Actions column and click the edit icon to edit an existing MAC authentication profile.

   The Add MAC Authentication Profile dialog box opens.

   NOTE: To modify an existing profile, edit the applicable fields and then click **Done**. To delete a MAC authentication profile, hover over the Actions column for the profile and click the **X** icon.

2. Complete the following fields.

Field	Description
Profile	Enter a name for the MAC authentication profile.
Max auth failure	Maximum number of authentication failures allowed before the station is denied access. Default is 1, and the range is 1-5.
Quiet period	Interval (in seconds) to wait before attempting the retry after the failed authentication. Default is 30 sec, and the range is 1-65535 sec.

Field	Description
Server retry count	Maximum number of retries that can be made on each server in a server group. If a server is not available, after the specified number of retries, Orchestrator attempts to access the next server in the server group. Default is 3, and the range is 1-5.
Server group retry period	Timeout duration (in seconds). If the appliance cannot reach the server within the specified period, the session times out. Default is 60 sec, and the range is 60-65535 sec.
Reauthentication	Select this check box to force the appliance to perform a reauthentication within the configured reauthentication interval.
Max reauthentication	Maximum number of reauthentication attempts. Default is 3, and the range is 1-10.
Reauthentication interval	Interval (in seconds) between reauthentication attempts. Default is 86400 sec, and the range is 60-864000 sec.
Message authenticator	Select this check box to require all MAC authentication requests to include the Message-Authenticator attribute.

3. Click **Done**.

Continue to Server for server and server groups configuration.

Server

Use the Server tab to add or edit the servers and server groups you want to use to authenticate the stations that are attempting to log in to the network.

Server Fields

1. On the Network Access Control (NAC) dialog box, click **Server**.
2. Click **Add** to add a new server, or hover over the Actions column and click the edit icon to edit an existing server.

   NOTE: To modify an existing server, edit the applicable fields and then click **Done**. To delete a server, hover over the Actions column for the server and click the **X** icon.
3. Complete the following fields.

Field	Description
ID	Unique identifier of the server.
Server name	Enter a name for the server.
Host	Enter the FQDN, IPv4 address, or IPv6 address of the RADIUS server.

Field	Description
Key	Pre-shared key of the authentication server. This key is shared between the Mobility Conductor and the server. The maximum length is 128 characters.
Auth port	Server port on the server.
Interface for source IP address	IP address of the RADIUS server. This option allows the user to configure the interface to reach the RADIUS server.
Source segment	Segment name of the interface configured to reach the server.

4. Click **Done**.

Server Groups Fields

You can create groups of servers. If one server is not reachable based on the server retry count configured on the 802.1x/MAC tab, the appliance will try to reach another server.

1. Click **Add** to add a new server group, or hover over the Actions column and click the edit icon to edit an existing server group.

   The Server Group dialog box opens.

   NOTE: To modify an existing server group, edit the applicable fields and then click **Done**. To delete a server group, hover over the Actions column for the server group and click the **X** icon.

2. Complete the following fields.

Field	Description
ID	Unique identifier of the server group.
Server group name	Enter a name for the server group.
Servers	Servers in the server group. Click the cell to select servers from the list.

3. Click **Done**.

Continue to AAA Profile.

AAA Profile

Use the AAA Profile tab to add or edit AAA profiles to map the 802.1x and MAC authentication profiles to a server group you want to use to authenticate stations. This profile is used for dynamic authorization—for example, when a station needs to be reauthenticated or when the existing session is disconnected. After you create a AAA profile, you assign the profile to an interface label.

1. On the Network Access Control (NAC) dialog box, click **AAA Profile**.
2. Click **Add** to add a AAA profile, or hover over the Actions column and click the edit icon to edit an existing AAA profile.

   The Edit AAA Profile dialog box opens.

   NOTE: To modify an existing AAA profile, edit the applicable fields and then click **Done**. To delete a AAA profile, hover over the Actions column for the profile and click the **X** icon.
3. Complete the following fields.

Field	Description
Profile	Enter a name for the AAA profile.
DA Enable	Select this check box to enable Dynamic Authorization functionality.
DA Server	If you enabled DA, select the server to use for Dynamic Authorization.
802.1x Auth Profile	Select the name of the 802.1x authentication profile.
802.1x Default Role	Select the default role assigned to 802.1x stations.
802.1x Auth Server Group	Select the server group used for 802.1x authentication.
MAC Auth Profile	Select the name of the MAC authentication profile.
MAC Default Role	Select the role assigned to the station for MAC stations.
MAC Auth Server Group	Select the name of the server group used for MAC authentication.
Enable Survivability	Select this check box to enable authentication survivability. When enabled, if the authentication server becomes unavailable, stations that were already successfully authenticated retain access because the session information is stored in the local cache of the EdgeConnect appliance. Access for cached stations is based on the configured survivability duration (timer). Authentication Survivability is only applicable for MAC authentication. 802.1x enabled stations do not survive the authentication when an upstream RADIUS server becomes unreachable.
Survivability Duration	Maximum period of time (in hours) that a cached station can stay connected when all RADIUS servers are unreachable. After the survivability duration expires, the station must reauthenticate. Default is 24 Hrs, and the range is 1-168 Hrs. Survivability Duration must be set equal to or greater than the reauthorization interval if reauthorization is enabled (802.1x or MAC authentication profile).

Field	Description
Enable Deny List	Select this check box to add stations to the deny list if they have excessive authentication failures and are denied access. When a station MAC address is added to the deny list, it is prevented from attempting authentication for a configured period of time. While on the deny list, any new authentication requests from the MAC address are silently dropped.
Deny List Duration	Maximum period of time (in minutes) that a station remains on the deny list. After the deny list duration expires, the station is automatically removed from the list and can retry authentication. Default is 15 Min, and the range is 1-1440 Min.
Deny List Threshold Count	The Deny List Threshold Count controls how aggressively the system identifies and temporarily blocks stations (clients) that are experiencing excessive authentication failures. This setting works in conjunction with the Max auth failure value configured in your 802.1x or MAC authentication profile to determine the total number of authentication attempts a station is allowed before it is added to the deny list. Default is 3, and the range is 1–50. Formula: [Deny List Threshold Count] × [Max auth failure] = Total allowed authentication attempts Example: If the Deny List Threshold Count is set to 2 and the Max auth failure is set to 3 in the MAC or 802.1x authentication profile, the station is allowed a total of 6 authentication attempts before being denied. [2] × [3] = 6 total allowed authentication attempts

4. Click **Done**.

Continue to Apply Policies.

Apply Policies

Use the Apply Policies tab to modify the policies that are assigned to each interface label. Stations plugged into the LAN port with the assigned interface label will be authenticated using the policy you select.

Each LAN interface label defined in your Orchestrator deployment is assigned the default authentication policy. The default authentication policy is set to “trusted.” When authentication is set to “trusted”, no authentication is required to access the network.

1. On the Network Access Control (NAC) dialog box, click **Apply policies**.

2. Click **Modify**.

   The Interfaces dialog box opens. All available LAN interfaces are listed.

3. Hover over the Actions column and click the edit icon for the interface label that you are assigning a policy to.

   The Apply Policy dialog box opens.

4. Select a AAA profile.

   NOTE: If AAA profile is set to *none*, the Auth type is automatically set to *Trusted*.

5. Select an authentication method from the Auth type menu.

   - **Trusted**: Select **Trusted** if no authentication is required.
   - **Both**: Select **Both** to first attempt 802.1x authentication and then fall back to MAC authentication.
   - **802.1x**: Select **802.1x** if the port only supports 802.1x authentication.
   - **Mac**: Select **Mac** if the port only supports MAC authentication.

6. To close the Apply Policy dialog box, click **Done**.
7. To close the Interfaces dialog box, click **Save**.
8. To close the Network Access Control (NAC) dialog box and finish enabling NAC, click **Save**.

   You will receive a green status message if your policy was successfully applied.

Delete a Policy

To delete a policy from a LAN interface, hover over the Actions column for the LAN interface and click the **X** icon. The NAC security settings for this LAN interface return to the default values.

Secure Web Services

Configuration > Overlays & Security > Security > Secure Web Services

EdgeConnect appliances integrate with cloud-based OpenText™ Threat Intelligence (BrightCloud), a cybersecurity solution that provides real-time blocking of web-based threats. Specifically, EdgeConnect SD-WAN integrates the OpenText Threat Intelligence (BrightCloud) services, leveraging its Web Classification, URL Reputation and IP Reputation capabilities to enable real-time security enforcement and content filtering.

- **Web Classification and Reputation:** Dynamically assesses URL and domain risk, enabling you to enforce your web policies by granting users access to reputable websites while blocking them from inappropriate or malicious ones. You can set matching traffic criteria by:
  - One or more URL addresses

  - Web category: strict, moderate, or custom, each defined as a subset of site categories provided by OpenText Threat Intelligence (BrightCloud) (including high-risk categories)
  - Web reputation: high risk, suspicious, moderate risk, low risk, or trustworthy based on reputation scoring

- **IP Reputation:** Dynamically assesses the reputations of billions of IP addresses (both IPv4 and IPv6) and assigns reputation scores. You can set matching traffic criteria by bad IP reputation, which blocks high-risk IP addresses.

Users receive block notification messages when they attempt to access blocked pages as configured in your organization's policies.

To use Secure Web Services, your appliances must have one of the following valid licenses:

- Advanced Security Standard (EC-AS-Standard)
- Advanced Security Unlimited (EC-AS-UL)
- Dynamic Threat Defense (EC-DTD)

To configure licenses on your appliances, see Licenses. For assistance, contact HPE Aruba Networking support.

EdgeConnect appliances regularly synchronize with the most current intelligence data from cloud-based OpenText Threat Intelligence (BrightCloud). Before using Secure Web Services, ensure that the management service for Secure Web Services is configured appropriately on the Management Services tab. See Management Services.

The Secure Web Services tab displays status and configuration information for appliances selected in the appliance tree, as follows:

Field	Description
Appliance	Name of the appliance.
Status	Configuration status of Secure Web Services on the appliance. Possible statuses are: Web Classification & IP Reputation Web Classification IP Reputation Disabled
Licensed	Indicates whether the appliance is licensed to use Secure Web Services.
Query on miss	Indicates whether query on miss of URL classification will occur. When categorization of the queried URL is not available in the current URL-related threat intelligence database, the query is sent over the network to the cloud-based OpenText Threat Intelligence (BrightCloud) platform. The asynchronous response is then internally cached so that future queries for the URL category and reputation will return information successfully.

Field	Description
Incremental Update for Web	Indicates whether real-time updates to the appliance's URL classification database are enabled. These updates do not apply to EC-XS, EC-US, and EC-10104 appliance models.
Incremental Update for IP	Indicates whether real-time updates to the appliance's IP reputation database are enabled.
Unclassified URL Reputation	Risk category assigned to URLs that do not have a URL classification in the appliance's URL classification database. Possible risk categories are: High Risk (1 - 20) Suspicious (21 - 40) Moderate Risk (41 - 60) Low Risk (61 - 80) Trustworthy (81 - 100)
Web Traffic Blocked	Number of blocked web traffic flows that have occurred on the appliance over the last five minutes.
Events	Click the chart icon to view the following five charts showing statistics for the appliance: Total URL Lookup vs Total URL Success Total Flows Seen vs Denied Flows Total Network Queries Sent High Risk IPs High Risk URLs

To configure Secure Web Services for an appliance, click the edit icon associated with the appliance listed in the table. The Secure Web Services dialog box opens.

NOTE: The appliance must be licensed to use Secure Web Services. If not licensed, the edit icon is not available.

Secure Web Services Dialog Box

Use the Secure Web Services dialog box to enable and configure Secure Web Services on the appliance.

1. To enable Secure Web Services configuration on this appliance, move the **Secure Web Services Config** toggle to the right.
2. Complete the following configuration fields as appropriate.

Field	Description
Enable URL Classification	Move the toggle to the right to enable the URL classification feature on the appliance.
Incremental Update *(for URL Classification)*	Select this check box to enable real-time updates to the appliance's URL classification database.

Field	Description
Query on miss	Select this check box to enable queries when a URL classification is missed. When categorization of the queried URL is not available in the current URL-related threat intelligence database, the query is sent over the network to the cloud-based OpenText Threat Intelligence (BrightCloud) platform. The asynchronous response is then internally cached so that future queries for the URL category and reputation will return information successfully.
Unclassified URL Reputation	From the drop-down list, select a risk category to assign to URLs that do not have a URL classification in the appliance's URL classification database. Risk categories are: High Risk (1 - 20) Suspicious (21 - 40) Moderate Risk (41 - 60) Low Risk (61 - 80) Trustworthy (81 - 100) The default risk category is High Risk.
Enable IP Reputation	Move the toggle to the right to enable the IP reputation feature on the appliance.
Incremental Update *(for IP Reputation)*	Indicates whether real-time updates to the appliance's IP reputation database are enabled.

3. Click **Save**.

Cipher Settings

Configuration > Overlays & Security > Security > Cipher Settings

Cipher Profile settings allow you to restrict the use of certain ciphers for the Orchestrator and EdgeConnect OS (ECOS) services that use cryptography, as outlined in the following table.

Cryptographic Service	System	Applicability
TLS	Orchestrator and ECOS	Orchestrator acts as a client or a server in various TLS connection use cases.
Tunnels	ECOS	EdgeConnect IPSec tunnels
SSH	ECOS	EdgeConnect CLI
Certificate	Orchestrator and ECOS	Orchestrator web server, EdgeConnect web server, Syslog client, and IKE-based IPSec tunnels

Cryptographic Service	System	Applicability
SNMP	ECOS	The EdgeConnect SNMP interface supports Management Information Base II (MIB-II) as described in RFC 1213. This interface is used to receive notifications or traps from the appliance whenever an alarm is raised or cleared. SNMP is not used for configuration.
NTP	ECOS	Secure NTP is available starting with ECOS Release 9.4.2.
Cluster	ECOS	The cluster feature was introduced in Orchestrator Release 9.5.2; security configuration is performed in the Cipher Profile settings in Orchestrator Release 9.6.0 and later.

Cipher restrictions are applied to cryptographic services using a Cipher Profile. Some cryptographic services apply to both Orchestrator and EdgeConnect OS (ECOS) and some apply only to ECOS.

Cipher Profile settings determine which ciphers, algorithms, or parameters are allowed to be used by services requiring cryptographic functions. Changing the Cipher Profile settings may or may not change the ciphers used in existing connections and configurations.

IMPORTANT: For cryptographic services that require direct configuration, such as ECOS IPSec tunnels, you must manually change the configuration of the service before activating Cipher Profile changes.

This topic describes the following: - Cipher Settings - Service Impacts of Activating Cipher Profiles - Preconfigured Cipher Profiles - Cipher Settings Tab - Edit a Cloned Cipher Profile - TLS Settings - Validate TLS Settings on Self-Hosted/On-Prem Orchestrators - Validate TLS Settings on Orchestrator-as-a-Service - Validate TLS Settings on ECOS - Tunnels (ECOS) Settings - Avoid Dropped Tunnels When Changing Cipher Profile Tunnel Settings - Troubleshoot Dropped Tunnels When Changing Cipher Profile Tunnel Settings - SSH (ECOS) Settings - Configure SSH on ECOS - Certificate Settings - Configure Certificates on Orchestrator and ECOS - SNMP (ECOS) Settings - Configure SNMP on ECOS - NTP (ECOS) Settings - Configure NTP on ECOS - Set Up NTP Servers Using the Orchestrator Template - Configure Secure NTP via the CLI - Cluster Settings - Configure Clusters on Orchestrator

Service Impacts of Activating Cipher Profiles

The following table describes the impact to Orchestrator and ECOS for each cryptographic service when a Cipher Profile is activated, and actions the user must take to modify the service configuration.

Cryptographic Service	Impact to Orchestrator	Impact to ECOS
TLS	Changing the Cipher Profile changes the service configuration. You must validate the profile changes before applying them using the Validate button. After changing the Cipher Profile, existing TLS sessions remain up. Self-Hosted/On Prem Orchestrators: Perform a *service gms restart* using the Orchestrator CLI, at which point existing TLS sessions are dropped. After the gms service is restarted, TLS sessions are automatically reestablished with the new Cipher Profile settings. Orchestrator-as-a-Service: Changing the Cipher Profile only affects services for which Orchestrator is client, such as API connections to third-party services. You will need to open a TAC case to restart the Orchestrator as a Service (OaaS) to have TLS Cipher Profile changes take effect.	Changing the Cipher Profile changes the service configuration. You must validate the profile changes before applying them using the Validate button. After changing the Cipher Profile, TLS sessions are dropped and reestablished.

Cryptographic Service	Impact to Orchestrator	Impact to ECOS
Tunnels	Not applicable	Changing the Cipher Profile DOES NOT change the service configuration for any use case. You must perform tunnel configurations before changing the Cipher Profile, and you must explicitly validate the service configuration against the intended (to be activated) Cipher Profile. **IMPORTANT:** Existing tunnels, established with settings that will be disallowed in the intended Cipher Profile, will be DROPPED when the new Cipher Profile is activated.
SSH	For self-hosted/on-prem Orchestrators, SSH configuration is performed at the Linux level via the SSH config file; Cipher Profile settings do not apply. SSH does not apply to Orchestrator-as-a-Service.	Changing the Cipher Profile changes the service configuration. ECOS SSH configuration is orchestrated via the Cipher Profile settings. There is no need to configure via CLI.
Certificate	Existing end entity certificates configured with settings not allowed in the new/intended Cipher Profile will be invalid and not usable. You must validate the profile changes before applying them using the Validate button.	Existing end entity certificates configured with settings not allowed in the new/intended Cipher Profile will be invalid and not usable. You must validate the profile changes before applying them using the Validate button.

Cryptographic Service	Impact to Orchestrator	Impact to ECOS
SNMP	Not applicable	Starting with Orchestrator Release 9.5.4, ECOS SNMP configuration is performed via Orchestrator templates or via the Cipher Profile settings.
NTP	Not applicable	ECOS NTP server configuration is performed via Orchestrator templates. NTP secure algorithm parameters are configured via CLI and must be validated by the user before changing the service configuration.
Cluster	Not applicable	Changing the Cipher Profile changes the service configuration. The cluster feature was introduced in Orchestrator Release 9.5.2 with fixed cryptographic settings. Security configuration is performed in the Cipher Profile settings starting with Orchestrator Release 9.6.0.

NOTE: The Validate action (button) provides feedback on the compatibility of HTTPS server certificates indicating if the new Cipher Profile will work with existing TLS connections. Whether the Validate function returns a pass or fail, you can still execute the profile change. The Validate action DOES NOT verify other services such as tunnels.

For more information on EdgeConnect SD-WAN Cipher Profiles and cryptographic algorithms, see the Security Algorithms and Cipher Profiles PDF.

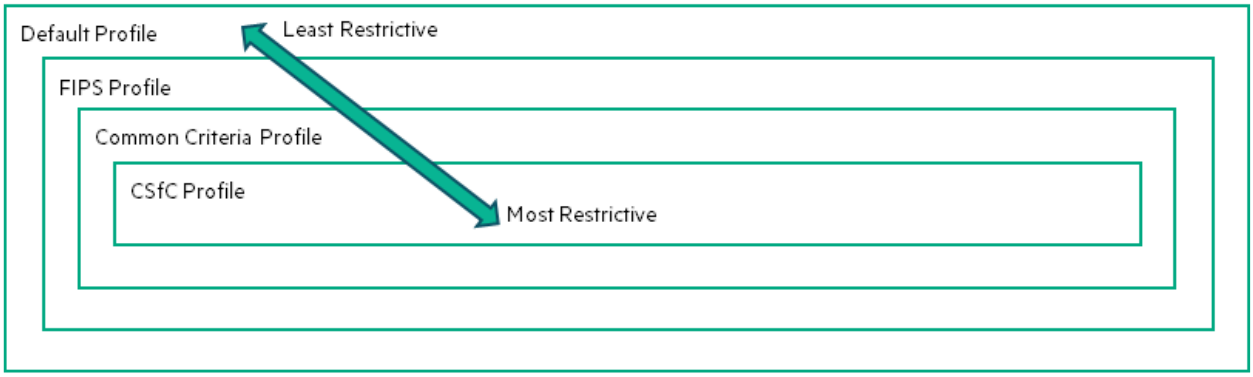
Preconfigured Cipher Profiles

Orchestrator provides four preconfigured Cipher Profiles.

- Default – This profile reflects the default values for new and upgraded Orchestrators with the least restrictive cipher algorithm options.
- FIPS – This profile applies cipher restrictions that meet FIPS requirements, as outlined in FIPS 140-2 (Federal Information Processing Standards Publication 140-2, Security Requirements for Cryptographic Modules).

- Common Criteria – This profile applies cipher restrictions that meet Common Criteria requirements outlined in the international standard The Common Criteria for Information Technology Security Evaluation (ISO/IEC 15408). More information on the standard is available at Common Criteria Portal.
- CSfC – This profile applies cipher restrictions that meet Commercial Solutions for Classified (CSfC) requirements, as outlined in the NSA document Commercial Solutions for Classified (CSfC) Selections for Internet Protocol Security (IPsec) Virtual Private Network (VPN) Gateways.

The following figure shows the comparative restrictiveness of the preconfigured Cipher Profiles.



The preconfigured Cipher Profiles cannot be edited. You can create a custom profile by cloning one of the preconfigured profiles and then editing it to match your requirements. You should clone the preconfigured profile that is the closest to your required enterprise security policy.

NOTE: As of Orchestrator Releases 9.4.4 and 9.5.4, the FIPS and Common Criteria Profiles are identical.

Cipher Settings Tab

Use the Cipher Settings tab to configure Cipher Profiles. The following table describes the fields displayed on this tab for each Cipher Profile.

Field	Description
Profile Name	The profile name. For cloned profiles, this name is entered by the user when they clone a profile.
Status	Shows the status of the Cipher Profile (Inactive or Active).

Field	Description
Comment	For the four preconfigured profiles, this is a description of the profile and indicates if it contains cipher restrictions to meet specific requirements. For cloned profiles, this shows the text the user enters in the Comments field for the profile.
Actions	Hover to view available actions for this row (View, Validate, Clone, Edit, and Delete). Edit and Delete are not available for the four preconfigured profiles because they cannot be edited or deleted.

You can use the Cipher Settings tab to view, clone, edit, validate and then activate, and delete Cipher Profiles.

- To view the cipher restrictions associated with a Cipher Profile, hover over the Actions column and click **View**. This only allows you to view the profile. It does not allow you to edit the profile.
- To clone a Cipher Profile, hover over the Actions column and click **Clone**.
- To edit a cloned Cipher Profile, hover over the Actions column and click **Edit**. For more information on cloning and editing a Cipher Profile, see Edit a Cloned Cipher Profile.
- To validate a Cipher Profile, hover over the Actions column and click **Validate**.
  - The Validate <profile name> dialog box opens.
  - What is validated: Only Orchestrator and ECOS end-entity HTTPS Server Certificates are checked for compatibility against the profile.
  - What is not validated: All other services.
  - When validation is complete, the Activate button on the dialog box is enabled. To apply the Cipher Profile, click **Activate**, enter the confirmation code, and click **Activate**. Only one Cipher Profile can be active at a time.

  IMPORTANT: For services that require direct configuration, such as ECOS IPSec tunnels, you must manually change the service configuration before activating any Cipher Profile changes.

  NOTE: The Validate action (button) provides feedback on the compatibility of HTTPS server certificates indicating if the new Cipher Profile will work with existing TLS connections. Whether the Validate function returns a pass or fail indication the user can still execute the profile change. The Validate action DOES NOT verify other services such as tunnels.
- To delete a cloned Cipher Profile, hover over the Actions column and click **Delete**. You cannot delete the four preconfigured Cipher Profiles.

Edit a Cloned Cipher Profile

To edit the cipher settings for a cloned Cipher Profile:

1. Navigate to **Configuration > Overlays & Security > Security > Cipher Settings**.
2. Hover over the Actions column for the cloned Cipher Profile you want to edit and click **Edit**.

   The Cipher Profiles dialog box opens.
3. Click the cryptographic service for which you want to edit cipher settings: TLS, Tunnels (ECOS), SSH (ECOS), Certificate, SNMP (ECOS), NTP (ECOS), Cluster Settings.
4. Edit the cipher settings for the cryptographic service. Click the following links to learn more about the cipher settings for each cryptographic service:
   - TLS Settings
   - Tunnels (ECOS) Settings
   - SSH (ECOS) Settings
   - Certificate Settings
   - SNMP (ECOS) Settings
   - NTP (ECOS) Settings
   - Cluster Settings
5. Click **Save**.

TLS Settings

TLS settings apply to both Orchestrator and EdgeConnect. There are two sections for the TLS profile settings, Cipher Suites and Advanced Properties. In the Cipher Suites section, you choose algorithms available for the following services: key exchange, encryption, and hash function. These building blocks form the composite cipher suites, which are shown in the Enabled and Disabled sections on the following figure. When editing the TLS settings, you can enable or disable the algorithms or the composite cipher suites.

To select an algorithm for a service, drag it from the Disabled to the Enabled column. When you enable an algorithm, any cipher suites that contain those algorithms move from Disabled to Enabled. The order that cipher suites appear in the Enabled column does not indicate the order in which the system (TLS server) offers them; all cipher suites that appear in the Enabled column are available for TLS.

The following table describes the TLS settings in the Cipher Suites section.

Setting	Description	Options
Key exchange algorithm	The algorithms available for Orchestrator to use to establish secure TLS connections.	NONE ECDHE_ECDSA ECDHE_RSA DHE_RSA RSA (disabled by default)
Encryption algorithm	The algorithms available for Orchestrator to use to encrypt data being sent during a TLS session.	AES_128_GCM AES_256_GCM AES_128_CCM AES_128_CCM_8 (disabled by default) AES_256_CBC (disabled by default) AES_256_CCM (disabled by default) AES_128_CBC (disabled by default) AES_256_CCM_8 (disabled by default)
Hash function	The secure hash algorithms available for Orchestrator to use for data integrity	SHA256 SHA384 SHA (disabled by default).

The following table describes the TLS settings in the Advanced Properties section.

Setting	Description	Options
Min DH Key length	The minimum length of the Diffie-Hellman key for Orchestrator when establishing secure TLS connections.	2048 (default) 3072 4096
Enable Session resumption	Enable/disable session resumption for TLS. Select or clear the check box. If enabled, an existing TLS secure session can be “reused.” This allows for faster session establishment. TLS session resumption is not allowed in CSfC.	Enabled or disabled

Setting	Description	Options
Named Curves	To customize, click **Custom** and select the check boxes for the named curves to have available.	Custom options: secp224r1 secp256r1 (default) secp384r1 secp521r1 ffdhe2048 ffdhe3072 ffdhe4096 ffdhe6144 ffdhe8192

Validate TLS Settings on Self-Hosted/On-Prem Orchestrators

For self-hosted/on-premises Orchestrators, changing the TLS Cipher Profile settings changes the service configuration for all Orchestrator TLS connections. The user must validate proposed changes before activation. When connecting to a third-party service, it is important to understand the cipher options that the third-party service offers and ensure the TLS Cipher Profile is in alignment.

When the TLS Cipher Profile settings are changed, existing TLS sessions remain up. The user must perform `service gms restart` using the Orchestrator CLI, at which point existing TLS sessions are dropped.

After the gms service is restarted, TLS sessions are automatically reestablished with the new Cipher Profile settings.

Validate TLS Settings on Orchestrator-as-a-Service

For OaaS, changing the TLS Cipher Profile settings only affects services for which Orchestrator is the client, such as API connections to third-party services and HTTPS or Syslog remote receiver connections.

The user must validate proposed changes before activation. When connecting to a third-party service, it is important to understand the cipher options that the third-party service offers and ensure the TLS Cipher Profile is in alignment.

Users need to open a TAC case to restart their OaaS to have TLS Cipher Profile changes take effect. After TAC restarts the OaaS, TLS sessions are automatically reestablished with the new Cipher Profile settings.

Validate TLS Settings on ECOS

Changing the TLS Cipher Profile settings changes the service for all ECOS TLS connections. The user must validate proposed changes before activation. After Cipher Profile changes are activated, ECOS resets all TLS connections automatically.

Tunnels (ECOS) Settings

The Cipher Profile settings for Tunnels apply to all tunnel types:

- Orchestrated underlay tunnels
  - IPSec UDP (IPSec phase only)
  - IKE-based, IPSec underlay tunnels
- Orchestrated IKE-based, IPsec passthrough tunnels to select SSE services (HPE SSE, Zscaler, Netskope, etc.)
- Unified Fabric, HPE ANW Central tunnels (requires Orchestrator release 9.5.4 or later)
- Service Orchestration constructed, EdgeConnect-to-third party passthrough tunnels
- Manually constructed, EdgeConnect-to-EdgeConnect tunnels
- Manually constructed, EdgeConnect-to-third party passthrough tunnels

For Cipher Profile settings for Tunnels, there are two sections, IKE and IPSEC. The IKE section only applies to IKE-based IPSec tunnels and does not apply to IPSec UDP tunnels. To enable an algorithm, select its associated check box. All tunnel algorithms are enabled in the default profile.

Setting	Description	Options
IKE Minimum Pre-Shared Key (PSK) length	The required minimum length of the IKE PSK (bytes). Enter a numeral in the range of 8-64.	8-64 bytes (8 bytes is the default.)
IKE Authentication Algorithms	The authentication algorithms available for the IKE security association (SA).	SHA-1 (default) SHA2-256 SHA2-384 SHA2-512
IKE Encryption Algorithms	The encryption algorithms available for the IKE security association (SA).	AES-CBC-256 AES-CBC-128 AES-GCM-128 AES-GCM-256 (Auto is the default.)

Setting	Description	Options
IKE Diffie-Hellman Group	The Diffie-Helman groups available for IKE security association (SA).	1 2 5 14 (default) 15 16 17 18 19 20 21 26 31
IKE Pseudo Random Function	The secure hash algorithms available for EdgeConnect IKE-based IPSec tunnels to use for pseudo random function.	SHA2-256 (default when AES-GCM-128 is selected for Encryption/Authentication) SHA2-384 (default when AES-GCM-256 is selected for Encryption/Authentication) SHA2-512
IPSEC Authentication Algorithms	The authentication algorithms available for EdgeConnect IPSec tunnels to use for IPSec security association (SA).	SHA-1 (default) SHA2-256 SHA2-384 SHA2-512 AES-GMAC-128 AES-GMAC-256
IPSEC Encryption Algorithms	The encryption algorithms available for EdgeConnect IPSec tunnels to use for IPSec security association (SA).	AES-CBC-256 (default) AES-CBC-128 AES-GCM-128 AES-GCM-256

Setting	Description	Options
IPSEC Perfect Forward Secrecy	The Diffie-Hellman groups available for EdgeConnect IPSec tunnels to use for IPSec security association (SA) negotiation.	1 2 5 14 (default) 15 16 17 18 19 20 21 26 31

NOTE: Changing the Tunnel Cipher Profile settings DOES NOT change the service configuration for any ECOS Tunnels.

Tunnel configurations must be performed prior to changing the Cipher Profile. The user must explicitly validate the service configuration against the intended (to be activated) Cipher Profile. ALL pre-existing tunnels, established with settings that will be disallowed in the intended (to be activated) Cipher Profile, will be DROPPED when the new Cipher Profile is activated.

Avoid Dropped Tunnels When Changing Cipher Profile Tunnel Settings

To avoid dropped tunnels when you change the Cipher Profile Tunnel settings, perform the following steps. These instructions reflect an Orchestrator that is starting with all default values, and the Default Cipher Profile is active.

1. Review the settings for the preconfigured Cipher Profiles (Default, FIPS, Common Criteria, and CSfC) to determine which profile has settings that are closest to your enterprise target security policy.

   NOTE: The Default, FIPS, and Common Criteria profiles have all algorithm options enabled for Tunnels (ECOS) settings. This is subject to change in future releases.
2. If none of the preconfigured profiles match the target security policy, clone the profile that is closest to the target.
3. Note any algorithms you plan to enable or disable for the Tunnels service.
4. Edit the cloned profile and make the necessary changes, but do not activate it yet.
5. After editing the cloned profile, you must implement the desired tunnel setting changes for all tunnel types. Click the following links for instructions on implementing tunnel settings changes for each tunnel type.

   - Implement Tunnel Setting Changes for Orchestrated Underlay Tunnels (all labels)

- Implement Tunnel Setting Changes for Orchestrated IKE-based, IPSec Passthrough Tunnels to Select SSE Services (HPE SSE, Zscaler, Netskope, etc.)
- Implement Tunnel Setting Changes for Unified Fabric, HPE ANW Central Tunnels
- Implement Tunnel Setting Changes for Manually Constructed, EdgeConnect-to-EdgeConnect Tunnels
- Implement Tunnel Setting Changes for Manually Constructed, EdgeConnect-to-Third Party Passthrough Tunnels

IMPORTANT: To avoid dropped tunnels, you must implement tunnel setting changes before you activate the Cipher Profile.

Implement Tunnel Setting Changes for Orchestrated Underlay Tunnels for ALL Labels

In Orchestrator, navigate to **Orchestrator > Orchestrator Server > Tools > Tunnel Settings**.

Orchestrated IPSec UDP Underlay Tunnels

- IPSec UDP is the default for Orchestrated Underlay Tunnels
- Only IPSec phase applies

General tab: No changes are required if you will continue to use IPSec UDP tunnels.

IPSec tab: There are two algorithms that correspond to Cipher Profile tunnel settings: Authentication algorithm and Encryption algorithm. Ensure that both settings use algorithms allowed in the Cipher Profile you plan to activate. The following figure shows the out-of-the-box default values. If for example, the new Cipher Profile disallows SH1 and CBC-128, you can change to the recommended setting of AES-GCM-256.

Orchestrated IKE-based IPSec Underlay Tunnels

Customers who require FIPS, Common Criteria, and/or CSfC compliance, can use standard, IKE-based IPsec tunnels.

NOTE: When configuring the IKE and IPSec encryption algorithms for FIPS and Common Criteria Compliance, the strength of the IKE algorithms must be greater than or equal to the IPSec algorithms.

The following table shows an example of a supported and valid configuration.

Setting	IKE-phase value	IPsec-phase value
Authentication Algorithm	n/a	n/a
Encryption Algorithm	AES-256-GCM-16	AES-256-GCM-16
Diffie-Hellman Group	DH 19	DH 19

The following table shows an example of a configuration that is unsupported and invalid because the IKE phase setting is cryptographically weaker than the IPSec phase setting.

Setting	IKE-phase value	IPsec-phase value
Authentication Algorithm	n/a	n/a
Encryption Algorithm	AES-128-GCM-16	AES-256-GCM-16
Diffie-Hellman Group	DH 18	DH 14

General tab: To select IKE-based IPSec tunnels for a specific label, from the Mode menu select **IPSec**. No other changes are needed on the General tab, unless the IPSec Suite B option is selected (only available for IKE-based IPSec Tunnels).

NOTE: GCM-256 is recommended. GMAC options cover authentication only and do not provide encryption.

IKE tab: For IKE-based IPSec tunnels, there are three algorithms that correspond to Cipher Profile tunnel settings: Authentication algorithm, Encryption algorithm, and Diffie-Hellman group. Ensure these settings use algorithms allowed in the Cipher Profile you plan to activate.

IPSec tab: For IKE-based IPSec tunnels, there are three algorithms that correspond to Cipher Profile tunnel settings: Authentication algorithm, Encryption algorithm, and Perfect forward secrecy group. Ensure these settings use algorithms allowed in the Cipher Profile you plan to activate.

Implement Tunnel Setting Changes for Orchestrated IKE-based, IPSec Passthrough Tunnels to Select SSE Services (HPE SSE, Zscaler, Netskope, etc.)

Tunnels are orchestrated for specific SSE services such as HPE SSE, Zscaler, and Netskope, and for Infrastructure as a Service (IaaS) cloud services AWS and Azure. Service Orchestration is used to automate the integration of third-party service providers without an API.

For each service, navigate to the tunnel settings for that service. Perform the actions outlined in the following table on the algorithms for each subsection (General, IKE, and IPSec) to ensure the settings use algorithms that are enabled in the Cipher Profile you plan to activate.

Tab	Action
General tab	No changes are needed.
IKE tab	For IKE-based IPSec tunnels, there are three algorithms that correspond to Cipher Profile tunnel settings: Authentication algorithm Encryption algorithm Diffie-Helman (DH) group Ensure these settings use algorithms that are enabled in the Cipher Profile you plan to activate.

Tab	Action
IPSec tab	For IKE-based IPSec tunnels, there are three algorithms that correspond to Cipher Profile tunnel settings: Authentication algorithm Encryption algorithm Perfect forward secrecy group Ensure these settings use algorithms that are enabled in the Cipher Profile you plan to activate.

You must also verify the allowable range of algorithm parameters for any partner service. Changing the Cipher Profile settings affects the EdgeConnect end of the tunnel, and you must verify that any changes are compatible with the partner service.

Follow the instructions for each service that you partner with to verify the changes you made to the Cipher Profile settings are compatible with that service.

- AWS Network Manager
  1. In Orchestrator, navigate to **Configuration > Cloud Services > AWS Network Manager**.
  2. Click **Tunnel settings**.
- HPE SSE
  1. In Orchestrator, navigate to **Configuration > Cloud Services > HPE SSE**.
  2. Click **Tunnel settings**.
- Microsoft Azure Virtual WAN
  1. In Orchestrator, navigate to **Configuration > Cloud Services > Azure Network Manager**.
  2. Click **On-prem Gateways**.
  3. Click **Tunnel Settings**.
- Zscaler Internet Access
  1. In Orchestrator, navigate to **Configuration > Cloud Services > Zscaler Internet Access**.
  2. Click **Tunnel Settings**.
- Netskope
  1. In Orchestrator, navigate to **Configuration > Cloud Services > Netskope**.
  2. Click **Tunnel Settings**.
- Service Orchestration Constructed EdgeConnect-to-Third Party Passthrough Tunnels

1. For services created within Service Orchestration, navigate to **Configuration > Cloud Services > Service Orchestration**.
2. Click the tab for the service (if configured).
3. Click **Tunnel Settings**.

Implement Tunnel Setting Changes for Unified Fabric, HPE ANW Central Tunnels

Unified Fabric, which supports tunnels between EdgeConnect and SD-Branch was introduced in Orchestrator Release 9.5.2.

Cipher Profiles were introduced in Orchestrator 9.5.4, however the Cipher Profile tunnel settings are not communicated to HPE Aruba Networking Central Overlay Tunnel Orchestrator (OTO). With Orchestrator and ECOS Release 9.5.4+, the following algorithms are used by HPE Aruba Networking Central and must be enabled in the Cipher Profile tunnel settings:

- Authentication: SHA1
- Encryption: AES-256-CBC

Starting with Release 9.6.0, ECOS sends the authentication and encryption methods to HPE Aruba Networking Central OTO. When a subsequent change is made to the Cipher Profile tunnel settings, for example to disallow SHA1, ECOS sends a resync message to OTO.

Implement Tunnel Setting Changes for Manually Constructed, EdgeConnect-to-EdgeConnect Tunnels

1. In Orchestrator, navigate to **Configuration > Networking > Tunnels > Tunnels**.
2. In the appliance tree, select the appliance.

   NOTE: To modify an Underlay or Passthrough tunnel, click **Underlay** or **Passthrough**, then hover over the Actions column of the tunnel and click **Edit**. The Modify Tunnel dialog box opens.
3. Click **Underlay**.
4. Click **Add Tunnel**.

   The Add Tunnel dialog box opens.
5. Perform the actions in the following table for each subsection (General, IKE, and IPSec) to ensure the settings use algorithms that are enabled in the Cipher Profile you plan to activate.

Tab	Action
General tab	No changes are needed.

Tab	Action
IKE tab	For IKE-based IPSec tunnels, there are three algorithms that correspond to Cipher Profile tunnel settings: Authentication algorithm Encryption algorithm Diffie-Helman (DH) group Ensure these settings use algorithms that are enabled in the Cipher Profile you plan to activate.
IPSec tab	For IKE-based IPSec tunnels, there are three algorithms that correspond to Cipher Profile tunnel settings: Authentication algorithm Encryption algorithm Perfect forward secrecy group Ensure these settings use algorithms that are enabled in the Cipher Profile you plan to activate.

Implement Tunnel Setting Changes for Manually Constructed, EdgeConnect-to-Third Party Passthrough Tunnels

1. In Orchestrator, navigate to **Configuration > Networking > Tunnels > Tunnels**.
2. In the appliance tree, select the appliance.

   NOTE: To modify an Underlay or Passthrough tunnel, click **Underlay** or **Passthrough**, then hover over the Actions column of the tunnel and click **Edit**. The Modify Tunnel dialog box opens.
3. Click **Passthrough**.
4. Click **Add Tunnel**.

   The Add Passthrough Tunnel dialog box opens.
5. Perform the actions in the following table for each subsection (General, IKE, and IPSec) to ensure the settings use algorithms that are enabled in the Cipher Profile you plan to activate.

Tab	Action
General tab	No changes are needed.

Tab	Action
IKE tab	For IKE-based IPSec tunnels, there are three algorithms that correspond to Cipher Profile tunnel settings: Authentication algorithm Encryption algorithm Diffie-Helman (DH) group Ensure these settings use algorithms that are enabled in the Cipher Profile you plan to activate.
IPSec tab	For IKE-based IPSec tunnels, there are three algorithms that correspond to Cipher Profile tunnel settings: Authentication algorithm Encryption algorithm Perfect forward secrecy group Ensure these settings use algorithms that are enabled in the Cipher Profile you plan to activate.

Troubleshoot Dropped Tunnels When Changing Cipher Profile Tunnel Settings

If you do not follow the recommended order of operations when changing Cipher Profile settings, tunnels that use algorithms that are disallowed by the active Cipher Profile will drop. If this happens, you can expect to see the following.

Alarms

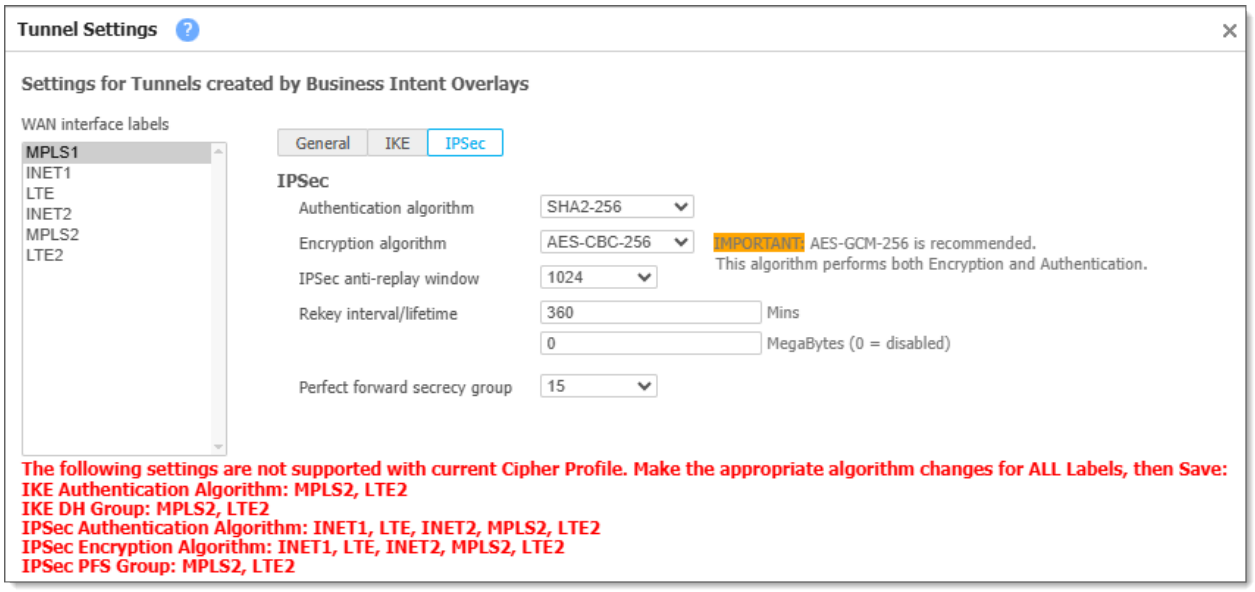
If tunnels drop due to disallowed Cipher Profile settings, Orchestrator generates an alarm for each tunnel that drops. To view the alarm description and recommended actions, navigate to **Monitoring > Summary > Alarms**.

Tunnels Tab Status

On the Tunnels tab, tunnels that dropped due to disallowed ciphers show a status of “down - bad” in the Status column.

Tunnel Settings Error Messages

If tunnels drop due to disallowed Cipher Profile settings, you need to manually correct the tunnel settings (**Orchestrator > Orchestrator Server > Tools > Tunnel Settings**). When attempting to save corrections for a single label, error messages appear on the dialog box for all nonconforming labels whether there are tunnels built for those labels or not, as shown in the following figure. You must correct the tunnel settings for each label to ensure they comply with the active Cipher Profile before you can click **Save**.



SSH (ECOS) Settings

To make an algorithm available for SSH service, select its associated check box.

Setting	Description	Options
Key Exchange Algorithms	The algorithms available for EdgeConnect to use to establish a secure connection to the appliance CLI.	diffie-hellman-group1-sha1 diffie-hellman-group14-sha1 diffie-hellman-group14-sha256 diffie-hellman-group16-sha512 diffie-hellman-group18-sha512 diffie-hellman-group-exchange-sha1 diffie-hellman-group-exchange-sha256 ecdh-sha2-nistp256 ecdh-sha2-nistp384 ecdh-sha2-nistp521 curve25519-sha256 curve25519-sha256@libssh.org
Encryption Algorithms	The algorithms available for EdgeConnect to use to encrypt data being sent to the appliance CLI.	aes128-cbc aes192-cbc aes256-cbc aes128-ctr aes192-ctr aes256-ctr aes128-gcm@openssh.com aes256-gcm@openssh.com

Setting	Description	Options
HMAC Algorithms	The Hash-based Message Authentication Code (HMAC) algorithms available for EdgeConnect to use when connecting to the appliance CLI.	hmac-sha1 hmac-sha1-96 hmac-sha2-256 hmac-sha2-512
Host Key Algorithms		ecdsa-sha2-nistp256-cert-v01@openssh.com ecdsa-sha2-nistp384-cert-v01@openssh.com ecdsa-sha2-nistp521-cert-v01@openssh.com ssh-ed25519-cert-v01@openssh.com rsa-sha2-512-cert-v01@openssh.com rsa-sha2-256-cert-v01@openssh.com ssh-rsa-cert-v01@openssh.com ecdsa-sha2-nistp256 ecdsa-sha2-nistp384 ecdsa-sha2-nistp521 ssh-ed25519 rsa-sha2-512 rsa-sha2-256 ssh-rsa
CA Sign Algorithms		ecdsa-sha2-nistp256 ecdsa-sha2-nistp384 ecdsa-sha2-nistp521 ssh-ed25519 rsa-sha2-512 rsa-sha2-256 ssh-rsa
Public Key Algorithms		ecdsa-sha2-nistp256-cert-v01@openssh.com ecdsa-sha2-nistp384-cert-v01@openssh.com ecdsa-sha2-nistp521-cert-v01@openssh.com ssh-ed25519-cert-v01@openssh.com rsa-sha2-512-cert-v01@openssh.com rsa-sha2-256-cert-v01@openssh.com ssh-rsa-cert-v01@openssh.com ecdsa-sha2-nistp256 ecdsa-sha2-nistp384 ecdsa-sha2-nistp521 ssh-ed25519 rsa-sha2-512 rsa-sha2-256 ssh-rsa

Configure SSH on ECOS

ECOS SSH configuration is orchestrated via the Cipher Profile settings. While still an option, there is no need to configure via CLI.

Certificate Settings

The certificate settings apply to end-entity certificates for both Orchestrator and EdgeConnect.

Setting	Description	Options
Fail connection on OCSP undetermined	When a certificate is uploaded in Orchestrator, the Online Certificate Status Protocol (OCSP) is run by Orchestrator and EdgeConnect to verify the status of the certificate. If communication cannot be established with the OCSP server, then the revocation check is ignored, but the connection does not necessarily fail. When this setting is enabled, if communication cannot be established with the OCSP server then the connection will fail.	Enable or disable
Min RSA Key length	Indicates minimum RSA key length for certificates.	2048 (default) 3072 4096
Min EC Key length		224 256 (default) 384 521
ECC Curves		Custom options: secp224r1 (included in default profile) secp256r1 (included in default profile) secp384r1 (included in default profile) secp521r1 (included in default profile)

Setting	Description	Options
HMAC Algorithms		Custom options: sha1 (included in default profile) sha256 (included in default profile) sha384 (included in default profile) sha512 (included in default profile)

NOTE: If ECOS is using a Certificate with key length 2048, it can talk to other TLS servers using a longer key length of 3072 (CSfC-compliant) certificates.

When applying a Cipher Profile, Orchestrator checks the ECOS TLS client to determine if it has a valid certificate. If ECOS does not have a valid certificate, the Cipher Profile will NOT be applied, and Orchestrator will raise an alarm. During the time of Zero Touch Provisioning (initial deployment), the ECOS TLS client will accept a server certificate, for example key length of 3072, even if ECOS does not yet have the CSfC Cipher Profile activated. This will enable initial communications for ECOS to Cloud Portal and ECOS to Orchestrator.

Configure Certificates on Orchestrator and ECOS

For information on configuring certificates, see End Entity Certificates Tab and End Entity Certificate Validation at the Time of Upload.

SNMP (ECOS) Settings

To make an algorithm available for an SNMP service to use, select its associated check box.

Setting	Description	Options
Encryption Algorithms	The encryption algorithms available for EdgeConnect to use to for SNMP service.	AES-CBC-128 AES-CBC-256
Hash Algorithms	The secure hash algorithms available for EdgeConnect to use for SNMP service.	SHA-1 SHA2-256 SHA2-384 SHA2-512

Configure SNMP on ECOS

Starting with Orchestrator Release 9.5.4, SNMP configuration is performed via Orchestrator templates or via the Cipher Profile settings.

- If Cipher Profile settings change, and any SNMP configuration uses privacy algorithms that are no longer valid, the system overwrites the hash algorithms to use one of the values allowed in the Cipher Profile.

- Similarly, if the SNMP configuration is changed to set hash and privacy algorithms, the system performs a check at the ECOS end to ensure the values are WRT allowed values and if they are not, Orchestrator overwrites the values.

NTP (ECOS) Settings

For NTP, you can select the algorithms available for EdgeConnect to use for network time protocol (SHA2-265, SHA2-384, and SHA2-512). To make an algorithm available for the NTP service to use, select its associated check box.

Configure NTP on ECOS

EdgeConnect OS supports NTP and manual time setting. NTP server can be set up for all appliances in the SD-WAN fabric using the Orchestrator template. However, Secure NTP requires configuration via the CLI.

Set Up NTP Servers Using the Orchestrator Template

Set up the NTP servers for all EdgeConnect appliances in the SD-WAN fabric using the Orchestrator Date/Time Template. For more information, see Date/Time Setting.

Configure Secure NTP via the CLI

Secure NTP Parameters

- Key id can be between 1-255.
- Algorithm is recommended to be SHA384.
- Pre-shared key can be any string. There is no restriction on string length or character composition.

Secure NTP Procedure

1. Authenticate to the EdgeConnect CLI/SSH.
2. Execute *enable*
3. Execute *configure terminal*
4. Add PSKs to the appliance:

   Execute *ntp key <Key id> <Algo> <pre shared key>*

   Example: *ntp key 10 sha384 uwefh8239hngli28*
5. Select one of the keys to securely connect:

   NOTE: You can only select keys from the list of available keys.

   To see the list of available keys:

   Execute *show ntp key*

   Execute *ntp select-key <Key id>*

   Example: *ntp select-key 10*

6. Restart the ntpd service:

   Execute: `pm process ntpd restart`

Optional CLI commands for secure NTP:

- Delete a Key: Execute `no ntp key <key id>`
- Displays all the ntp keys: Execute `show ntp key`
- Displays a specific key: Execute `show ntp key <key id>`

Manual NTP Server Setup

To manually set the NTP server for a single EdgeConnect appliance:

1. Execute `ntp server <IP address> version 4`
2. Execute `ntp enable`
3. Execute `write memory`

Manually Change Appliance Time or Date

1. Authenticate to the EdgeConnect CLI/SSH.
2. Execute `enable`
3. Execute `configure terminal`
4. Execute `clock set hh:mm:ss yyyy/mm/dd`
5. Execute `write memory`

Cluster Settings

To make an algorithm available for the Cluster service to use, select its associated check box.

Setting	Description	Options
Encryption Algorithm	The encryption algorithms available for Orchestrator and EdgeConnect to use for Clusters.	AES-CBC-256 AES-CBC-128 AES-GCM-128 AES-GCM-256
Authenticat Algorithm	The authentication algorithms available for Orchestrator and EdgeConnect to use for Clusters.	SHA-1 SHA2-256 SHA2-384 SHA2-512 AES-GMAC-128 AES-GMAC-256

Configure Clusters on Orchestrator

For information on configuring Clusters, see Cluster Profiles and Clusters.

NOTE: Cluster security is only configured via Cipher Profile settings starting with Orchestrator Release 9.6.0.

SSL WAN Optimization Certificates

Configuration > Overlays & Security > SSL for WAN Optimization > SSL WAN Optimization Certificates

EdgeConnects provide deduplication for SSL-encrypted WAN traffic by supporting the use of SSL WAN Optimization certificates and other keys. An EdgeConnect decrypts SSL data using the configured certificates and keys, optimizes data, and transmits data over an IPSec tunnel. The peer EdgeConnect uses configured SSL WAN Optimization certificates to re-encrypt data before transmission.

- Peers that exchange and optimize SSL traffic must use the same certificate and key.
- For the SSL WAN Optimization certificates to function:
  - Tunnels must be in IPSec or IPSec UDP mode for both directions of traffic.
  - TCP acceleration and SSL acceleration must be enabled in the Optimization Policy.

The SSL WAN Optimization Certificates tab summarizes the SSL WAN Optimization certificates installed on appliances for decrypting non-SaaS traffic.

TIP: For a historical matrix of EdgeConnect and Orchestrator security algorithms, view the Security Algorithms PDF.

SSL WAN Optimization Certificates Dialog Box

Use the SSL WAN Optimization Certificates dialog box when the server is part of your enterprise network and has its own enterprise SSL certificates and key pairs.

Use this dialog box to load the certificate and key directly into the appliance.

- You can add either a PFX certificate (generally, for Microsoft servers) or a PEM certificate. If the PFX Certificate File check box is not selected, the default certificate is PEM.
- If the key file has an encrypted key, enter the passphrase needed to decrypt it.

Before installing the certificates:

- Configure the tunnels bilaterally for IPSec mode.

  To do so, navigate to **Configuration > Networking > Tunnels > Tunnels**, and then click the edit icon for the appropriate tunnel. The Tunnels dialog box opens. Click the edit icon for the appropriate tunnel, and then set the Mode field to **IPSec**.

- Verify that TCP acceleration and SSL acceleration are enabled.

  To do so, navigate to **Configuration > Templates & Policies > Policies > Optimization Policies**, and then review the Set Actions.

SSL WAN Optimization CA Certificates Tab

Configuration > Overlays & Security > SSL for WAN Optimization > SSL WAN Optimization CA Certificates

This tab lists any installed **Certificate Authorities (CA)** that the browser uses to validate up the chain to the root CA.

If the enterprise certificate that you used for signing substitute certificates is subordinate to higher level **Certificate Authorities (CA)**, you must add those CA certificates. If the browser cannot validate up the chain to the root CA, it will warn you that it cannot trust the certificate.

TIP: For a historical matrix of EdgeConnect and Orchestrator security algorithms, view the Security Algorithms PDF.

SSL WAN Optimization CA Certificates Edit Row

If the enterprise CA certificate you use for signing substitute certificates is subordinate to higher level **Certificate Authorities (CA)**, you must add those CA certificates here.

Those same CA certificates must also be present in the browser. If the browser cannot validate up the chain to the root CA, it will warn you that it cannot trust the certificate.

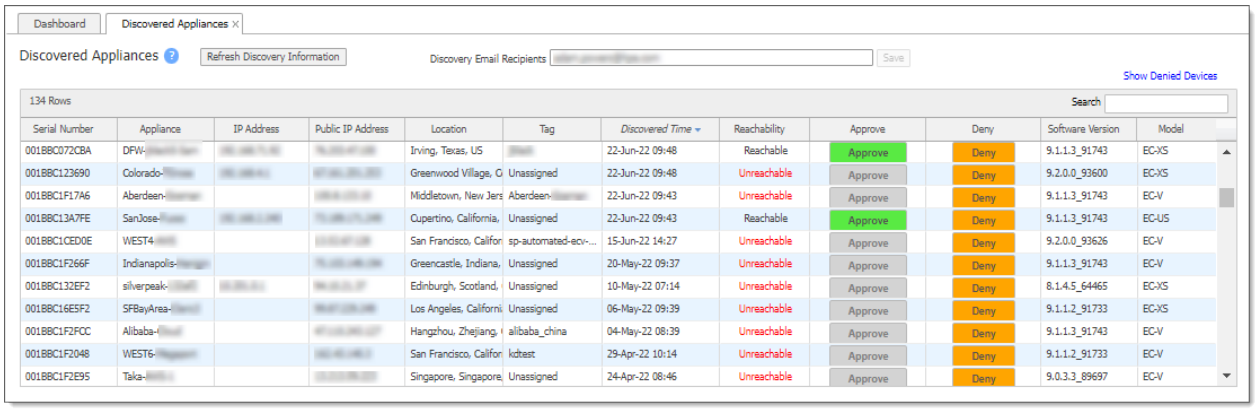
- Use this page to directly load the CA certificate into the appliance.
  - You can add either a PFX certificate (generally, for Microsoft servers) or a PEM certificate.
  - The default is PEM when **PFX Certificate File** is deselected.
- EdgeConnect supports:
  - X509 Privacy Enhanced Mail (PEM), Personal Information Exchange (PFX), and RSA key 1024-bit and 2048-bit certificate formats.
  - SAN (Subject Alternative Name) certificates. SAN certificates enable sharing of a single certificate across multiple servers and services.

TIP: For a historical matrix of EdgeConnect and Orchestrator security algorithms, view the Security Algorithms PDF.

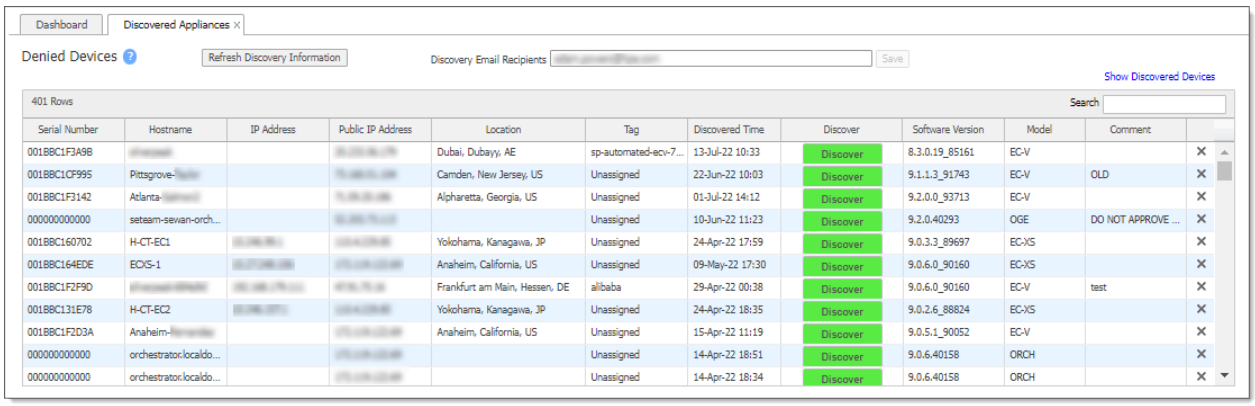
Discovered Appliances

Configuration > Overlays & Security > Discovery > Discovered Appliances

This tab lists each appliance that Orchestrator discovers.



- To enable Orchestrator to manage an appliance after you verify its credentials, click **Approve**.
- If the appliance does not belong in your network, click **Deny**. If you want to include it later, click **Show Denied Devices**, locate it in the table, and click **Discover**.

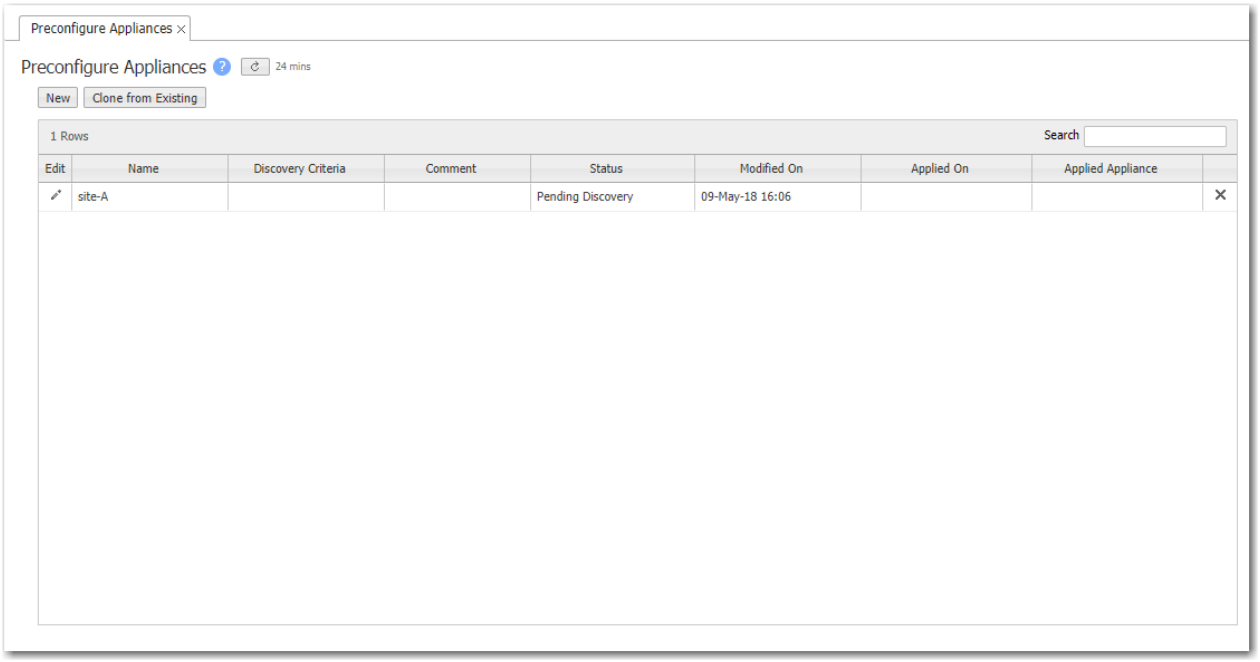


- As a security measure to prevent unauthorized management of your network, any Orchestrator with your Account Name and Account Key must be approved by the originally deployed Orchestrator.
- If you enabled Air-Gap mode, the **Air-Gap appliances awaiting discovery** button may appear. For more information on discovering appliances and assigning serial numbers to appliances in Air-Gap mode, see HPE Aruba Networking Orchestrator Air-Gap User Guide.

Preconfigure Appliances

Configuration > Overlays & Security > Discovery > Preconfiguration

Use this page to prepopulate flat data files that are matched with appliances as you add them to your network.

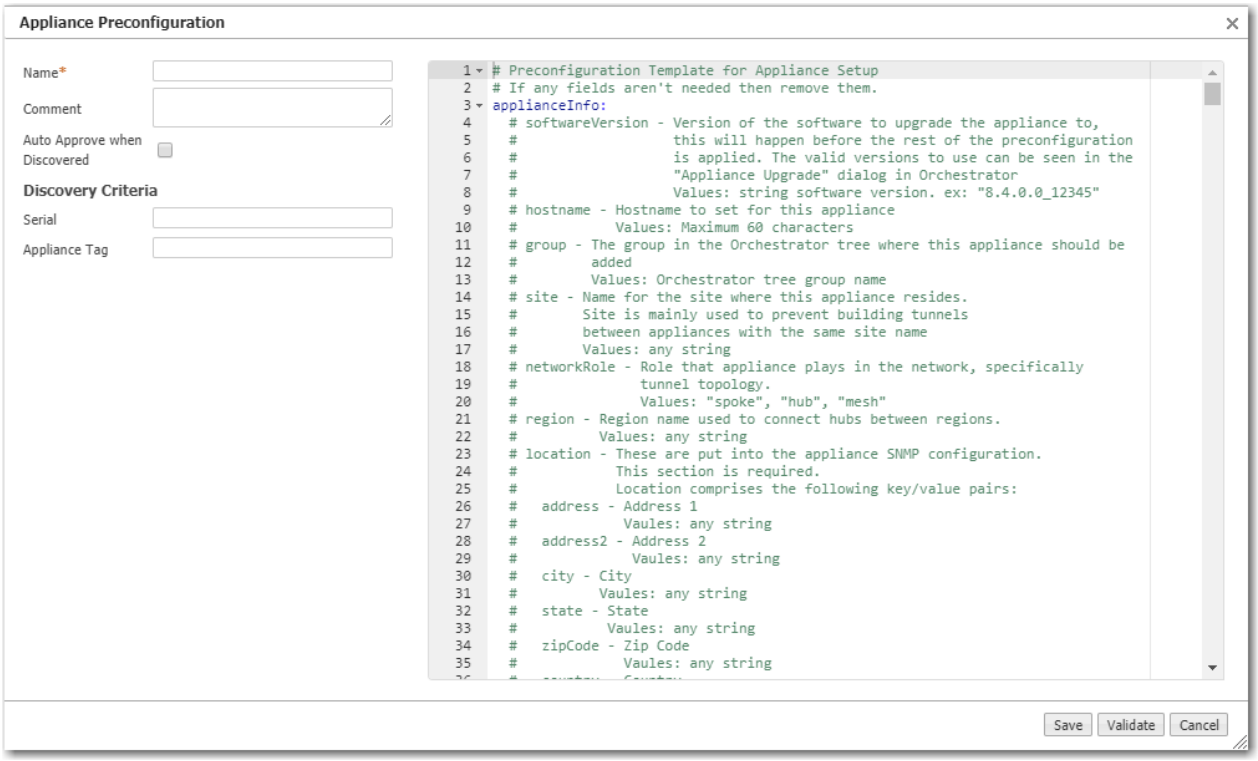


The information in the files is a combination of items found in the Appliance Configuration Wizard, along with site-specific information such as BGP, OSPF, IP SLA rules, VRRP, interfaces, and addressing.

You can create a new file or clone (and rename) an existing one. Make any changes with the built-in editor.

After the appliance is discovered and approved, software upgrade and configuration push are done automatically.

New or Clone



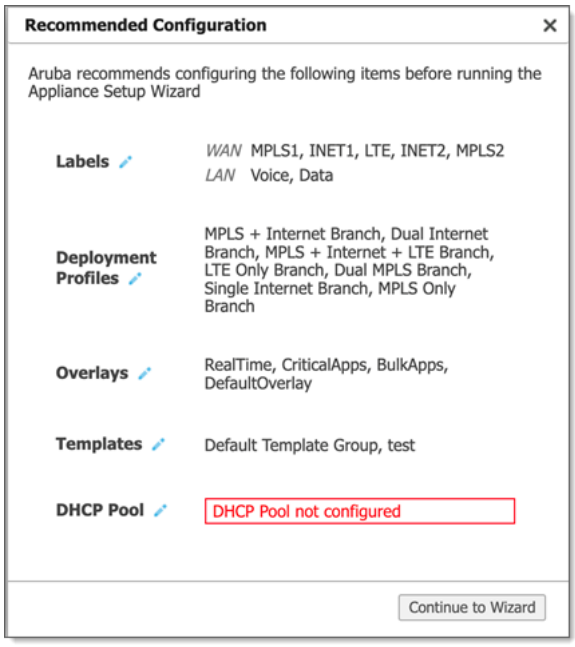
Field	Description
Name	Assigns a name to the preconfiguration file.
Comment	Optional descriptive field.
Auto Approve when Discovered	When **selected**, Orchestrator finds the appliance that matches the Discovery Criteria and automatically loads it without needing user intervention. When **deselected**, the user will be prompted to manually approve the association of the preconfiguration file to the appliance.
Serial	Serial number associated with the appliance that is to receive this configuration.
Appliance Tag	Free-form text or unique identifier that an administrator can associate with the appliance. Available as a discovery criteria for EC-Vs.

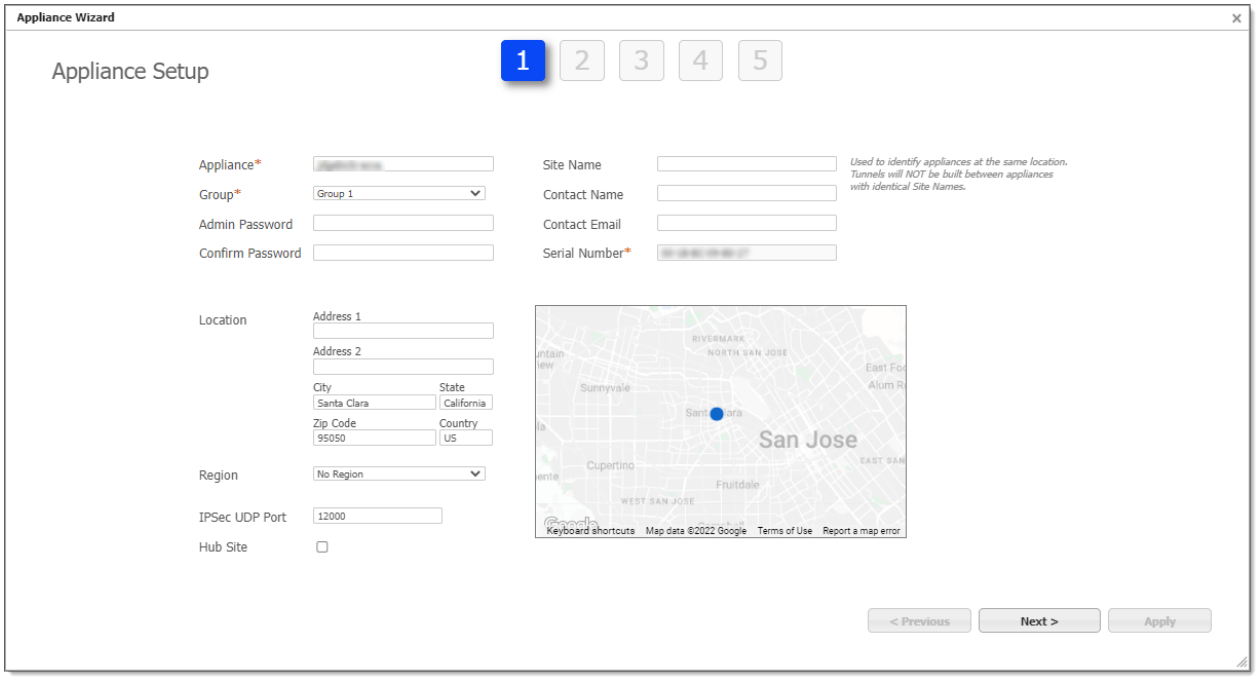
Appliance Wizard

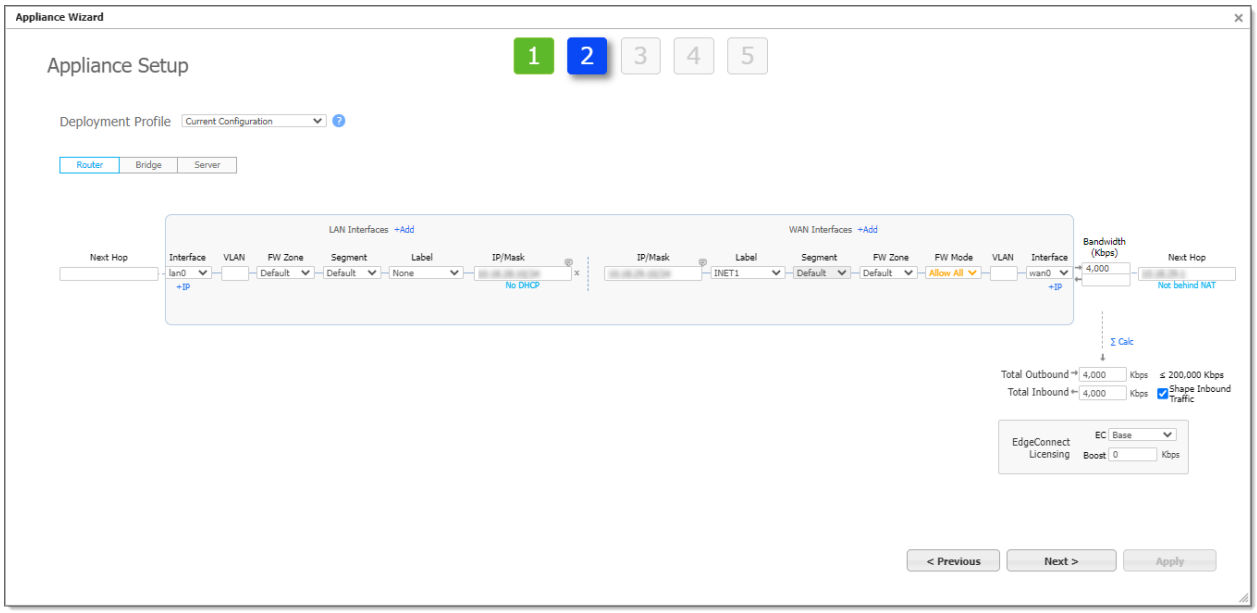
Configuration > Overlays & Security > Discovery > Configuration Wizard

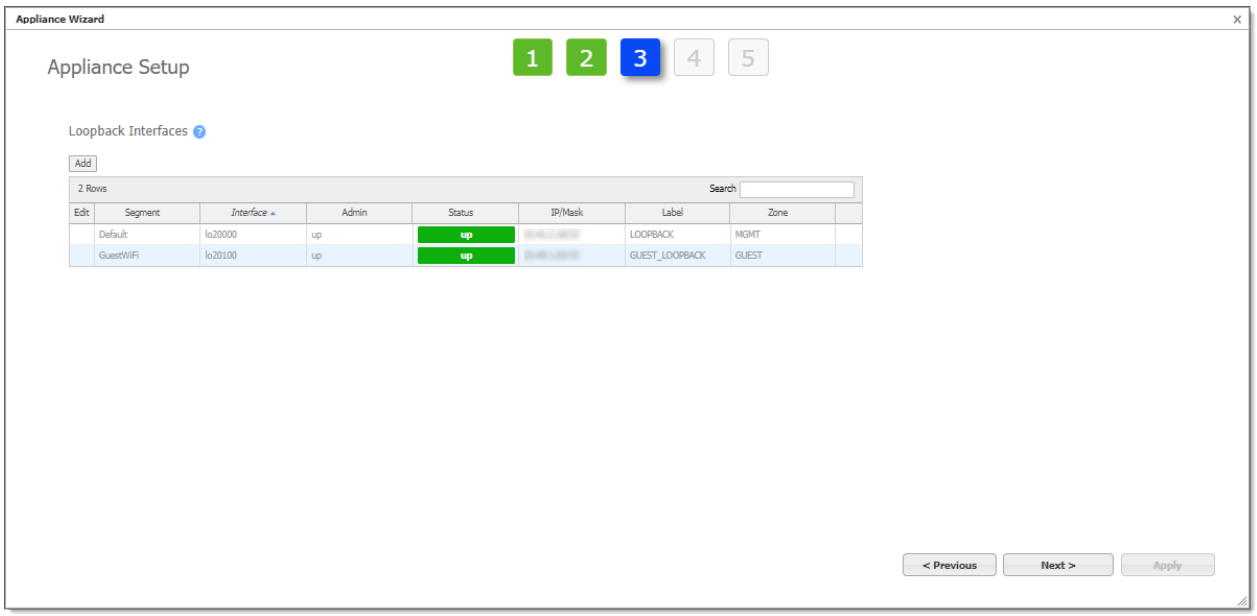
Use the Appliance Wizard to set up a newly added appliance or to reconfigure an appliance that is already in your network.

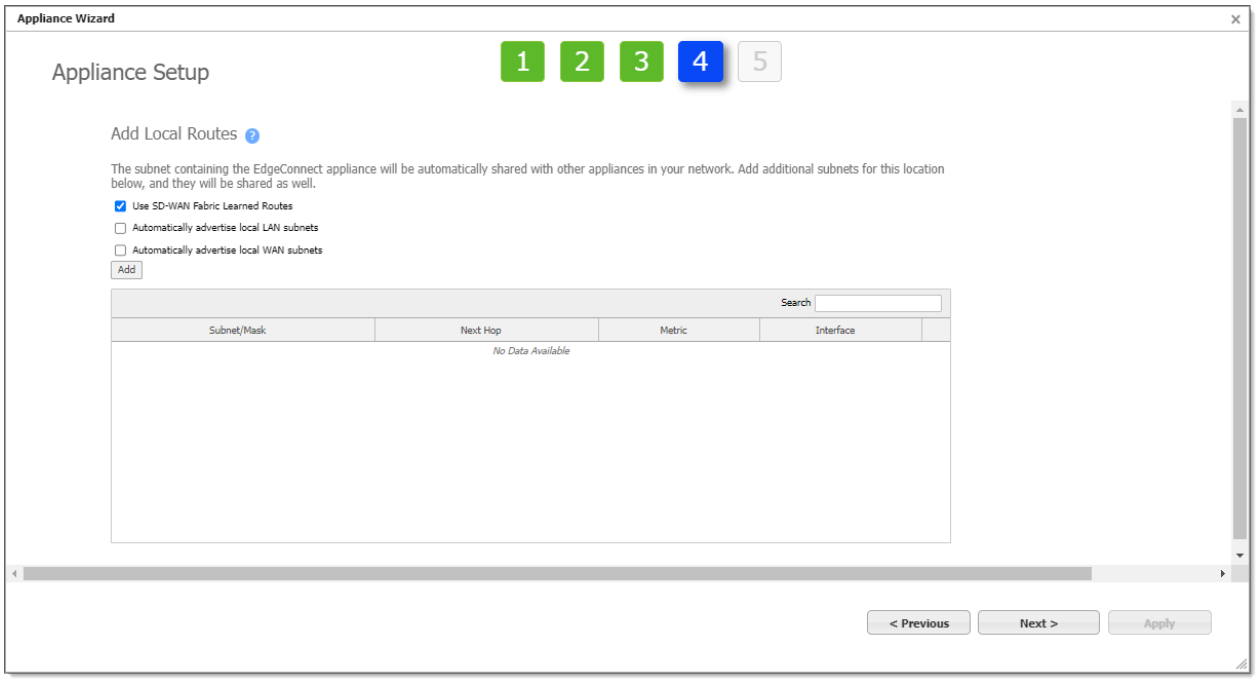
NOTE: Orchestrator assumes you will push many of the same configuration items to each appliance. To that end, it surveys the templates and Overlay prerequisite items and displays the **Recommended Configuration** list, showing what comprehensive items you have and have not yet configured.

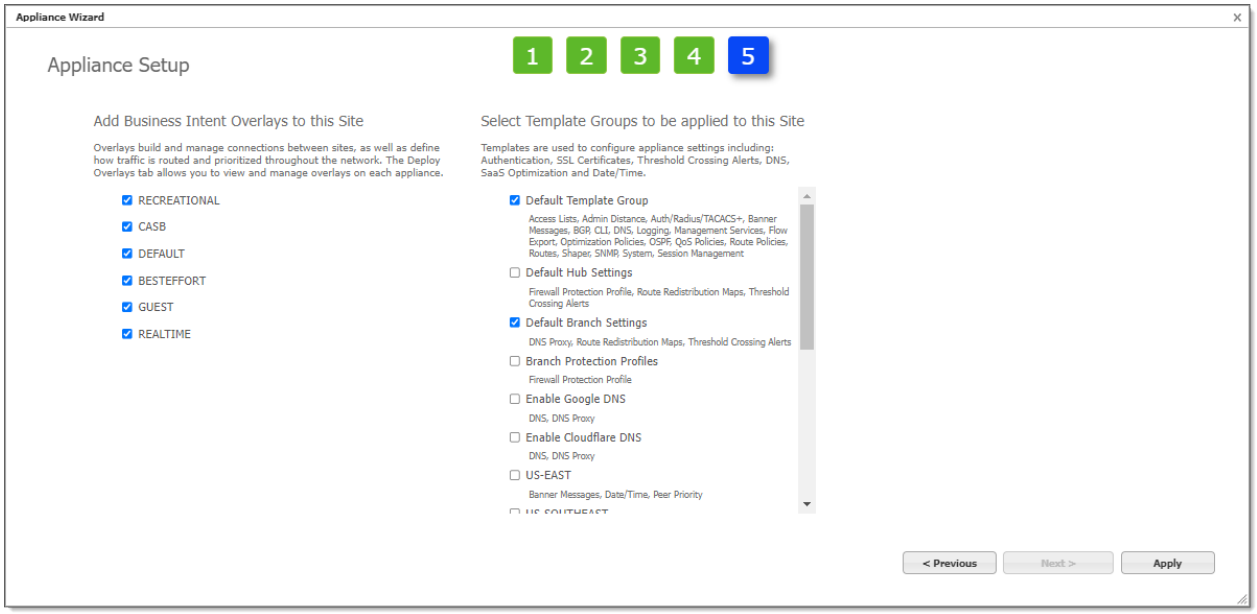





Appliance Wizard
Appliance Setup
1 2 3 4 5
Deployment Profile
Current Configuration
Router
Bridge
Server
LAN Interfaces +Add
WAN Interfaces +Add
Next Hop
Interface
VLAN
FW Zone
Segment
Label
IP/Mask
IP/Mask
Label
Segment
FW Zone
FW Mode
VLAN
Interface
Bandwidth (Kbps)
Next Hop
lan0
Default
Default
None
No DHCP
INET1
Default
Default
Allow All
wan0
4,000
Not behind NAT
+IP
Σ Calc
Total Outbound
Total Inbound
4,000
Kbps
≤ 200,000 Kbps
Shape Inbound Traffic
EdgeConnect Licensing
EC
Base
Boost
0
Kbps
< Previous
Next >
Apply


Appliance Wizard
Appliance Setup
1 2 3 4 5
Loopback Interfaces
Add
2 Rows
Search
Edit
Segment
Interface
Admin
Status
IP/Mask
Label
Zone
Default
lo20000
up
up
LOOPBACK
MGMT
GuestWiFi
lo20100
up
up
GUEST_LOOPBACK
GUEST
< Previous
Next >
Apply





EC-Enterprise Licenses

Configuration > Overlays & Security > Licensing > Licenses

- This page lists the appliance model, serial number, appliance name, feature licenses, and license terms for the appliances selected in the appliance tree.
- You can add, edit, or revoke EdgeConnect (EC) licenses from an appliance.

- A license summary including the number of used licenses and total number of available licenses is displayed above the table. The expiration date of the Boost license and each feature license is also listed.

NOTE: EdgeConnect stops passing traffic when a license expires.

Assign a License to an Appliance

1. In the appliance tree, select one or more appliances to display in the table.
2. Do one of the following:
   - To assign licenses to one appliance, click the **Edit** icon next to that appliance.
   - To assign licenses in bulk (to all appliances in the table), click **Assign Licenses to Appliances**.

     NOTE: To assign licenses in bulk, all appliances must be on the same software version.

   The **Assign Licenses to Appliances** dialog box opens.
3. Complete the following elements as needed:

Field	Description
EC	Select the **Add/Replace** check box, and then select the EC size from the list: **Mini**, **Base**, **Base + Plus**, **50 Mbps**, **200 Mbps**, **500 Mbps**, **1 Gbps**, **2 Gbps**, or **Unlimited**.
Boost	Select the **Add/Replace** check box, and then enter the amount of Boost to apply to the EC.
Feature licenses	1. To add a feature license, select the **Add/Replace** check box. 2. If required, select a license option from the list and specify a quantity, such as amount of bandwidth.

4. To revoke a license or Boost, select the **Revoke** check box next to the license or Boost you want to revoke.

   NOTE: If you revoke an EC license from an appliance, Silver Peak will revoke the Boost license and all feature licenses from that appliance.

   NOTE: You must revoke the license from an appliance before you can RMA it. For more information on how to RMA an appliance, see RMA Wizard.
5. Click **Apply**.

Licenses

This page lists the appliance model, serial number, appliance name, and feature licenses for the appliances selected in the appliance tree. You can add, edit, or revoke EdgeConnect (EC) licenses from an appliance.

- To configure an EC-Enterprise license, go to EC-Enterprise Licenses.
- To configure an EC-Metered license, go to EC-Metered Licenses.

EC-Enterprise Licenses

Configuration > Overlays & Security > Licensing > Licenses

Use this tab to view and administer licenses for all EdgeConnect appliances.

To view a summary of license information, click **Summary**. The table displays the following information.

Field	Description
Account	The Orchestrator account that the license subscription is associated with.
EC License	The license subscription type.
Licenses Available	The number of individual appliance licenses that are not assigned to an appliance.
License Usage	The number of individual appliance licenses that are assigned to an appliance and in use.
License Count	The total number of individual appliance licenses available with the license subscription.
License End	The expiration date for the license subscription.

To view license information by appliance, click **Appliances**. You can assign, edit, or revoke licenses for appliances on this page. You can filter the list of appliances as follows:

- EdgeConnect: click to show only appliances that have an EdgeConnect license applied.
- WAN Optimization: click to show only appliances that have a WAN Optimization license applied.
- Advanced Security: click to show only appliances that have an advanced security license applied.

Licensing Subscription Tiers

EC-Enterprise licensing subscription tiers are Foundation, Advanced, and On Prem, and there are some optional feature add-on licenses. To access the full scope of certain features, Advanced or On-Prem licensing is needed. If you do not have the appropriate licensing to access a feature, Orchestrator indicates this using onscreen messages.

The following information outlines the features that vary by licensing tier and describes the optional feature add-on licenses.

Foundation (Cloud Orchestrator)

- Bandwidth tier licenses allowed: 100 Mbps, 1 Gbps, and Unlimited
- Up to three Business Intent Overlays (RealTime, Critical, and Default)

- Up to 2 VRF segments (Default + one custom)
- Limited QoS
- Standard stats retention (Cloud Orchestrator)
- AppExpress (Monitor only)
- Multi-Regional topologies – 4 max regions (Default + three custom), 4 hosts /region

Advanced (Cloud Orchestrator)

- Bandwidth tier licenses allowed: 20 Mbps, 50 Mbps, 100 Mbps, 200 Mbps, 500 Mbps, 1 Gbps, 2 Gbps, and Unlimited
- Up to seven BIOs (RealTime, Critical, BulkApps, Default, and three custom)
- Maximum available allowed VRF segments
- Enhanced stats retention (Cloud Orchestrator)
- AppExpress (Monitor and steer)
- Multi-Regional topologies – maximum available allowed

On-Prem

- Up to seven BIOs (RealTime, Critical, BulkApps, Default, and three custom)
- Maximum available allowed VRF segments
- Custom stats retention
- AppExpress (Monitor and steer)
- Multi-Regional topologies – maximum available allowed

Optional Feature Add-On Licenses

- WAN Optimization
  - Capabilities include latency Mitigation and data reduction, minimizes the effects of latency on application performance, and improves application response times across the WAN.
  - Per-megabit-per-second, per-month.
- Advanced Security/Dynamic Threat Defense (DTD)
  - Intrusion detection and prevention security features (IDS/IPS)
  - Adaptive DDoS (baseline learning, auto rate limit, smart burst)
  - SYN Cookie
  - Web security intelligence
  - This license must be purchased for each EdgeConnect SD-WAN appliance.

Assign a License to an Appliance

1. On the Licenses tab, click **Appliances**.
2. Do one of the following:
   - To assign licenses to one appliance, click the edit icon next to that appliance.
   - To assign licenses in bulk (to all appliances in the table), click **Assign Licenses to Appliances**.

     NOTE: To assign licenses in bulk, all appliances must be on the same software version.

   The Assign Licenses to Appliances dialog box opens.
3. Complete the following elements as needed:

Field	Description
EC	Select the **Add/Replace** check box, and then select the EC size from the menu (**50 Mbps**, **200 Mbps**, **500 Mbps**, **1 Gbps**, **2 Gbps**, **Unlimited**, **Base**, or **Base + Plus**.
WAN Optimization	Select the **Add/Replace** check box, and then enter the amount of WAN Optimization to apply to the EC.
Advanced Security	To add an advanced security/DTD license, select the **Add/Replace** check box, and then select the type of advanced security/DTD license from the menu (**EC-AS** or **EC-UL**).
Reassign Appliance to Account	This option appears if you have registered a secondary Cloud Portal account. It allows you to assign appliances to either of your registered Cloud Portal accounts.

4. To revoke a license from an appliance, select the **Revoke** check box for the license type (EC, WAN Optimization, or Advanced Security).

   NOTE: If you revoke an EC license from an appliance, any WAN Optimization, advanced licenses, and legacy feature licenses are also revoked from that appliance.

   NOTE: You must revoke the license from an appliance before you can RMA it. For more information on how to RMA an appliance, see RMA Wizard.
5. Click **Apply**.

EC-Metered Licenses

Configuration > *Overlays & Security* > *Licensing* > *Licenses*

To filter the list, click one of the following buttons:

Button	Description
EC-Metered License	Display the EC-metered licenses for all appliances selected in the appliance tree. To filter the list, click one of the following buttons: **All** – Display all appliances. **WAN Optimization** – Display appliances with WAN Optimization licenses granted. **Feature license** – Display appliances with this feature license granted.
Bandwidth Usage Report	Display the bandwidth usage report for all appliances selected in the appliance tree. To aggregate the usage report, click **Summary**, **Appliance**, or **Daily**, and then select a month and year.
Feature License Usage Report	Display the feature license usage report for all appliances selected in the appliance tree. To aggregate the usage report, click **Summary**, **Appliance**, or **Daily**, and then select a month and year.

NOTE: EdgeConnect stops passing traffic when a license expires.

Assign a License to an Appliance

1. In the appliance tree, select one or more appliances to display in the table.
2. Do one of the following:
   - To assign licenses to one appliance, click the **Edit** icon next to that appliance.
   - To assign licenses in bulk (to all appliances in the table), click **Assign Licenses to Appliances**.

     NOTE: To assign licenses in bulk, all appliances must be on the same software version.

   The **Assign Licenses to Appliances** dialog box opens.
3. Complete the following elements as needed:

Field	Description
EC	Select the **Add/Replace** check box to apply the EC-metered license.
WAN Optimization	Select the **Add/Replace** check box, and then enter the amount of WAN Optimization to apply to the EC.

Field	Description
Feature licenses	1. To add a feature license, select the **Add/Replace** check box. 2. If required, select a license option from the list and specify a quantity, such as amount of bandwidth.

4. To revoke a license or WAN Optimization, select the **Revoke** check box next to the license or WAN Optimization you want to revoke.

   NOTE: If you revoke an EC license from an appliance, Silver Peak will revoke the WAN Optimization license and all feature licenses from that appliance.

   NOTE: You must revoke the license from an appliance before you can RMA it. For more information on how to RMA an appliance, see RMA Wizard.

5. Click **Apply**.

Bandwidth Usage Report

This page lists the maximum outbound bandwidth usage, maximum inbound bandwidth usage, and WAN Optimization bandwidth for the account.

To aggregate the usage report, click **Summary**, **Appliance**, or **Daily**, and then select a month and year.

Feature License Usage Report

This page lists the feature license usage report for the account.

To aggregate the usage report, click **Summary**, **Appliance**, or **Daily**, and then select a month and year.

EC-Metered Licenses

Configuration > *Overlays & Security* > *Licensing* > *Licenses*

To filter the list, click one of the following buttons:

Button	Description
EC-Metered License	Display the EC-metered licenses for all appliances selected in the appliance tree. To filter the list, click one of the following buttons: **All** – Display all appliances. **Boost** – Display appliances with Boost licenses granted. **Feature license** – Display appliances with this feature license granted.
Bandwidth Usage Report	Display the bandwidth usage report for all appliances selected in the appliance tree. To aggregate the usage report, click **Summary**, **Appliance**, or **Daily**, and then select a month and year.
Feature License Usage Report	Display the feature license usage report for all appliances selected in the appliance tree. To aggregate the usage report, click **Summary**, **Appliance**, or **Daily**, and then select a month and year.

- This page lists the appliance model, serial number, appliance name, and feature licenses for the appliances selected in the appliance tree.
- You can add, edit, or revoke EdgeConnect (EC) licenses from an appliance.

NOTE: EdgeConnect stops passing traffic when a license expires.

Assign a License to an Appliance

1. In the appliance tree, select one or more appliances to display in the table.
2. Do one of the following:
   - To assign licenses to one appliance, click the **Edit** icon next to that appliance.
   - To assign licenses in bulk (to all appliances in the table), click **Assign Licenses to Appliances**.

     NOTE: To assign licenses in bulk, all appliances must be on the same software version.

   The **Assign Licenses to Appliances** dialog box opens.
3. Complete the following elements as needed:

Field	Description
EC	Select the **Add/Replace** check box to apply the EC-metered license.

Field	Description
Boost	Select the **Add/Replace** check box, and then enter the amount of Boost to apply to the EC.
Feature licenses	1. To add a feature license, select the **Add/Replace** check box. 2. If required, select a license option from the list and specify a quantity, such as amount of bandwidth.

4. To revoke a license or Boost, select the **Revoke** check box next to the license or Boost you want to revoke.

   NOTE: If you revoke an EC license from an appliance, Silver Peak will revoke the Boost license and all feature licenses from that appliance.

   NOTE: You must revoke the license from an appliance before you can RMA it. For more information on how to RMA an appliance, see RMA Wizard.

5. Click **Apply**.

Bandwidth Usage Report

This page lists the maximum outbound bandwidth usage, maximum inbound bandwidth usage, and Boost bandwidth for the account.

To aggregate the usage report, click **Summary**, **Appliance**, or **Daily**, and then select a month and year.

Feature License Usage Report

This page lists the feature license usage report for the account.

To aggregate the usage report, click **Summary**, **Appliance**, or **Daily**, and then select a month and year.

Cloud Portal

Configuration > *Overlays & Security* > *Licensing* > *Cloud Portal*

Orchestrator > *Orchestrator Server* > *Licensing* > *Cloud Portal*

The Cloud Portal dialog box is used to register cloud-based features and services, such as SaaS optimization and EdgeConnect.

NOTE: Orchestrator 9.5.2 and later support IPv6 and IPv4. To support IPv6, the Cloud Portal URL changed to portal2.silverpeak.cloud. All EdgeConnect appliances on the network must be at least on ECOS 9.5.2.x or a later release to use this option. If you have EdgeConnect appliances running on releases earlier than ECOS 9.5.2.x or you have an environment with a mix of releases (some appliances are on ECOS 9.5.2.x and later releases and some are on releases earlier than ECOS 9.5.2.x), you must use portal.silverpeak.cloud. See the Orchestrator 9.5.2 release notes for more details.

NOTE: If you enabled Air-Gap mode, you cannot access the Cloud Portal. You will log in to the Air-Gap Portal to generate a new account key, which you will provide on this dialog box. Follow the instructions for enabling Air-Gap and accessing the Air-Gap Portal here: HPE Aruba Networking Orchestrator Air-Gap User Guide.

- When you purchase one of these services, an Account Name and instructions to obtain your Account Key are sent to you. You will use these to register your appliances.
- Use of these services requires that your appliances can access the Cloud Portal via the Internet.
- You can require hardware appliances to be provisioned with the Account Name and Account Key in order to be discovered.
- You can register a secondary account using the **Secondary Accounts** button.
- If you subscribed to an AAS License, the license type is displayed. The AAS license subscription determines which features are available in Orchestrator. The license token is refreshed from Cloud Portal every 24 hours.
- The Orchestrator UUID (universally unique identifier) is available under the Registration heading on the Cloud Portal dialog box.

Secondary Accounts

Secondary accounts allow you to manage multiple license end dates across a single SD-WAN, and can be used in the following situations:

- You want to add an account that has a different end date from the existing primary account in your Orchestrator.
- You want to merge accounts from one Orchestrator to another Orchestrator and have one SD-WAN fabric to manage.

NOTES:
- Hardware assets must co-reside in the same secondary account where the license resides.
- It is recommended to co-terminate licenses when possible for ease of management.

The primary account is where Orchestrator resides, and secondary accounts are associated to Orchestrator through the registration process. Hardware must reside in the same account as the software licenses associated with that hardware. You can register any number of secondary accounts under an Orchestrator supporting the multi-account feature.

As with a primary account, an Account Name and instructions to obtain your Account Key are sent to you. You will use these to register your appliances to the secondary account. To add a secondary account click **Secondary Accounts**, then click **+Add**, enter the account information, and click **Add**. To register the added secondary account, click **Save**.

Any secondary account changes, such as adding, updating, or deleting accounts, will trigger the registration process for all secondary accounts, even for those not modified directly in Orchestrator. This will be evident when a spinning icon is displayed next to “Registered” after saving any changes. After the registration process is complete, the spinning icon will be replaced with registration information that Cloud Portal returns to Orchestrator.

Configuration > Networking

The options under **Configuration > Networking** focus on configuring components of your network, such as deployments, interfaces, loopback, virtual tunnel interfaces (VTIs), and DHCP. Other options are related to configuring routes and tunnels.

Starting in Orchestrator 9.7.0, all DHCP configuration and monitoring features are consolidated in the DHCP Manager.

Deployment Tab

Configuration > Networking > Deployment

This tab provides summary and detailed views of the selected appliance's deployment settings.

To change an appliance's deployment settings, click the **Edit** icon next to the name of the desired appliance.

The following table describes the fields on the Summary view of this tab.

Field	Description
Appliance	Name of the deployed appliance.
HA	Name of the appliance with which this appliance is paired for EdgeHA.
Mode	Indicates the deployment mode for the appliance: **Inline Router:** Uses separate LAN and WAN interfaces to route data traffic. **Bridge:** Uses a virtual interface, **bvi**, created by binding the WAN and LAN interfaces. **Server:** Both management and data traffic use the mgmt0 interface.
Outbound Bandwidth	Deployment's total outbound bandwidth in Kbps.
Inbound Bandwidth	Deployment's total inbound bandwidth in Kbps.
WAN Labels Used	Identify the service, such as **MPLS** or **Internet**.
LAN Labels Used	Identify the data, such as **data**, **VoIP**, or **replication**.
Segment	Names of the segments used for this appliance deployment.
License	License tier granted for this appliance deployment.
Details	Select the **information** icon to view further deployment details of an appliance.

The following table describes the fields on the Details view of this tab.

Field	Description
Appliance	Name of the deployed appliance.
Interface	Name of the LAN or WAN interface. There are three types of interfaces: **Physical interface:** The physical or virtual NIC on the appliance (for example, wan0 and lan0) **Sub-interface:** The logical IP interface associated with any physical interface via a VLAN ID (for example, wan0.100 and lan0.200) **IP alias interface:** The logical IP interface associated with any physical interface or sub-interface (for example, wan0:1, lan0:1, and lan0.200:1).
Label	Label mapped to the interface. LAN labels refer to traffic type, such as VoIP, data, or replication. WAN labels refer to the service or connection type, such as MPLS, internet, or Verizon.
Zone	Firewall zone applied to the interface.
Segment	Name of the segment used for this interface.
IP/Mask	Interface's IP address and subnet mask.
WAN/LAN Side	Indicates that the interface is WAN-side or LAN-side.
Next Hop	Deployment interface's next hop router address.
Public IP	Public IP address.
Inbound	Interface's inbound bandwidth in Kbps.
Outbound	Interface's outbound bandwidth in Kbps.
NAT	Indicates whether the appliance is behind a NAT-ed interface.
Firewall Mode	Indicates the firewall mode for the appliance's WAN-side interface: **Allow All** – Permits unrestricted communication. **Stateful** – Only allows communication from the LAN-side to the WAN-side. Used if the interface is behind the WAN edge router. **Stateful+SNAT** – Applies Source NAT to outgoing traffic. Used if the interface is directly connected to the Internet. **Harden** – For traffic inbound from the WAN, the appliance accepts __*only__* IPSec tunnel packets that terminate on an EdgeConnect appliance. For traffic outbound to the WAN, the appliance __*only__* allows IPSec tunnel packets and management traffic that terminate on an EdgeConnect appliance.
DHCP	Indicates whether the interface's IP address is obtained from the DHCP server.

Field	Description
HA Interface	Indicates whether the interface is part of an EdgeHA link.
License	License tier granted for this appliance deployment.
Comment	Additional information for this deployment interface.

Deployment Dialog Box

The three deployment modes are **Bridge**, **Router**, and **Server**.

WARNING: ALWAYS use Router mode unless you have a legacy, WAN Optimization–specific use case and are well-acquainted with the requirements of Bridge or Server mode deployments.

NOTE: Customers with a large number of staged appliances offline but not in maintenance mode experience Orchestrator slowness when saving new appliance deployments or making other changes to appliance configurations from Orchestrator.

Enable EdgeHA

EdgeHA mode is a high availability cluster configuration that provides appliance redundancy by pairing two EdgeConnect devices together.

When you configure two EdgeConnect appliances in EdgeHA mode, the resilient cluster acts as a single logical system for orchestrated WAN functions. It extends the robust SD-WAN multipathing capabilities, such as Business Intent Overlays, seamlessly across the two devices as though they were one entity.

With EdgeHA mode, a WAN uplink is physically plugged into a single one of the EdgeConnect appliances but is available to both in the cluster. For WAN connections that perform NAT (for example, a consumer-grade Broadband Internet connection), it means that only a single Public IP needs to be provisioned in order for both EdgeConnect devices in the EdgeHA cluster to be able to build Business Intent Overlays using that transport resource. The same is true for orchestrated tunnels to third-party cloud services, such as Zscaler and AWS Transit Gateway.

NOTE: EdgeHA mode provides clustering for WAN-side functions only. You must select and configure an appropriate LAN-side redundancy mechanism for a given business location. Available options are VRRP+IP SLA, BGP, and OSPF.

To enable EdgeHA:

1. Select the **EdgeHA** check box.
2. Configure the interfaces (LAN-side and WAN-side) on both EdgeConnect devices to reflect the WAN connections that are plugged into each one of the respective appliances.

   NOTE: Both EdgeConnect devices will be able to leverage all WAN connections regardless of which chassis they are physically plugged into. It is, however, important to match the interface configuration displayed on the Deployment dialog box to the actual chassis the WAN connection is physically and directly connected to.

3. Select the physical ports on the respective EdgeConnect appliances that you will connect to each other using an Ethernet cable (RJ-45 twisted pair or SR optical fiber).

   NOTE: You can choose any LAN or WAN port combination for this HA Link that is available on the respective EdgeConnect chassis. You must match the media type and speed for both ends of the HA link. (For example, 1 Gigabit-Ethernet RJ-45 to RJ-45 or 10 Gigabit-Ethernet multimode fiber LC-connector-to-LC-connector). Also, note that you cannot use MGMT ports for the HA Link; only LAN or WAN ports.

IPSec over UDP Tunnel Configuration

For both EdgeConnect appliances in a high availability cluster to be able to share a common transport connection, you must set the tunnel type to IPSec over UDP mode. This is the default tunnel mode for all deployments running ECOS 8.1.6/Orchestrator 8.2 or later.

NOTE: For SD-WAN fabrics upgraded from earlier releases, see Tunnel Settings in Orchestrator (**Orchestrator > Orchestrator Server > Tools > Tunnel Settings**) to change to IPSec over UDP mode.

You must configure the same site name for both appliances in the EdgeHA pair so that Orchestrator assigns a unique IPSec UPD port number for each appliance.

LAN-side High Availability

Typically, in a branch site deployment, you will choose to configure the cluster with VRRP+IP SLA to modify priority and subnet sharing metrics based on VRRP and WAN interface status. For more advanced deployments with Layer 3 routers or switching on the LAN side, BGP or OSPF can be configured.

LAN-side Monitoring

The IP SLA feature should be configured to monitor the LAN-side VRRP state in order to automatically disable subnet sharing from that appliance in the case of a LAN link failure.

For more information, refer to the IP SLA configuration guide.

Map Labels to Interfaces

- On the **LAN** side, labels are optional. You can use them as match criteria for Business Intent Overlay ACLs, such as __*data__*, __*VoIP__*, or __*replication__*.
- On the **WAN** side, labels identify the link type, such as __*MPLS__* or __*Internet__*. These labels are mandatory. Orchestrator uses them to build Business Intent Overlay policies.
- To create or manage a global pool of labels, either:
  - Navigate to **Configuration > Overlays & Security > Deployment Profiles**, click the **Edit** icon next to Label, and make the appropriate changes, or
  - Navigate to **Configuration > Overlays & Security > Interface Labels**) and make the appropriate changes.
- The change you make to a label propagates automatically. For example, it renames tunnels that use that labeled interface.

LAN and WAN Interface Configuration: Segment, Firewall Zone, and Role

EdgeConnect Segmentation (VRF) provides orchestrated Layer 3 segmentation, Zone-Based Firewall, Role, and an end-to-end Intrusion Detection System (IDS) across the SD-WAN fabric.

Segment and FW Zone policies are global in scope. They are managed on the **Configuration > Networking > Routing > Routing Segmentation (VRF)** tab.

Roles are managed on the **Orchestrator > Orchestrator Server > Users & Administration > Role Based Access Control** tab.

Use the Deployment dialog box to assign segments, zones, and roles to LAN interfaces. Zones and roles can be assigned to WAN interfaces (but not segments). Segment, FW Zone, and Role fields are initially set to the system-generated default. You can select a different segment, firewall zone, and role from the drop-down lists.

You can specify an Interface Role for a flow when other identity pipelines, such as Radius snooping, NACD, and VXLAN/GPID, are not present.

For example, a VLAN for unauthenticated IoT devices could have the Interface Role set to "IoT." This Interface Role is transported across the fabric and is used to reconstruct the VXLAN GPID role on the other end of the fabric.

NOTE: The segment for WAN interfaces cannot be changed.

LAN–side Configuration: DHCP and Router Advertisements

- By default, the LAN IP does not act as a DHCP Server. Based on your configuration, you can set the interface to act as a DHCP relay server when the appliance is in Router mode.
- The global defaults are set in **Configuration > Networking > DHCP > DHCP Server Defaults** and pre-populate this page. The other choices are **No DHCP/No RA** and having the appliance act as a **DHCP/BOOTP Relay**.
- To customize an individual interface on the Deployment screen, click the DHCP-related link under the IP/Mask field. The DHCP Settings / Router Advertisements dialog box opens.
- From **Configuration > Networking > DHCP Manager**, you can open the same DHCP Settings / Router Advertisements dialog box when exactly one appliance is selected in the appliance tree. Click **DHCP Settings**, then select the deployed LAN-side interface.
- Only one DHCP mode is allowed per interface without any static alias interfaces on the WAN-side.
- Before you can configure DHCP Relay, you must navigate to Management Services and select an interface for DHCP Relay. See Management Services for more information.

If the LAN interface has an IPv4 IP address, click **V4** to display the DHCP configuration settings. See V4.

If the LAN interface has an IPv6 IP address, click **V6** to display the Router Advertisement settings. See V6.

V4

The following tables describe the various DHCP settings you can configure for LAN interfaces that have IPv4 IP addresses.

DHCP Server

Setting	Description
Subnet Mask	Mask that specifies the default number of IP addresses reserved for any subnet. For example, entering **24** reserves 256 IP addresses.
IP Range	You can designate one or more IP address ranges available for use. Specify **Start IP** and **End IP** addresses. To add another IP address range, click **Add**. **IMPORTANT:** Multiple IP ranges cannot overlap.
Default lease,Maximum lease	Specify, in seconds, how long an interface can keep a DHCP–assigned IP address.
Gateway IP	Specifies the IP address for the gateway to use.
DNS server(s)	Specifies the associated Domain Name System servers.
NTP server(s)	Specifies the associated Network Time Protocol servers.
NetBIOS name server(s)	Used for Windows (SMB) type sharing and messaging. It resolves the names when you are mapping a drive or connecting to a printer.
NetBIOS node type	**NetBIOS node type** of a networked computer relates to how it resolves NetBIOS names to IP addresses. There are four node types: **B**-node = 0x01 Broadcast **P**-node = 0x02 Peer (WINS only) **M**-node = 0x04 Mixed (broadcast, then WINS) **H**-node = 0x08 Hybrid (WINS, then broadcast)
DHCP failover	Enables DHCP failover. To set it up, click the **Failover Settings** link.

DHCP/BOOTP Relay

Setting	Description
Destination DHCP/BOOTP Server	IP address of the DHCP server assigning the IP addresses. This setting applies to the local interface only.

Setting	Description
Common DHCP server for all segments	Select this check box to set the default values for all segments. **HINT:** You can reset the defaults in Management Services by setting the DHCP Relay interface to “any” and then selecting an interface label again. However, this might impact service. Or, you can manually reset the defaults by selecting the following values: Option 82 = enabled, Option 82 Policy = append, and select the following sub options: 1, 5, 10, 11, 151, and 152.
Distinct DHCP server per segment	Select this option to override the DHCP relay configuration set in the Manages Services tab with the settings you select in this dialog box.
Enable Option 82	When selected, inserts additional information into the packet header to identify the client’s point of attachment. This setting applies to all LAN-side interfaces on this appliance. **IMPORTANT:** Changing this setting will modify Option 82 settings on all LAN-side interfaces that are enabled as DHCP Relay.
Option 82 Policy	Tells the relay what to do with the hex string it receives. The choices are **append**, **replace**, **forward**, and **discard**. This setting applies to all LAN-side interfaces on this appliance. **IMPORTANT:** Changing this setting will modify Option 82 settings on all LAN-side interfaces that are enabled as DHCP Relay.

Setting	Description
Sub Options	Select one or more of the following: **1** - Agent Circuit ID: Provides information about the interface or circuit through which the DHCP request was received. **5** - Link selection: Specifies the IP address used by the DHCP server to determine the appropriate subnet for addressing the DHCP client. **10** - Client Unicast/Broadcast Indication flag: Indicates whether the DHCP relay received the client packet as a unicast or broadcast packet. **11** - Server ID Override: Allows the DHCP relay agent to act as a proxy for the DHCP server to process unicast lease renewals. **150** - Link selection (Cisco proprietary): Provides information about a segment or VPN that is necessary to allocate an address to a DHCP client on that segment. **151** - VRF name/VPN ID **152** - VRF name/VPN ID Control Sub-Option OR Server ID Override (Cisco proprietary): Indicates whether the DHCP server supports sub option 151 (VRF Name/VPN ID). If this option is present in the reply from the server, the server does not support option 151.

V6

The following table describe the various router advertisement settings you can configure for LAN interfaces that have IPv6 IP addresses. The LAN clients can use these options to autoconfigure IPv6 addresses and to learn default gateway addresses.

NOTE: DHCP for IPv6 is not supported.

Setting	Description
Enable Router Advertisements	Specifies whether the router should send RA messages.
Managed Flag	Select this option to instruct IPv6 hosts to use DHCPv6 to obtain their IPv6 addresses in addition to any other configuration information.

Setting	Description
Other Flag	Select this option to instruct IPv6 hosts to use DHCPv6 to obtain additional configuration information, such as DNS server addresses and other network parameters.
Link MTU	Set the maximum transmission unit (MTU) size that can be transmitted without fragmentation. This helps ensure that all hosts on the network use the same MTU, avoiding issues related to packet fragmentation and reassembly.
Max Interval	Specify the maximum interval in seconds between unsolicited RA messages. This helps to control the frequency of RA messages.
Min Interval	Specify the minimum interval in seconds between unsolicited RA messages. This helps to control the frequency of RA messages.
Current Hop Limit	Set the default hop limit for IPv6 packets sent by hosts on the network. Hosts use this value to configure their own hop limit for outgoing packets.
Default Router Preference	Select **High**, **Medium**, or **Low** to set the preference level of the router for use as a default router. Hosts use this value to prioritize multiple routers on the same link.
Default Router Lifetime	Specify the lifetime in seconds of the default route that is advertised by the router. The hosts use this value to determine how long the router should be used as the default gateway.
Reachable Time	Specify the time in milliseconds that an IPv6 host considers a neighbor reachable after receiving a confirmation. This value maintains accurate and timely reachability information in the neighbor cache.
Retrans Timer	Specify the time in milliseconds between retransmissions of neighbor solicitation messages. This value reduces the frequency of retries when attempting to discover or confirm the reachability of neighbors on the network.

Add a Router Advertisement Prefix

Click **Add** and complete the following fields.

Considerations

- RA can be configured only on LAN side interfaces.
- Users can configure RA only on IPv6 configured interfaces.
- DHCPv4 server and RA cannot be configured on the same interface at the same time.
- DHCPv4 relay and RA cannot be configured on the same interface at the same time.

- RA, DHCPv4 Server, and DHCPv4 Relay cannot be enabled if there is an alias interface configured for the main/primary interface.
- A maximum of 10 prefixes can be configured in the RA configurations per interface.

Setting	Description
Prefix-id	The ID assigned to the prefix.
Prefix	The IPv6 prefix to advertise to hosts on the network. Hosts use this prefix to configure their IPv6 addresses and determine the network portion of the IP addresses.
Autonomous flag	Select whether the prefix can be used by hosts for SLAAC. When set to true, hosts can use the prefix to generate their own IPv6 addresses.
Onlink flag	Specifies whether the prefix is on-link, which affects how hosts handle routing for addresses within the prefix. If set to true, hosts assume that addresses within the prefix can be reached directly on the local network segment.
Valid Lifetime	Specify the duration in seconds for which the advertised prefix is valid.
Preferred Lifetime	Specify the duration in seconds (relative to the time the packet is sent) that addresses generated from the prefix via stateless address auto-configuration remain preferred.

WAN–side Configuration

Interface mode: Orchestrator release 9.5.2 and later supports IPv4 and IPv6 to meet the increasing demand for IP addresses. This feature also allows you to deploy appliances that support both IPv4 and IPv6 for a dual stack solution.

NOTE: The WAN interfaces of factory deployed appliances with ECOS version 9.5.2 and later support both IPv4 and IPv6 addresses for Zero Touch Provisioning (ZTP). Appliances with previous versions of ECOS only support IPv4.

Select one of the following options for each WAN interface:

- **Static** – Configure the IPv4 or IPv6 address manually on the interface.
- **DHCPv4** – The IPv4 address is configured dynamically by DHCPv4 on the interface.
- **DHCPv6** – The IPv6 address is configured dynamically by DHCPv6 on the interface.
- **SLAAC** – The IPv6 address is configured dynamically by SLAAC (State Less Address Auto Configuration) on the interface.
- **DHCPv4 + DHCPv6** – This is a dual-stack mode that supports the simultaneous use of IPv4 and IPv6 addresses. The IPv6 address is configured dynamically by DHCP on the primary interface and the IPv4 address is configured dynamically by DHCP on the alias interface. All properties except Label are inherited from the primary interface and cannot be edited.

- **DHCPv4 + SLAAC** – This is a dual-stack mode that supports the simultaneous use of IPv4 and IPv6 addresses. The IPv6 address is configured dynamically by SLAAC on the primary interface and the IPv4 address is configured dynamically by DHCP on the alias interface. All properties except Label are inherited from the primary interface and cannot be edited.

NOTE: Only one DHCP configuration option is allowed for each physical interface.

Firewall Zone: Zone-based firewall policies are configured globally on the Orchestrator. A zone is applied to an **Interface**. By default, traffic is allowed between interfaces labeled with the same zone. Any traffic between interfaces with different zones is dropped. You can create exception rules (Security Policies) to allow traffic between interfaces with different zones.

Firewall Mode: Four options are available at each WAN interface:

- **Allow All** permits unrestricted communication.

  WARNING: Use this option with extreme caution and only if the interface is behind a WAN edge firewall.

- **Stateful** __*only__* allows communication from the LAN-side to the WAN-side.

  Use this option if the interface is behind a WAN edge router.

- **Stateful with SNAT** applies Source NAT to outgoing traffic.

  Use this option if the interface is connected directly to the Internet and you want to enable local internet breakout.

- **Harden**

  - For traffic inbound from the WAN, the appliance accepts __*only__* IPSec tunnel packets that terminate on an EdgeConnect appliance.
  - For traffic outbound to the WAN, the appliance __*only__* allows IPSec tunnel packets and management traffic that terminate on an EdgeConnect appliance.

VLAN Settings: With Orchestrator release 9.5.2 and later, you can assign multiple IP aliases to the same VLAN interface ID. Prior to Orchestrator release 9.5.2, Orchestrator only allowed you to assign one unique IP address to a VLAN ID. This new feature also allows you to assign VLAN IDs to interfaces in a dual stack solution with SLAAC.

Sub-interfaces behave the same as physical interfaces.

NOTE: If you modify the IP address or subnet mask of a sub-interface, all sub-interfaces and IP aliases with the same ID will be deleted and added back. The label, segment, and zone will remain unchanged. There will be a brief outage of all IPs while the interfaces are deleted and added back.

IP aliases and sub-interfaces will appear in all tabs where interfaces are listed or selectable (for example, the Dynamic table on the Interfaces tab). The first sub-interface will be the main interface. Additional IP aliases will be assigned with incremented interface numbers. For example, the main sub-interface will be wan0.120 and additional IP aliases will be assigned wan0.120:1, wan0.120:2, and so on.

NAT Settings: To change the NAT setting, click the NAT-related link under the Next Hop field on the WAN side. The NAT Settings dialog box opens.

Select one of the following options:

- If the appliance is behind a NAT-ed interface, select **NAT**.
- If the appliance is not behind a NAT-ed interface, select **Not behind a NAT**.
- To assign a destination IP address for tunnels being built from the network to this WAN interface, select the last option and enter the IP address.

Shaping: You can limit bandwidth selectively on each WAN interface.

- **Total Outbound** bandwidth is always enabled.
- **Total Inbound** bandwidth is disabled by default and can be enabled by clicking the Edit icon for a specific interface on the Shaper tab.

EdgeConnect Licensing: Only visible on EdgeConnect appliances.

- You can change the bandwidth allotted for this appliance by selecting the appropriate option from the **EC** drop-down list. Your options are based on the licensing you have purchased.
- If you have purchased a pool of WAN Optimization for your network, you can allocate a portion of it in the **WAN Opt** field on the Deployment dialog box. You can also direct allocations to specific types of traffic in the Business Intent Overlays.
- To view the licensing and distribution of EdgeConnect and WAN Optimization bandwidth for your appliances, navigate to the **Configuration > Overlays & Security > Licensing > Licenses** tab.

BONDING

- EdgeConnect supports etherchannel bonding of multiple physical interfaces of the same media type into a single virtual interface. For example, **wan0** plus **wan1** bond to form **bwan0**. This increases throughput on a very high-end appliance and/or provides interface-level redundancy.
- For bonding on a virtual appliance, you would need to configure the host instead of the appliance. For example, on a VMware ESXi host, you would configure NIC teaming to get the equivalent of etherchannel bonding.
- Whether you use a physical or a virtual appliance, etherchannel must also be configured on the directly connected switch/router. Refer to the switch or router user documentation for configuring interface bonding.

Interfaces Tab

Configuration > Networking > Interfaces

The Interfaces tab lists the interfaces for appliances selected in the appliance tree. Fields available in the table depends on the selected filter button that is located immediately above the table.

The All button displays all hardware and dynamic interfaces for the selected appliances.

Descriptions of the fields on this tab (All, Hardware, Dynamic, USB LTE) follow:

Field	Description
Appliance	Name of the appliance for the interface.
Name	Name of the interface.
Status	Status of the interface (up or down).
LACP Status	Link aggregation status of the interface (up or down). This status applies only to bonded interfaces (blan0, blan1, bwan0, and bwan1). This field is displayed only if the channel group is in LACP mode.
IP Address/Mask	IP address for the interface.
Public IP	Public IP address for the interface.
Segment	Name of the configured segment being used.
DHCP	Indicates whether this interface's IP address is obtained from the DHCP server. Displays as **Yes**, **No**, **No data** (not configured), or **Invalid data** (error condition).
Speed	Current interface speed state and setting.
Duplex	Current interface duplex state and setting.
MTU	Maximum number of packets being transmitted.
MAC Address	MAC address applied to the interface.
SNMP IfIndex	Index number of the network interface.
Type	Port type (RJ45 or SFP). Drop-down box appears on interfaces with Combo ports. RJ45 is the default for Combo port. **NOTE:** Combo ports apply to hardware interfaces only. Combo ports are not applicable to VLANs (virtual interfaces), so in those cases, this field appears blank.

- Best practice is to assign static IP addresses to management interfaces to preserve their reachability.
- **Duplex** should never display as **half duplex** after auto-negotiation. If it does, performance issues and dropped connections will occur on the appliance. To resolve, check the cabling on the appliance and the ports on the adjacent switch or router.
- To directly change interface parameters for a particular appliance, click the corresponding edit icon, which opens the Interfaces dialog box for the appliance.
- To change the IP address for a **lan** or **wan** interface, either use the Appliance Manager **Configuration > System & Networking > Deployment** page or the CLI (Command Line Interface).
- To change the IP address for **mgmt0**, either use the Appliance Manager **Administration > Basic Settings > Hostname/IP** page or the CLI.

Descriptions of the fields on this tab (PoE) follow:

Field	Description
Appliance	Name of the appliance for the interface.
Name	Name of the interface.
PoE Config	Power mode of the PoE interface.
PoE Status	Enable status of the interface.
Priority	Priority status of the interface (dynamic or static).
PoE Class	Power class of the interface.
Voltage (V)	Total voltage available on the appliance.
Current (A)	Total current available on the appliance.
Power (W)	Total power available on the appliance.

Terminology

Interface	Description
blan	Bonded LAN interfaces (as in **lan0** + **lan1**).
bvi0	Bridge Virtual Interface. When the appliance is deployed in-line (Bridge mode), it is the routed interface that represents the bridging of **wan0** and **lan0**.
bwan	Bonded WAN interfaces (as in **wan0** + **wan1**).
tlan	10-Gbps fiber LAN interface.
twan	10-Gbps fiber WAN interface.

PoE (Power Over Ethernet)

This table appears only for interfaces that support PoE.

Field	Description
Edit	Click to open the PoE Configuation Dialog box.
Name	Name of the interface.
PoE Config	Enabled status of the interface.
PoE Status	Detection status. Settings include Searching, Delivering, or Disabled.
Priority	Priority status. Priority determines the interfaces that receive power when power is over-subscribed on the appliance. Setting of 0 denotes highest priority.
PoE Class	Power class of the interface.
Voltage (V)	Total voltage available on the appliance.

Field	Description
Current (A)	Total current available on the appliance.
Power (W)	Total power available on the appliance.

Interfaces Dialog Box

The dialog box that appears depends on the selected filter button located immediately above the table. - When all, Hardware, Dynamic, or USB LTE is selected, the Interfaces dialog box appears. Use this dialog box to change configurations for the specified interface.

- When PoE is selected, the PoE Configuration dialog box appears. Use this dialog box to modify the configuration interfaces on the specified appliance.

The All Interfaces button displays all interfaces for the appliance, including both assigned and unassigned hardware interfaces. MAC addresses indicate assigned interfaces.

Descriptions of the fields on this dialog box follow:

Hardware

Field	Description
Name	Name of the interface.
Admin	Admin status of the interface (up or down). Click this field to change the value.
Status	Status of the interface (up or down).
IP Address/Mask	IP address for the interface. If the address is blue, you can click it to open the Deployment dialog box, from which you can change IP addresses/masks.
Public IP	Public IP address for the interface.
Segment	Name of the configured segment being used.
Speed (Mbps) / Duplex	Current interface speed and duplex settings. **auto** means auto-negotiation, which is the process by which terminating devices automatically negotiate for maximum bandwidth.
State	Current interface speed and duplex states.
MTU	Maximum number of packets being transmitted. Click this field to change the value.
MAC	MAC address applied to the interface. To unassign the MAC address, click the field and select **Unassigned**.
SNMP IfIndex	Index number of the network interface.
Type	Port type (RJ45 or SFP). Drop-down box appears on interfaces with Combo ports. RJ45 is the default for Combo port.

Dynamic

Field	Description
Name	Name of the interface.
Status	Status of the interface (up or down).
LACP Status	Link aggregation status of the interface (up or down). This status applies only to bonded interfaces (blan0, blan1, bwan0, and bwan1). This field is displayed only if the channel group is in LACP mode.
IP Address/Mask	IP address for the interface. If the address is blue, you can click it to open the Deployment dialog box, from which you can change IP addresses/masks.
Segment	Name of the configured segment being used.
MTU	Maximum number of packets being transmitted.
SNMP IfIndex	Index number of the network interface.

PoE Configuration Dialog Box

The PoE Configuration dialog box modifies the configuration on interfaces located on the specified appliance. Configuration parameters include:

Power Mode: Specifies power management and priority parameters for appliance PoE interfaces. The following mode options are available: - Dynamic, without pre-allocation and no priority - Static, no priority - Dynamic with pre-allocation and no priority - Dynamic with pre-allocation and priority - Static with priority - Dynamic without pre-allocation and with priority

Individual Power Mode settings include:

Dynamic Mode: Power budget calculations are performed based on power that Powered Devices are currently consuming.

Static Mode: Power budget calculations are performed based on the class of the Powered Devices. Power consumption by Powered Devices is not considered.

Pre-Allocation: When connecting a new Powered Device, power budget calculations are based on the maximum power consumption of the class. If the sum of the current power allocated and the maximum power of the class is outside this budget, the device is not powered. When Pre-Allocation is disabled, the device is allowed to connect and make decisions based on consumption.

Port Limits: Individual ports can have maximum power draw limits set. Ports that exceed the power limit are be denied power.

Interface: The name of the interface modified by corresponding parameter settings.

Enable PoE: Specifies the PoE enable status of the correlated interface.

Automatic Bandwidth Detection

Starting with Orchestrator 9.7.0, the Interfaces tab includes support for manual bandwidth measurement on WAN interfaces through Automatic Bandwidth Detection. In this release, bandwidth testing is manual and on demand. Scheduled or automatic testing is not included.

Prerequisites

Before you can run a speed test on a WAN interface, ensure the following conditions are met:

- The WAN interface has a valid IP address configured on the Deployments tab (Configuration > Networking > Deployments).
- The WAN interface has internet access. Automatic Bandwidth Detection uses Cloudflare-based speed test services to measure bandwidth, and the test cannot complete without internet connectivity.

Running a Speed Test

Automatic Bandwidth Detection allows you to manually measure the available download and upload bandwidth of a WAN interface. This helps validate actual link performance against expected or provisioned circuit capacity.

Speed tests are supported only on WAN-side interfaces that have valid IP addresses configured on the Deployments tab. They are not available for LAN or non-WAN interfaces. Speed test controls appear in the Speed Status column only for supported WAN interfaces on appliances running a supported software version.

Speed test controls can be run in the Speed Test column when any filter other than PoE is selected, on the interface dialog box, and in the deployment workflow. During initial appliance installation, you can run a speed test from the Deployments tab to measure circuit capacity and use the results when configuring interface bandwidth values.

To run a speed test:

1. Navigate to **Configuration > Networking > Interfaces**.
2. Locate the WAN interface you want to test.
3. In the Speed Status column, click **Run speed test**.

   While the test is running, Orchestrator displays a spinner in the Speed Status column and disables the run button on all other interfaces. Only one speed test can run at a time.

   IMPORTANT: Speed tests leverage Cloudflare-based speed test services. Test traffic runs passthrough unshaped and is not subject to the EdgeConnect QoS engine, traffic shaper, or configured license bandwidth. Because the test runs wide open to measure actual WAN capacity, it can temporarily affect other user traffic on the link.

   This behavior is intentional. If shaping or license enforcement were applied, the test could report only the configured limit instead of the available bandwidth supplied by the WAN connection. Use the result to determine whether interface bandwidth settings or appliance licensing should be adjusted.

   A speed test usually completes within about a minute, depending on link conditions.

4. To stop a test before it completes, in the Speed Status column, click **Stop speed test**.

   When a test is stopped manually, the Speed Status column displays *Stopped*.

   If the test does not complete within 6 minutes, it times out and the Speed Status column displays *Failed*.

If a speed test does not complete successfully, the Speed Status column shows *Failed* in red, and Orchestrator displays an error message describing the failure.

NOTE: Speed tests are user-initiated and run on demand. Automatic or scheduled testing is not included in this release. Users with read-only RBAC permissions cannot initiate or stop speed tests.

Speed Test Results

When a speed test completes, the Speed Status column displays the latest measured download and upload bandwidth values inline. To view detailed results, hover over the interface row to display a tooltip with up to 5 of the most recent successful test results. Each entry in the tooltip shows the following:

Column	Description
Download	Measured download bandwidth (for example, 95.20 Mbps).
Upload	Measured upload bandwidth (for example, 48.50 Mbps).
Latency	Measured round-trip latency (for example, 12 ms).
Last Checked	Date and time the test completed.

These values represent the available bandwidth detected during the test interval and can be used to assess how closely link performance aligns with expected circuit capacity, configured bandwidth values, and appliance licensing. Speed test results persist across tab refreshes.

NOTE: A speed test reports an instantaneous, on-demand measurement of available bandwidth at the time of the test. It is not a continuously learned or averaged value.

Automatic Bandwidth Detection — Scope and Limitations

Automatic Bandwidth Detection Phase 1 includes the characteristics in the following table.

Characteristic	Details
Compatibility	Requires Orchestrator 9.7.0 or later and EdgeConnect OS 9.7.0.0 or later on the appliance. Supported only on WAN interfaces with internet access.

Characteristic	Details
User interface only	There are no CLI enhancements associated with this feature.
REST API	Automatic Bandwidth Detection is driven entirely from the Orchestrator UI. This release does not add any public, customer-facing REST APIs for the feature.
React-based UI implementation	The feature is implemented in the React.js Orchestrator UI. This release does not include support in the legacy UI for this feature.
Manual operation	The user explicitly initiates speed tests. This release does not include continuous or automated detection.
Concurrency	Only one speed test can run at a time.
Alerts and notifications	This release does not include alarms, notifications, or automated reporting based on speed test results.
Speed test service	Uses Cloudflare-based speed test services to measure WAN bandwidth.
Traffic treatment	Speed test traffic runs passthrough unshaped and is not subject to the EdgeConnect QoS engine, traffic shaper, or configured license bandwidth during the test.
UI locations	Speed test controls can be run in the Speed Test column when any filter other than PoE is selected, on the interface dialog box, and in the deployment workflow.

VRRP Tab

Configuration > Networking > VRRP

This tab summarizes the configuration and state for appliances deployed with Virtual Router Redundancy Protocol (VRRP).

VRRP specifies an election protocol that dynamically assigns responsibility for a virtual router to one of the VRRP routers on a LAN. The VRRP router controlling the IP addresses associated with a virtual router is called the Master and forwards packets sent to these IP addresses. The election process provides dynamic failover in the forwarding responsibility should the Master become unavailable. This allows end hosts to use any virtual router IP addresses on the LAN as the default first-hop router. The advantage gained from using VRRP is a higher availability default path without the need to configure dynamic routing protocols such as BGP or OSPF.

VRRP Configuration Considerations

When configuring VRRP, observe the following restrictions:

- If you set the VRRP virtual IP (VIP) to a subnet that is different than that of the LAN physical interfaces, do not use static routes on the LAN side.
- By default, EdgeHA operates within the IPv4 link-local address range. If you configure the LAN interfaces to use the same range, ensure that there are no duplicate IP addresses.
- If the LAN physical interfaces are set to the link-local subnet 169.254.0.0/16, make sure that this subnet is not shared via route-map filtering.
- DHCP server, DHCP relay, or other management services on the VRRP VIP with a different subnet are not supported.
- VRRP convergence is optimized when the Site/Cluster Tracker option is enabled.
- EdgeConnect appliances do not support the implementation of different max advertise intervals for different routers in the same VRRP group (IPv4 and IPv6). You must configure this interval to be the same on all routers in the same VRRP group, whether they are considered master or backup. When advertisement timers are different, devices can declare as master irrespective of their priority.

VRRP Edit Row

From the list of appliances, click the edit icon to display the **VRRP** screen. Click **Add VRRP** to add a VRRP instance. Use the information in the following table to assist you in configuring a VRRP. Click **Save** to deploy your VRRP configuration.

VRRP Settings

Field	Description
Group ID	The value is assigned to a group of routing devices. The group most commonly includes two appliances but depending on the deployment could contain one or more appliances and a router (or L3 switch), or more than two appliances. The valid range is 1 to 255.
Interface	Choose an interface that VRRP will use for peering from a list of configured system interfaces.
Version	Select the VRRP version that applies for your system: **2** – Supports only IPv4. **3** – Supports IPv4 and IPv6; does not support authentication strings.
State	The VRRP instance has three states: **Backup** – Instance is in VRRP backup state. **Init** – Instance is initializing, it is disabled, or the interface is down. **Master** – Instance is the current VRRP master.
Admin	Select **up** (enable) or **down** (disable).
Virtual IP	IP address of the VRRP instance. Configure the VRRP VIP on a different subnet than the LAN physical interfaces; this maximizes the number of available LAN IP addresses. VRRP instances can run between two or more appliances, or appliances and routers. The VRRP VIP subnet is shared with the peers.
Hold Down	The number of seconds a higher-priority backup router waits before preempting the primary router after it has just started up. It is best practice to configure a hold time so that routing protocols converge and tunnels come back up before preemption occurs. The default value is **60 sec** to account for the default *quiescent tunnel keep alive time*. The minimum value is 1 second.
Advertisement Timer	The time interval between sent advertisements. For version 2, the time is measured in seconds, and the default is **1 sec**. For version 3, the time is measured in centi-seconds, and the default is **1 centi-sec**. 1 centi-sec = 10ms.
Priority Config	The greater the number, the higher the priority. The appliance with the higher priority is the VRRP Master.
Priority State	The current VRRP priority, which can be influenced by IP SLA rules.

Field	Description
Preemption	Leave this selected/enabled so that after a failure, the appliance with the highest priority comes back online and again assumes primary responsibility.
Site/Cluster Tracking	Enabling this option optimizes traffic convergence, which is particularly useful for traffic originating from a hub and destined for an HA site where this VRRP option is enabled. If it is not enabled, convergence may take up to sixty seconds (the default reclassification interval). When enabled, traffic is steered to the new VRRP primary device without waiting for route updates or the reclassification interval. The VRRP convergence happens quickly.
Authentication String	Clear text password for authenticating VRRP version 2 group members. You cannot use an authentication string if you are using VRRP version 3.
Description	Free-form text field where you can enter a description of the VRRP instance.
Details	Click the info icon in this column to view the following details about the VRRP instance. **Master IP** – The interface or local IP address of the current VRRP Master. **Virtual MAC Address** – MAC Address that the VRRP instance is using. On a hardware appliance, this is in 00-00-5E-00-01-{VRID} format. On virtual appliances, the VRRP instance uses the MAC Address assigned to the interface (for example, the MAC address that the hypervisor assigned to **wan0**). **State Uptime** – Time elapsed since the VRRP instance entered the state it is in. **Master State Transitions** – Number of times the VRRP instance went from Master to Backup and vice versa. A high number of transitions indicates a problematic VRRP configuration or unstable network. In this case, check the configuration of all local appliances and routers, and then review the log files. **IP Address Owner** – An EdgeConnect appliance cannot use one of its own IP addresses as the VRRP IP, so this will always be **No**.
Segment	The name of the segment, if enabled.

DNS Proxy

Configuration > Networking > DNS Proxy

The DNS (Domain Name Server) Proxy stores public IP addresses with their associated domain name. By default, Server A is used primarily as a private DNS to backhaul traffic and Server B is used to match all other domains that are not included under Server A. Server B is also used for public (cloud services) to breakout traffic. Other deployment modes include AppExpress Only and Hybrid.

The following table describes fields on this tab.

Field	Description
Appliance	Name of the appliance associated with DNS proxy.
Segment	Name of the segment applied to your appliances, if enabled.
DNS Proxy Enabled	Whether the DNS Proxy is enabled. Select **True** or **False**.
Deployment Mode	**Default** – Server A is used primarily as a private DNS to backhaul traffic and Server B is used to match all other domains that are not included under Server A. Server B is also used for public (cloud services) to breakout traffic. All UDP-based queries are proxied. **AppExpress Only** — Only AppExpress domains are proxied. **NOTE:** To deploy AppExpress Only mode, you must push an empty DNS template to the appliances before you apply AppExpress groups to the appliances. See Templates Overview. **Hybrid** — All AppExpress domains are proxied and all other domains are subject to the default mode for Server A and Server B. **NOTE:** If "No data available" is displayed, DNS proxy is disabled. No DNS queries are proxied.
Interface	Name of the interface associated with the DNS proxy.
Server A Addresses	IP addresses of Server A.
Server A Domains	Domain addresses of Server A.
Server A Caching	Whether you configured the server to be cached.
Server B Addresses	IP addresses of Server B.
Server B Domains	Domain addresses of Server B.
Server B Caching	Whether you configured the server to be cached.

Configure DNS Proxy Policies

Complete the following steps to configure and define your DNS Proxy policies.

NOTE: This feature is only configurable if you have loopback interfaces configured.

1. Choose whether you want to enable the DNS Proxy by selecting **ON** or **OFF**.
2. Select the name of the loopback interface or the LAN-side label associated with your DNS proxy.
3. Enter the IP addresses for Server A in the **Server A Addresses** field.
4. Choose whether you want caching to be **ON** or **OFF**. If selected, the domain name to the IP address mapping is cached. By default, caching is **ON**.
5. Enter the domain names of the Server A for the above IP addresses.
6. Enter Server B IP addresses in the **Server B Addresses field**. Server B will be used if there are no matches to the Server A domains.

NOTE: You can **Clear DNS Cache**. This will erase the domain name to the IP address mapping you had cached for both Server A and B.

Branch NAT

Configuration > Networking > Branch NAT

Branch NAT enables multiple sites with overlapping IP addresses to connect to a single SD-WAN fabric.

NOTE: Branch NAT configuration is not available in Appliance Manager.

You can use Branch NAT features for a variety of overlapping IP use cases.

- The most common scenario is when all branches share the same IP configuration. You can use NAT to translate overlapping local IPs within a branch into globally scoped and routable IPs within the fabric.
- You can also use routing segmentation in combination with NAT within your branch locations.

NOTE: Typically, the NAT feature should not be used for Local Breakout situations. Local Breakout relies on the Stateful+SNAT feature to automatically SNAT into the IP address of the WAN-side interface.

You can configure SNAT (Source Network Address Translation), DNAT (Destination Network Address Translation), destination TCP, and UDP port translation rules that apply to all LAN to WAN traffic in ingress and egress directions. Configured rules apply to all traffic, including tunneled traffic, internet-bound traffic, and other passthrough traffic. The following address translation options are supported:

- 1:1 source and destination IP address translation
- 1:1 subnet to subnet source and destination IP address translation
- Many to one IP source address translation
- NAT pools for translated source IP address

NOTE: You can NAT to and from non-EdgeConnect IP addresses, but NAT functionality does not apply to the EdgeConnect local IP addresses. Also, be aware that NAT between segments is not supported. This is handled by the Inter-Segment SNAT Exceptions feature. To navigate to this feature, enter "Inter-Segment SNAT Exceptions" in the search field at the top of the Orchestrator application. For feature details, see Inter-Segment SNAT Exceptions.

On the Branch NAT tab, click **NAT Rules** to view NAT rules configured for your network. To view configured NAT pools, click **NAT Pools**.

You can select a segment from the **Segment** drop-down list to filter NAT rules in the table. This list includes all segments configured on the Routing Segmentation (VRF) tab. To display all NAT rules for any segment, select **All**. The Segment filter is available only if routing segmentation is enabled.

To save the contents of the currently displayed table to a CSV file, click **Export**.

NAT Rules and Pools

To add and configure NAT rules and pools for an appliance, click the edit icon associated with the appliance listed in the table. The NAT dialog box opens.

NAT Rules

Each NAT rule has a directional indicator. Outbound rules are applied to traffic flows initiated from the LAN and destined to the WAN. Inbound rules are applied to traffic flows initiated from the WAN and destined to the LAN. NAT rules include all tunneled traffic, internet-bound traffic, and other passthrough traffic. Return traffic for a given flow does not require an additional rule. You must configure a destination IP address for each rule.

NOTE: You must disable advertisements of local, static routes on the LAN side at the site so that routes are completely unique. In addition, you must configure announce-only static routes for your NAT pools and advertise them to the WAN by allowing the routes in your "Redistribute routes to SD-WAN fabric" route map. For details, see Routes Tab.

To add and configure a NAT rule:

1. On the NAT dialog box, click **Add Rule**.

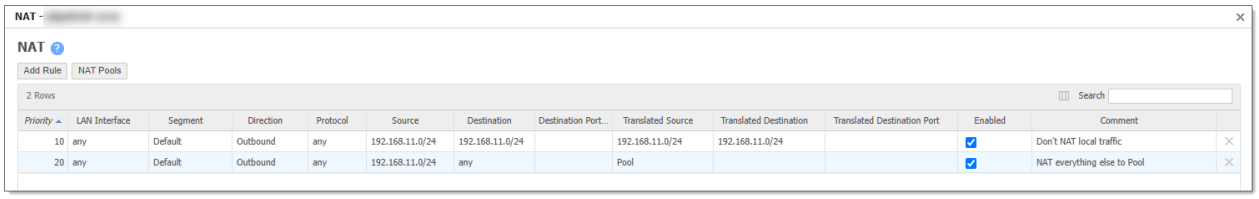
   A row is added to the table.
2. Complete the following fields as appropriate.

Field	Description
Priority	Order in which the rules are applied. The lower the priority, the higher probability your NAT rule will be applied.
LAN Interface	LAN interface the NAT rule uses. You can select a specific LAN interface or "any". This is configurable for an outbound NAT rule only.

Field	Description
Segment	Name of the segment to apply to this NAT rule. This field is available only if Routing Segmentation is enabled and if the appliance version supports NAT with segmentation. If you select a specific LAN interface for this NAT rule, this field will be automatically populated with the segment for that interface inferred from the Deployment dialog box. In this case, you cannot edit this field. You can select a segment for any NAT rule (inbound or outbound) if the LAN interface is set to “any”. For inbound rules, you must specify the destination segment the rule belongs to. For outbound rules, specify the source segment the rule belongs to. Deleting a segment from the Routing Segmentation (VRF) tab will automatically remove all associated branch NAT rules.
Direction	Select the direction of traffic: Outbound (LAN to Fabric or WAN-side) Inbound (Fabric or WAN-side to LAN)
Protocol	Type of protocol used for each NAT. By default, this is set to “any”.
Source	Original source IP address of the IP packet. You can specify a source IP address or “any”.
Destination	IP address of the LAN/WAN interface to which traffic is going. You can specify a destination IP address or “any”.
Destination Port	Original destination port to match in the flow.
Translated Source	Translated source IP address or NAT pool when the NAT rule is applied. You can enter an IP address or select a NAT pool from the field’s drop-down list if you have set up one or more of them. **NOTE:** If you select a NAT pool, both the segment and direction for the NAT rule and the NAT pool must match.
Translated Destination	Translated destination IP address when the NAT rule is applied.
Translated Destination Port	Translated destination port when the NAT rule is applied.
Enabled	Select this check box to enable your customized NAT rule. Direction can be inbound or outbound.
Comment	Comment you can add for the NAT rule.

3. Click **Save**.

NOTE: You can create an explicit rule to not NAT by using the same IP subnet in the Source, Destination, Translated Source, and Translated Destination fields. In the example shown below, the first rule causes local traffic within the 192.168.11.0/24 range to not get NATed. The second rule, however, causes any other traffic from the 192.168.11.0/24 network to be source NATed into the NAT pool named “Pool”.



NAT Pools

To add and configure a NAT pool:

1. On the NAT dialog box, click **NAT Pools**.

   The NAT Pools dialog box opens.

2. Click **Add**.

   A row is added to the table.

3. Complete the following fields as appropriate.

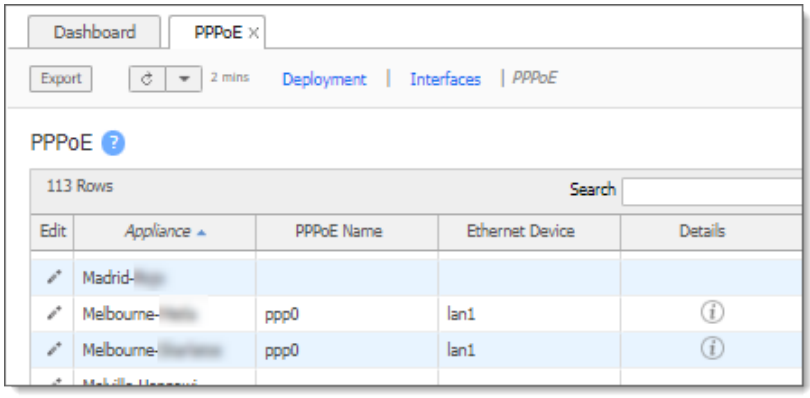
Field	Description
Name	Name of the NAT pool.
Direction	Direction of traffic (Inbound or Outbound). **NOTE:** The direction applied to the NAT rule and the associated translated NAT pool must be the same.
Segment	Name of the segment to apply to this NAT pool. This field is available only if Routing Segmentation is enabled and if the appliance version supports NAT with segmentation. **NOTE:** The segment applied to the NAT rule and the associated translated NAT pool must be the same.
Subnet	Subnet IP address. **NOTE:** You can set up overlapping subnet IP addresses for NAT pools, but the associated segments must be unique (that is, not the same segment name).
Translate Ports	Select this check box to enable source port address translation if the NAT pool is too small to accommodate multiple flows simultaneously with one-to-one IP address translation.

4. Click **Save**.

PPPoE Tab

Configuration > *Networking* > *PPPoE*

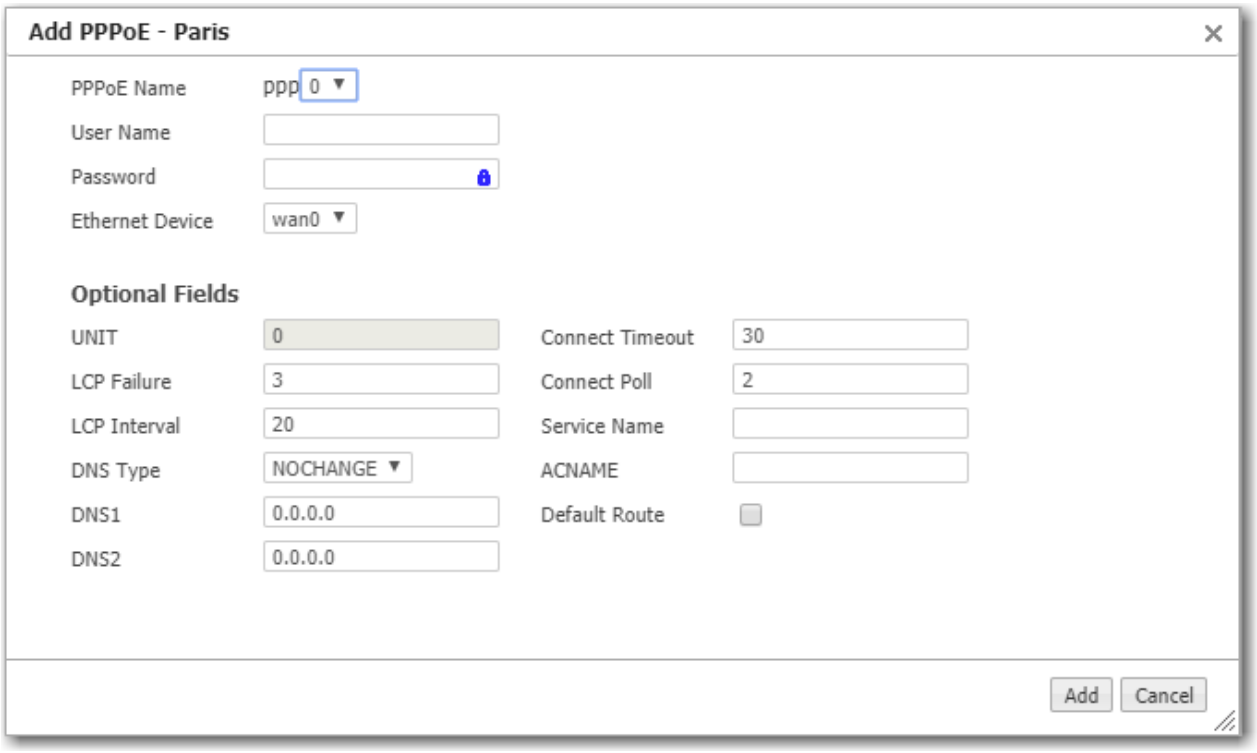
Point-to-Point Protocol over Ethernet (**PPPoE**) is a network protocol for encapsulating PPP frames inside Ethernet frames. It is used mainly with DSL services where individual users connect to a DSL modem over Ethernet.



When configuring a PPPoE connection, complete the following fields:

Field	Description
Ethernet Device	Specifies the physical interface to use for sending the protocol. Generally, this is a WAN-side interface.
Password	This is set up with your Internet Service Provider (ISP).
PPPoE Name	Name is **ppp** followed by a numerical suffix from **0** to **9**.
User Name	This is set up with your Internet Service Provider (ISP).

Generally, this is all the configuration required. If your ISP is fine-tuning the access, you might be asked to configure some of the **Optional Fields**, below.



Field	Description
ACNAME	Access Concentrator Name. Provided by the ISP.
Connect Poll	Specifies how many times to try to establish the link. The default value is **2**.
Connect Timeout	When trying to establish the link, this specifies how many seconds until the effort times out. The default value is **30** seconds.
Default Route	If the check box is selected, the connection uses the default gateway provided by the ISP.
DNS Type	This specifies the resolver to use: **NOCHANGE** – Do not accept or configure the ISP's Domain Name Server (DNS). Use the DNS configured on the **Administration > General Settings > Setup > DNS** tab. **SERVER** – Accept the ISP's DNS. This then overrides the EdgeConnect DNS configuration. **SPECIFY** – Use **DNS1** and **DNS2** to resolve domain names.
LCP Failure	*Link Control Protocol Failure*. Specifies the number of times the keep-alive can fail before the link goes down. The default value is **3**.
LCP Interval	The default value for this keep-alive interval is **20** seconds.

Field	Description
Service Name	Provided by the ISP.

USB LTE Tab

Configuration > Networking > USB LTE

Use this tab to configure and monitor USB LTE modems.

The USB LTE tab provides two views:

- Click **Config** to display configuration details associated with the local appliance. For field descriptions, see Enable or Edit a USB LTE Modem.
- Click **Status** to display the status of each appliance and additional details including:

Field	Description
RSSI	Signal Strength Indicator. See Understanding RSSI Values.
Phone Number	The phone number associated with the cellular device.
Carrier	The network provider for the USB LTE device.
Uplink Status	Status indicating whether the USB LTE modem is able to establish an uplink to the carrier network.
IMEI	International Mobile Equipment Identity (IMEI). The unique 15-digit number that identifies the USB LTE modem device.
IMSI	International Mobile Subscriber Identity (IMSI). The IMSI number on the USB LTE modem.
Model Number	The model number of the USB LTE modem.
Product ID	The product ID on the USB LTE modem.
Part Number	The part number associated with the USB LTE modem.
Serial Number	The serial number on the USB LTE modem.
Firmware Version	The firmware version on the USB LTE modem.

Understanding RSSI Values

RSSI	Signal Strength	Description
> -65 dBm	Excellent	Strong signal with maximum data speeds.
-65 dBm to -75 dBm	Good	Strong signal with good data speeds.

RSSI	Signal Strength	Description
-75 dBm to -85 dBm	Fair	Fair but useful.
-85 dBm to -95 dBm	Poor	Performance will drop drastically.
<= -95 dBm	No signal	Disconnection.

TIP: Select the **Hide Missing Interfaces** check box in the top-right corner of the tab to filter the list to display only those appliances assigned an interface. Clear the **Hide Missing Interfaces** check box to display all interfaces.

Technical Tips and Best Practices

- You can plug the USB LTE modem into any USB port on the EdgeConnect.
- You can attach a maximum of one USB LTE modem to each appliance.
- Aruba offers an extender cable for better placement of the USB LTE modem.
- HPE Aruba Networking does not recommend the use of the USB LTE modem with a USB hub.
- EC-V and EC-US do not support USB LTE modems.
- If you encounter issues, admin down and admin up the `cell0` interface.

Enable or Edit a USB LTE Modem

To enable or edit a USB LTE modem:

1. If the USB LTE modem is not plugged into an appliance, plug the modem into the appliance now.

   The USB LTE modem will be in the Down state.
2. Navigate to **Configuration > Networking > USB LTE**.
3. Click the edit icon next to the appliance on which you want to enable or edit a USB LTE modem.

   The Edit LTE Interface dialog box opens.
4. Configure the following elements as needed:

Field	Description
Interface	The name of the cell interface. The appliance shows cell0, cell1, cell2, and cell3. Only cell0 is used for the USB LTE modem. You must select and configure the cell interface name on the Deployment page. FW Model must be set to Stateful. Click **Deployment** in the tab header to configure the cell interface.
Admin	Indicates whether the interface is Up or Down. The USB LTE modem ships in the Down state.
Generation	Set the cellular modem to either **auto**, **4G**, or **3G**.
APN	The Access Point Name for the interface. Consult your network provider for a list of possible values. If a specific APN is not required, select **default**.
Authentication	Select **None**, **CHAP**, or **PAP**. If you select PAP or CHAP, Credentials fields display, where you must provide the Username and Password.

5. Click **Save**.

Loopback Interfaces

Configuration > Networking > Loopback Interfaces

The loopback feature enhances reliability and security by enabling you to access your network using a single static IP address. If one interface goes down, you can access all interfaces through the single static IP address.

To add a loopback interface to your network:

1. Navigate to **Configuration > Networking > Loopback Interfaces**.

   The Loopback tab opens.
2. Click the edit icon next to the appliance to which you want to add a loopback interface.

   The Loopback Interfaces dialog box opens.
3. Click **Add**.

   The Add Interface dialog box opens.
4. Configure the following elements as needed:

Field	Description
Segment	Name of the segment, if enabled.
Interface	Name of the loopback interface.

Field	Description
IP/Mask	IP address for the loopback interface.
Admin	Select whether the admin status is up or down.
Label	Label of the loopback interface.
Zone	Zone you want to apply to your loopback interface.

5. Click **Add**.

Loopback Orchestration

Configuration > Networking > Loopback Orchestration

Use the Loopback Orchestration tab to create a pool of loopback addresses from which Orchestrator can automatically create loopback interfaces. You can assign IP addresses from the pool to each appliance in the network.

IMPORTANT: Loopbacks must be configured consistently on all appliances. Loopbacks are used for IP SLA source IP, DNS Proxy source IP, and other management functions, such as TACACS, NetFlow, and SNMP. You must create a loopback label and management zone before you start setting up loopback orchestration. Be sure to configure a unique LAN-side label (such as “Loopback”) for the orchestrated loopback interfaces. If you use the built-in zone-based firewall, you should also configure a dedicated firewall zone such as “Management”.

You can select a segment in the **Segment** drop-down list to filter loopback interfaces in the table. This list includes all segments configured on the Routing Segmentation (VRF) tab. Select **All** to display all loopback interfaces.

NOTE: If a loopback interface exists for the Global segment, and an appliance is in a region not represented by any other loopback interface, the appliance will have loopback based on the Global loopback pool.

The following table describes the fields for each loopback interface listed on the Loopback Orchestration tab.

Field	Description
Segment	Segment associated with the loopback interface. Each loopback interface is associated with a specific segment.
Region	Region associated with the loopback interface. By default, this is set to Global, which includes all regions. You can create multiple loopback interfaces for the same segment, but with different regions.
Label	Label of the LAN-side interface being used. The default setting is NONE.
Zone	Firewall zone associated with the loopback interface. By default, this is set to the system-provided Default firewall zone.

Field	Description
Role	Role selection enables specifying the source role for a flow when other identity pipelines, such as Radius snooping, NACD, and VXLAN/GPID, are not present.
Management IP	Indicates whether management applications running on the appliance should use this loopback interface. **NOTE:** If you configure multiple loopback interfaces for different regions, you can specify different management interfaces for each region defined. Otherwise, if you configure multiple loopback interfaces for the same region or no regions, only one management interface is allowed.
Loopback Pool	Subnet IP address and mask used for the loopback pool. To change the loopback pool subnet IP, see Change the Subnet IP for a Loopback Pool below. You can configure one loopback pool per region and per segment.
Allocated / Total	Number of loopback IP addresses allocated from the pool out of the total number of IP addresses in the pool.
Deleted	Number of loopback IP addresses deleted from the loopback pool. The Reclaim link is displayed if the number is greater than zero (0). To return deleted loopback IP addresses to their original pools so they can be used again, see Reclaim Deleted Loopback IP Addresses below. **NOTE:** You must use Appliance Manager to delete an interface from an appliance.

Additional information:

- When enabling regional loopbacks, Orchestrator assigns each loopback a sequentially incremented loopback number. For example, if you have two regions (EAST and WEST) with IP addresses 10.1.1.0/24 and 10.2.2.0/24 respectively, Orchestrator creates lo20001 for EAST and lo20002 for WEST. Features such as BGP refer to explicit loopback names. You need to adjust them after Orchestrator assigns regional loopbacks to the appliance.
- It is recommended that you select the Management IP check box for the loopback associated with the Default segment. Only one management IP can be selected per appliance. This option causes the EdgeConnect to prefer the management IP as source IP for many management functions. It also causes the Orchestrator tree to display the loopback that has been selected using the Management IP check box.
- If a loopback interface exists for the Global segment, and an appliance is in a region not represented by any other loopback interface, the appliance will have loopback based on the Global loopback pool.

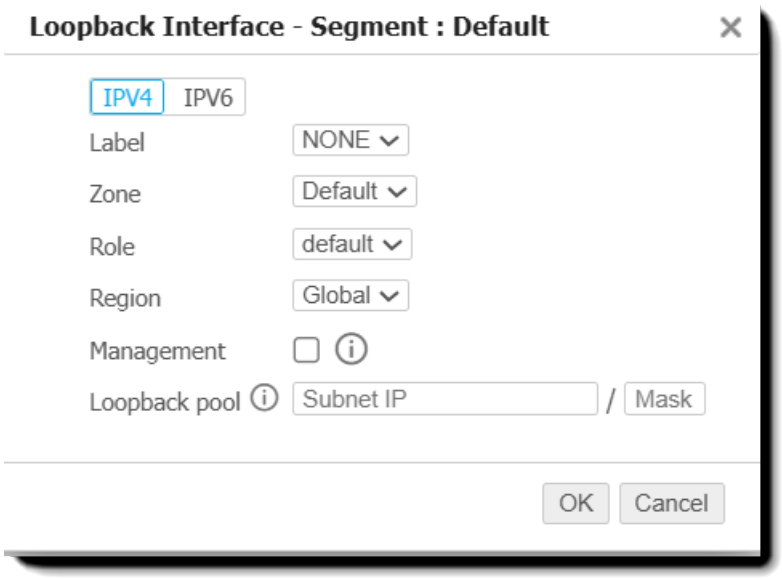
- If a loopback interface does not exist for the Global segment, and an appliance is in a region not represented by any other loopback interface, the appliance will not have loopback because a Global loopback pool does not exist.
- When an appliance is removed from the Orchestrator, the associated loopback address is retired and moved to Deleted status. If the appliance is later re-added to the Orchestrator, the original loopback IP is restored to the appliance.
- If you configure regional loopbacks, and then an appliance is moved from one region to another region, the loopback for the appliance is retired and moved to Deleted status. If the appliance is later moved back into the original region, the original loopback is restored.

Create a Loopback Interface

To create a loopback interface:

1. From the **Segment** drop-down list, select the segment to which you want to apply the new loopback interface.
2. Click **+Add**.

   The Loopback Interface dialog box opens.



   The selections you make on this dialog box are applied to the segment selected on the Loopback Orchestration tab.

   NOTE: You can create more than one loopback interface for a segment. However, you can create only one unique loopback range per segment.
3. Select the IP version: **IPv4** or **IPv6**.
4. Optionally, you can do the following:

- Select a label from the **Label** drop-down list. This list includes all LAN-side interface labels configured on the Interface Labels tab.

NOTE: Use a dedicated "Loopback" label. Do not re-use an existing LAN-side label. If your system does not already have a "Loopback" or similar label, you should create one.

- Select a firewall zone from the **Zone** drop-down list if you want the loopback interface to be part of a specific firewall zone. This list includes all zones configured in the Firewall Zone Definition dialog box. If you do not select one, the system-provided Default zone will be used.
- Select a region from the **Region** drop-down list. This list includes all regions configured on the Regions tab. If you do not select one, the system-provided Global pseudo-region will be used. Roles are managed on the *Orchestrator* > *Orchestrator Server* > *Users & Administration* > *Role Based Access Control* tab.
- Select a role from the **Role** drop-down list. Role selection enables specifying the source role for a flow when other identity pipelines, such as Radius snooping, NACD, and VXLAN/GPID, are not present.

5. For IPv4 only, select the **Management** check box if you want management applications running on the appliance to use the loopback interface. If selected, management applications running on the appliance prefer this loopback interface. Choosing at least one orchestrated loopback as "Management IP" is highly recommended. This setting also causes the Orchestrator appliance tree to display the Management IP when you hover over the appliance.
6. Specify the Loopback pool Subnet and Mask. All Loopback interfaces in this segment draw IPs from this pool.
7. Click **OK** to close the Loopback Interface dialog box, and then click **Save** on the Loopback Orchestration tab page to save your changes.

   The new loopback interface is added to the table on the Loopback Orchestration tab.

Change the Subnet IP for a Loopback Pool

To change the loopback pool subnet IP for a loopback interface listed in the table on the Loopback Orchestration tab:

1. Click the link in the Loopback Pool column for the appropriate loopback interface.

   The Loopback Pool dialog box opens.
2. In the **Subnet IP** field, enter the subnet IP address you want to use for the loopback pool.
3. Click **Update**.

   NOTE: Orchestrator immediately begins to reassign loopback IPs to all appliances. If the appliances were using the existing loopback to communicate with Orchestrator, you might see the appliance disconnect and then reconnect to the Orchestrator. It is recommended that you make loopback changes within a maintenance window.

Reclaim Deleted Loopback IP Addresses

The Orchestrator tries to maintain the loopback to appliance mapping even after an appliance is removed from Orchestrator or moved to a different region. When an appliance is removed, its associated loopback is retired and moved to Deleted status. If your loopback pool becomes depleted, you can reclaim deleted loopback IPs by returning them to their original pools so they can be used again for new appliances.

To reclaim a deleted loopback IP address:

1. Click the **Reclaim** link displayed in the Deleted column.

   The Reclaim Deleted Loopback IPs dialog box opens.

2. To return a listed loopback IP address to the original pool, click the **Reclaim** link in the Reclaim column.

   NOTE: Alternatively, you can click **Reclaim All** to return all listed loopback IP addresses to their original pools.

   A confirmation dialog box opens, asking whether you want to return the deleted loopback IP to the pool.

3. Click **Reclaim**.

Virtual Tunnel Interfaces (VTI)

Configuration > Networking > Virtual Tunnel Interfaces (VTI)

Use the Virtual Tunnel Interfaces (VTI) tab to manage VTIs for your appliances. A VTI is a tunneling protocol that does not require a static mapping of IPSec sessions to a physical interface. The tunnel endpoint is associated with a tunnel interface that enables a constant, secure, and stable connection throughout your network. With this feature, you can establish IPSec UDP SD-WAN tunnels nested within third-party IPSec tunnels (“tunnel-in-tunnel”). IPv6 is supported for outer tunnel; inner tunnel mode must be IPSEC_UDP.

NOTE: Tunnel-in-tunnel capabilities require that your appliances are running ECOS 9.5.x.x or later.

- The table on this tab lists all appliances in your network or those selected in the appliance tree.
- The columns are populated if a VTI has been set up for the appliance.
- An appliance will be listed more than once if multiple VTIs have been set up for the appliance.

To create a VTI or modify an existing one for an appliance, click any edit icon associated with the appliance. The VTI dialog box opens.

VTI Dialog Box

The VTI dialog box lists VTIs associated with the appliance if any exist. You can use this dialog box to set up:

- One or more VTIs for each tunnel interface.
- IP/Mask aliases for IP/Masks configured for the VTI.
- Automatic distribution of WAN-side VTIs into BGP.

Use the following procedures to configure a VTI with an associated tunnel in Orchestrator.

Add a VTI

1. Click **Add**.

   The Add VTI Interface dialog box opens.
2. Complete the following fields as appropriate.

Field	Description
Segment	Segment to associate with the VTI. This field is enabled only if Routing Segmentation is enabled. Select **Default** from the drop-down list for the system-supplied default segment or one of the other listed segments, which reflect custom segments defined using Routing Segmentation.
Interface	Identifier of the VTI. For example, to assign vti200 as the identifier, enter 200 in the field. **NOTE:** IDs in the range 20000 to 30000 are reserved for Orchestrator.
Interface Type	Interface type (lan or wan). Select the appropriate type from the drop-down list. The default setting is wan. **NOTE:** The interface type must match the associated passthrough tunnel configuration. If the passthrough tunnel is configured with a WAN-side interface label, select wan. If the passtrough tunnel is configured with a LAN-side interface label, select lan.
IP/Mask	Root IP address and subnet mask of the VTI. If these fields are empty, the default setting of 0.0.0.0/0 is used.
Admin	Select whether the interface is up or down. The default setting is up.
Status	Status of the VTI tunnel.

Field	Description
IP/Mask Alias	Alias IP address and subnet mask of the VTI. To display these fields, click the **+IP** link. Configuring an IP/Mask alias is optional. The default setting is 0.0.0.0/0. Deleting an IP/Mask alias resets it to the default setting. **NOTE:** This field is displayed only if the appliance is running ECOS 9.5.x.x or later. The IP/Mask alias is associated with the IP/Mask configured for the VTI. It assumes the zone of the associated VTI. You can configure any routable address for the IP address alias. **NOTE:** Appliances automatically source SD-WAN tunnel packets from IP address aliases when they are configured. The underlying tunnel will be sourced from the alias IP address.
Passthrough Tunnel	Passthrough tunnel to associate with the VTI. Select the appropriate tunnel from the drop-down list.
Label	Interface label to associate with the VTI. The default setting is NONE. To apply a label to the VTI, select one from the drop-down list. The list includes all LAN or WAN interface labels depending on your selection in the Interface Type field. Labels are defined by using the Interface Labels dialog box.
(Alias) Label	Interface label to associate as an alias with the VTI. When you click the **+IP** link, the Label field moves from the right of the IP/Mask fields to the right of the IP/Mask Alias fields. **NOTE:** This field and the +IP link are displayed only if the appliance is running ECOS 9.5.x.x or later. If you want to apply an alias label to the VTI, select an appropriate label from the drop-down list. The list includes all LAN or WAN interface labels depending on your selection in the Interface Type field. Labels are defined by using the Interface Labels dialog box. **NOTE:** Only one label can be specified for either the IP/Mask or IP/Mask alias (not both). Also, you cannot assign the same label to any other WAN-side interface.

Field	Description
Auto Distribute	Indicates whether the WAN-side VTI will be automatically redistributed into BGP. This feature is enabled by default. **NOTE:** This check box is displayed only if the appliance is running ECOS 9.5.x.x or later. This setting provides a per-VTI exception that ignores the “Automatically advertise local WAN subnets” setting on the Update Segment dialog box accessed from the Routes template. When VTI is provisioned, BGP must advertise WAN-side VTI configurations.
Zone	Firewall zone to associate with the VTI. The drop-down list includes all zones configured in the Firewall Zone Definition dialog box. Select **Default** for the system-supplied firewall zone.

3. To change the default NAT setting (Not behind NAT), click the **Not behind NAT** link.

   NOTE: The NAT-related link is displayed only if the appliance is running ECOS 9.5.x.x or later.

   The NAT Settings dialog box opens.

   1. Select **NAT** if the appliance is behind a NAT-ed interface or select the last option and enter an IP address to assign a destination IP for tunnels being built from the network to this VTI interface.
   2. Click **OK**.
4. On the Add VTI Interface dialog box, click **Add**.
5. On the VTI dialog box, click **Save**.

Edit a VTI

1. Click the edit icon associated with the VTI you want to modify.

   The Edit VTI Interface dialog box opens.
2. Complete the fields as appropriate. The fields are described above.
3. To change the default NAT setting (Not behind NAT), click the NAT-related link.

   NOTE: The NAT-related link is displayed only if the appliance is running ECOS 9.5.x.x or later.

   The NAT Settings dialog box opens.

   1. As appropriate, select **Not behind a NAT**, **NAT**, or select the last option and enter an IP address to assign a destination IP for tunnels being built from the network to this VTI interface.
   2. Click **OK**.
4. On the Edit VTI Interface dialog box, click **Update**.
5. On the VTI dialog box, click **Save**.

Delete a VTI

To delete a VTI listed in the table, click the corresponding delete icon (**X**) in the last column.

DHCP Manager

Configuration > Networking > DHCP Manager

Starting in Orchestrator 9.7.0, SD-WAN introduces a centralized DHCP Manager that significantly enhances how administrators configure, monitor, and operate DHCP services across appliances. These enhancements are designed to support distributed DHCP architectures, improve operational visibility, and reduce the effort required to manage DHCP scopes, leases, reservations, and failover.

All DHCP-related configuration and monitoring features are consolidated under a single workflow in the Orchestrator UI.

NOTE: For appliances running versions earlier than 9.7.0, DHCP settings must be configured locally on the appliance UI. See DHCP Server Defaults, DHCP Leases, DHCP Failover, and DHCP Failover State. DHCP Manager requires Orchestrator and appliances running 9.7.0 or later. Older appliances display no supported data in DHCP Manager views.

The DHCP Manager includes the following subtabs, each addressing a distinct aspect of DHCP management:

- Address pool
- Lease
- Scope
- Reservation
- Service
- Profiler
- Failover state
- Relay

Address pool

Address pools define the IP subnets and ranges from which DHCP leases are allocated.

Key capabilities include:

- Support for both directly connected and non-directly connected subnets. Non-directly connected subnets enable DHCP services for remote LAN segments without requiring a physical interface per subnet.
- Multiple IP ranges per subnet (up to 10 ranges per pool).
- Optional scope naming for easier identification.
- Lease duration controls (default and maximum).
- Gateway, DNS, NTP, and NetBIOS configuration per pool.

Address pools can be local or remote.

Pool type	Description	Configuration path / behavior
Local pool	DHCP pool for a directly connected LAN-side subnet where EdgeConnect is in the client broadcast domain.	Select one appliance, click **DHCP Settings**, select a deployed LAN-side interface, and configure DHCP settings. The same dialog box is available from the Deployment page.
Remote pool	DHCP pool not bound to a local interface. EdgeConnect responds as a DHCP server when DHCP relay/IP helper sends requests to an EdgeConnect LAN-side IP.	Select one appliance and click **Add Remote Pool**. Pool selection is based on the relay *giaddr* (gateway IP address).

IMPORTANT: In 9.7.0, remote pools do not start DHCP service by themselves. Enable DHCP server on at least one LAN-side interface before relying on remote pools.

NOTE: Address pools can be created, edited, or deleted only when a single appliance is selected in the appliance tree.

NOTE: The scope utilization threshold is not configured per pool. It is a global setting (accessible from the System Information tab in the Orchestrator UI) that applies to both directly connected and non-directly connected address pools.

Lease

The Lease subtab provides centralized visibility and operational control of active DHCP leases. The subtab displays only active leases.

Capabilities include:

- Viewing all active leases for a selected appliance set.
- Filtering by IP/Subnet, MAC address, or interface.
- Deleting:
  - Selected leases
  - All leases matching a filter
  - All leases on a single appliance

These functions support faster troubleshooting and IP reclamation during reconfiguration or device replacement scenarios.

DHCP lease statuses

The different lease statuses provide visibility into the state of IP address assignments and help you manage the DHCP pool effectively. IP address status designations are as follows:

- **free**: The IP address is available in the DHCP pool and has not been assigned to any client. Typically, the **Primary** DHCP Failover server allocates **free** leases.
- **backup**: The IP address is not assigned to a client by the DHCP Failover peer servers and can be assigned by any DHCP Failover peer server. This occurs during failover scenarios or when the active server is unavailable. Typically, the **Secondary** DHCP Failover server allocates **backup** leases.
- **active**: The IP address is assigned to a client, and the lease is currently held by the active DHCP server.
- **abandoned**: The IP address was previously assigned to a client, but the lease has expired, and the client has not renewed it. The IP address is not immediately available for reuse and may go through a waiting period before being returned to the free pool.

Each DHCP server of a failover pair acts as a backup for the DHCP leases of its peer. For example, the DHCP server of EdgeConnect-1 hands an IP lease to a client 172.23.2.50. This IP address shows as active on both EdgeConnect-1 and EdgeConnect-2 DHCP Failover peer servers.

Scope

The Scope subtab provides a real-time summary of DHCP scope utilization, including:

- Subnet and scope name.
- Total, allocated, and available IP addresses.
- Utilization percentage per scope.

This view enables administrators to identify potential scope exhaustion issues and plan capacity adjustments before service is impacted.

Reservation

The Reservation subtab allows centralized management of static IP assignments.

Capabilities include:

- Add, edit, and delete reservations.
- Validate MAC and IP address formats.
- Prevent duplicate IP usage (duplicate MAC addresses are allowed).
- Search reservations by device name, MAC, or IP address.

NOTE: Reservations can only be modified when a single appliance is selected.

NOTE: Reservations can only be made for IP addresses outside the configured IP pool ranges. In addition, a reserved IP address must not belong to the lease file, and it must belong to one of the subnets (local or remote) configured by the user.

Service

The Service subtab displays the current status (**Up** / **Down**) of the DHCP and Relay services on each appliance. It is read-only and does not provide control over the service lifecycle.

NOTE: If DHCP server status is down, remote pools do not serve clients until DHCP server is enabled on a LAN-side interface.

Profiler

The Profiler subtab supports simultaneous DHCP server operation and DHCP Discovery forwarding, enabling device profiling workflows (for example, integration with ClearPass). This feature is commonly referred to as *IP Helper*.

IMPORTANT: The Profiler subtab is not a replacement for DHCP/BOOTP Relay mode. DHCP/-BOOTP Relay is used when a given interface does not serve DHCP—it instead forwards all client requests for a given interface to a remote DHCP server. The Profiler is used when the appliance is the DHCP server and additionally needs to forward DHCP Discovery traffic to a NAC or profiler system (such as ClearPass) for device fingerprinting. This is referred to as *ClearPass Device Insight* (CPDI). These are distinct features and can operate simultaneously. Segmentation (VRF) must be enabled before the Profiler subtab is configurable.

Capabilities include:

- Configure one or more IP helper addresses (maximum: 2).
- Map IP helpers to segments (VRFs).
- Validate helper/segment pairings.
- Enforce segmentation prerequisites.

This enables forwarding DHCP discovery traffic to a profiler for device classification while the appliance continues to assign addresses locally.

Failover state

The Failover state subtab enables real-time monitoring of failover health and synchronization:

- Monitor real-time failover health and synchronization status.
- View partner state, timestamps, and failover group details.

This subtab displays the DHCP failover peer states of each server for troubleshooting purposes.

Failover group name support allows failover to work when peers use different LAN-side interface names (for example, lan0 on one appliance and lan1 on the other).

NOTE: Disabling EdgeHA for RMA or other changes invalidates the DHCP failover configuration, and some values—such as My IP and Peer IP—can be deleted.

DHCP failover state fields

Field	Description
Appliance	Name of the EdgeConnect appliance that is part of the DHCP failover configuration.
Failover Group Name	Failover group name that is the same for all the tagged and untagged interfaces corresponding to one physical interface.

Field	Description
My State	Failover endpoint state of the selected primary appliance. The states are: **Normal**, **Communications-Interrupted**, **Partner-Down**, **Recover**, **Recover-Wait**, **Recover-Done**, **Unknown**.
My State Time	Date and time when the selected appliance's DHCP server entered the specified state in the table.
Partner State	Failover endpoint state of the partner appliance. The states are: **Normal**, **Communications-Interrupted**, **Partner-Down**, **Recover**, **Recover-Wait**, **Recover-Done**, **Unknown**.
Partner State Time	Date and time when the partner appliance entered the specified state in the table.
MCLT	Maximum client lead time: the maximum amount of time that one server can extend a lease for a client's binding beyond the time known by the partner.

DHCP failover state descriptions

- **Normal:** Both EdgeConnect DHCP servers are exchanging keepalives correctly and DHCP bindings are synchronized.
- **Communications-Interrupted:** Each EdgeConnect is unable to exchange DHCP failover messages and independently services DHCP requests. In this state, each EdgeConnect assumes that the partner EdgeConnect is not down, but only unable to exchange failover messages.
- **Partner-Down:** The local EdgeConnect has determined that its failover partner is down and has taken over responsibility for serving all DHCP requests for the shared address pools. The local server continues to operate independently until the partner returns and the two servers resynchronize.
- **Recover:** The local EdgeConnect is trying to establish full synchronization of DHCP bindings with the previously down EdgeConnect. In this state, EdgeConnect does not serve any DHCP requests until the DHCP bindings are fully synchronized between them.
- **Recover-Wait:** The DHCP server is waiting for a period equal to the Maximum Client Lead Time (MCLT) before transitioning to the Recover-Done state. During this time, the server is preparing to synchronize its lease database with its peer server so as to avoid any inconsistencies.
- **Recover-Done:** EdgeConnects enter this state after recovering from a communication failure and completing the synchronization of DHCP bindings. EdgeConnects transition to this state after they recover state but prior to the normal state.
- **Unknown:** This indicates the peer state is unreachable or has not received any responses from the peer for the amount of Max Response Delay timer.

Relay

The Relay subtab centralizes DHCP relay configuration:

- Select relay interfaces.

- Configure destination DHCP servers.
- Enable Option 82 and define relay policies.
- Support per-segment DHCP relay behavior.

DHCP relay forwards client DHCP packets to a DHCP server, receives the server responses, and returns the responses to the clients.

Failover Configuration

DHCP failover is configured through the DHCP Settings dialog box for an interface, not through a dedicated DHCP Manager subtab. The DHCP Manager provides the Failover state subtab for monitoring failover health and synchronization.

DHCP failover configuration supports:

- Creating and managing DHCP failover relationships.
- Assigning a failover relationship for all non-directly connected subnets.
- Assigning Primary and Secondary roles.
- Configuring peer IPs, ports, MCLT, and load balancing (SPLIT) parameters.

For full details, see DHCP Failover.

Monitoring and Alarms

The DHCP Manager introduces proactive monitoring capabilities:

- Configurable scope utilization thresholds (default 90%).
- Alarms when utilization crosses defined high-water marks.
- Central visibility into DHCP health across appliances.

This supports early detection of scope exhaustion and service degradation risks. See Alarms for a full list of DHCP-related alarms.

Release and Compatibility Notes

- DHCP Manager is available starting with Orchestrator and appliance version 9.7.0.
- No new external packages are required.
- Existing configurations are preserved and upgraded automatically where applicable.

DHCP settings Integration

The DHCP settings subtab in DHCP Manager opens the same DHCP configuration dialog box used from the Deployment page.

1. Select exactly one appliance in the tree.
2. Click **DHCP settings**.
3. Select the deployed LAN-side interface.
4. Configure DHCP server, relay, and related interface-level settings.
5. Click **Save**.

DHCP Pseudo LAN Client Isolation

A DHCP-based pseudo LAN Client Isolation pattern can use a /32 subnet mask option so clients ARP for the default gateway and traffic is routed through EdgeConnect for policy evaluation.

CAUTION: This pattern is not universal segmentation. Some clients might not behave correctly with /32 mask behavior.

ARP Visibility

The web UI provides ARP table visibility and supports clearing individual or all ARP entries.

Logging and Troubleshooting Notes

- There is no new dedicated DHCP CLI for this enhancement set.
- In ECOS 9.6.0.0 and later, feature-level debug can capture DORA (Discover, Offer, Request, Acknowledge) DHCP events in audit logs.
- Formal top-level DHCP logging is not available as a separate feature.

Known Limitations

- DHCPv6 server is not available in this feature set.
- DHCP configuration and failover are not orchestrated automatically. Configuration remains per appliance.
- The UI does not directly start, stop, or restart the DHCP service.

DHCP Server Defaults

Configuration > Networking > DHCP Server Defaults

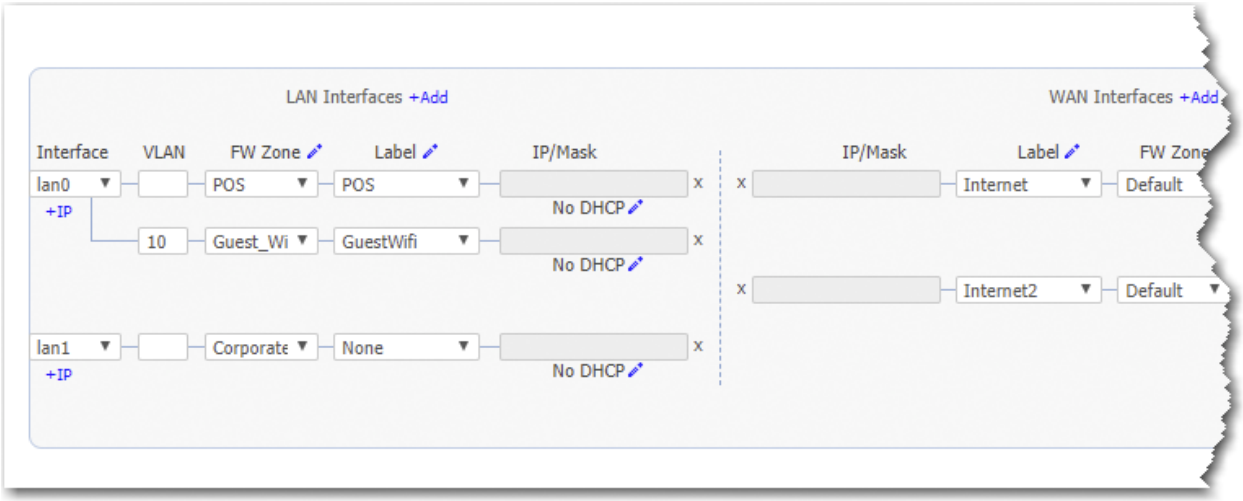
NOTE: Starting in Orchestrator 9.7.0, DHCP configuration and monitoring are centralized in the DHCP Manager. The content on this page applies to appliances running versions earlier than 9.7.0.

You can reduce your workload by using this tab to configure global defaults for Dynamic Host Configuration Protocol (DHCP).

- These defaults apply to the LAN interfaces in **Deployment Profiles** that specify **Router** mode.
- There are three choices:
  - **No DHCP/No RA**.
  - Each LAN interface acts as a **DHCP Server**.

    NOTE: If you enable DHCP Failover, you must use the same interface label for each physical interface.
  - The EdgeConnect appliance acts as a **DHCP/BOOTP Relay** between a DHCP server at a data center and clients needing an IP address.

- On the **Configuration > Overlays & Security > Deployment Profiles** tab, the selected default displays consistently under each LAN–side IP/Mask field.



  For any LAN–side interface, you can override the global default by clicking the DHCP-related link under the IP/Mask field and changing the values or selection.

- Changes you save to the global default only apply to new configurations.
- To view or revise the list of reserved subnets, select **Monitoring**.
- Before you can configure DHCP, you must navigate to Management Services and select an interface for DHCP Relay. DHCP relay forwards DHCP messages between clients and servers that are not on the same local network or subnet. This is particularly useful in larger networks where a single DHCP server needs to manage IP address assignments across multiple subnets. See Management Services for more information.

If the LAN interface has an IPv4 IP address, click **V4** to display the DHCP configuration settings. See V4.

If the LAN interface has an IPv6 IP address, click **V6** to display the Router Advertisement settings. See V6.

DHCP Settings / Router Advertisements

V4

The following tables describe the various DHCP settings you can configure for LAN interfaces that have IPv4 IP addresses.

DHCP Server

Field	Description
DHCP Pool Subnet/Mask	Enter the DHCP pool subnet and mask IP addresses.

Field	Description
Subnet Mask	Mask that specifies the default number of IP addresses reserved for any subnet. For example, entering **24** reserves 256 IP addresses.
Exclude first N addresses	Specifies how many IP addresses are not available at the beginning of the subnet's range.
Exclude last N addresses	Specifies how many IP addresses are not available at the end of the subnet's range.
Default lease, Maximum lease	Specify, in hours, how long an interface can keep a DHCP–assigned IP address.
Default gateway	Indicates whether the default gateway is being used.
DNS server(s)	Specifies the associated Domain Name System servers.
NTP server(s)	Specifies the associated Network Time Protocol servers.
NetBIOS name server(s)	Used for Windows (SMB) type sharing and messaging. It resolves the names when you are mapping a drive or connecting to a printer.
NetBIOS node type	NetBIOS node type of a networked computer relates to how it resolves NetBIOS names to IP addresses. There are four node types: **B**-node – 0x01 Broadcast **P**-node – 0x02 Peer (WINS only) **M**-node – 0x04 Mixed (broadcast, then WINS) **H**-node – 0x08 Hybrid (WINS, then broadcast)
DHCP failover	Enables DHCP failover. To set it up, click the **Failover Settings** link.

DHCP/BOOTP Relay

Field	Description
Destination DHCP/BOOTP Server	IP address of the DHCP server assigning the IP addresses.

Field	Description
Common DHCP server for all segments	Select this option to use one DHCP server for all VRF segments. Because the DHCP server must be segment-aware and use options and sub-options accordingly when Common mode is used, it is recommended to use Distinct mode in most cases. **HINT:** You can reset the defaults in Management Services by setting the DHCP Relay interface to “any” and then selecting an interface label again. However, this might impact service. Or, you can manually reset the defaults by selecting the following values: Option 82 = enabled, Option 82 Policy = append, and select the following sub-options: 1, 5, 10, 11, 151, and 152.
Distinct DHCP server per segment	Select this option to override the DHCP relay configuration set in the Management Services tab with the settings you select in this dialog box. Distinct mode uses a distinct DHCP server for each segment. **NOTE:** This is the default setting, and is recommended for most scenarios.
Enable Option 82	When selected, inserts additional information into the packet header to identify the client’s point of attachment. This setting applies to all LAN-side interfaces on this appliance. **IMPORTANT:** Changing this setting will modify Option 82 settings on all LAN-side interfaces that are enabled as DHCP Relay.
Option 82 Policy	Tells the relay what to do with the hex string it receives. The choices are **append**, **replace**, **forward**, and **discard**. This setting applies to all LAN-side interfaces on this appliance. **IMPORTANT:** Changing this setting will modify Option 82 settings on all LAN-side interfaces that are enabled as DHCP Relay.

Field	Description
Sub Options	Select one or more of the following: **1** - Agent Circuit ID: Provides information about the interface or circuit through which the DHCP request was received. **5** - Link selection: Specifies the IP address used by the DHCP server to determine the appropriate subnet for addressing the DHCP client. **10** - Client Unicast/Broadcast Indication flag: Indicates whether the DHCP relay received the client packet as a unicast or broadcast packet. **11** - Server ID Override: Allows the DHCP relay agent to act as a proxy for the DHCP server to process unicast lease renewals. **150** - Link selection (Cisco proprietary): Provides information about a segment or VPN that is necessary to allocate an address to a DHCP client on that segment. **151** - VRF name/VPN ID **152** - VRF name/VPN ID Control Sub-Option OR Server ID Override (Cisco proprietary): Indicates whether the DHCP server supports sub option 151 (VRF Name/VPN ID). If this option is present in the reply from the server, the server does not support option 151.

V6

The following table describe the various router advertisement settings you can configure for LAN interfaces that have IPv6 IP addresses. The LAN interface will use these options to auto-configure IPv6 addresses and to learn default gateway addresses.

Setting	Description
Enable Router Advertisements	Specifies whether the router should send RA messages.
Managed Flag	Select this option to instruct IPv6 hosts to use DHCPv6 to obtain their IPv6 addresses in addition to any other configuration information.

Setting	Description
Other Flag	Select this option to instruct IPv6 hosts to use DHCPv6 to obtain additional configuration information, such as DNS server addresses and other network parameters.
Link MTU	Set the maximum transmission unit (MTU) size that can be transmitted without fragmentation. This helps ensure that all hosts on the network use the same MTU, avoiding issues related to packet fragmentation and reassembly.
Max Interval	Specify the maximum interval in seconds between unsolicited RA messages. This helps to control the frequency of RA messages.
Min Interval	Specify the minimum interval in seconds between unsolicited RA messages. This helps to control the frequency of RA messages.
Current Hop Limit	Set the default hop limit for IPv6 packets sent by hosts on the network. Hosts use this value to configure their own hop limit for outgoing packets.
Default Router Preference	Select **High**, **Medium**, or **Low** to set the preference level of the router for use as a default router. Hosts use this value to prioritize multiple routers on the same link.
Default Router Lifetime	Specify the lifetime in seconds of the default route that is advertised by the router. The hosts use this value to determine how long the router should be used as the default gateway.
Reachable Time	Specify the time in milliseconds that an IPv6 host considers a neighbor reachable after receiving a confirmation. This value maintains accurate and timely reachability information in the neighbor cache.
Retrans Timer	Specify the time in milliseconds between retransmissions of neighbor solicitation messages. This value reduces the frequency of retries when attempting to discover or confirm the reachability of neighbors on the network.

Add a Router Advertisement Prefix

Click **Add** and complete the following fields. You can create up to 10 prefixes.

Setting	Description
Prefix-id	The ID assigned to the prefix.
Prefix	Select whether the prefix can be used by hosts for SLAAC. When set to true, hosts can use the prefix to generate their own IPv6 addresses.

Setting	Description
Autonomous flag	Select whether the prefix can be used by hosts for SLAAC. When set to true, hosts can use the prefix to generate their own IPv6 addresses.
Onlink flag	Specifies whether the prefix is on-link, which affects how hosts handle routing for addresses within the prefix. If set to true, hosts assume that addresses within the prefix can be reached directly on the local network segment.
Valid Lifetime	Specify the duration in seconds for which the advertised prefix is valid.
Preferred Lifetime	Specify the duration in seconds (relative to the time the packet is sent) that addresses generated from the prefix via stateless address auto-configuration remain preferred.

DHCP Leases

Configuration > Networking > DHCP Leases

NOTE: Starting in Orchestrator 9.7.0, lease visibility and management are available on the Lease subtab of the DHCP Manager tab. The content on this page applies to appliances running versions earlier than 9.7.0.

This page lists the IP addresses that are currently being leased from the DHCP pool.

DHCP Lease Statuses

The different lease statuses provide visibility into the state of IP address assignments and help you manage the DHCP pool effectively. IP address status designations are as follows:

- **free**: The IP address is available in the DHCP pool and has not been assigned to any client. Typically, the **Primary** DHCP Failover server allocates **free** leases.
- **backup**: The IP address is not assigned to a client by the DHCP Failover peer servers and can be assigned by any DHCP Failover peer server. This occurs during failover scenarios or when the active server is unavailable. Typically, the **Secondary** DHCP Failover server allocates **backup** leases.
- **active**: The IP address is assigned to a client, and the lease is currently held by the active DHCP server.
- **abandoned**: The IP address was previously assigned to a client, but the lease has expired, and the client has not renewed it. The IP address is not immediately available for reuse and may go through a waiting period before being returned to the free pool.

Each DHCP server of a failover pair acts as a backup for the DHCP leases of its peer. For example, the DHCP server of EdgeConnect-1 hands an IP lease to a client 172.23.2.50. This IP address shows as active on both EdgeConnect-1 and EdgeConnect-2 DHCP Failover peer servers.

DHCP Failover

NOTE: Starting in Orchestrator 9.7.0, DHCP failover is configured in the DHCP Settings dialog box and monitored on the Failover state subtab of the DHCP Manager tab. The content on this page applies to appliances running versions earlier than 9.7.0.

DHCP Failover performs seamless failover, real-time database synchronization, and load balancing that ensures high availability of the DHCP service. DHCP Failover is available for IPv4 addressing. Besides providing more robust DHCP services, DHCP Failover avoids duplicate addressing when two EdgeConnect gateways at the same site enable their DHCP servers.

DHCP Failover server peers are configured on two EdgeConnects, one on each. Configure consistent subnet IP pool ranges, peer IP addresses, port numbers, and other failover settings across the peer pairs for deterministic behavior of the DHCP service. Settings are not compared and validated between the failover peers. DHCP Failover settings examples are provided below.

DHCP Failover can be configured in a Deployment Profile or directly on a LAN interface of an EdgeConnect. The procedure below provides the steps for configuring DHCP Failover directly on EdgeConnect LAN interfaces.

Preparation

Before starting this process, record the IP addresses for the EdgeConnect LAN interfaces that you plan to use for the DHCP Failover servers. Use the same interface label for each physical interface.

NOTE: DHCP failover is time sensitive. For DHCP failover to function properly, both the Primary and Secondary DHCP servers must be configured to use the same NTP server.

Procedure

On the DHCP Failover dialog box of each EdgeConnect, configure the settings listed below to provision your DHCP failover servers.

1. Navigate to the DHCP Failover configuration screen.
    - Go to the **Configuration > Interfaces** screen.
    - Click the **edit icon** of the interface you will use. This opens the **Interfaces** configuration screen.
    - Click the **IP address** of the **Hardware** interface you will use (for example, lan0). This opens the **Deployment** screen.
    - Click the **NoDHCP/No RA** link to open the **DHCP Settings / Router Advertisements** screen.
    - Configure or update the DHCP settings as needed.
    - Select the **DHCP Failover** check box.
    - Click the **Failover Settings** link to open the **DHCP Failover** settings screen.
2. Configure the **DHCP Failover** settings listed in the table below. Click **Save** to save your settings. Upon saving the configuration, the associated EdgeConnect LAN interface(s) show that they are in use for a DHCP server. Repeat this procedure to configure the DHCP Failover server peer on the other EdgeConnect.

Field	Use	Comments
Role	Select Primary or Secondary.	Differentiates the DHCP Failover peer servers, particularly when the optional SPLIT option is used to determine the percentage of DHCP requests each peer DHCP Failover server handles.
My IP	IP address of the EdgeConnect LAN interface you assign to the DHCP server.	Must be on the same Layer 3 subnet as the Peer IP below. If you have multiple VLANs/Sub-Interfaces configured under one physical interface (for example, LAN0), choose any of the interface IP addresses as My IP and the corresponding EdgeConnect LAN interface IP of the peer as the Peer IP.
My Port	TCP port number of the LAN interface.	TCP port 647 is the default port EdgeConnects use to establish a peer DHCP server state and synchronize DHCP leases between them.
Peer Port	TCP port number of the peer DHCP server.	TCP port 647 is the default port EdgeConnects use to establish a peer DHCP server state and synchronize DHCP leases between them.

Field	Use	Comments
MCLT	Optional. The Maximum Client Lead Time (MCLT) default is 60 minutes. This field cannot be zero.	This setting only impacts DHCP renewal requests when one of the DHCP servers is down. It defines the maximum amount of time that an EdgeConnect can extend a lease for a DHCP client during a failure of a peer DHCP server that issued the original lease. Using a low value causes the DHCP renewal clients to look for lease renewals too frequently. The default of 60 minutes (3600sec) typically allows enough time for the failed DHCP server to have recovered.
SPLIT	Optional. Controls the percentage of DHCP requests that each EdgeConnect handles.	The range is 0 to 255, with 128 being the default value. A SPLIT value of 128 distributes the load evenly between the two EdgeConnects. For example, a SPLIT value of 64 for the Primary DHCP Failover server would result in it handling 75% of the client requests and the secondary server handling 25%.
Max Response Delay	Optional. Determines how long the EdgeConnect delays its response to a client request, such as a lease request or renewal, if it hasn't received a timely acknowledgment from its failover partner.	By waiting for a defined period, the EdgeConnect can ensure that it does not prematurely take over leases or assignments that the failover partner might still be handling. This setting provides a buffer to accommodate temporary network issues or delays in communication between EdgeConnects.

Field	Use	Comments
Max Unacked Updates	Defines the maximum number of update messages (BNDUPD) that one EdgeConnect can send to its failover peer without receiving an acknowledgment (BNDACK) before it takes action to mitigate potential issues.	Limiting the number of unacknowledged updates helps ensure that both EdgeConnects maintain consistent lease information, reducing the risk of IP address conflicts.
Load Balance Max	Optional. Allows you to configure the number of seconds to wait for a cutoff, after which load balancing is disabled.	The cutoff is based on the number of seconds since the client sent its first DHCP DISCOVER or DHCP REQUEST message. This only works with clients that correctly implement the SECS field (most clients support this). Set this in the range of 3 to 5. The result is that when one of the EdgeConnects is responding to failover keepalives but not responding to actual DHCP requests from clients, the peer EdgeConnect takes over its client load automatically as the clients retry.

NOTE: In a DHCP Failover configuration, if failover is checked but the 'my IP' and 'peer IP' are not provided, and you try to remove the pair, the system complains that 'My IP' and 'Peer IP' need to be added because failover is checked. In this case, uncheck failover, save the changes, and then remove the pair.

DHCP Failover Settings Examples

The DHCP failover settings below are for the DHCP servers of two EdgeConnects, each with identical failover settings and correct Role, My IP, and Peer IP settings. You can configure one or more failover groups. The examples provided below are for a single failover group and for two failover groups.

Each failover group must be configured on a per physical interface basis. For example, if you select the DHCP Server for any of the LAN0 subnets and configure the Failover Settings, then those failover settings are retained for all the DHCP servers under LAN0 interfaces. If LAN0 has VLAN2301 and VLAN2302 on it, they would have the same DHCP failover settings. Likewise, a LAN2 interface DHCP server has its own failover settings. Use the examples below as a reference for how to configure DHCP failover settings. You can also use a pre-configuration template to standardize DHCP failover settings.

Single DHCP Failover Group Example

- **EdgeConnect 1 DHCP Server**
  - Interface: lan0 (IP/Mask: 172.23.2.3/25)
  - Role: Primary
  - MY IP: 172.23.2.3
  - Peer IP: 172.23.2.4
  - Port: 647
  - MCLT: 3600
  - SPLIT: 128
  - Max Response Delay: 10
  - Max Unacked Updates: 10
  - Load Balance Max: 5
- **EdgeConnect 2 DHCP Server**
  - Interface: lan0 (IP/Mask: 172.23.2.4/25)
  - Role: Secondary
  - MY IP: 172.23.2.4
  - Peer IP: 172.23.2.3
  - Port: 647
  - MCLT: 3600
  - SPLIT: 128
  - Max Response Delay: 10
  - Max Unacked Updates: 10
  - Load Balance Max: 5

Two DHCP Failover Groups Example

- **EdgeConnect 1 DHCP Server LAN0 Group**
  - Interface: lan0 (IP/Mask: 172.23.2.3/25)
  - Role: Primary
  - MY IP: 172.23.2.3
  - Peer IP: 172.23.2.4
  - Port: 647
  - MCLT: 3600
  - SPLIT: 128
  - Max Response Delay: 10
  - Max Unacked Updates: 10
  - Load Balance Max: 5

- **EdgeConnect 2 DHCP Server LAN0 Group**
  - Interface: lan0 (IP/Mask: 172.23.2.4/25)
  - Role: Secondary
  - MY IP: 172.23.2.4
  - Peer IP: 172.23.2.3
  - Port: 647
  - MCLT: 3600
  - SPLIT: 128
  - Max Response Delay: 10
  - Max Unacked Updates: 10
  - Load Balance Max: 5
- **EdgeConnect 1 DHCP Server LAN2 Group**
  - Interface: lan2 (IP/Mask: 172.23.3.3/25)
  - Role: Primary
  - MY IP: 172.23.3.3
  - Peer IP: 172.23.3.4
  - Port: 647
  - MCLT: 3600
  - SPLIT: 128
  - Max Response Delay: 10
  - Max Unacked Updates: 10
  - Load Balance Max: 5
- **EdgeConnect 2 DHCP Server LAN2 Group**
  - Interface: lan2 (IP/Mask: 172.23.3.4/25)
  - Role: Secondary
  - DHCP Failover: Secondary
  - MY IP: 172.23.3.4
  - Peer IP: 172.23.3.3
  - Port: 647
  - MCLT: 3600
  - SPLIT: 128
  - Max Response Delay: 10
  - Max Unacked Updates: 10
  - Load Balance Max: 5

DHCP Failover Fundamentals

EdgeConnect **DHCP Failover** settings under **DHCP Server Configuration** enable you to provision DHCP Failover server peers in your network. The topics below explain how DHCP Failover servers operate.

DHCP Failover Primary or Secondary Roles

In the DHCP Failover settings screen, **Role** specifies how the pool of IP addresses is managed across the DHCP **Primary** and **Secondary** servers.

NOTE: The DHCP server Failover role does not refer to which server is active, and which is backup.

The pool of IP addresses the DHCP Failover servers manage is specified by the Layer 3 subnet that both the primary and secondary DHCP servers must be configured to use. That is, the interface subnet mask specifies the pool of addresses the DHCP Failover servers manage. For example, if a DHCP Failover server is configured to use lan0 with its IP/Mask set to 172.23.2.3/25, the pool would have 126 addresses (128 minus two, one for the network address and one for the broadcast address).

DHCP Active and Backup Server Behavior

During failover events, observe the following distinctions between the backup and active server designations:

- **Active-Backup Pair**: In DHCP Failover mode, you are deploying two DHCP servers, designated as “active” and “backup.” The Primary DHCP server does not always service the requests. The active server handles all DHCP requests and maintains the DHCP lease database.
- **SPLIT:** The configuration option that specifies the percentage split of the DHCP requests that each EdgeConnect handles. The SPLIT values are 0 to 255. The default of 128 distributes 50% of the load to each DHCP Failover server.
- **Lease Assignment and Failover**: The active server assigns IP addresses and other configuration parameters to clients. If the active server fails, the backup server seamlessly takes over, ensuring uninterrupted DHCP service.
- **Lease Database Synchronization**: The active server continuously synchronizes its DHCP lease database with the backup server. This ensures that the backup server has the latest information and can take over without any lease conflicts or disruptions in case of a failover.

DHCP Database Synchronization

DHCP Failover synchronization does not rely on an Edge-HA link. In DHCP Failover mode, DHCP servers synchronize their databases in these ways.

- **Real-time Updates**: The active DHCP server sends real-time updates to the backup server whenever a lease is assigned, renewed, released, or expires. This ensures that both servers have the most up-to-date information.
- **Full Synchronization**: Periodically, a full synchronization of the DHCP lease database is performed between the active and backup servers. This acts as a safeguard to ensure complete consistency in case any updates were missed during real-time synchronization.

FAQs

In a split-brain scenario, would both EdgeConnects provide IP addresses for DHCP requests?

A split-brain scenario is when the DHCP Failover servers are not communicating and synchronizing their data. Once the DHCP Failover servers detect the peer is unavailable, both serve all clients. New clients get DHCP offers from both. The client chooses one. Client DHCP lease renewal is different. If the EdgeConnect that provided the previous DHCP lease is not responding, the renewal fails for its clients. If the client eventually sends out new DHCP Requests (DORA process Discover-Offer-Request-Ack), an active EdgeConnect responds.

In a split-brain scenario, could an EdgeConnect assign a duplicate IP to a second device?

No. The EdgeConnect DHCP-Server uses a standards-based implementation that provides a mechanism to avoid duplicate IP addressing scenarios. EdgeConnects do a standard ARP request prior to issuing DHCP requests, thus preventing assignment of duplicate IP addresses.

What happens if one sub-interface is disabled?

If a disabled sub-interface whose IP address is *not* configured as **My IP** under the failover settings, the DHCP server does not provide DHCP services for that sub-interface range. If a disabled sub-interface whose IP address *is* configured as **My IP** under the failover settings, then DHCP failover is triggered for all the DHCP servers configured on that interface. This causes the local DHCP server to be unreachable from the DHCP peer perspective. The peer DHCP server then serves all new clients and honors the DHCP renewals that were issued by the previous DHCP server that is now unreachable. Likewise, if the physical interface goes down or the peer IP is unreachable, then the DHCP failover trigger happens.

DHCP Failover State

Configuration > Networking > DHCP Failover State

NOTE: Starting in Orchestrator 9.7.0, failover state monitoring is available on the Failover state subtab of the DHCP Manager tab. The content on this page applies to appliances running versions earlier than 9.7.0.

EdgeConnect appliances can act as a DHCP server for clients on the LAN side. DHCP failover allows redundancy by creating failover groups when two appliances are combined in an HA configuration. DHCP failover also provides stability if one EdgeConnect appliance dies by allowing the other EdgeConnect HA pair to take over as the DHCP server. To do so, the primary and secondary servers must be completely synchronized so that each server can rely on the other if one fails.

This tab displays the DHCP failover peer states of each server for troubleshooting purposes.

NOTE: Disabling EdgeHA for RMA or other changes invalidates the DHCP failover configuration, and some values—such as My IP and Peer IP—can be deleted.

DHCP Failover State Fields

Field	Description
Appliance	Name of the EdgeConnect appliance that is part of the DHCP failover configuration.
Failover Group Name	Failover group name that is the same for all the tagged and untagged interfaces corresponding to one physical interface.
My State	Failover endpoint state of the selected primary appliance. The states are: **Normal**, **Communications-Interrupted**, **Partner-Down**, **Recover**, **Recover-Wait**, **Recover-Done**, **Unknown**.
My State Time	Date and time when the selected appliance's DHCP server entered the specified state in the table.
Partner State	Failover endpoint state of the partner appliance. The states are: **Normal**, **Communications-Interrupted**, **Partner-Down**, **Recover**, **Recover-Wait**, **Recover-Done**, **Unknown**.
Partner State Time	Date and time when the partner appliance entered the specified state in the table.
MCLT	Maximum client lead time: the maximum amount of time that one server can extend a lease for a client's binding beyond the time known by the partner.

DHCP Failover State Descriptions

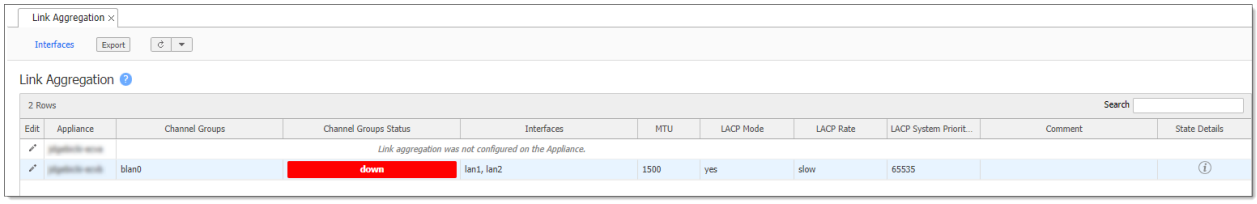
- **Normal:** Both EdgeConnect DHCP servers are exchanging keepalives correctly and DHCP bindings are synchronized.
- **Communications-Interrupted:** Each EdgeConnect is unable to exchange DHCP failover messages and independently services DHCP requests. In this state, each EdgeConnect assumes that the partner EdgeConnect is not down, but only unable to exchange failover messages.
- **Partner-Down:** The local EdgeConnect has determined that its failover partner is down and has taken over responsibility for serving all DHCP requests for the shared address pools. The local server continues to operate independently until the partner returns and the two servers resynchronize.
- **Recover:** The local EdgeConnect is trying to establish full synchronization of DHCP bindings with the previously down EdgeConnect. In this state, EdgeConnect does not serve any DHCP requests until the DHCP bindings are fully synchronized between them.
- **Recover-Wait:** The DHCP server is waiting for a period equal to the Maximum Client Lead Time (MCLT) before transitioning to the Recover-Done state. During this time, the server is preparing to synchronize its lease database with its peer server so as to avoid any inconsistencies.

- **Recover-Done:** EdgeConnects enter this state after recovering from a communication failure and completing the synchronization of DHCP bindings. EdgeConnects transition to this state after they recover state but prior to the normal state.
- **Unknown:** This indicates the peer state is unreachable or has not received any responses from the peer for the amount of Max Response Delay timer.

Link Aggregation

Configuration > Networking > Link Aggregation

The Link Aggregation tab displays channel group and link aggregation details for appliances selected in the appliance tree.



Link aggregation combines data from multiple physical or virtual interfaces into a channel group, which provides a single high-speed link. Configuring link aggregation adds failover redundancy to the interfaces in the channel group.

IMPORTANT: If you aggregate interfaces that are currently in use, those interfaces are removed from deployment before aggregation occurs. When attempting to apply channel group additions or changes on the Link Aggregation dialog box, a confirmation dialog box opens that gives you the choice to proceed with aggregating the interfaces or to cancel your link aggregation changes.

The table on the Link Aggregation tab displays the following information:

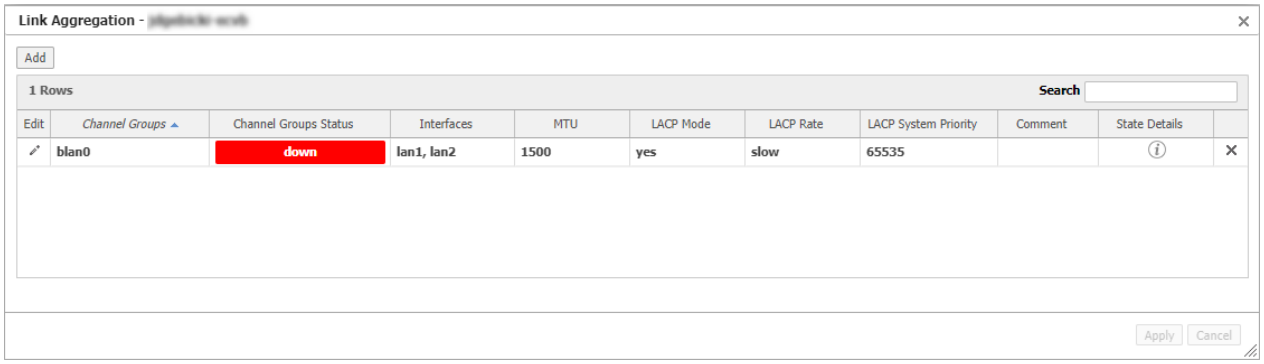
Column	Description
Appliance	Name of the EdgeConnect appliance.
Channel Groups	If any channel groups are configured for the appliance, the names of the channel groups are listed in this column. **NOTE:** You can create up to four channel groups (bonded interfaces) for each appliance; two for the LAN side (blan0, blan1) and two for the WAN side (bwan0, bwan1).
Channel Groups Status	Status of the channel group (up or down). For dynamic mode (LACP mode), this status combines the LACP and link statuses. For static mode, it combines the link statuses of the underlying ports.

Column	Description
Interfaces	Physical or virtual interfaces included in the channel group. A channel group consists of one, two, three, or four interfaces. **NOTE:** An interface can be part of only one channel group.
MTU	Maximum transmission unit (MTU) size (in bytes) configured for the channel group. The configured MTU overrides any existing MTU settings when the channel group is deployed. The default size is 1500 bytes.
LACP Mode	Indicates whether Link Aggregation Control Protocol (LACP) is enabled for the channel group (yes or no).
LACP Rate	Affects the timeout and the rate at which the LACP partner (switch) is requested to send LACPDU packets. For slow, one packet per 30 seconds, and timeout after 90 seconds. This is the default rate. For fast, one packet per second, and timeout after three seconds.
LACP System Priority	Priority number used to break ties with the LACP partner. This value can be set from 1 to 65535 with the lowest number having the highest priority. The default value is 65535.
Comment	Additional information about the channel group.
State Details	Provides status information on the channel group, including details about the channel group (bonded interface) state and port (interface) states. Click the info icon to open a dialog box that displays this status information. The Channel group state tab on the dialog box includes three status indicators: Link status, LACP status, and Channel group status. (LACP status is displayed only in dynamic mode [LACP mode].) The Channel group status reflects the Link status and LACP status. If either is down, the Channel group status will be down. The LACP status reflects the LACP statuses on the Port states tab. If the LACP statuses of all interfaces are down, the LACP status on the Channel group state tab will be down. To refresh the status information on this dialog box, click **Refresh**. The State Details icon is also displayed in the table on the Link Aggregation dialog box. The same dialog box opens if you click it there.

Configure Link Aggregation

To add, change, or delete channel groups for an appliance, click the edit icon in the applicable table row on the Link Aggregation tab.

The Link Aggregation dialog box opens.



The Channel Groups Status column displays the current status of the channel group (up, down, or pending). Pending status indicates that the link aggregation configuration has not yet been applied, and the state of the link aggregation is not known at this time.

Add or Modify a Channel Group

To add a channel group, perform the following steps:

NOTE: To modify a channel group, click the edit icon next to the channel group. The Modify link aggregation dialog box opens. Change the fields as described below, and then click **Apply**. You cannot modify an existing channel group name, but you can change the other settings.

1. Click **Add** on the Link Aggregation dialog box.

   The Add link aggregation dialog box opens.



2. Complete the following fields:

Column	Description
Channel Group (interface) name	Select a name for the channel group from the drop-down list (**blan0**, **blan1**, **bwan0**, or **bwan1**).
Interfaces to be grouped	From the drop-down lists, select one, two, three, or four interfaces to include in the channel group. **IMPORTANT:** If you aggregate interfaces that are currently in use, those interfaces are removed from deployment before aggregation occurs. When you click the Apply button on the Link Aggregation dialog box, a confirmation dialog box opens that gives you the choice to proceed with aggregating the interfaces or to cancel your link aggregation changes.
MTU	Specify the MTU size (in bytes) to be applied to all interfaces in the group. The default size is 1500 bytes.
LACP mode	Select this check box to enable LACP for the channel group. By default, this check box is not selected.
LACP rate	Select **slow** or **fast** from the drop-down list. The default is slow. This field is available only if LACP mode is selected.
LACP priority	Specify a priority number from 1 to 65535. Priority number is used to break ties with the LACP partner. The lower the number, the higher the priority. The default is 65535. This field is available only if LACP mode is selected.
Comment	(Optional) Provide additional information about the channel group.

3. Click **Add**.

Delete a Channel Group

To delete a channel group listed in the table on the Link Aggregation dialog box, click the corresponding delete icon (**X**) in the last column.

Apply Your Changes

To apply your link aggregation configurations:

1. On the Link Aggregation dialog box, click **Apply**.

   A confirmation dialog box opens.

   IMPORTANT: If you aggregate interfaces that are currently in use, those interfaces are removed from deployment before aggregation occurs. This confirmation dialog box gives you the choice to proceed with aggregating the interfaces or to cancel your link aggregation changes.

2. Click **Aggregate Interfaces** to proceed. Otherwise, click **Cancel**.

Cluster Profiles

Configuration > Networking > Cluster Profiles

On this tab you can view all cluster profiles that you have configured. Cluster profiles allow you to manage multiple EdgeConnect appliances as a cluster and orchestrate flow redirection within a cluster. A cluster profile contains configuration settings that are applied to a cluster. Each peer in a cluster then inherits those same settings. For more information on clusters see, Clusters. Cluster profiles also do the following:

- Synchronizes user role information between cluster peers.
- Synchronizes bandwidth availability between two EdgeHA appliances.

Orchestrator comes with a built-in cluster profile called "Profile [EdgeHA]" and all EdgeHA appliances are automatically configured using this profile. Flow redirection is disabled for this profile because EdgeHA does not support Active/Active deployments.

The following table describes the information displayed on the Cluster Profiles tab for each cluster profile.

Field	Description
Name	The name of the cluster profile.
Interface/label	The interface label that is used for inter-cluster synchronization when this cluster profile is applied to a cluster.
Flow redirection	Indicates if flow redirection is enabled or disabled for the cluster profile.
Wait time	Indicates how long the system waits for an update from other cluster peers before redirecting a flow when this cluster profile is applied to a cluster.
User session sync	Indicates if user session synchronization is enabled or disabled for the cluster profile.
Secure	Indicates if encrypted control messages are enabled for the cluster profile.

Create a Cluster Profile

1. From the Cluster Profiles tab, click **+Add**.

   The Cluster Profile dialog box opens.

2. Enter the following details for the cluster profile.

Field	Description
Name	Enter a name for the cluster profile.
Interface/label	Select a LAN-side label to use for inter-cluster synchronization. This interface is used for flow redirection and user session synchronization within a cluster. It is recommended to create a LAN-side label for this purpose, such as a label called “CLUSTER”.
Flow redirection	Click to enable flow redirection for peers in a cluster.
Wait time	Enter a value in milliseconds. This is used in conjunction with flow redirection to control how long the system waits for an update from other cluster peers before redirecting a flow.
User session sync	Click to enable user session synchronization between cluster peers. This synchronizes user role information between peers in a cluster.
Secure	Click to enable encrypted control messages between cluster peers.

3. Click **Save**.

For information on applying cluster profiles to clusters, navigate to **Configuration > Networking > Clusters**, or see Clusters.

Clusters

Configuration > Networking > Clusters

On this tab you can view all clusters and apply profiles to clusters. There is also a link to the Cluster Profiles tab where you can create cluster profiles. Clusters are created using Site/Cluster Names, and the Site/Cluster name becomes the identifier for a cluster. All appliances with the same Site/Cluster Name are part of the same cluster. Tunnels are not formed between appliances with the same Site/Cluster Name.

- If you plan to create clusters and you have previously used the Site setting to prevent tunnels from forming between EdgeConnect appliances at different sites, you need to use tunnel exceptions to prevent tunnels from forming between the appliances.
- Starting with release 9.5, Site/Cluster Names are used to create Locations in HPE SSE.

IMPORTANT: You need to configure a Site/Cluster Name for EdgeHA pairs. The Site/Cluster Name must match precisely for each appliance in the pair.

IMPORTANT: Starting with release 9.5, Site Names are called Site/Cluster Names.

You can apply cluster profiles to clusters. Cluster profiles do the following:

- Orchestrate flow redirection within a cluster.

  For flow redirection guidelines and limitations, see TCP Accelerated Flows.
- Synchronizes user role information between cluster peers.

- Synchronizes bandwidth availability between two EdgeHA appliances.

For more information on cluster profiles, see Cluster Profiles.

The following table describes the information displayed on the Clusters tab for each cluster.

Field	Description
Cluster	The name of the cluster.
Appliances	The appliances in the cluster.
Cluster Profile	The cluster profile that is currently applied to the cluster.
Config Status	Indicates whether Orchestrator has fully pushed the cluster profile that contains all of the cluster configuration settings to all of the appliances in the cluster or not (**In Sync** or **Out of Sync**). This status does not convey information about reachability between cluster peers. That information is available in the Cluster Config dialog box, which you can access by clicking the information icon for a cluster.
Peer	The number of peers in the cluster.
Interface	The interface label that is used for inter-cluster synchronization.
User Session Sync	Indicates if user session synchronization is enabled or disabled for the cluster.
Flow Redirection	Indicates if flow redirection is enabled or disabled for the cluster.
Monitor	Click the chart icon to view Flow Redirection stats and charts for the cluster.
Alarms	Indicates if there are any cluster-specific alarms. If there are active alarms, you can click the linked alarm information. The Alarms tab opens and is filtered to show only those alarms and the appliances affected by the alarms.

Cluster Config Dialog Box

To view details about a cluster, click the information icon in the table for a cluster. The Cluster Config dialog box opens. It shows configuration details for the cluster profile that is applied to the cluster and reachability status between the peers in the cluster. The following table describes the fields shown on the Cluster Config dialog box.

Field	Description
Profile	Shows the configuration settings of the cluster profile applied to the cluster. **Name** - The name of the cluster profile. **Interface/label** - The interface label that is used for inter-cluster synchronization when this cluster profile is applied to a cluster. **Flow redirection** - Indicates whether flow redirection is enabled for the cluster or not. **Wait time** - This value (in milliseconds) is used in conjunction with flow redirection to control how long the system waits for an update from other cluster peers before redirecting a flow. **User session sync** - Indicates if user session synchronization is enabled or disabled for the cluster profile. **Status** - Indicates whether Orchestrator has pushed configuration settings in the cluster profile to all of the appliances in the cluster or not. If there are any active alarms for appliances in the cluster, those are listed. If you click the linked alarm information, the Alarms tab opens and is filtered to show only the active alarms for the cluster and the affected appliances.
Peers	Shows the following information about each peer in the cluster: **Appliance** - Name of the appliance. **IP** - IP address of the appliance. **Interface** - Interface label for the appliance. The reachability status between an appliance and each peer in the cluster (**Reachable** or **Unreachable**).

Add an Appliance to a Cluster

To add an appliance to a cluster, you need to assign the cluster name to the appliance. There are multiple ways to do this.

- You can assign a Site/Cluster name when setting up a new appliance using the Appliance Configuration Wizard (**Configuration > Overlays & Security > Discovery > Configuration Wizard**).
- You can assign a Site/Cluster name to a specific appliance by accessing the System Settings for the appliance from either the appliance tree or the System Information tab.

  1. In the appliance tree, locate an appliance and click the menu button, then click **System Information**. Or, on the System Information tab (**Administration > Software > Upgrade > System Information**), click the edit icon for an appliance.

     The System Information dialog box opens.
  2. Click **System Settings**.
  3. To add an appliance to a cluster, in the Site/Cluster name field enter the name of the cluster to add the appliance to.

4. To remove an appliance from a cluster, in the Site/Cluster name field delete the name of the cluster to remove the appliance from it.
5. Click **Save**.

- You can assign a Site/Cluster name to multiple appliances on the Cluster tab or on the Tunnels tab.
  1. From either the Cluster tab or the Tunnels tab, click **Sites/Clusters**.

     The Appliance Site/Cluster Info dialog box opens.
  2. To add an appliance to a cluster, find the appliance in the list. Click in the Site/Cluster column and enter the name of the cluster to add the appliance to.
  3. To remove an appliance from a cluster, find the appliance in the list. Then click in the Site/Cluster column and delete the name of the cluster to remove the appliance from it.
  4. Click **Apply**.

Applying or Removing a Cluster Profile

1. In the appliance tree, select the appliances to which you want to apply or remove a cluster profile.
2. Click **Apply Cluster Profiles**.

   The Apply Cluster Profiles dialog box opens.
3. Select the profile from the Cluster Profile menu.
4. To apply the cluster profile to the selected appliances, select the **Add** check box. To remove the cluster profile from the selected appliances, select the **Remove** check box.
5. Click **Apply**.

Bridge Groups

Configuration > *Networking* > *Bridge Groups*

The Bridge Groups tab displays details about configured bridge groups for appliances selected in the appliance tree. With a bridge group you can create a bridged (switched) interface while in inline router mode on the LAN side of an EdgeConnect OS. Doing this achieves the following:

- You can define bridging between multiple EdgeConnect SD-WAN LAN side physical interfaces, similar to switching between ports on a switch.
- Creates a BVI (Bridged Virtual Interface), so that the devices in the bridge group can use the IP address of the BVI as a default gateway to other IP networks.

In Orchestrator, BVI interfaces are referred to as “slan” and there are four predefined interfaces that you can use for a bridge group (slan0, slan1, slan2, and slan3). You can configure the BVI with the same parameters available on physical or sub-interfaces, such as segmentation, firewall zones, labels, DHCP server and relay, VRRP, BGP, OSPF, multicast, and branch NAT. If a bridge group is configured for an appliance, it is available to add as a LAN interface on the Deployment tab (**Configuration > Networking > Deployment**).

The following items are not currently supported with bridge groups:

- VXLAN
- Spanning Tree Protocol (STP).

  WARNING: EdgeConnect provides no direct layer 2 loop prevention. The LAN segment must be loop free.
- VLAN tagged BVI interfaces or sub-interfaces.
- Network Access Control (NAC)

  NOTE: All bridged unicast traffic is subject to policy and will appear in the flow table.
- Virtual EdgeConnects are not supported.
- Bonded/LACP interfaces cannot be used as bridge ports.

The table on the Bridge Groups tab displays the following information:

Column	Description
Appliance	Name of the EdgeConnect appliance.
Bridge Groups	If any bridge groups are configured for the appliance, the names of the bridge groups are listed in this column.
Bridge Groups Status	Status of the bridge group (up, down, or pending). **up** – the BVI interface is in service. **down** – the BVI interface is out of service. **pending** – the BVI interface configuration is not applied yet.
Interfaces	Physical or virtual interfaces included in the bridge group. A bridge group can include up to 16 interfaces.
MTU	Maximum transmission unit (MTU) size (in bytes) configured for the bridge group. The configured MTU overrides any existing MTU settings when the bridge group is deployed. The default size is 1500 bytes.
Admin Status	Indicates if the BVI is active (up) or inactive (down). The default is up.
Comment	Additional information about the bridge group.

Configure Bridge Groups

To add, change, or delete bridge groups for an appliance, click the edit icon in the applicable table row on the Bridge Groups tab. The Bridge Groups dialog box opens. Interfaces to be added to a bridge group must not be in use on the Deployment tab.

Add or Modify a Bridge Group

To add a bridge group, perform the following steps:

NOTE: To modify a bridge group, click the edit icon next to the bridge group. The Modify bridge group dialog box opens. Change the fields as described below, and then click **Apply**. You cannot modify an existing bridge group name, but you can change the other settings.

1. Click **Add** on the Bridge Groups dialog box.

   The Add bridge group dialog box opens.

2. Complete the following fields:

Field	Description
Bridge Group (interface) name	Select a name for the bridge group from the drop-down list (**slan0**, **slan1**, **slan2**, or **slan3**).
Interfaces to be grouped	Drag interfaces that you want to include in the bridge group from the __Available box to the **Grouped** box. You can include up to 16 interfaces.
MTU	Specify the MTU size (in bytes) to be applied to all interfaces in the group. The default size is 1500 bytes.
Admin status	Select **up** or **down**. The default is **up**. Selecting **down** takes the BVI out of service.
Comment	(Optional) Provide additional information about the bridge group.

3. Click **Add**.

After you have added the bridge group (slan interface), you must navigate to the Deployments tab (**Configuration > Networking > Deployment**) and configure the slan interface for use.

Delete a Bridge Group

To delete a bridge group listed in the table on the Bridge Groups dialog box, click the corresponding delete icon (**X**) in the last column.

WARNING: Deleting a bridge group removes the associated Deployment tab configuration.

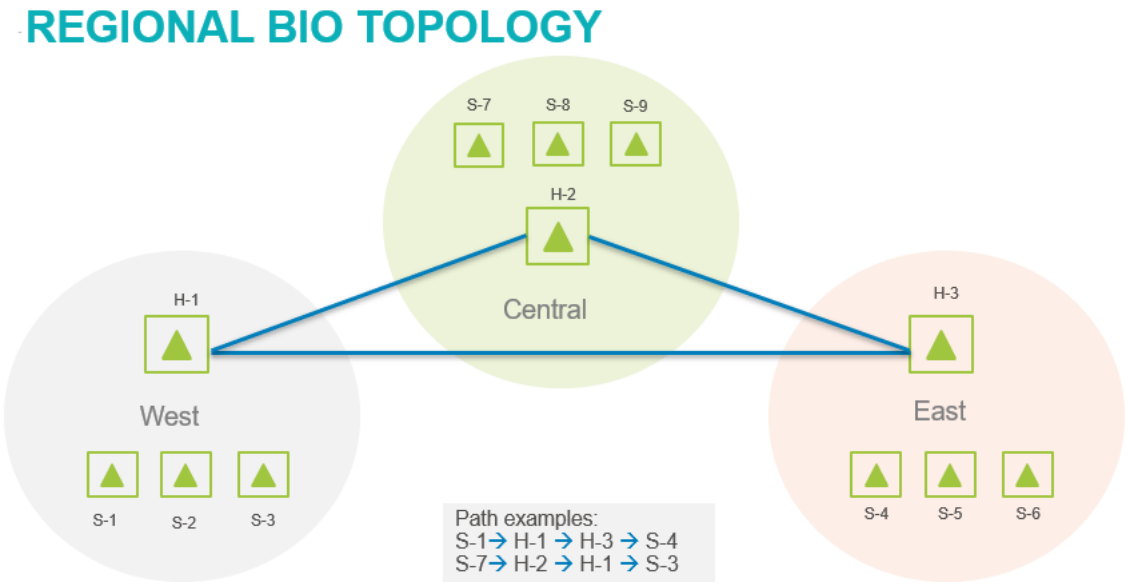
Regions

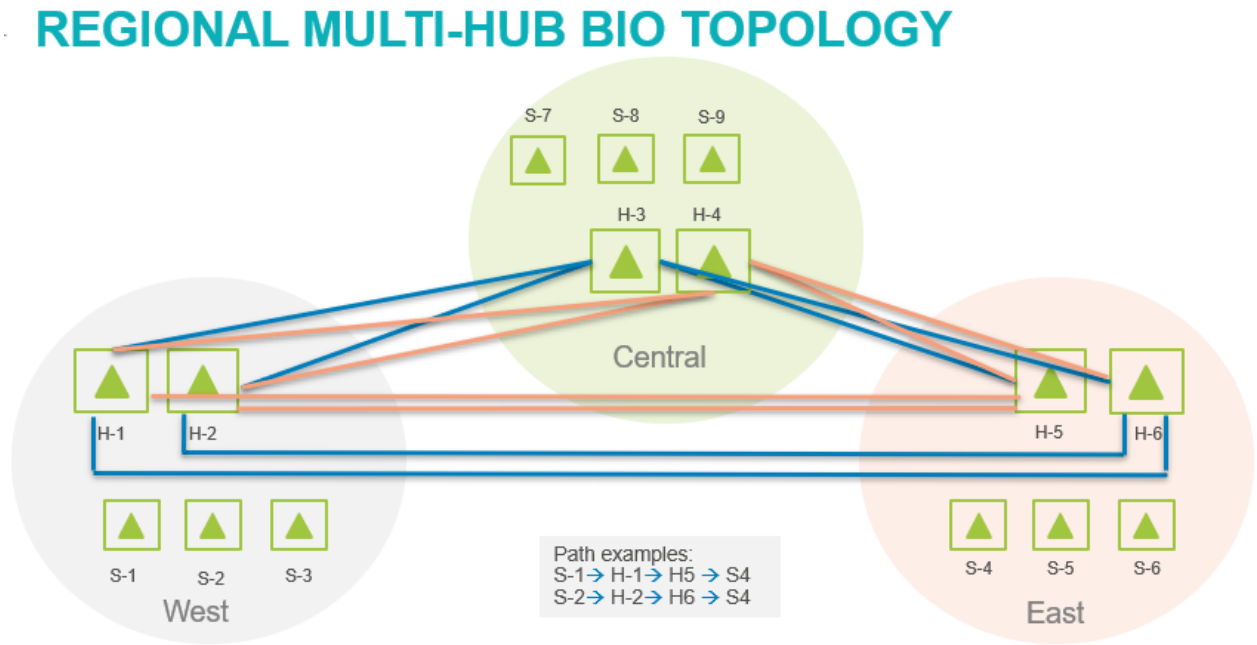
Configuration > Overlays & Security > Regions

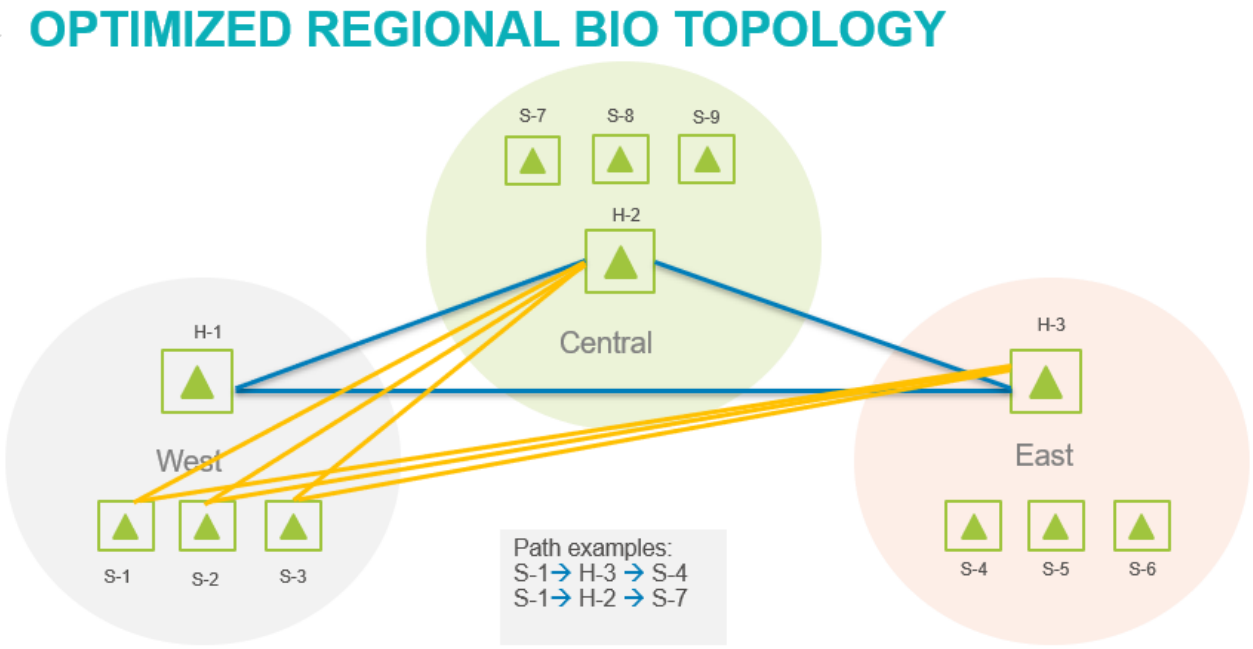
Use this tab to add or remove regions from the SD-WAN fabric and configure regional routing. The regions within your SD-WAN fabric can represent geographical regions, administrative regions, or a set of sites in the network that have common business goals.

Regional Routing

When activated, regional routing lets you manage your SD-WAN fabric by regions. This involves intra-region and inter-region route distribution across the SD-WAN fabric. The regions within your network can represent geographical regions, administrative regions, or a set of sites in the network that have common business goals. You can provide a different Business Intent Overlay (BIO)for each region by enabling regional routing and customizing BIOs per region. The following diagrams show examples of different regional network topologies you can build by activating regional routing.







To activate regional routing, click the **Enable Regional Routing** link. The Regional Routing dialog box opens. Select the **Enable regional routing** check box.

NOTE: With regional routing activated, any appliances configured as hubs will serve as hubs for all overlays applied to those appliances.

View Status

Click **View Status** to view the status of the added or updated appliances to regions.

Edit Regions

Complete the following steps to add a region or edit existing regions that you want to add to your overlays.

1. Click **Edit Regions**.
2. Click **New Region**.
3. Enter the name of your new region in the **Region Configuration** dialog box.
4. Click **Save**.

You can also edit an existing region.

1. Click the **Edit** icon next to the region you want to edit.
2. Enter the region name.
3. Click **Save**.

Navigate to the **Business Intent Overlay** tab to make further customizations to your regions and overlays.

Regional Routing

Configuration > Networking > Routing > Regional Routing

Use the Regional Routing dialog box to activate regional routing in Orchestrator. Select the **Enable regional routing** check box.

NOTE: With regional routing activated, any appliances configured as hubs will serve as hubs for all overlays applied to those appliances.

Use the Regions tab to manage your SD-WAN fabric by regions. For details and examples of regional network topologies, see Regions.

Routing Segmentation

Configuration > Networking > Routing > Routing Segmentation (VRF)

Use this tab to enable and disable routing segmentation across your network and apply unique configuration to your segments. Routing segmentation allows for the configuration of VRF (Virtual Routing and Forwarding)–style Layer 3 segmentation in your SD-WAN deployments. Note the following before configuring routing segmentation in Orchestrator:

- You must upgrade all EdgeConnect appliances and Orchestrator to version 9.0.
- All EdgeConnects must be configured to Inline Router mode.
- If a new appliance has been added to your network, or if an existing appliance has been replaced, you need to upgrade the appliance software to the appropriate version running in the network.
- After upgrading, segmentation is **disabled** by default. You will have to enable it on this tab.
- Regardless of whether segmentation is enabled or disabled, a **Default** segment is automatically created when you upgrade to 9.0.
- The system-generated Default segment cannot be deleted.
- After you enable routing segmentation, all existing configuration across your network is associated with the Default segment.

Add a New Segment

Before adding a segment, you must enable segmentation by moving the toggle at the top of the page. If Routing Segmentation is not enabled, you cannot make any modifications to the Default segment or add any new segments.

To add a new segment, click **+Add Segment** and enter a **Segment Name**. You can make further specifications by clicking the edit icon or by selecting the **+Add** icon in any of the columns in the table.

Segment Configuration

You can uniquely configure your segments by specifying the following on this page:

- Overlays & Breakout Policies
- Firewall Zone Policies
- Inter-Segment Routing & DNAT
- Inter-Segment SNAT
- Loopback

NOTE: Inter-Segment Routing & DNAT and **Inter-Segment SNAT** are applicable only if you are using different segments.

The following sections provide more details.

Overlays & Breakout Policies for Segments

Use this dialog box to configure overlays and breakout policies for your segments. This configuration determines the overlays used by each segment when traffic is originating from that segment and sent over the SD-WAN fabric to other sites. This configuration is also used when traffic breaks out locally to the Internet and Cloud Services using the Preferred Policy Order on the **Business Intent Overlay** (BIO) tab. For traffic to match what is on the specified BIO tab, ensure the following two conditions are true:

- BIO must include the defined segment policy
- The BIO match criteria must match the new flow

The overlays are arranged by priority defined in the **Match** field in the **Overlay Configuration** dialog box on the **BIO** page. You can specify if you want to include or skip the segment for each overlay by clicking **Include** or **Skip** icon in the table cell. By default, all overlays are included for all configured segments.

Include and Skip

If you want to skip an overlay, click the enabled **Include** icon and **Skip** appears grayed out. The segment will not be applied to the specified overlay. Click **Skip** again to include the segment; it will turn back to green. If an overlay is set to Skip, traffic will not match that overlay and moves to the next prioritized BIO. Additionally, if no BIOs match, traffic is dropped.

TIP: If overlay is set to **Skip**, Flow Details on the **Flows** tab displays the list of skipped overlays.

Firewall Zone Policies

Use this dialog box to enable and associate firewall zones to your segments. With segmentation enabled, firewall zone security policies are orchestrated and there is no need for Firewall Security Templates. After migration, deactivate the Security Policies Template in all Template Groups. If left active, the template will override any default-default segment security policies configured on this dialog box.

Before you begin Firewall Zone configuration, note the following:

- Review your existing security policies.
- Create a new security templates group with the new firewall zoning policies that only includes zones associated with LAN and WAN interfaces.
- Delete all rules in your previous Security Policy Template on the **Apply Template Groups** tab.
- Ensure you have selected the **Replace** option in the previous Security Policy Template.
- Save the previously used Security Policy Template. This deletes the security policy rules on your appliances.

Complete the following steps to set a rule or policy to your firewall zones within your segment.

1. Select the cell of the segment you want to update in the Matrix View. The **From Zone To Zone** dialog box opens.

   NOTE: If you are already in Table View, click **Add Rule**.
2. Enter the Source Segment in the **Source Segment** field. This is the segment that the firewall is starting from.
3. Enter the Destination Segment in the **Destination Segment** field. This is the segment where the firewall is going to.
4. Select **Add Rule**.
5. Complete the content in the table.

Field	Description
Priority	Enter the priority amount.
Match Criteria	Click the edit icon in this column to modify and create the match criteria for each zone.
Action	Select **Allow** or **Deny** to determine whether this zone will apply the selected segment.
Enabled	Select the check box to enable or clear it to disable.
Logging	Determines the filter for the zone-based firewall drop logging levels. You can select one of the following levels to apply: **None**, **Emergency**, **Alert**, **Critical**, **Error**, **Warning**, **Notice**, **Info**, or **Debug**.
Tag	Use tags to categorize or identify the purpose of a rule.
Comment	Any additional details about the firewall zone.

6. Click **Save**. The Save Segment Firewall Zone Policies dialog box opens.
7. Enter a comment (optional) in the **Audit Log Comment** field, and then click **Save**. Any text entered in the **Audit Log Comment** field appears on the Audit Logs tab.

NOTE: Firewall zones are unique to each segment. For example, the default zone in Segment X will not be the same default zone in Segment Y.

Inter-Segment Routing & DNAT Exceptions

Use this tab to configure inter-segment routing and DNAT rules when traffic is crossing between segments.

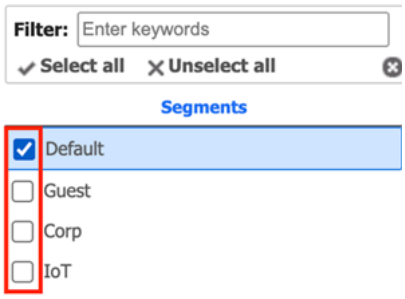
Starting with Orchestrator release 9.5.1, you can configure rules that allow multiple source segments to connect to one subnet destination. This configuration will form a group of rules. Source segments connected to the same subnet destination must be grouped in one rule. For example, if you select both “Guest” and “IoT” as the Source Segment for a subnet destination, you cannot add another rule that contains either “Guest” or “IoT” for that same destination.

Field	Description
Source Segment	Name of the segment that traffic is initiating from. You can select multiple source segments to create a group of rules.
Matches Destination IP	IP address that matches the destination segment IP address, before DNAT. The IP address is included in the defined policy match criteria.
Send to Segment	Name of the segment the packets are translated to from the matched destination IP address. This is included in the set criteria. Click in the cell to display the multi-selector, and then select or clear segments.
Translated Destination IP	IP address of the DNAT IP address when the segment is translated. **NOTE:** If DNAT is not needed, this field is empty.
Enabled	Indicates whether inter-segment DNAT is enabled or disabled within your segment. You can enable or disable multiple rules.
Comment	Any additional information.

Add a Rule

1. Click the **+Add** link under the Inter-Segment Routing & DNAT column to open the Inter-Segment Routing & DNAT dialog box.
2. Click **+Add Rule** to add a new rule.

   NOTE: To edit a rule that is part of a group of rules, you must delete the existing rule from the grouped rule by clearing the segment from the Source Segment list. Click in the Source Segment cell to display the multi-selector, as shown in the following screen capture.



3. Click in any cell to provide the details for the new rule (see field descriptions above).
4. Click **Save** to create the new rule or click **Cancel** to close the dialog box without making any changes.

NOTE: Inter-segment routing & DNAT rules are orchestrated globally to all appliances from this tab. To review rules on individual appliances, click **Inter-Segment Routing & DNAT Exceptions** and select the appliance in the tree. It is best practice to use only the globally orchestrated rules and avoid using local exceptions per appliance.

Delete a Rule

1. Click the corresponding delete icon (**X**). If the rule is a grouped rule, each rule that contains the same source segment will be deleted also. Deleting one rule, could result in multiple rules being deleted.
2. Click **Save**.

Inter-Segment SNAT Exceptions

This tab enables you to enable source network address translation to your segments.

NOTE: The default setting for SNAT is enabled for inter-segment traffic.

Field	Description
Source	Name of the segment that the SNAT is starting from.
Destination	Name of the segment that SNAT is translated to.
SNAT	Whether SNAT is enabled or disabled.

Loopback

Click **+Add** and you are redirected to the **Loopback Orchestration** tab. Select the segment you want to apply a loopback interface from the table, and then click **+Add Loopback Interface**.

Appliances

This column represents the amount of appliances the selected segment is enabled on.

Comment

Click the cell in the **Comment** column to add a comment including any additional information for that particular segment.

Delete a Segment

WARNING: Segmentation involves drastic changes to your physical network. Deleting segments can be service affecting. Carefully read this section before deleting any of your segments.

Deleting a segment removes all the segmentation configuration from all the appliances within your network. When you delete a segment, Orchestrator automatically deletes the following:

- The segment's association with the overlay and break-out policies
- The intra-segment and inter-segment firewall zone policies
- The inter-segment routing & DNAT rules
- The inter-segment SNAT rule
- The loopback interfaces associated with the segment
- The VTI interfaces associated with the segment
- All the interface and VLAN interfaces

Manual Tasks to Complete Before Deleting a Segment

The following configuration is disassociated from the segment and you need to manually delete the following:

- Any manual created tunnels
- BGP peers in the segment
- Internal subnet table rules
- Overlay ACL rules associated to the deleted segment

To delete a segment, click the **X** in the last column in the table. A Delete Routing Segment warning appears. Click **Delete** or **Cancel**.

Disable a Segment

To disable routing segmentation across your network, you need to delete all configured segments in the network, except the default segment (which cannot be deleted). After all the segments are deleted, navigate to this tab and move the toggle at the top of the page to disable.

Management Services

Configuration > *Networking* > *Routing* > *Management Services*

Use the Management Services tab to configure management services, regardless of whether routing segmentation is enabled.

- When enabled, management services are functional in the associated segment based on the selected interface.
- When disabled, all interfaces are available for configuration.

NOTE: Management services still function if routing segmentation is disabled in Orchestrator. In this case, you can use only the default configuration; that is, any interface with the Default segment.

Beginning with version 9.0, Orchestrator provides two tabs for configuring management services:

- **Management Routes:** Use this tab to configure static routes for management services traffic from an appliance (egress traffic).
- **Management Services:** Use this tab to specify the source IP address of the interface used for each management service.

While it is recommended that you now use the Management Services tab to configure services, you can continue to use the Management Routes tab if you are not required to specify source IP addresses for management services.

The Management Services tab displays the following fields:

Field	Description
Appliance	Name of the appliance selected.
Management Service	Management service used by your appliance.
Interface for Source IP Address	IP address of the interface used by the management service. By default, management services are configured to use any source IP address. You can modify the interface for the Source IP address by updating this field for the corresponding management service.
Source Segment	Name of the associated segment applied to the management service when your source IP address is selected.

Click the edit icon associated with the management service you want to configure.

Management Services Dialog Box

To configure a management service listed in this dialog box:

1. Click twice in the **Interface for Source IP Address** field associated with that service.

   A drop-down list of all interfaces configured for your appliance appears.

2. Select an interface.

   The Source Segment field updates automatically with the associated segment.

3. Click **Save**.

If the Interface for Source IP Address field is set to “any”, there is no control over which source IP address will be used for management services egress packets. Depending on the route lookup, the corresponding source IP configured in the Management Routes table is used as the source IP of the packet. If the Source IP is not configured (0.0.0.0) in the Management Routes table for the selected route, the egress interface’s IP address is used as the source IP address.

Descriptions of management service behaviors follow:

Service	Behavior
HTTP(S), Cloud Portal, Orchestrator, Secure Web Services	These services use the selected interface's Interface for Source IP Address as the source address to establish reachability and WebSocket connections to the Cloud Portal, Orchestrator, and Secure Web Services. HTTP/HTTPS uses the Interface for Source IP Address for connection as well. Secure Web Services uses this interface to connect with cloud-based OpenText™ Threat Intelligence (BrightCloud). Interfaces available in the Interface for Source IP Address drop-down list are in the Default segment; non-Default segment interfaces are not listed. If you do not select an interface, "any" is used by default, which means that any reachable interface in the Default segment is used to connect with OpenText Threat Intelligence (BrightCloud). The selected interface is used to periodically download the most current intelligence data to the URL classification and IP reputation databases on the appliances. It is also used to send cloud queries over the TLS Server to the OpenText Threat Intelligence (BrightCloud) platform. **WARNING:** Changing the source interface setting for "HTTP(S), Cloud Portal, Orchestrator, Secure Web Services" impacts robust appliance reachability to the Cloud Portal, Orchestrator, and Secure Web Services. If you change it, you could lose management plane connectivity. It is highly recommended that you leave this setting set to "any".
DHCP Relay, NTP, Other VRF mgmt Apps, NetFlow, RADIUS/TACACS+, SNMP, SSH, and Syslog	Each of these management services use Interface for Source IP Address as the source IP address. The source interface configured from the management route table is ignored if the Interface for Source IP Address is not "any". **NOTE:** Currently, EST is the only service included in "Other VRF mgmt Apps".

Inter-Segment Routing and DNAT Exceptions

Use this tab to configure inter-segment routing and Destination NAT (DNAT) rules when traffic is crossing between segments.

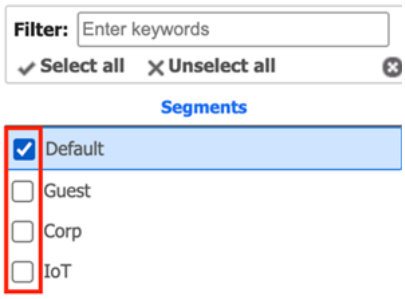
Starting with Orchestrator release 9.5.1, you can configure rules that allow multiple source segments to connect to one subnet destination. This configuration will form a group of rules. Source segments connected to the same subnet destination must be grouped in one rule. For example, if you select both “Guest” and “IoT” as the Source Segment for a subnet destination, you cannot add another rule that contains either “Guest” or “IoT” for that same destination.

Field	Description
Source Segment	Name of the segment that traffic is initiating from. You can select multiple source segments to create a group of rules.
Matches Destination IP	IP address that matches the destination segment IP address, before DNAT. The IP address is included in the defined policy match criteria. If you enter IPv6 in the Matches Destinations field, “No Translation” is the only option for Translated Destination.
Send to Segment	Name of the segment the packets are translated to from the matched destination IP address. This is included in the set criteria. Click in the cell to display the multi-selector, and then select or clear segments.
Translated Destination IP	IP address of the DNAT IP address when the segment is translated. **NOTE:** If DNAT is not needed, this field is empty.
Enabled	Indicates whether inter-segment DNAT is enabled or disabled within your segment. You can enable or disable multiple rules.
Comment	Any additional information.

Add a Rule

1. Click the edit icon to open the **Inter-Segment Routing & DNAT** dialog box.
2. Click **+Add Rule** to add a row to the table.

   NOTE: To edit a rule that is part of a group of rules, you must delete the existing rule from the grouped rule by clearing the segment from the Source Segment list. Click in the Source Segment cell to display the multi-selector, as shown in the following screen capture.



3. Click in any cell to provide the details for the new rule (see field descriptions above).
4. Click **Save** to create the new rule or click **Cancel** to close the dialog box without making any changes.

NOTE: Inter-segment routing & DNAT rules are orchestrated globally to all appliances from this tab. To review rules on individual appliances, click **Inter-Segment Routing & DNAT Exceptions** and select the appliance in the tree. It is best practice to use only the globally orchestrated rules and avoid using local exceptions per appliance.

Delete a Rule

1. Click the corresponding delete icon (**X**). If the rule is a grouped rule, each rule that contains the same source segment will be deleted also. Deleting one rule, could result in multiple rules being deleted.
2. Click **Save**.

Inter-Segment SNAT Exceptions

Configuration > Networking > Routing > Inter-Segment SNAT Exceptions

Use this tab to enable source network address translation to your segments. Select an appliance or group of appliances in the Orchestrator appliance tree to apply your Source NAT (SNAT) exceptions.

NOTE: The default setting for S-NAT is enabled for inter-segment traffic.

Field	Description
Appliance	Name of the segment that the SNAT exception is being applied to.
Source	Name of the segment that the SNAT is starting from.
Destination	Name of the segment that the SNAT is translated to and going to.
IPv4 SNAT	Indicates whether SNAT is enabled or disabled for the specified segment.
Comment	Any additional information.

BGP Tab

Configuration > Networking > Routing > BGP

On this tab, you can configure **BGP (Border Gateway Protocol)** for appliances and add their BGP peers (also known as BGP "neighbors"). You can also add and modify peer-based advertisement and redistribution rules. EdgeConnect has the following behaviors relative to **communities**:

- EdgeConnect will propagate any existing communities and can be configured to assign and append new communities, remove, or replace all communities for routes advertised to peers.
- Appliances can display up to ten communities per route.
- Appliances subnet-share communities with their EdgeConnect peers.
- Appliances advertise communities to remote peers, if learned from EdgeConnect peers.
- Appliances advertise communities to BGP neighbors.
- All BGP-learned subnets also appear in the appliance Routes table, displayed on the Routes configuration page. In addition, any AS Path or BGP Community information learned with a particular subnet will also be displayed with that subnet entry in the table.
- BGP route updates are not refreshed unless the peer specifically asks for it. To update the BGP routes, go to the **Peers** table and select **Soft Reset** in the desired row.
- BGP Equal-cost multi-path (ECMP) is supported for eBGP and iBGP. Multiple next-hops will be installed for the same prefix if all BGP path attributes are the same, enabling BGP to load balance egress traffic across multiple peers.
- When EBGP or IBGP metric values are not specified, the system defaults to the following settings:

Received MED Path Attributes	Expected Final Metric Value
MED > 0	MED value
MED = 0 or No MED attributes	Admin distance + 50

- A maximum of 64 BGP peers and 64 OSPF neighbors is supported per appliance, with 200 next-hops supported per interface.
- A small set of community numbers are used as internal communities that represent the source domain of a particular route:

Value	Description
100	Locally configured
101	Subnet shared (learned from another appliance)
102	Local BGP

Value	Description
103	Remote BGP (learned from another appliance)
104	Local OSPF
105	Remote OSPF (learned from another appliance)

These internal community values only use the appliance's local ASN in the ASN portion of the community. When the ASN portion of an attached community exactly matches the local ASN and the community portion exactly matches one of these internal values, they are flagged as internal communities only and stripped when advertising the route to BGP peers.

Click the **Summary** button on the BGP tab to display configuration details associated with the local appliance, such as its local AS number and router ID. Click the icon in the **BGP State Details** column to display a summary, including the number of routes learned and advertised via BGP by this appliance.

Click the **Peers** button on the BGP tab to display information about all configured peers for the appliances selected in the appliance tree. Click the icon in the **Peer Details** column to display the connection status of each peer that is configured for the appliance.

Filter by Segment

To filter the rows displayed in the BGP table by segment:

- Select **Default** from the **Segment** drop-down list to display for the system-supplied default segment, or
- Select one of the other listed segments, which reflect the custom segments defined using Routing Segmentation (Configuration > Networking > Routing > Routing Segmentation (VRF)).

Select **All** to display for all segments, which is the default setting.

The table below describes the fields displayed for the BGP configuration.

Field	Description
Appliance	Name of the appliance.
Segment	Name of the segment being used, if enabled.
Peer IP	IP address of the EdgeConnect peer. IPv4 and IPv6 are both supported. **NOTE:** When using VXLAN with BGP, only IPv4 addresses are supported.
Local Interface	A list of the interfaces that can be chosen: **Any**, **lan0**, **wan0**, or **wan1**.
Peer ASN	Peer's Autonomous System Number.

Field	Description
Peer State	State of the peer. A peer state of Established indicates that full adjacency has been established and routes can be advertised to and learned from that peer.
Soft Reset	Click the **Soft Reset** button to manually request a route update from the BGP peer without resetting the session. **NOTE:** Before you perform a soft reset, ensure that Soft Reconfiguration is disabled for this BGP peer.
Soft Reconfiguration	Indicates whether Soft Reconfiguration is enabled for this BGP peer.
Established Time	Final peer state that indicates neighbor connection as complete.
Type	Governs what kinds of routes the appliance is allowed to advertise to this BGP peer. These routes are itemized as Route Export Policies.
Inbound Route Map	Route map being used for the inbound traffic.
Outbound Route Map	Route map being used for the outbound traffic.
Local Preference	Local preference is the first attribute an EdgeConnect appliance looks at to determine which route towards a certain destination is the “best” one. This value is not exchanged between external BGP routers. Local preference is a discretionary BGP attribute. Default value is 100. The path with the highest local preference is preferred.
MED	*Multi Exit Discriminator.* When BGP chooses the best route to reach a certain destination, it first looks at the local preference and AS path attributes. When the local preference and AS path length are the same for two or more routes towards a certain prefix, the Multi Exit Discriminator (MED) attribute is chosen. With MED, the lowest value is preferred. **NOTE:** If you configured the Metric Delta parameter in an earlier version of our software, this value has been translated into a MED value.
Input Metric	Metric that is advertised with the route when shared.
Enable Imports	Allows the learning of routes from this specific BGP peer.
AS Prepend Count	Learned path from an external prepend between a remote BGP site to local BGP peers.
Next-Hop-Self	Advertised route connected to a CE router that an EdgeConnect appliance learns from the eBGP with a PE router.
Override ASN	Indicates whether routes are advertised to the BGP peer where the BGP peer’s own ASN is in the AS-Path.
Keep Alive Timer	Interval, in seconds, between keep alive signals to a peer.

Field	Description
Hold Timer	When availability to a peer is lost, this specifies how long to wait before dropping the session.
BFD	Indicates whether BFD is enabled for the BGP peer. This field is set to N/A if BFD is not supported on the appliance.
Adjacency	Indicates the adjacency of the BGP peer (Single-Hop or Multi-Hop). This field is set to N/A if BFD is not supported on the appliance.
Peer Details	Additional details about the peer or its state.

To edit the BGP configuration for one of the listed appliances, click the edit icon in the left column of the table.

BGP Information

Use this window to enable BGP for your appliances and to configure BGP peers. Complete the following steps to start BGP configuration.

1. Move the toggle to **Enable BGP**.
2. Complete the following fields.

Field	Description
Autonomous system number (ASN)	Configure this number as needed for your network.
Router ID	This router identifier is the IPv4 address by which the remote peer can identify this appliance for purposes of BGP.
Route Target	The BGP route target that identifies the routes associated with a segment. The route target must be in the following form: 2bytesASN:4bytesID or 4byteASN:2byteID (for example, 65001:1200). **NOTE:** Route target is available for all segments.
Import:Export	Select this check box to specify a unique import and export value.

Field	Description
Graceful restart	Enable receiver-side graceful restart capability. EdgeConnect retains routes learned from the peer and continues to use it for forwarding (if possible) if/when a BGP peer goes down. The retained routes are considered stale routes. They will be deleted and replaced with newly received routes. **Max Restart Time** – Specifies the maximum time (in seconds) to wait for a Graceful Restart capable peer to come back after a peer restart or peer session failure. **Stale Path Time** – Specifies maximum time (in seconds) following a peer restart that EdgeConnect waits before removing stale routes associated with that peer.
AS path propagate	Select this check box to enable this appliance to send the full AS path, associated with a prefix to other routers and appliances, avoiding routing loops. This will provide the learned path from an external prepend between a remote BGP site to local BGP peers.
Log BGP update messages	Select this check box to enable logging of BGP peer messages on the segment. This feature provides detailed logging of BGP update messages exchanged between BGP peers. It captures and logs every detail for each update, whether sent or received, including key attributes such AS Path list, routes, next-hop, origin, MED, and so on as they are shared with or received from a BGP peer.

3. Under the Common settings for all segments header, complete the following fields:

Field	Description
Max route updates per peer	The maximum number of route updates for each peer running route loop detection. The default value is 10, and the range is 5–100.
Detection interval	The interval, in minutes, at which route advertisement loop detection runs. A route advertisement loop occurs when the same route is being advertised, removed, and re-advertised repeatedly within a short time period. The default value is 15, and the range is 1–60.

To add a BGP peer, select **Add**. The Add Peer dialog box opens.

Add Peer

Complete the following fields to add a BGP peer.

Field	Description
Peer IP	IP address of the EdgeConnect peer. IPv4 and IPv6 are both supported.
Peer Adjacency	To specify the adjacency of the BGP peer, click **Single-Hop** or **Multi-Hop**. Single-Hop is the default selection.
EVPN Peer	Select this check box to enable the extended BGP technology for casting VXLAN information. Enabling EVPN Peer disables IPv4/IPv6 unicast routing for this peer. Only one address family is supported per BGP peer. Only IPv4 is available when you enable VXLAN. **NOTE:** EVPN Peer is only displayed for the default segment (VRF ID = 0). **NOTE:** This field is only available if you configured VXLAN using the VXLAN template or by configuring it on the VXLAN tab and you selected a VXLAN tunnel endpoint (VTEP) source interface. For more information on configuring VXLAN and a VTEP source interface, see VXLAN Tab or VXLAN Template.
Local Interface	You can specify the source address or interface for a specific BGP peer. Select the interface from the drop-down list: **any**, **lan0**, **wan0**, or **wan1**. If you selected the EVPN Peer check box, this value is automatically set to the VTEP source interface you set when you configured VXLAN.
Peer ASN	Replace all ASNs in the AS-Path of routes advertised to this peer with the appliance ASN.
Override ASN	Select this check box to advertise routes to the BGP peer where the BGP peer's own ASN is in the AS-Path. All instances of the BGP peer ASN are replaced with the local ASN of the appliance in all routes advertised to the BGP peer.
Peer Type	Select the type of peer from the drop-down list: **Branch** or **PE-router**. If you selected the EVPN Peer check box, this value is automatically set to Branch.
Admin Status	Select whether you want the Admin Status **UP** or **DOWN**.

Field	Description
Soft Reconfiguration	Select this check box to prevent the appliance from sending a route-refresh message to the BGP peer when a policy is changed. When enabled, the appliance will apply policy changes against BGP peer learned routes stored in memory. **NOTE:** To request a route update from the peer, click the **Soft Reset** button for the peer on the BGP tab. Before you perform a soft reset, ensure that Soft Reconfiguration is disabled for this BGP peer.
Next-Hop-Self	Select this check box to enable the next-hop-self.
Inbound route map	Route map for inbound traffic. Select the edit icon to load or configure inbound route maps.
Outbound route map	Route map for outbound traffic. Select the edit icon to load or configure outbound route maps.
Outbound route map	Route map for outbound traffic. Select the edit icon to load or configure outbound route maps.
BFD	Select this check box to enable BFD for the BGP peer. This field is not displayed if BFD is not supported on the appliance. **NOTE:** Before you select this check box, enable and configure BFD from the BFD tab.
Keep Alive Timer	Interval, in seconds, between keep alive signals to a peer.
Hold Timer	Specified time to wait before dropping the session when the reachability to a peer is lost.
Enable MD5 Password	Select this check box to add a password to authenticate the TCP session with the peer.

NOTE: Multiple address families are not supported on a single BGP peer. You must configure a separate BGP peer for each address family. IPv6 peers are only capable of exchanging IPv6 routes. IPv4 peers are only capable of exchanging IPv4 routes.

BGP Inbound and Outbound Route Redistribution Maps

Route Maps are policies applied to IP routes during redistribution between routing protocols. They have **Match Criteria** and **Set Actions** that allow for filtering routes or modifying metrics and attributes for routes that meet the criteria defined in the match statement. Route-map rules follow a top-down order based on the sequence number defined for each entry.

EdgeConnect Enterprise supports applying Route Maps inbound from and outbound to BGP peers and outbound to OSPF neighbors and the SD-WAN Fabric. It is best practice to use Orchestrator to apply Route Maps using templates.

You can specify up to 20 BGP inbound route maps, 20 BGP outbound route maps, and 128 rules per route map.

You can specify up to 6 pipe separated prefixes for each rule applied to a route map as shown in this example: *172.168.24.0/21 ge 21 le 21|172.168.40.0/21 ge 21 le 21.*

You can add, delete, rename, or clone route maps using this window. You can add rules to your route map by clicking **Add Rule**. A route map without any enabled rules is treated as a default deny all.

Prefix Match Criteria

The default for prefix match criteria is exact-match + greater-than. Both the specified prefix and any subnets of that prefix will be matched, up to a length of 32 for IPv4 or 128 for IPv6 (subnet sharing route maps only).

Less-than-or-equal-to (LE) and greater-than-or-equal-to (GE) clauses can also be applied to specify the inclusion of certain subnets.

To match a default-route, deny 0.0.0.0/1, deny 128.0.0.0/1, and then permit any.

GE Clause

If a GE clause is applied, the rule will also include all prefixes that have a prefix length greater than or equal to the GE value and less than or equal to 32 or 128 (for IPv6).

Example: A.B.C.D/X GE Y

In this example the following will be included:

- The exact match to A.B.C.D/X
- All the prefixes that belong to the subnet A.B.C.D/X that have a length greater than or equal to Y and less than or equal to 32

For example, you have a route map entry of 192.168.0.0/16 GE 24, and a peer advertises the following prefixes:

- 192.168.0.0/16
- 192.168.1.0/24
- 192.168.2.0/25
- 192.168.2.128/25

192.168.0.0/16 – The prefix length is 16, which is not greater than or equal to 24. This route does not match.

192.168.1.0/24 – The prefix length is 24, which is equal to the specified value. Therefore, this route matches the condition.

192.168.2.0/25 – The prefix length is 25, which is greater than or equal to 24. Therefore, this route matches the condition.

192.168.2.128/25 – The prefix length is 25, which is greater than or equal to 24. Therefore, this route also matches the condition.

So, the routes that would be matched by the given route map entry are:

- 192.168.1.0/24
- 192.168.2.0/25
- 192.168.2.128/25

LE Clause

If an LE clause is applied, the rule will also include all prefixes that have a prefix length less than or equal to the LE value.

Example: A.B.C.D/X LE Y

In this example the following will be included:

- The exact match to A.B.C.D/X
- All the prefixes that belong to the subnet A.B.C.D/X that have a length greater than or equal to X and less than or equal to 32
- All the prefixes that belong to the subnet A.B.C.D/X that have a length less than or equal to Y

For example, you have a route map entry of 192.168.0.0/16 LE 24, and a peer advertises the following prefixes:

- 192.168.0.0/16
- 192.168.1.0/24
- 192.168.2.0/25
- 192.168.2.128/25

192.168.0.0/16 – The prefix length is 16, which is less than or equal to 24. This route matches.

192.168.1.0/24 – The prefix length is 24, which is equal to the specified value. This route matches.

192.168.2.0/25 – The prefix length is 25, which is not less than or equal to 24. This route does not match.

192.168.2.128/25 – The prefix length is 25, which is not less than or equal to 24. This route does not match.

So, the routes that would be matched by the given route map entry are:

- 192.168.0.0/16
- 192.168.1.0/24

Combining LE and GE Clauses

Example: A.B.C.D/X LE Y GE Z

In this example the following will be included:

- The exact match to A.B.C.D/X

- All the prefixes that belong to the subnet A.B.C.D/X that have a length less than or equal to Y
- All the prefixes that belong to the subnet A.B.C.D/X that have a length greater than or equal to Z and less than or equal to 32

Exact Match

If both GE and LE clauses are specified and are equal, the rule will result in an exact match.

Example: A.B.C.D/X LE Y GE Y

In this example, the following will be included:

- The exact match to A.B.C.D/X
- The exact match to the subnet A.B.C.D/X that has a length equal to Y

You can specify the following fields in each rule for the selected route map.

Priority (Inbound and Outbound)

Field	Description
Priority	If you are using Orchestrator templates to add rules, Orchestrator will delete all entries from **1000 – 9999** before applying its policies. You can create rules with higher priority than Orchestrator rules (**1 – 999**) and rules with lower priority (**10000 – 19999** and **25000 – 65534**). **NOTE:** The priority range from **20000** to **24999** is reserved for Orchestrator. When adding a rule, the priority is incremented by 10 from the previous rule. The priority can be changed, but this default behavior helps to ensure you can insert new rules without having to change subsequent priorities.

Select Match Criteria (Inbound)

Source Protocol	Complete the Following Fields (based on protocol selected)
BGP	Prefix + optional LE/GE parameters BGP Communities

Select Match Criteria (Outbound)

Source Protocol	Complete the Following Fields (based on protocol selected)
Local/Static	Prefix + optional LE/GE parameters

Source Protocol	Complete the Following Fields (based on protocol selected)
SD-WAN (Local/Static)	Prefix + optional LE/GE parameters BGP Communities
BGP	Prefix + optional LE/GE parameters BGP Communities
OSPF	Prefix + optional LE/GE parameters OSPF Tag
SD-WAN (BGP)	Prefix + optional LE/GE parameters BGP Communities
SD-WAN (OSPF)	Prefix + optional LE/GE parameters OSPF Tag BGP Communities
SD-WAN (CFGSET)	Prefix + optional LE/GE parameters BGP Communities
SD-WAN (RIP)	Prefix + optional LE/GE parameters BGP Communities
SD-WAN (OVERLAY)	Prefix + optional LE/GE parameters BGP Communities
SD-WAN (IAPVPN)	Prefix + optional LE/GE parameters BGP Communities
OAP-BGP	Prefix + optional LE/GE parameters BGP Communities
OAP-OSPF	Prefix + optional LE/GE parameters OSPF Tag
OAP-CFGSET	Prefix + optional LE/GE parameters
OAP-RIP	Prefix + optional LE/GE parameters
OAP-OVERLAY	Prefix + optional LE/GE parameters
OAP-IAPVPN	Prefix + optional LE/GE parameters
OAP-STATIC	Prefix + optional LE/GE parameters
OAP-DIRECT	Prefix + optional LE/GE parameters

Source Protocol	Complete the Following Fields (based on protocol selected)

Set Actions (Inbound and Outbound)

Field	Description
Permit	Enable or disable. This setting allows or denies the route map.
BGP Local Preference	Best BGP destination. The default value is 100.
Metric	Metric for the route.
BGP Communities	Label of extra information that is added to one or more prefixes advertised to BGP neighbors. There are three options for how this information is added: **Append** – Click to add this information to the prefix when the route is advertised to BGP neighbors. You can append up to six BGP Communities in inbound and outband route maps using BGP route map rules. The maximum number of communities supported for an advertised or received BGP route is 11. For the No-Export community, enter 655535:65281 in the BGP Communities field of the Add Rule page. For the No-Advertise community, enter 65535:65282 in the BGP Communities field of the Add Rule page. **Override** – Click to replace the communities in the route with the community specified. **Remove** – Click to remove this information from the prefix when the route is advertised to BGP neighbors.
Nexthop	Advertised route connected to a CE router that an EdgeConnect appliance learns from the eBGP with a PE router.
ASN Prepend Count	Original route path that was used. **NOTE:** This field is displayed only for the Outbound redistribution map.
Comment	Comment you want to include.

The following table describes the redistribution commands supported in the BGP routing protocol.

Command	Redistribution Support
Match prefix	Yes
Set metric	Yes
Set tag	Yes

BGP ASN Global Pool

Configuration > Networking > Routing > BGP ASN Global Pool

Use this dialog box to configure the ASN Range to assign Autonomous System Numbers (ASNs) for new appliances. Note the following before configuration:

- ASNs are applied only to new appliances. The ASNs configured in this dialog box do not impact or change any previous or manually configured ASNs.
- ASN Range is configured for Default Segment and cannot be changed.
- ASN Orchestration assigns the same ASN to EdgeHA appliances.
- ASN Orchestration assigns the same ASN to appliances with same site name.
- If you use Orchestrator ASNs, do not use ASN 65534 and 1. These numbers are reserved.

Enter the start and end ranges for the ASNs. Click the **+Add Reserved ASN** to exclude any ASNs from being applied to an appliance. You can reassign ASNs manually by using the BGP tab.

Routes Tab

Configuration > Networking > Routing > Routes

Each appliance builds a route table with entries that are added automatically by the system, manually by a user, or learned from a routing protocol (SD-WAN Fabric Subnet Sharing, BGP, or OSPF). This tab allows you to view all routes for all appliances.

You can filter the type of routes displayed by clicking **All**, **Local / Static**, **SD-WAN Fabric**, **BGP**, **OSPF**, or **OAP**.

This tab also links to the following tabs: BGP, OSPF, BFD, Peer Priority, Admin Distance, and Multicast. Clicking **Enable Subnet Sharing with System Templates** opens the Templates tab and launches the Add/Edit Template dialog box.

Route Maps

Route Maps are policies applied to IP routes during redistribution between routing protocols. They have **Match Criteria** and **Set Actions** that allow for filtering routes or modifying metrics and attributes for routes that meet the criteria defined in the match statement. Route-map rules follow a top-down order based on the sequence number defined for each entry.

EdgeConnect Enterprise supports applying Route Maps inbound from and outbound to BGP peers and outbound to OSPF neighbors and the SD-WAN Fabric. It is best practice to use Orchestrator to apply Route Maps using templates.

Route mapping is supported for the following protocols and the direction of those protocols:

- Local, static to SD-WAN fabric
- BGP, OSPF to SD-WAN fabric
- SD-WAN fabric to BGP Outbound peers
- Local, BGP, OSPF to BGP outbound peers
- Local BGP Peers to EdgeConnect BGP sessions

The following table lists the routing protocols and the associated commands that are supported.

Command	Redistribution Support	BGP	OSPF	SD-WAN	Local/Static
Match prefix	Yes	Yes	Yes	Yes	Yes
Set metric	Yes	Yes	Yes	Yes	Yes
Set tag	Yes	Yes	Yes	Yes	Yes

You can specify up to 20 route maps per protocol per direction, 128 rules per route map, and six prefixes per rule. A route map without enabled rules is treated as a default deny all.

Additionally, if a route map is not selected for BGP, OSPF, or SD-WAN redistribution points this is also considered a deny-all. You must select a route map to advertise routs via one of the protocol intersections.

Filter by Subnet

To filter the routes displayed in the Routes table by subnet, enter the subnet in the **Filter by subnet** field, and then click **Apply**.

Filter by Segment

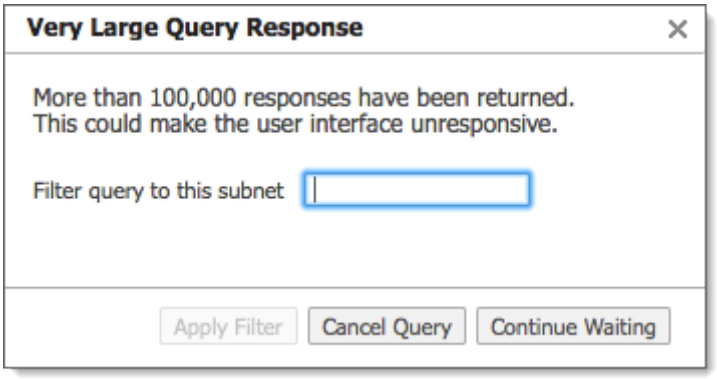
To filter the routes displayed in the Routes table by segment:

- Select **Default** from the **Segment** drop-down list to display for the system-supplied default segment, or
- Select one of the other listed segments, which reflect the custom segments defined using Routing Segmentation (Configuration > Networking > Routing > Routing Segmentation (VRF)).

Select **All** to display for all segments, which is the default setting.

If the number of the routes filtered exceeds 500,000, a **Very Large Query Response** pop-up will display. To mitigate this issue, you can filter by subnet and/or segment, cancel the query, or continue waiting.

NOTE: If the number of the routes filtered is greater than 500,000 the following pop-up will display.



Segment

The segments you have configured on the Routing Segmentation tab are listed in the Segment field. After you specify the segment, the Routes table displays only the routes belonging to that segment.

The following information is displayed for each route listed in the table:

Field	Description
Appliance	Name of the appliance.
Segment	Routes displayed belonging to this segment.
Subnet/Mask	Actual subnet to be shared or learned.
Next Hop	Next hop IP address for the route. A maximum of 200 next-hops are supported per logical interface.
Interface	Interface for outgoing traffic. Display only.
Zone	Firewall zone associated with the route.
State	Shows whether the route is up or down.
VXLAN	Indicates if static VXLAN is configured (yes or no)
VNI	Indicates the Virtual Network Identifier for the route.
VTEP peer MAC	Indicates the MAC address for the VXLAN Tunnel End Point peer.
Metric	Metric of the subnet. Value must be between 0 and 100. When a peer has more than one tunnel with a matching subnet (for example, in a high availability deployment), it chooses the tunnel with the lower numerical value.

Field	Description
Type	Indicates one of the following route types: **Auto (System)** – Automatically added subnets of interfaces on this appliance. **Auto (SaaS)** – Automatically added subnets from SaaS services. **Added by user** – Subnets manually added or configured on this appliance. **SP: Hostname** – Subnets added by exchanging information with peer appliances. If the peer has learned the subnet from a remote BGP or OSPF peer, that information is appended. **<BGP peer Type>: <BGP peer ip>** – Subnets added by exchanging information with local BGP peers. **OSPF: OSPF neighbor IP** – Subnets added by exchanging information with local OSPF peers. **<BGP peer Type> EVPN: <BGP peer ip>** – Subnets added by exchanging information with local EVPN enabled BGP peers. **OAP<device id>(STATIC)** – Static subnets learned from ORO (Overlay Route Orchestrator). **OAP<device id>(DIRECT)** – Direct (connected) subnets learned from ORO. **OAP<device id>(OSPF)** – Subnets added by exchanging information with an OAP (Overlay Agent Protocol) OSPF neighbor. **OAP<device id>(EBGP)** – Subnets added by ORO exchanging routing information with a router outside the company-wide network. **OAP<device id>(IBGP)** – Subnets added by ORO exchanging routing information with a router inside the company-wide network. **OAP<device id>(BGP)** – Subnets added by exchanging information with an OAP BGP peer in an external network. **OAP<device id>(IAP-VPN)** – Instant Access Point subnets learned from ORO. **OAP<device id>(OVERLAY)** – Subnets added by ORO. **OAP<device id>(RIP)** – Routing Information Protocol subnets learned from ORO. **OAP<device id>(CFGSET)** – BGW (branch gateway) subnets learned from ORO.

Field	Description
Nexthop Tunnel	Next hop tunnel for the route.
Label	Shows the label applied to the route and the current state (Up or Down) of the IP SLA rule(s) associated with the route label. The label can cause the route to go DOWN depending on the status of associated IP SLA rules. A route label can be associated with one or more IP SLA rules. Depending on the IP SLA Condition set on the label, the route is marked DOWN if *any* of the IP SLA rules are DOWN or only when *all* IP SLA rules associated with the label are DOWN. See Route Labels.
Region	Indicates the physical region where the appliance is located.
Advertise OSPF Tag to SD-WAN Fabric	Shows whether an OSPF tag for the route is advertised to the SD-WAN fabric (yes or no).
SD-WAN Site ID	Numeric ID used for advertising the route in the SD-WAN fabric.
Subnet Message Type	Indicates the highest subnet version supported by the route.
Additional Info	Indicates any tags for restricting route lookups: **Tag FROM LAN** – Used to restrict route lookups to traffic arriving on a LAN–side interface. **Tag FROM WAN** – Used to restrict route lookups to traffic arriving on a WAN–side interface. **NOTE:** If the route is a BGP route and EVPN is enabled for the route, only the import route target (labeled “Route Target”) is displayed. The export route target is not displayed.
Comment	Any additional information you would like to include.

To edit a route, select the edit icon in the Routes table.

Route Table Lookup Criteria

Each Route table has lookup criteria that is used in the following order:

- Longest prefix match
- Admin distance of source protocol (lower is better)

- Metric (lower is better)
- Peer Priority (lower is better; default value is 1024)
- BGP AS_PATH (if present; shorter is better)
- Flow-based ECMP

If there are two or more routes that match all the above criteria, use multiple routes.

System-Generated Default Routes

Orchestrator may display **system-generated default routes** (for example, *0.0.0.0/0*) in route tables with metric values in the 1250 range (such as 1250, 1251, or 1253).

These routes are automatically created by EdgeConnect as part of its internal SD-WAN forwarding and path-selection logic.

Characteristics of system-generated routes

System-generated default routes have the following characteristics:

- They are created automatically by EdgeConnect and are not user-configured.
- They are used internally by the system to ensure deterministic routing and path selection.
- They may appear alongside user-configured and learned routes in the Orchestrator Routes view.
- They are not advertised externally (for example, through BGP or OSPF).
- They are not configurable, redistributable, or tunable using route maps or policy.

Routes displayed in Orchestrator can include:

- System-generated routes
- User-configured static or local routes
- Learned routes (for example, via Fabric, BGP, or OSPF)

About metric values

Metric values associated with system-generated default routes (including 1250-range values) are:

- Assigned internally by EdgeConnect
- Not user-configurable
- Not documented for interpretation or tuning

These values are implementation details and should not be interpreted as traditional routing metrics or correlated with specific overlays, tunnels, or WAN labels.

Interpretation guidance

When system-generated default routes appear in the Routes table:

- Treat them as internal routing artifacts, not user-managed routes.
- Do not attempt to modify or suppress them through configuration.

- In most cases, they can be observed but safely ignored, unless troubleshooting internal routing behavior under the guidance of Aruba support.

Admin Distance Configuration

You can configure the admin distance by using the Admin Distance template on the **Templates** tab. The default settings in this template determine the most reliable route with the use of admin distance. See the table below for the various default admin distances per route type.

Route Type	Default Admin Distance
Local	1
Subnet Shared - Static Routes	10
Subnet Shared - BGP Remote	15
Subnet Shared CFGSET	15
Subnet Shared IAPVPN	15
Subnet Shared - OSPF Remote	15
Subnet Shared Overlay	15
Subnet Shared RIP	15
eBGP	20
OAP BGP	25
OAP CFGSET	25
OAP Direct	25
OAP IAPVPN	25
OAP OSPF	25
OAP Overlay	25
OAP RIP	25
OAP Static	25
OSPF	110
iBGP	200

Navigate to the **BGP** and **OSPF** tabs for more information about applying or configuring your route maps.

Edit or Add Routes

The following table describes the elements in the Routes dialog box. They represent various features you can apply to your route.

Field	Description
Automatically advertise local LAN subnets	Indicates whether the system-created LAN subnets of your appliance should be advertised to your peers.
Automatically advertise local WAN subnets	Indicates whether the system-created local WAN subnets of your appliance should be advertised to your peers.
Metric for automatically added routes	Metric assigned to subnets of interfaces on this appliance. Specify a value from 0 to 100. The default value is 50. When a peer has more than one tunnel with a matching subnet (for example, in a high-availability deployment), it chooses the tunnel with the lower metric value.
Redistribute routes to SD-WAN fabric	Route redistribution map for the SD-WAN fabric. Click the edit icon next to this field and specify the appropriate route redistribution map.
Filter routes from SD-WAN fabric with matching local ASN	Indicates whether to filter routes from the SD-WAN fabric with matching local Autonomous System Number (ASN).
Include BGP local ASN to routes sent to SD-WAN fabric	Indicates whether all routes must carry local ASN over subnet sharing to remote EdgeConnect peers.
Tag BGP communities to routes	Send the specified communities with routes that are advertised to both SD-WAN fabric peers and BGP peers, if the routes are learned from any of the following source protocols: Local/Static SD-WAN (Local/Static) SD-WAN (BGP) SD-WAN (OSPF) If you select this option, enter the BGP communities you want to be tagged in the field.

Field	Description
Communities	BGP communities to share. A community must be a combination of two numbers (0 to 65535) separated by a colon. For multiple communities, use a comma to separate them. You can have up to nine communities per route shared with subnet sharing. *Subnet sharing* is the protocol used to exchange routes between EdgeConnect appliances across the SD-WAN fabric.
Use SD-WAN fabric learned routes	Indicates whether to use SD-WAN fabric learned routes.
Enable Equal Cost Multi Path (ECMP)	Indicates whether you want to enable Equal Cost Multi-Path routing support.
Segment pruning	Controls whether route updates for a segment (VRF) are exchanged across an SD-WAN tunnel only when that segment is configured on both peer appliances. Available settings: **OFF** — Disables pruning; route updates for all segments are exchanged with all peers. **auto** — The default setting. Pruning is applied automatically based on the segments configured on each peer appliance. **Manual** — Enables pruning for a user-specified list of segments. When **Manual** is selected, the **Segment allow list** field appears. Enter up to 64 segments whose route updates should be exchanged across SD-WAN tunnels. The Default segment (VRF 0) is always included and cannot be removed.
Segment allow list	Available only when **Segment pruning** is set to **Manual**. Specifies the segments whose route updates are exchanged across SD-WAN tunnels. Accepts up to 64 segments. The Default segment (VRF 0) is always included and cannot be removed.

Add Routes

Use the Add Routes dialog box to add a user-defined route to an appliance's route table.

1. In the Routes dialog box, click **Add Routes**.

   The Add Route dialog box opens.

2. Configure the following elements as needed.

Field	Description
Subnet/Mask	Subnet IP address and mask (for example, 4.4.4.4/32).
VXLAN	If you are configuring a VXLAN network, select this check box and complete any additional fields. See VXLAN Tab.
Next Hop	The Next Hop IP address for the route. If you specify a Next Hop, you cannot select a zone for the route. Leaving the Next Hop option blank signals to the system that this is an advertise-only route. These routes are not used for forwarding traffic within the EdgeConnect appliance itself but are disseminated to neighboring devices for routing purposes. Advertise-only routes allow for the injection of a default route in use cases such as a default-originate route for BGP and OSPF. They also allow the summarization of routes advertised to SD-WAN fabric peers or OSPF and BGP neighbors. Advertise-only routes prevent overloading the EdgeConnect appliance with traffic it is not meant to handle directly. (Optional)
Interface	Interface for outgoing traffic. Click in the field and select the appropriate interface. If you specify an interface, you cannot select a zone for the route. (Optional)
Zone	Firewall zone to apply to the route. Select the appropriate firewall zone from the drop-down list. Initially, this field is set to Default. If you specify a next hop or an interface, you cannot select a zone for the route; the field automatically sets to None and cannot be changed. (Optional)
Metric	Metric for the subnet. Specify a value from 0 to 100. When a peer has more than one tunnel with a matching subnet (for example, in a high-availability deployment), it chooses the tunnel with the lower metric value. The default value is 50.
Tag	Tag for restricting route lookups. It is primarily used to filter routes from being redistributed in a routing loop. Select one of the following options from the drop-down list: **ANY** – Allows route lookups for traffic arriving on a LAN-side or WAN-side interface. **FROM_LAN** – Restricts route lookups to traffic arriving on a LAN-side interface. **FROM_WAN** – Restricts route lookups to traffic arriving on a WAN-side interface.

Field	Description
Label	Select a route label to associate with the route. A route label can be associated with one or more IP SLA rules, and can cause the route to go DOWN depending on the status of the IP SLA rules. Depending on the IP SLA Condition set for the route label, the route is marked DOWN if *any* of the IP SLA rules are DOWN or only when *all* IP SLA rules associated with the label are DOWN. See Route Labels.
Comments	Additional information you want to provide about this route. (Optional)

3. Click **Add**.

Segment / VRF Pruning

When routing segmentation is enabled, each segment (VRF) maintains a separate routing table and independently exchanges route updates with peer appliances across SD-WAN tunnels. Without pruning, an appliance would send route updates for all locally configured segments to every peer appliance—regardless of whether the peer has those segments configured. Exchanging route updates for segments not present on a peer appliance adds unnecessary control-plane overhead to the SD-WAN fabric.

Segment / VRF pruning limits route update exchange to segments that are configured on both ends of an SD-WAN tunnel. When a segment is not configured on a peer appliance, route updates for that segment are not sent to that peer across the tunnel. Segments that are configured on both endpoints continue to exchange route updates normally.

Segment / VRF pruning operates dynamically. As segments are added to or removed from peer appliances, the set of segments exchanging route updates across each SD-WAN tunnel is updated accordingly, without requiring manual reconfiguration.

The **Segment pruning** field in the Edit or Add Routes dialog box provides the following settings:

- **OFF**—Disables pruning. Route updates for all segments are sent to all peer appliances regardless of which segments are configured on each peer.
- **auto**—The default setting. Pruning is applied automatically based on the segments configured on each peer appliance.
- **Manual**—Enables pruning for a user-specified list of segments. The **Segment allow list** field appears and accepts up to 64 segments. The Default segment (VRF 0) is always included and cannot be removed from the list.

NOTE: Segment / VRF pruning requires EdgeConnect OS 9.7.0.x or later. Orchestrator automatically excludes the pruning configuration when applying templates to appliances running earlier versions.

NOTE: Segment / VRF pruning applies only when routing segmentation is enabled. For information about enabling segmentation and configuring segments, see Routing Segmentation.

Import Subnets

Before you import route details from a CSV file into a selected appliance, ensure that each row in the CSV file contains values for the following fields in the exact order specified with commas to separate the values:

- Subnet
- Mask Length
- Metric
- Is Local *(no longer used; leave this value blank)*
- Advertise to Silver Peak Peers *(no longer used; leave this value blank)*
- Advertise to BGP Peers *(no longer used; leave this value blank)*
- Next Hop
- Advertise to OSPF Neighbors *(no longer used; leave this value blank)*
- Interface Name
- Segment
- Zone

The following lines illustrate what two rows in a CSV import file might look like:

```
10.1.0.0,16,50,,,,10.1.0.1,,lan0,Default
10.2.0.0,16,50,,,,,,,
```

NOTE: Do not include a header row in the CSV file. Also, do not add spaces after commas in rows.

Complete the following steps to import route details from a CSV file into a selected appliance.

1. Click **Choose File**.
2. Locate and select the CSV file on your local machine, and then click **Open**.
3. Click **Import**.

   Orchestrator imports the information from the selected file and the Routes table displays new or updated route details.

Export Subnets

To save the contents of the Routes table to a CSV file, click **Export**.

SD-WAN Fabric Route Redistribution Maps

Route Maps are policies applied to IP routes during redistribution between routing protocols. They have **Match Criteria** and **Set Actions** that allow for filtering routes or modifying metrics and attributes for routes that meet the criteria defined in the match statement. Route-map rules follow a top-down order based on the sequence number defined for each entry.

EdgeConnect Enterprise supports applying Route Maps inbound from and outbound to BGP peers and outbound to OSPF neighbors and the SD-WAN Fabric. It is best practice to use Orchestrator to apply Route Maps using templates.

You can specify up to 20 SD-WAN route maps and 128 rules per route map.

You can specify up to 6 pipe separated prefixes for each rule applied to a route map as shown in this example: *172.168.24.0/21 ge 21 le 21|172.168.40.0/21 ge 21 le 21.*

You can add, delete, rename, or clone route maps using this window. You can add rules to your route map by clicking **Add Rule**. A route map without any enabled rules is treated as a default deny all.

Prefix Match Criteria

The default for prefix match criteria is exact-match + greater-than. Both the specified prefix and any subnets of that prefix will be matched, up to a length of 32 for IPv4 or 128 for IPv6 (subnet sharing route maps only).

Less-than-or-equal-to (LE) and greater-than-or-equal-to (GE) clauses can also be applied to specify the inclusion of certain subnets.

The previous workaround approach of denying 0.0.0.0/1 and 128.0.0.0/1 and then permitting any was used in systems that did not support direct matching of 0.0.0.0/0. In the current HPE Aruba Networking SD-WAN Orchestrator, you can simply permit 0.0.0.0/0 to match the default route directly.

GE Clause

If a GE clause is applied, the rule will also include all prefixes that have a prefix length greater than or equal to the GE value and less than or equal to 32 or 128 (for IPv6).

Example: A.B.C.D/X GE Y

In this example the following will be included:

- The exact match to A.B.C.D/X
- All the prefixes that belong to the subnet A.B.C.D/X that have a length greater than or equal to Y and less than or equal to 32

For example, you have a route map entry of 192.168.0.0/16 GE 24, and a peer advertises the following prefixes:

- 192.168.0.0/16
- 192.168.1.0/24
- 192.168.2.0/25
- 192.168.2.128/25

192.168.0.0/16 – The prefix length is 16, which is not greater than or equal to 24. This route does not match.

192.168.1.0/24 – The prefix length is 24, which is equal to the specified value. Therefore, this route matches the condition.

192.168.2.0/25 – The prefix length is 25, which is greater than or equal to 24. Therefore, this route matches the condition.

192.168.2.128/25 – The prefix length is 25, which is greater than or equal to 24. Therefore, this route also matches the condition.

So, the routes that would be matched by the given route map entry are:

- 192.168.1.0/24
- 192.168.2.0/25
- 192.168.2.128/25

LE Clause

If an LE clause is applied, the rule will also include all prefixes that have a prefix length less than or equal to the LE value.

Example: A.B.C.D/X LE Y

In this example the following will be included:

- The exact match to A.B.C.D/X
- All the prefixes that belong to the subnet A.B.C.D/X that have a length greater than or equal to X and less than or equal to 32
- All the prefixes that belong to the subnet A.B.C.D/X that have a length less than or equal to Y

For example, you have a route map entry of 192.168.0.0/16 LE 24, and a peer advertises the following prefixes:

- 192.168.0.0/16
- 192.168.1.0/24
- 192.168.2.0/25
- 192.168.2.128/25

192.168.0.0/16 – The prefix length is 16, which is less than or equal to 24. This route matches.

192.168.1.0/24 – The prefix length is 24, which is equal to the specified value. This route matches.

192.168.2.0/25 – The prefix length is 25, which is not less than or equal to 24. This route does not match.

192.168.2.128/25 – The prefix length is 25, which is not less than or equal to 24. This route does not match.

So, the routes that would be matched by the given route map entry are:

- 192.168.0.0/16
- 192.168.1.0/24

Combining LE and GE Clauses

Example: A.B.C.D/X LE Y GE Z

In this example the following will be included:

- The exact match to A.B.C.D/X
- All the prefixes that belong to the subnet A.B.C.D/X that have a length less than or equal to Y
- All the prefixes that belong to the subnet A.B.C.D/X that have a length greater than or equal to Z and less than or equal to 32

Exact Match

If both GE and LE clauses are specified and are equal, the rule will result in an exact match.

Example: A.B.C.D/X LE Y GE Y

In this example, the following will be included:

- The exact match to A.B.C.D/X
- The exact match to the subnet A.B.C.D/X that has a length equal to Y

You can specify the following fields in each rule for the selected route map.

Priority

Field	Description
Priority	If you are using Orchestrator templates to add rules, Orchestrator will delete all entries from **1000 – 9999** before applying its policies. You can create rules with higher priority than Orchestrator rules (**1 – 999**) and rules with lower priority (**10000 – 19999** and **25000 – 65534**). **NOTE:** The priority range from **20000** to **24999** is reserved for Orchestrator. When adding a rule, the priority is incremented by 10 from the previous rule. The priority can be changed, but this default behavior helps to ensure you can insert new rules without having to change subsequent priorities.

Select Match Criteria

Source Protocol	Complete the Following Fields (based on protocol selected)
Local/Static	Prefix + optional LE/GE parameters
BGP	Prefix + optional LE/GE parameters BGP Communities

Source Protocol	Complete the Following Fields (based on protocol selected)
OSPF	Prefix + optional LE/GE parameters OSPF Tag
ANY	Prefix + optional LE/GE parameters OSPF Tag BGP Communities
OAP-BGP	Prefix + optional LE/GE parameters BGP Communities
OAP-OSPF	Prefix + optional LE/GE parameters OSPF Tag
OAP-CFGSET	Prefix + optional LE/GE parameters
OAP-RIP	Prefix + optional LE/GE parameters
OAP-OVERLAY	Prefix + optional LE/GE parameters
OAP-IAPVPN	Prefix + optional LE/GE parameters
OAP-STATIC	Prefix + optional LE/GE parameters
OAP-DIRECT	Prefix + optional LE/GE parameters

NOTE: The above fields in the right column will change depending on the source protocol chosen.

Set Actions

Field	Description
Permit	Enable or disable. This setting allows or denies the route map.
OSPF Tag	Value of OSPF tag to set in routing information sent to destination. **NOTE:** This field is displayed only if Source Protocol is set to OSPF or OAP OSPF.
Metric	Metric for the route.
Comment	Comment you want to include.

OSPF Tab

Configuration > Networking > Routing > OSPF

This tab manages **OSPF (Open Shortest Path First)** on LAN and WAN interfaces.

OSPF learns routes from routing peers, and then subnet shares them with EdgeConnect peers and/or BGP neighbors.

A route tag is applied to a route to better identify the source of the network it originated from. It is primarily used to filter routes from being redistributed in a routing loop.

A maximum of 64 OSPF neighbors and 64 BGP peers is supported per appliance, with 200 next-hops supported per interface.

- For BGP, only 64 peers can be added. For OSPF, more than 64 neighbors can be added, though an error will be logged.
- If more than 64 OSPF neighbors are added, the active OSPF neighbors are chosen in a deterministic manner. All OSPF neighbors that are added are queried in a sorted order using segment ID as the primary index and the neighbor IP address as the secondary index. For example, if there are 65 OSPF neighbors, the peer in the highest segment and with the highest IP address will be the one that is always dropped. It will not drop a random OSPF neighbor.
- Also, if there are 60 OSPF neighbors in the default segment, which always has ID:0, and 10 OSPF neighbors in segment 1, the 60 neighbors in the default segment will always be included, as well as the 4 neighbors in segment 1 with the lowest IP addresses.

Filter by Segment

To filter the rows displayed in the OSPF table by segment:

- Select **Default** from the **Segment** drop-down list to display for the system-supplied default segment, or
- Select one of the other listed segments, which reflect the custom segments defined using Routing Segmentation (Configuration > Networking > Routing > Routing Segmentation (VRF)).

Select **All** to display for all segments, which is the default setting

The table below describes the fields displayed for the OSPF configuration.

Field	Description
Appliance	Name of the appliance.
Segment	Name of the segment being used, if enabled.
Enable	[Route Metric] Cost associated with a route. The higher the value, the less preferred.
Router ID	This router identifier is the IPv4 address by which the remote peer can identify this appliance for purposes of OSPF.
Redistribute Routes to OSPF	Redistribution map being used to redistribute routes to OSPF.
Details	Additional details about the route.

Select the edit icon in the OSPF table to edit and enable OSPF.

OSPF Edit Row

Use this dialog box to manage **OSPF (Open Shortest Path First)** on LAN and WAN interfaces.

OSPF learns routes from routing peers, and then subnet shares them with EdgeConnect peers and/or BGP neighbors.

Field	Description
Enable OSPF	When enabled, the appliance has access to use the OSPF protocol.
Router ID	IPv4 address of the router that the remote peer uses to identify the appliance for purposes of OSPF.
Redistribute routes to OSPF	Redistributing routes into OSPF from other routing protocols or from **static** will cause these routes to become OSPF external routes. Select the edit icon to the left of this field and select the OSPF route redistribution maps you would like to select.
Opaque LSA support	Enable for acknowledgment of opaque LSAs through OSPF protocol.

To add an area to an OSPF route, click **Add** in the **Areas** section. To add an additional interface to an OSPF route, click **Add** in the **Interfaces** section.

NOTE: The BFD field in the Interfaces table on the OSPF dialog box is set to N/A if BFD is not supported on the appliance.

To configure or modify an OSPF route map, select the edit icon next to the Redistribute routes to OSPF field.

Add Area

Complete the following fields to add an area to OSPF.

Field	Description
Area ID	Number of the area in which to locate the interface. The Area ID is the same for all interfaces.
Area Type	**Standard** – The most common area type for OSPF. This area type allows and propagates all OSPF LSA types. **NSSA** – Not-So-Stubby Area type.

Field	Description
Redistribution	Check box only available when NSSA is selected for Area Type. It indicates if the router should originate NSSA LSAs (Link State Advertisements) into the area for external routes or not. If the check box is selected the router will suppress NSSA LSAs. This check box is selected by default.
Summary	Check box only available when NSSA is selected for Area Type. It controls the import of summary LSAs into NSSA areas. If this is not selected, the router does not originate or propagate summary LSAs into the NSSA areas. This check box is selected by default.
Translator Role	Only available when NSSA is selected for Area Type. It indicates the ability of the Border router to perform NSSA translation of Type-7 LSAs into Type-5 LSAs.There are two options: **Candidate** – Allows the translator to participate in election. **Always Translate** – Always translates Type-7 LSAs into Type-5 LSAs.

Add Interface

Complete the following fields to add an interface to OSPF.

Field	Description
Interface	Indicates whether a Backup Designated Router (BDR) is specified for the Designated Router (DR). Options are **Yes** or **No**.
Area ID	Number of the area in which to locate the interface. The Area ID is the same for all interfaces. It can be an integer between 0 and 4294967295, or it can take a form similar to an IP address, A.B.C.D.
Cost	The cost of an interface in OSPF is an indication of the overhead required to send packets across a certain interface. It is used in the OSPF path calculation to determine link preference.
Priority	Router priority. (If two or more best routes are subnet shared, peer priority is used as the tiebreaker.)
Admin Status	Indicates whether the interface is set to admin **UP** or **DOWN**.
Hello Interval	Specifies the length of time, in seconds, between the hello packets that a router sends on an OSPF interface.
Dead Interval	Number of seconds that a router's Hello packets have not been seen before its neighbors declare the OSPF router down.

Field	Description
Transmit Delay	Number of seconds required to transmit a link state update packet. Valid values are 1 to 65535.
Retransmit Interval	Amount of time (in seconds) the router will wait to send retransmissions if the router receives no acknowledgment.
BFD	Select this check box to enable BFD for the OSPF interface. This field is not displayed if BFD is not supported on the appliance. **NOTE:** Before you select this check box, enable and configure BFD from the BFD tab.
Authentication	**None** – No authentication. **Text** – Simple password authentication allows a password (key) to be configured per area. **MD5** – Message Digest authentication is a cryptographic authentication. A key (password) and key-id are configured on each router. The router uses an algorithm based on the OSPF packet, the key, and the key-id to generate a “message digest” that gets appended to the packet.
Comment	Any information you want to include for your own use.

OSPF Route Redistribution Maps

Route Maps are policies applied to IP routes during redistribution between routing protocols. They have **Match Criteria** and **Set Actions** that allow for filtering routes or modifying metrics and attributes for routes that meet the criteria defined in the match statement. Route-map rules follow a top-down order based on the sequence number defined for each entry.

EdgeConnect Enterprise supports applying Route Maps inbound from and outbound to BGP peers and outbound to OSPF neighbors and the SD-WAN Fabric. It is best practice to use Orchestrator to apply Route Maps using templates.

You can specify up to 20 OSPF route maps and 128 rules per route map.

You can specify up to 6 prefixes for each rule applied to a route map.

You can add, delete, rename, or clone route maps using this window. You can add rules to your route map by clicking **Add Rule**. A route map without any enabled rules is treated as a default deny all.

Prefix Match Criteria

The default for prefix match criteria is exact-match + greater-than. Both the specified prefix and any subnets of that prefix will be matched, up to a length of 32 for IPv4 or 128 for IPv6 (subnet sharing route maps only).

Less-than-or-equal-to (LE) and greater-than-or-equal-to (GE) clauses can also be applied to specify the inclusion of certain subnets.

To match a default-route, deny 0.0.0.0/1, deny 128.0.0.0/1, and then permit any.

GE Clause

If a GE clause is applied, the rule will also include all prefixes that have a prefix length greater than or equal to the GE value and less than or equal to 32 or 128 (for IPv6).

Example: A.B.C.D/X GE Y

In this example the following will be included:

- The exact match to A.B.C.D/X
- All the prefixes that belong to the subnet A.B.C.D/X that have a length greater than or equal to Y and less than or equal to 32

For example, you have a route map entry of 192.168.0.0/16 GE 24, and a peer advertises the following prefixes:

- 192.168.0.0/16
- 192.168.1.0/24
- 192.168.2.0/25
- 192.168.2.128/25

192.168.0.0/16 – The prefix length is 16, which is not greater than or equal to 24. This route does not match.

192.168.1.0/24 – The prefix length is 24, which is equal to the specified value. Therefore, this route matches the condition.

192.168.2.0/25 – The prefix length is 25, which is greater than or equal to 24. Therefore, this route matches the condition.

192.168.2.128/25 – The prefix length is 25, which is greater than or equal to 24. Therefore, this route also matches the condition.

So, the routes that would be matched by the given route map entry are:

- 192.168.1.0/24
- 192.168.2.0/25
- 192.168.2.128/25

LE Clause

If an LE clause is applied, the rule will also include all prefixes that have a prefix length less than or equal to the LE value.

Example: A.B.C.D/X LE Y

In this example the following will be included:

- The exact match to A.B.C.D/X
- All the prefixes that belong to the subnet A.B.C.D/X that have a length greater than or equal to X and less than or equal to 32
- All the prefixes that belong to the subnet A.B.C.D/X that have a length less than or equal to Y

For example, you have a route map entry of 192.168.0.0/16 LE 24, and a peer advertises the following prefixes:

- 192.168.0.0/16
- 192.168.1.0/24
- 192.168.2.0/25
- 192.168.2.128/25

192.168.0.0/16 – The prefix length is 16, which is less than or equal to 24. This route matches.

192.168.1.0/24 – The prefix length is 24, which is equal to the specified value. This route matches.

192.168.2.0/25 – The prefix length is 25, which is not less than or equal to 24. This route does not match.

192.168.2.128/25 – The prefix length is 25, which is not less than or equal to 24. This route does not match.

So, the routes that would be matched by the given route map entry are:

- 192.168.0.0/16
- 192.168.1.0/24

Combining LE and GE Clauses

Example: A.B.C.D/X LE Y GE Z

In this example the following will be included:

- The exact match to A.B.C.D/X
- All the prefixes that belong to the subnet A.B.C.D/X that have a length less than or equal to Y
- All the prefixes that belong to the subnet A.B.C.D/X that have a length greater than or equal to Z and less than or equal to 32

Exact Match

If both GE and LE clauses are specified and are equal, the rule will result in an exact match.

Example: A.B.C.D/X LE Y GE Y

In this example, the following will be included:

- The exact match to A.B.C.D/X

- The exact match to the subnet A.B.C.D/X that has a length equal to Y

You can specify the following fields in each rule for the selected route map.

Priority

Field	Description
Priority	If you are using Orchestrator templates to add rules, Orchestrator will delete all entries from **1000 – 9999** before applying its policies. You can create rules with higher priority than Orchestrator rules (**1 – 999**) and rules with lower priority (**10000 – 19999** and **25000 – 65534**). **NOTE:** The priority range from **20000** to **24999** is reserved for Orchestrator. When adding a rule, the priority is incremented by 10 from the previous rule. The priority can be changed, but this default behavior helps to ensure you can insert new rules without having to change subsequent priorities.

Select Match Criteria

Source Protocol	Complete the Following Fields (based on protocol selected)
Local/Static	Prefix + optional LE/GE parameters
BGP	Prefix + optional LE/GE parameters BGP Communities
SD-WAN Routes	Prefix + optional LE/GE parameters BGP Communities OSPF Tag
OAP-BGP	Prefix + optional LE/GE parameters BGP Communities
OAP-OSPF	Prefix + optional LE/GE parameters OSPF Tag
OAP-CFGSET	Prefix + optional LE/GE parameters
OAP-RIP	Prefix + optional LE/GE parameters
OAP-OVERLAY	Prefix + optional LE/GE parameters
OAP-IAPVPN	Prefix + optional LE/GE parameters
OAP-STATIC	Prefix + optional LE/GE parameters

Source Protocol	Complete the Following Fields (based on protocol selected)
OAP-DIRECT	Prefix + optional LE/GE parameters

NOTE: The above fields in the right column will change depending on the source protocol chosen.

Set Actions

Field	Description
Permit	Enable or disable. This setting allows or denies the route map.
OSPF Tag	Value of OSPF tag to set in routing information sent to destination.
OSPF Metric Type	Filters redistributed routes to OSPF.
Metric	Metric for the route.
Comment	Comment you want to include.

BFD Tab

Configuration > Networking > Routing > BFD

Bidirectional Forwarding Detection (BFD) is a networking protocol that detects faults between devices. The EdgeConnect appliance supports BFD for both BGP and OSPF.

- Single and multi-hop BFD configurations are supported.
- BFD asynchronous mode is supported.
- BFD can be configured for up to 20 segments with a maximum of 100 simultaneous BFD sessions across all segments.

Configuring BFD for BGP or OSPF is a two-step process:

1. Click the edit icon for an appliance listed in the BFD table on the BFD tab, and then enable and configure BFD on the BFD dialog box. For details, see BFD Dialog Box.
2. Enable BFD for each BGP peer or OSPF interface.
    1. For BGP, navigate to **Configuration > Networking > Routing > BGP**. Click the edit icon for an appliance listed in the BGP table, and then click **Add** to add a BGP peer or click the edit icon for an existing BGP peer listed in the BGP Peers table. Select the **BFD** check box, make other changes as appropriate, and then click **Add** or **Save**.
    2. For OSPF, navigate to **Configuration > Networking > Routing > OSPF**. Click the edit icon for an appliance listed in the OSPF table, and then click **Add** in the Interfaces area to add an interface or click the edit icon for an existing interface listed in the Interfaces table. Select the **BFD** check box, make other changes as appropriate, and then click **Add** or **Save**.

The BFD tab provides two views of BFD information:

- Click the **Summary** button on the BFD tab to display configuration details associated with the local appliance. For field descriptions, see BFD Dialog Box.
- Click the **Sessions** button to display currently active BFD sessions. BFD establishes a session between two endpoints over a particular link. If more than one link exists between two systems, multiple BFD sessions can be established to monitor each of them.

Filter by Segment

To filter the rows displayed in the BFD table by segment:

- Select **Default** from the **Segment** drop-down list to display for the system-supplied default segment, or
- Select one of the other listed segments, which reflect the custom segments defined using Routing Segmentation (Configuration > Networking > Routing > Routing Segmentation (VRF)).

Select **All** to display for all segments, which is the default setting.

The following table describes the fields displayed in the Sessions view of the BFD tab.

Field	Description
Appliance	Name of the appliance.
Segment	Name of the segment. This field displays only if Routing Segmentation is enabled.
Local address	IP address of the local endpoint.
Remote address	IP address of the remote endpoint.
Local Interface	Name of the local interface (lan0, wan0, wan1).
State	Session state (UP or DOWN).
Uptime	Session up time.
Details	Additional details about the BFD session

BFD Dialog Box

Use this dialog box to enable and configure BFD for your appliances, as follows:

1. Move the toggle to **Enable BFD**.
2. Complete the following fields.

Field	Description
Min Tx Interval	Minimum transmit interval in milliseconds (ms). Specify a value from 300 to 5000. The default setting is 300.
Min Rx Interval	Minimum receive interval in milliseconds (ms). Specify a value from 300 to 5000. The default setting is 300.
Detection Multiplier	Detection time multiplier. In BFD, the detection time is the transmit interval multiplied by the detection multiplier. If BFD data is not received within the detection time, a failure occurs. Specify a value from 3 to 10. The default setting is 3.

Multicast

Configuration > Networking > Routing > Multicast

Orchestrator supports multicast routing, a method of sending data from a single IP address to a larger group of recipients.

- Up to 60 multicast routes are supported, including up to 30 (S,G) and 30 (*,G) groups.
- Up to 1,000 sites can participate in the same multicast stream. HPE Networking recommends up to 200 sites.
- All versions of IGMP are supported (IGMPv1, IGMPv2, and IGMPv3).
- Multicast routing is supported in Inline Router mode only.

Multicast and Segments

In Orchestrator 9.7.1 and later, you can enable multicast on up to eight segments (VRFs) independently. This allows multicast traffic to operate within multiple isolated routing domains while maintaining end-to-end segmentation. Multicast routing between segments is not supported.

IMPORTANT: Multicast traffic is not routed between segments. Sources and receivers must reside within the same multicast-enabled segment (VRF). To support multicast in multiple segments, configure multicast independently for each segment.

- Each multicast-enabled segment requires its own multicast configuration.
- Each segment can use a different Rendezvous Point (RP) and Allowed Group configuration.

Before multicast can be enabled, a segment must be created and interfaces must be assigned to that segment.

TIP: Beginning with Orchestrator 9.7.0, the Multicast template enables you to configure multicast settings across multiple appliances. HPE Networking recommends using the template to ensure configuration consistency and simplify multicast deployments. See Multicast Template.

Multicast Configuration Workflow Overview

Deploying multicast in a multi-segment environment requires more than simply enabling multicast on an interface. Before multicast traffic can flow, you must create the required segments (VRFs), associate interfaces with those segments, configure a Rendezvous Point (RP), and enable multicast on the interfaces that will participate in multicast routing.

The following workflow provides a high-level overview of the recommended deployment process and helps illustrate the relationship between segments, interfaces, multicast configuration, and multicast routing.

Configuring multicast in a multi-segment environment typically involves the following tasks:

1. Configure segments.
    - Define the required segments (VRFs).
    - Assign LAN interfaces to the appropriate segments.
    - Optionally, create multicast address groups that can be used to restrict multicast traffic.
2. Enable multicast on each segment.
    - Enable multicast for the segment.
    - Configure the Rendezvous Point (RP) IP address.
    - Optionally, configure allowed multicast address groups.
3. Configure a Rendezvous Point (RP) for the segment.
    - Optionally create a loopback interface to serve as the RP.
    - Assign the RP IP address.
    - For high-availability deployments, configure a secondary RP and the appropriate failover mechanism.
4. Associate interfaces with the multicast-enabled segment.
    - Verify that the interface is bound to the correct segment (VRF).
    - Add one or more LAN interfaces to the multicast-enabled segment.
    - Enable Internet Group Management Protocol (IGMP) and Protocol Independent Multicast (PIM) as required.
    - Configure any required PIM parameters.
5. Verify multicast operation.
    - Verify IGMP group membership.
    - Verify PIM neighbors.
    - Verify multicast routes.
    - Verify multicast traffic flows.

Multicast Tab Views

Orchestrator provides four views of multicast status, each accessible by one of the corresponding buttons at the top of the Multicast tab: **Summary**, **Interfaces**, **Neighbors**, and **Routes**.

Descriptions of fields on the Summary view follow:

Field	Description
Appliance	Name of the appliance associated with the multicast configuration.
Enable	Indicates whether multicast is enabled.
Rendezvous Point IP	IP address of the Rendezvous Point (RP) used by PIM Sparse Mode. The RP acts as a meeting point for multicast sources and receivers during multicast tree enablement.
Allowed Group	Only addresses included in the specified multicast address group are allowed. If no address group is specified, any IP address can multicast. The message *Feature is not supported for the appliance* displays in this field if the appliance does not support the Allowed Group feature.

Click the edit icon to enable or disable multicast, add an interface for multicast, or edit an existing interface.

Segment Filter

Use the **Segment** drop-down list at the top of the tab to filter the display by a specific segment. Select **All Segments** to view multicast data across all segments or select a specific segment to view only its multicast configuration.

For Orchestrator 9.7.1 and later, multicast information can be filtered by segment. Depending on the selected view, the information displayed reflects either all configured segments or the selected segment.

NOTE: For appliance versions 9.6.0.0 and earlier, only the Default segment is displayed.

Multicast Dialog Box

From the Summary, Interfaces, Neighbors, or Routes view on the Multicast tab:

1. Click the edit icon next to the appliance for which you want to set up multicast.

   The Multicast dialog box opens.

2. Move the **Enable multicast** toggle to the right to enable multicast for this appliance with its SD-WAN fabric peers.

3. In the **Rendezvous point IP address** field, enter the appropriate IP address.

   NOTE: When deploying an EdgeConnect appliance as the Rendezvous Point (RP), you can use a loopback interface address as the RP address. Using a loopback interface provides a stable RP address that remains available regardless of the state of individual physical interfaces.

   IMPORTANT: Only static Rendezvous Point (RP) configurations are supported. Bootstrap Router (BSR) and Anycast RP configurations are not supported.

4. In the **Allowed group** field, select an available address group from the drop-down list or enter a new address group. All IP addresses included in the specified address group will be allowed to multicast. This field is not displayed if the appliance does not support the Allowed Group feature. Address group names can include alphanumeric characters, periods, underscores, or hyphens.

NOTE: Multicast cannot be enabled on more than eight segments. Attempting to do so will result in an error message.

IMPORTANT: The address group you specify in the Address Group field must be valid. If you enter a new address group, ensure that you also create and set it up on the Address Groups tab (Configuration > Templates & Policies > ACLs > Address Groups). If the new address group remains invalid, no IP addresses will be allowed to multicast.

Interfaces

Field	Description
Interface	Name of the interface you want to connect.
PIM enabled	Indicates whether Protocol Independent Multicast (PIM) is enabled on the interface. When enabled, the interface can participate in multicast routing and exchange multicast routing information with neighboring routers.
IGMP enabled	Indicates whether Internet Group Management Protocol (IGMP) is enabled on the interface. When enabled, the interface can process multicast group membership messages from connected hosts and support multicast traffic delivery.
DR priority	Specifies the priority used during Designated Router (DR) elections. Interfaces with higher priority values are preferred when electing the DR.
DR router IP	Displays the IP address of the Designated Router (DR) elected for the multicast-enabled segment.
Actions	Hover to view available actions for this row (edit, delete, and so on).

To add an interface:

1. Click **Add**.

   The Add Interface dialog box opens.
2. Select the desired interface from the **Interface** drop-down list.
3. Select the **Enable PIM** check box if you want to enable it.
4. Select the **Enable IGMP** check box if you want to enable it.
5. In the **DR priority** field, specify the priority to use during Designated Router (DR) elections on the interface.
6. In the **Hello interval** field, select the number of seconds (1-300) for this interval. PIM messages are sent to the interface each time this interval is reached. The default value is 30.

7. In the **Join / Prune interval** field, select the number of seconds (30-180) that you want to send PIM messages to inform upstream routers about multicast group membership to join or prune interfaces. The default value is 30.
8. Click **Done**.
9. Click **Save**.

Neighbors

Field	Description
Interface	Name of the interface associated with the multicast neighbor.
PIM enabled	Indicates whether the neighboring router is participating in Protocol Independent Multicast (PIM) on the selected interface.
IGMP enabled	Indicates whether Internet Group Management Protocol (IGMP) is enabled on the neighboring multicast interface.
DR priority	Determines the interface's priority in Designated Router (DR) elections. Higher values increase the likelihood that the router will be elected as the DR.
DR router IP	IP address of the neighbor.

Routes

Field	Description
Source	Transmitter of the multicast data.
Segment	SD-WAN segment (VRF) for which multicast is enabled and managed. Multicast traffic remains within the selected segment.
Group	IP address of the multicast group.
Incoming interface/Peer	Interface that receives inbound traffic.
Outgoing interfaces/Peers	Interfaces or peers to which multicast traffic is forwarded.

On the Multicast tab, you can click the **Export CSV** icon to export a spreadsheet of the multicast report. You can also click the refresh button to update information displayed on the tab.

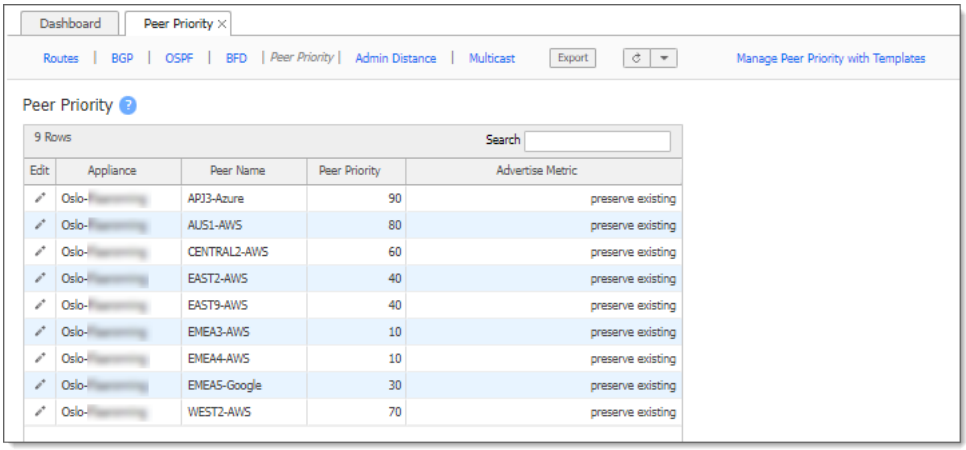
Peer Priority Tab

Configuration > *Networking* > *Routing* > *Peer Priority*

When an appliance receives a **Subnet** with the same **Metric** from multiple remote/peer appliances, it uses the Peer Priority list as a tie-breaker.

- If a **Peer Priority** is not configured, the appliance randomly distributes flows among multiple peers.
- The lower the number, the higher the peer's priority.

Click the edit icon to configure a peer and its peer priority.



NOTE: By default, the peer priority range starts at 1.

Peer Priority Edit Row

This dialog box displays a list of configured peers. The peer priority and advertise metric are displayed for each peer.

- Peer priority controls the peer to which traffic is sent when route ties occur. It acts similar to BGP's local preference.
- Advertise metric controls the return path of a flow back toward the local appliance. It adjusts the metric of all routes sent to Peer Name. Advertise metric announces different metrics to different fabric peers. It acts similar to BGP's Multi Exit Discriminator (MED). The default setting is *preserve existing* (do nothing).

Both peer priority and advertise metric impact all routes sent and received from Peer Name.

To add a peer:

1. Click **Add Peer**.
2. In the new row that is added to the table, enter the peer name, peer priority, and advertise metric.
3. To delete a peer, click the **X** in the far-right column of the peer's row.
4. When finished, click **Apply**.

Admin Distance Tab

Configuration > Networking > Routing > Admin Distance

This tab shows values associated with various types of **Admin Distance**. Admin Distance (AD) is the route preference value assigned to dynamic routes, static routes, and directly connected routes. When the appliance's Routes table has multiple routes to the same destination, the appliance uses the route with the lowest administrative distance.

The following table displays the values associated with various types of **Admin Distance**.

Field	Description
Appliance	Name of the appliance.
Local	Manually configured route, or one learned from locally connected subnets.
EBGP	External BGP: exchanging routing information with a router outside the company-wide network.
IBGP	Internal BGP: exchanging routing information with a router inside the company-wide network.
Subnet Shared - Static Routes	Route learned from an EdgeConnect peer.
Subnet Shared - BGP Remote	Route shared from an EdgeConnect peer in an external network.
OSPF	Route learned from an OSPF (Open Shortest Path First) neighbor.
Subnet Shared - OSPF Remote	Route learned from an EdgeConnect peer.
Subnet Shared CFGSET	ORO (Overlay Route Orchestrator) BGW (branch gateway) route learned from the SD-WAN fabric.
Subnet Shared RIP	ORO learned Routing Information Protocol route learned from the SD-WAN fabric.
Subnet Shared Overlay	ORO modified/added route learned from the SD-WAN fabric.
Subnet Shared IAPVPN	ORO Instant Access Point route learned from the SD-WAN fabric.
OAP BGP	Route learned from an OAP (Overlay Route Orchestrator) BGP peer in an external network.
OAP CFGSET	BGW (branch gateway) route learned from ORO.
OAP Direct	Direct (connected) route learned from ORO.
OAP IAPVPN	Instant Access Point route learned from ORO.
OAP OSPF	Route learned from an OAP OSPF neighbor.
OAP Overlay	ORO modified/added route.

Field	Description
OAP RIP	Routing Information Protocol route learned from ORO.
OAP Static	Static route learned from ORO.

To edit these fields, click the edit icon.

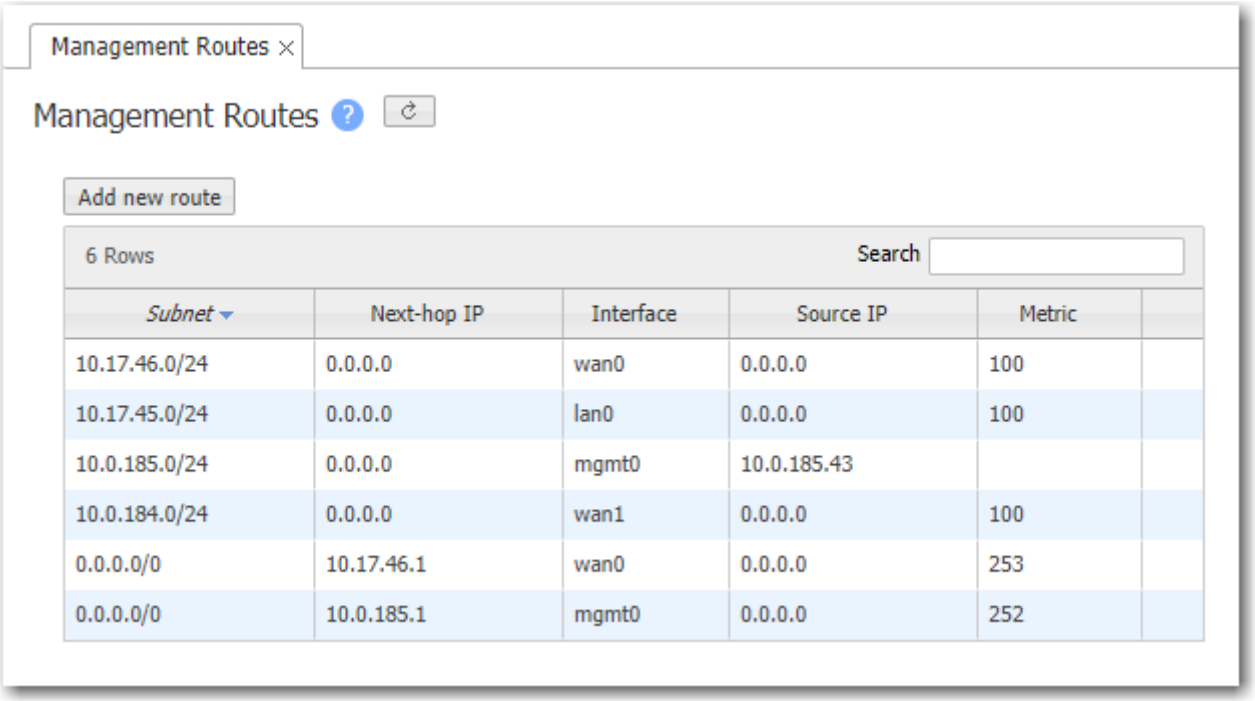
Admin Distance Edit Row

Use this dialog box to edit the admin distances for each type in the table. Click any cell in the Distance column to begin modifying the values. When finished, click **Save**.

Management Routes Tab

Configuration > *Networking* > *Routing* > *Management Routes*

Use this tab to configure **next-hops** for management interfaces.



- Management routes specify the **default gateways** and local IP subnets for the management interfaces.
- In a Dual-Homed Router Mode configuration, you might need to add a static management route for flow redirection between appliances paired for redundancy at the same site.

- The management routes table shows the configured static routes and any dynamically created routes. If you use **DHCP**, the appliance automatically creates appropriate dynamic routes. A user cannot delete or add dynamic routes.
- If the **Source IP** is listed as **0.0.0.0**, packets sent using this route use the **Interface**'s IP address as the Source IP address. If the **Source IP** lists a specific IP address, that IP address is used instead.

VXLAN Tab

Configuration > Networking > Routing > VXLAN

Use the VXLAN tab to specify Virtual Extensible Local Area Network (VXLAN) and Virtual Network Identifier (VNI) settings for routing segments already configured on HPE Aruba Networking CX switches or EdgeConnect appliances. VXLAN allows you to create multiple Layer 2 segments over a Layer 3 network. Each segment is identified by a 24-bit VNI that can support up to 16 million virtual networks.

VXLAN encapsulates Layer 2 Ethernet frames in Layer 3 UDP packets, enabling you to create virtualized Layer 2 subnets, or segments, that span physical Layer 3 networks. The entity that performs the encapsulation and decapsulation of packets is called a Virtual Tunnel Endpoint (VTEP). An EdgeConnect is a VTEP for WAN-to-LAN traffic. An HPE Aruba Networking CX switch is a VTEP for LAN-to-LAN traffic.

A VNI specifies a routing segment, a firewall zone, and a fallback role for a VXLAN instance. A VNI identifies different virtual networks in the data plane. A VNI is a 24-bit value in the VXLAN header and can support up to 16 million individual network segments. A VNI is like a VLAN ID but has a larger address space. A VNI maps the virtual network to a specific VXLAN segment. The VNI identifies the destination of the traffic in the VXLAN network. VNI is the basis for isolating different virtual networks from each other.

Once a VNI is configured for a segment or in a template, HPE Aruba Networking CX switches or EdgeConnect appliances automatically create a network virtual interface (NVE) as a VXLAN tunnel endpoint (VTEP). A VTEP encapsulates and decapsulates VXLAN packets. The only accepted peer is the NVE that is configured in BGP. Packets received with a VNI not mapped to a segment will be dropped.

EdgeConnect automatically binds the NVE to the VXLAN segment and specifies the source interface for the VXLAN tunnel - only loopback interfaces from the default segment are valid. If BGP EVPN Peer is enabled, the loopback interface you choose is automatically configured in the local interface field of the BGP EVPN Peer configuration. For more information on BGP EVPN Peer configuration, see the BGP tab.

The VXLAN packet tells BGP the target VTEP. BGP discovers the remote VXLAN tunnel endpoint address, advertises routes that are reachable over this tunnel as the forwarding next hop, and dynamically brings down this tunnel when reachability over the tunnel is no longer needed.

Prerequisites

Before you can assign a VNI to a VXLAN segment, you must configure the following settings:

- Segmentation must be enabled to support VXLAN. See the Routing Segmentation (VRF) tab.
- The IP routing on the BGP Layer 3 network that connects the EdgeConnect VTEPs must already be configured. This is necessary to enable VXLAN traffic to traverse the network. Therefore, only in-line router mode is supported.
- Currently, the EdgeConnect EVPN address family is only supported for BGP EVPN peers in the Default segment (VRF ID = 0).
- One or more loopback interfaces must already be available.
- VXLAN is only supported on LAN interfaces. Route-Targets must be defined, and BGP enabled for all segments, even if no BGP peers are configured in non-default segments.

Common Settings for all VNIs

Use this section of the **VXLAN** Tab to configure these common settings for all VNIs:

- **Destination UDP Port**: You can configure a custom destination UDP port for VXLAN. If not selected, the appliance uses the default port of 4789.
- **VTEP Source Interface**: Select a loopback interface from the list.

  NOTE: Only loopback interfaces are valid. The loopback interface you choose will automatically be configured in the local interface field of the BGP Peer configuration if EVPN Peer is enabled.

VNI Mappings

For this dialog box, use the steps below to map a VNI to a routing segment, a firewall zone, and a fallback role.

NOTE: All configured VNIs configured on an EdgeConnect are communicated to a BGP peer in a single VXLAN tunnel. While the VXLAN tab lists each VNI separately, it only reports the status of that single VXLAN tunnel in the first VNI on the list. Conversely, the Routes tab displays the routes for each VNI segment that is being communicated in the VXLAN tunnel.

Add

1. Click **Add** to create a new VNI for a segment.
2. Enter a value for the VNI segment. Valid values are 1-16777215.
3. Select the **Segment**, **Firewall Zone**, and **Fallback Role** (Don't Apply, Guest IOT, Untrusted).
4. Click **OK**.

Edit

1. Select an existing VNI from the list.

2. Click the **Edit** icon to modify an existing VNI.

   NOTE: In the Flows tab, enable the VNI Tx and VNI Rx columns to display the number of the VNI that received or sent the VXLAN traffic. Both values should match for every flow. If not, there might be a misconfiguration downstream from the EdgeConnect.

Role to GPID Mapping

Use the Roles dialog box to map a policy enforcement role to a VXLAN Group Policy Identifier (GPID). Mapping policy enforcement roles to a VXLAN GPID is optional. Policy enforcement role mapping to a GPID propagates globally across the SD-WAN Fabric. Enabling the identity-based policy enforcement capability of the HPE Aruba Networking SD-WAN solution in VXLAN segments provides a highly automated extensible way of enabling a zero-trust security architecture.

VXLAN Statistics Reporting

VXLAN statistics reporting provides real-time insights into the performance and health of VXLAN networks. The **VXLAN** tab Origin column shows whether the VXLAN tunnel was provisioned statically or dynamically. The Status column reports if the VXLAN tunnel us UP or DOWN.

NOTE: For the status of statically provisioned tunnels, view the VTEP Details dialog box by clicking the Details field of a tunnel. You can identify a tunnel in the VXLAN tab page list according to its UP or DOWN status.

VTEP Details

The **VTEP Details** dialog box provides the following real-time statistics:

- **VTEP Peer Details** Transmit and receive packet counts, byte counts, and error counts
  - IP address
  - MAC address
  - Status (UP/DOWN)
  - Uptime
  - Origin (DYNAMIC/STATIS)
  - Associated routes (for all VNIs on the EdgeConnect)
  - RX packets
  - TX packets
  - RX bytes
  - TX Bytes
  - RX drops
  - TX drops
- **Local VTEP Details**
  - VTEP source interface

  - VTEP source IP
  - VTEM source MAC address

ARP Tab

Configuration > Networking > Routing > ARP

NOTE: In Orchestrator 9.6.0 and later, this tab is labeled ARP. In earlier versions of the React-based Orchestrator UI, this tab was labeled Neighbors.

Starting in Orchestrator 9.7.0, the ARP tab provides centralized visibility into ARP and neighbor tables for EdgeConnect appliances. This visibility improves network monitoring and simplifies troubleshooting of Layer 2 and Layer 3 connectivity. From the ARP tab, you can view ARP entries, add static entries, and delete existing entries. These operations previously required the local appliance CLI.

The ARP tab displays a consolidated view of all Layer 2 ARP entries known to selected EdgeConnect appliances, including entries for management interfaces (`mgmt0` and `mgmt1`). The table displays the same information as the EdgeConnect CLI command `show neighbor`. You can select one or more appliances in the appliance tree to view their ARP entries in a single table.

NOTE: The ARP tab includes a shortcut link to the DHCP Manager tab, but the two are separate features. ARP entries are managed independently from their own location under Configuration > Networking > Routing > ARP.

NOTE: The ARP tab is available only in the React-based Orchestrator UI. It is not available in the local Appliance Manager. There is no corresponding Monitoring > ARP view; all ARP management functions are accessed from Configuration > Networking > Routing > ARP.

NOTE: The ARP tab requires EdgeConnect OS 9.6.0.0 or later. When an appliance running an earlier version is selected, the Add button is disabled and displays the tooltip “Appliance version does not support.”

ARP, Routes, and VRFs

Address Resolution Protocol (ARP) resolves Layer 3 IP addresses to Layer 2 MAC addresses for hosts on directly connected subnets. EdgeConnect appliances maintain per-interface ARP tables to support routing decisions.

Key relationships between ARP, routes, and VRFs:

- Each VRF maintains its own routing and ARP tables.
- Routes resolve their next hops through ARP. If the ARP entry for a next hop IP address is missing or stale, traffic through that next hop may fail.
- ARP entries are automatically refreshed when traffic is exchanged with a neighbor. Administrators can manually clear entries to force re-resolution.

ARP Table Columns

The following columns appear in the ARP table.

Column	Description
Appliance	Name of the EdgeConnect appliance.
IP address	Layer 3 IP address of the neighbor. Both IPv4 and IPv6 addresses are shown.
MAC address	Layer 2 MAC address associated with the IP address.
Interface index	Numeric index of the local interface on which the ARP entry was learned.
Interface name	Name of the local interface (for example, *wan0*, *lan0*, *mgmt0*, or *mgmt1*).
Type	Type of ARP entry. Dynamic or Static.

The equivalent CLI command is *show neighbor*. Use *show neighbor ipv4* or *show neighbor ipv6* to filter by address family.

NOTE: EdgeConnect supports up to 2,000 static ARP entries per appliance. This limit applies only to static entries; it does not limit the number of dynamically learned entries shown in the table. Use the inline search to filter by IP address, MAC address, or interface name.

NOTE: When adding a static ARP entry, enter the MAC address in colon-separated format, for example *00:11:22:33:44:55*.

NOTE: The ARP table includes entries for management interfaces (*mgmt0* and *mgmt1*). A MAC address may appear multiple times in the table—for example, when *mgmt0* is connected to the same subnet as a data path interface. This can be useful for verifying Layer 2 connectivity between management and data path networks.

Adding a Static ARP Entry

Use the Add button to create a new static ARP entry for the selected appliance.

NOTE: The Add button is available only when exactly one appliance is selected. If multiple appliances are selected, the Add button is disabled and displays the tooltip “Select only one appliance.”

To add a static ARP entry:

1. Navigate to **Configuration > Networking > Routing > ARP**.
2. Select a single appliance.
3. Click **Add**.
4. In the ARP Configuration dialog box, complete the following fields.

Field	Description
Interface name	Local interface to associate with this entry. Select from the drop-down menu. Required.
IP address	IP address of the neighbor. Required.

Field	Description
MAC address	MAC address to associate with the IP address. Required. Enter the address in colon-separated format, for example `00:11:22:33:44:55`.

5. Select **Save**.

Deleting ARP Entries

You can delete an individual ARP entry or clear all dynamic ARP entries at once.

Complete the following steps to delete a specific ARP entry.

1. Navigate to Configuration > Networking > Routing > ARP.
2. Select the appliance whose ARP entry you want to delete.
3. In the ARP table, locate the row to remove and select the corresponding delete icon (**X**) in the last column.
4. In the Delete ARP confirmation dialog box, review the message "Are you sure you want to delete the selected ARP?" and select **Delete**.

To clear all dynamic ARP entries:

1. Navigate to Configuration > Networking > Routing > ARP.
2. Select one or more appliances.
3. Select **Clear all dynamic types**.
4. In the confirmation dialog box, review the warning "Deleting dynamic types will cause a network outage. Are you sure you want to delete them?"
5. Select **Delete**.

CAUTION: Clearing all dynamic types removes all dynamically learned ARP entries (both IPv4 and IPv6) for the selected appliances, causing a temporary network disruption while the appliances relearn the entries. Static ARP entries are not affected.

Tunnels Tab

Configuration > Networking > Tunnels > Tunnels

EdgeConnect tunnels are the foundation of your SD-WAN fabric. This tab displays details about tunnels in your network. It includes the following four subtabs:

- Overlay – Displays SD-WAN bonded tunnels. Specifically, overlay tunnels consist of bonded underlay tunnels.
- Underlay – Displays tunnels that map to discrete transports.
- Passthrough – Displays third-party (IPSec) tunnels for service chaining to cloud security services, such as Zscaler and Netskope, and tunnels for local breakouts to trusted SaaS applications, such as Office 365.

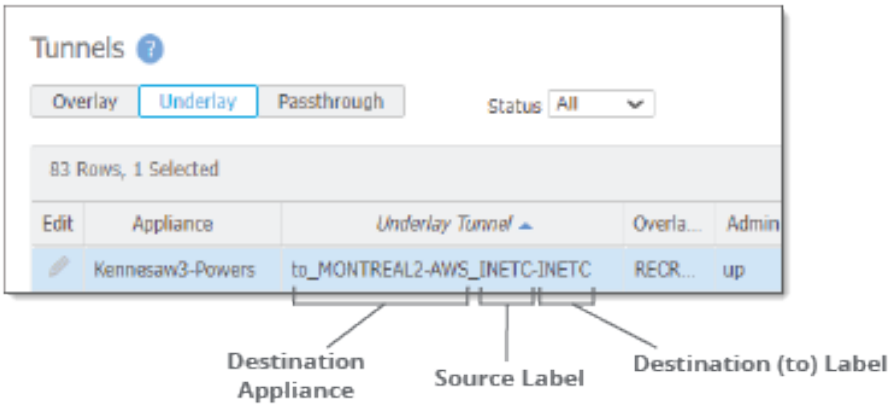
- HPE ANW Central – Displays tunnels orchestrated by HPE Aruba Networking Central.

In an SD-WAN network, Business Intent Overlays (BIOs) govern automatic tunnel creation and properties. This tab also provides the means to manually create IP Sec tunnels either between EdgeConnect appliances or from an EdgeConnect appliance to a third-party service.

NOTE: Manually created underlay tunnels cannot be used by BIOs.

Underlay Tunnel Naming

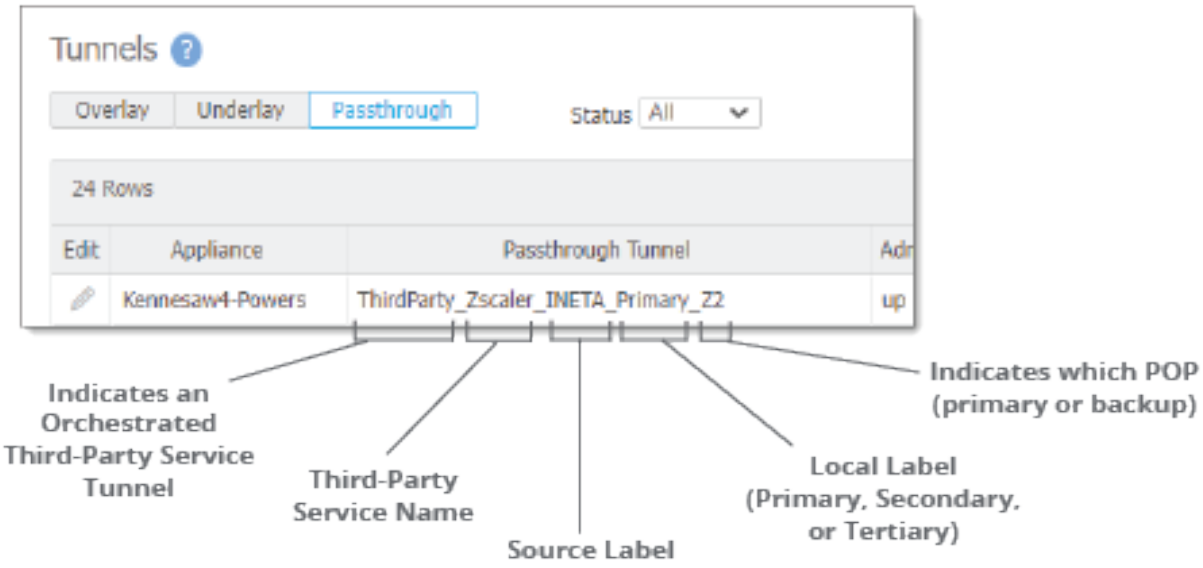
Underlay tunnels are uni-directional from the appliance listed in the Appliance column to the appliance listed after “to_” in the Underlay Tunnel column. The Underlay Tunnel column also includes the interface labels for the “from” and “to” sources.



Passthrough Tunnel Naming

Passthrough tunnels show “Passthrough” followed by the interface label for the source and the overlay (if applicable).

For orchestrated third-party tunnels, the Passthrough Tunnels column always shows “Third-Party” followed by the third-party service name, the interface label for the source, and the local label, which can be “Primary”, “Secondary”, or “Tertiary”. If there are POPs it indicates which POP after the local label. In the following example, it indicates which Zscaler POP where Z1 is the primary Zscaler POP and Z2 is the backup Zscaler POP.



Filter by Tunnel Status

To filter the rows displayed in the Tunnels table by tunnel status, select **Up** or **Down** from the Status drop-down list. Select **All** to display for all statuses, which is the default setting.

Subtab Field Descriptions

The following tables describe the fields displayed on the Overlay, Underlay, Passthrough, and HPE ANW Central subtabs. Field descriptions are not repeated if they appear on more than one subtab and have the same description.

Overlay Subtab

Field	Description
Appliance	Name of the appliance.
Overlay Tunnel	Designated overlay tunnel.
Overlay	Designated overlay to which the overlay tunnel is applied.
Admin Status	Indicates whether the tunnel has been set to admin **up** or **down**.

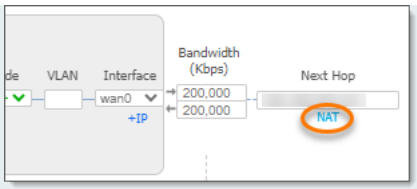
Field	Description
Status	Tunnel statuses are as follows: **down** – Tunnel is down. This can occur because the Admin Status of the tunnel is set to “down” (indicated when down status background is orange) or the tunnel cannot communicate with the appliance at the other end (indicated when down status background is red). Possible causes are: Lack of end-to-end connectivity / routability (test with *iperf*). Intermediate firewall is dropping the packets (open the firewall). Intermediate QoS policy (**be** packets are being starved; change control packet DSCP marking). Mismatched tunnel mode (udp / gre / ipsec / ipsec_udp). IPSec is misconfigured: (1) enabled on one side (see *show int tunnel configured*), or mismatched pre-shared key. **down - in progress** – Tunnel is down. Meanwhile, the appliance is exchanging control information with the appliance at the other end, trying to bring up the tunnel. **down - misconfigured** – The two appliances are configured with the same System ID (see *show system*). **up - active** – Tunnel is up and active. Traffic destined for this tunnel is being forwarded to the remote appliance. **up - idle** – Tunnel is up and active, but it has not had activity during the past five minutes, and it has slowed the rate of issuing keep-alive packets. **up - ip sla disabled** – Applies to passthrough tunnels only. Tunnel is up and has connectivity, but it is down because of a configured IP SLA consequent action. **up - reduced functionality** – Tunnel is up and active, but the two endpoint appliances are running mismatched software releases that provide no performance benefit. **UNKNOWN** – Tunnel status is unknown. This can occur because the appliance is unable to retrieve the current tunnel status. Try again later.

Field	Description
MTU	*Maximum Transmission Unit*. The largest possible unit of data that can be sent on a given physical medium. MTUs up to 9000 bytes are supported. **Auto** allows the tunnel MTU to be discovered automatically. It overrides the MTU setting.
Uptime	Length of time the tunnel has been up.
Underlay Tunnels	Designated underlay tunnels.
Live View	Click the chart icon to display a live view of the status of your selected tunnel. You can view by bandwidth, loss, jitter, latency, MOS, chart, traceroute, inbound or outbound, and lock the scale.
Historical Charts	Click the chart icon to display historical charts for the selected overlay and underlay tunnels.

Underlay Subtab

Field	Description
Segment	Name of the segment. This field displays only if Routing Segmentation is enabled.
Underlay Tunnel	Designated underlay tunnel.
Overlays	Overlays to which the tunnels for the appliance are applied.
Admin Status	Indicates whether the tunnel has been set to admin **up** or **down**. To change the admin status for the underlay tunnel, click the menu icon and select **Admin Up** or **Admin Down**.

Field	Description
Status	Click the status indicator in this field to display detailed troubleshooting information for this tunnel. HPE Networking Technical Assistance Center (TAC) can request that you capture and send this diagnostic information. Tunnel statuses are as follows: **down** – Tunnel is down. This can occur because the Admin Status of the tunnel is set to “down” (indicated when down status background is orange) or the tunnel cannot communicate with the appliance at the other end (indicated when down status background is red). Possible causes are: Lack of end-to-end connectivity / routability (test with *iperf*). Intermediate firewall is dropping the packets (open the firewall). Intermediate QoS policy (**be** packets are being starved; change control packet DSCP marking). Mismatched tunnel mode (udp / gre / ipsec / ipsec_udp). IPSec is misconfigured: (1) enabled on one side (see *show int tunnel configured*), or mismatched pre-shared key. **down - in progress** – Tunnel is down. Meanwhile, the appliance is exchanging control information with the appliance at the other end, trying to bring up the tunnel. **down - misconfigured** – The two appliances are configured with the same System ID (see *show system*). **up - active** – Tunnel is up and active. Traffic destined for this tunnel is being forwarded to the remote appliance. **up - idle** – Tunnel is up and active, but it has not had activity during the past five minutes, and it has slowed the rate of issuing keep-alive packets. **up - ip sla disabled** – Applies to passthrough tunnels only. Tunnel is up and has connectivity, but it is down because of a configured IP SLA consequent action. **up - reduced functionality** – Tunnel is up and active, but the two endpoint appliances are running mismatched software releases that provide no performance benefit. **UNKNOWN** – Tunnel status is unknown. This can occur because the appliance is unable to retrieve the current tunnel status. Try again later.

Field	Description
Local IP:Port	IP address and port number of the local appliance.
Remote IP:Port	Public IP address and port number of the remote-peer appliance. This represents the IP address of the EdgeConnect WAN-side interface learned by Orchestrator. Orchestrator uses this to create tunnels between remote sites. **NOTE:** If the NAT-related link under the Next Hop field on the WAN side of the appliance deployment is set to “NAT” (see the figure below), the appliance is behind a NAT-ed interface and the WAN Public IP address is auto discovered from the Cloud Portal.
	
Discovered IP:Port	Discovered IP address and port number, which represents the IP address and port contained in the NAT Discovery (NAT-D) packet sent at the beginning of tunnel setup. If this field displays “NONE:NONE”, the local appliance has not received a NAT-D packet from the remote appliance. This indicates connectivity issues between the locations. If the local appliance receives NAT-D packets, this field populates accordingly, and data path tunnel packets are being received as well. **NOTE:** If Remote IP:Port and Discovered IP:Port are different, they are shown in italics. This is informational only. No action is required by the administrator. Remote IP:Port is what was initially learned by Orchestrator. Discovered IP:Port is the accurate representation of the network in real time.
Max BW (Kbps)	Maximum bandwidth for the tunnel in kilobits per second (Kbps).
Mode	Indicates whether the tunnel protocol is **IPSec**, **IPSec UDP**, **IPSec OTO**, **UDP**, or **GRE**.
Advanced Options	Click the info icon to open the Tunnel Advanced Options dialog box, which displays details about the tunnel’s settings.
Traceroute	Click the chart icon to display a traceroute chart for the selected appliance.
Historical Charts	Click the chart icon to display historical charts for the selected underlay tunnel.

Passthrough Subtab

Field	Description
Passthrough Tunnel	Designated passthrough tunnel.
Charts	Click the chart icon to display historical charts for the selected passthrough tunnel.
Local IP	IP address of the local endpoint.
Remote IP	IP address of the remote endpoint.
Mode	Indicates whether the tunnel protocol is **GRE**, **IPSec**, or **No Encap**.
NAT	Indicates whether Network Address Translation (NAT) has been applied.
Peer/Service	Peer or service being used.

HPE ANW Central Subtab

Field	Description
From Link	The interface label on the local appliance.
To Site	The remote-peer appliance.
To Link	The interface label on the remote-peer appliance.
Mode	Indicates that the tunnel protocol is IPSec OTO.

Troubleshooting

1. *Have you created and applied the Overlay to all the appliances on which you are expecting tunnels to be built?*

   Verify this on the **Apply Overlays** tab.

2. *Are the appliances on which you are expecting the Overlays to be built using Release 8.0 or later?*

   View the active software releases on **Administration > Software > Upgrade > Software Versions**.

3. *Do you have at least one WAN Label selected as a Primary port in the Overlay Policy?*

   Verify this on the Business Intent Overlay tab in the **WAN Links & Bonding Policy** section.

4. *Are the same WAN labels selected in the Overlay assigned to the WAN interfaces on the appliances?*

   Verify that at least one of the *Primary* Labels selected in the Business Intent Overlay is identical to a Label assigned on the appliance's Deployment page. Tunnels are built between matching Labels on all appliances participating in the overlay.

5. *Do any two (or more) appliances have the same Site/Cluster Name?*

   We __*only__* assign the same Site/Cluster Name if we **do not** want those appliances to connect directly. To view the list of Site/Cluster Names, navigate to the **Configuration > Networking > Tunnels > Tunnels** tab, and then click **Sites/Clusters** at the top.

Tunnels Dialog Box

This dialog box enables you to add, modify, or delete manually created underlay and passthrough tunnels for an appliance. In an SD-WAN network, Business Intent Overlays (BIOs) govern automatic tunnel creation and properties.

The Tunnels dialog box includes the following three subtabs:

- Overlay – Displays SD-WAN bonded tunnels. Specifically, overlay tunnels consist of bonded underlay tunnels.
- Underlay – Displays tunnels that map to discrete transports.
- Passthrough – Displays third-party (IPSec) tunnels for service chaining to cloud security services, such as Zscaler and Netskope, and tunnels for local breakouts to trusted SaaS applications, such as Office 365.

About Authentication in IPSec Tunnels

Orchestrator and EdgeConnect Release 9.4 includes Public Key Infrastructure (PKI) use of RFC5280 X.509 certificates for IPSec peer authentication.

Both orchestrated and manually created EdgeConnect-to-EdgeConnect IPSec tunnels can use either certificate-based authentication using <RSA/ECDSA/both> X.509v3 certificates or PSK (pre-shared key). IPSec tunnel peer-authentication options are shown in the following table.

IPSec Tunnel Construction	EdgeConnect-to-EdgeConnect	EdgeConnect-to-3rd Party
Orchestrated IPSec_UDP (default mode)	Proprietary Authentication Not FIPS or Common Criteria approved	Not Applicable
Orchestrated IKE-based IPSec	x509v3 certificate OR Orchestrator-generated Pre-Shared Key (PSK), see Note 1	Orchestrator-generated Pre-Shared Key (PSK), see Note 2
Manually created IKE-based IPSec	x509v3 certificate OR User configured Pre-Shared Key (PSK), see Note 3	User configured Pre-Shared Key (PSK), see Note 3

NOTE 1: Orchestrator automatically creates pre-shared keys on orchestrated IKE-based tunnels with a length of 48 bytes. IPSec PSK is derived from Orchestrator's Random Number Generator (RNG) secure random implementation.

NOTE 2: Orchestrator automatically creates pre-shared keys on orchestrated IKE-based tunnels with a length of 36 characters. IPSec PSK is derived from a pseudo-randomly generated type-4 UUID (universal unique identifier). A 16-byte array is produced, which is then converted to a 36-character string.

NOTE 3: The pre-shared key must contain at least 8 characters, and cannot contain [ ] { } " # * characters. Max length is 64 characters. The default value is "silverpeak".

Add or Modify a Manually Created Underlay Tunnel

To add a manually created underlay tunnel, perform the following steps:

1. Navigate to **Configuration > Networking > Tunnels > Tunnels**.
2. Click the edit icon next to the appliance for which you want to add or modify a tunnel.

   The Tunnels dialog box opens.

   NOTE: To modify a tunnel, click the edit icon next to the tunnel. The Modify Tunnel dialog box opens. Change the fields as described below, and then click **Save**.
3. Click **Underlay**.
4. Click **Add Tunnel**.

   The Add Tunnel dialog box opens.
5. Complete the following fields as appropriate.

   __*Add Tunnel dialog box (for manually created underlay)__*

   The Add Tunnel dialog box displays a General tab. If you set the Mode field on this tab to IPSec, the IKE and IPSec tabs are also displayed.

   General tab (for manually created underlay)

   Access the following fields by clicking the General tab on the Add Tunnel dialog box.

   General

Field	Description
Alias	Alias name of the tunnel.
Mode	Indicates whether the tunnel protocol is **UDP**, **GRE**, **IPSec**, or **IPSec OTO**. If you select IPSec, you can specify the IKE version on the IKE tab. **NOTE:** If using IKE-based IPSec with IKEv2 you can leave this field set to Auto, otherwise it is recommended that you use the AES_256_GCM_16 algorithm, which performs both encryption and authentication, resulting in better performance.
IPSec Suite B Preset	This field is available only if the Mode field is set to IPSec. Select an IPSec Suite B preset if required by the security service (**GCM-128**, **GCM-256**, **GMAC-128**, or **GMAC-256**). The default setting is None. If IPSec Suite B Preset is set to None, no preset is selected, but GCM and GMAC algorithms are available to set independently. If an IPSec Suite B preset is selected, various settings on the IKE and IPSec tabs are configured automatically based on the selected preset.

Field	Description
Admin	Indicates whether the tunnel has been set to admin **up** or **down**.
Local IP / Interface	IP address or interface name of the local endpoint.
Auto discover MTU enabled	When enabled, allows the appliances to auto-negotiate the maximum tunnel bandwidth. Enabled by default.
MTU	Maximum Transmission Unit (MTU) is the largest possible unit of data that can be sent on a given physical medium. For example, the MTU of Ethernet is 1500 bytes. MTUs up to 9000 bytes are supported. Auto allows the tunnel MTU to be discovered automatically, and it overrides the MTU setting. This field is not available if the Auto discover MTU enabled check box is selected.
Auto max BW enabled	When enabled, allows the appliances to auto-negotiate the maximum tunnel bandwidth. Enabled by default.
Max BW Kbps	Maximum amount of bandwidth in kilobits per second. This field is not available if the Auto max BW enabled check box is selected.
UDP destination port	Used in UDP mode. Accept the default value unless the port is blocked by a firewall.
UDP flows	Used in UDP mode. Number of flows over which to distribute tunnel data.
Min BW Kbps	Minimum amount of bandwidth in kilobits per second.

Packet

NOTE: FEC settings do not apply when overlays are used. FEC settings only apply when routing directly to an underlay via Route Policy.

Field	Description
Reorder wait	Maximum time (in milliseconds) the appliance holds an out-of-order packet when attempting to reorder. **100** ms is the default value and should be adequate for most situations. FEC can introduce out-of-order packets if the reorder wait time is not set high enough.
FEC	Set Forward Error Correction (FEC) to **enable**, **disable**, or **auto**.
FEC ratio	When FEC is set to **auto**, FEC will range dynamically from off to 1:10 based on detected loss. The options are **1:1**, **1:2**, **1:5**, **1:10**, and **1:20**. This field is available only if FEC is set to enable.

Tunnel Health

Field	Description
Retry count	Number of failed keep-alive messages that are allowed before the appliance raises a tunnel-down alarm. Default value is 30; maximum value is 60.
DSCP	Determines the DSCP marking that the keep-alive messages should use.

FastFail Thresholds

NOTE: FastFail thresholds were used in a legacy application and should be ignored.

Field	Description
Fastfail enabled	When multiple tunnels are carrying data between two appliances, this feature determines how quickly to disqualify a tunnel from carrying data. The Fastfail connectivity detection algorithm for the wait time from receipt of last packet before declaring a **brownout** is: $Twait = Base + N * RTTavg$ where $Base$ is a value in milliseconds, and N is the multiplier of the average Round Trip Time over the past minute. For example, if: $Base = 200mSN = 2$ Then, $RTTavg = 50mS$ The appliance declares a tunnel to be in brownout if it does not see a reply packet from the remote end within 300 mS of receiving the most recent packet. In the Tunnel Advanced Options, $Base$ is expressed as **Fastfail wait-time base offset** (ms), and N is expressed as **Fastfail RTT multiplication factor**. **Fastfail enabled** – This option is triggered when a tunnel's keep-alive signal does not receive a reply. The options are **disable**, **enable**, and **continuous**. If the disqualified tunnel subsequently receives a keep-alive reply, its recovery is instantaneous. For disable, keep-alives are sent every second, and 30 seconds elapse before failover. In that time, all transmitted data is lost. For enable, keep-alives are sent every second, and a missed reply increases the rate at which keep-alives are sent from one per second to ten per second. Failover occurs after one second. For continuous, keep-alives are continuously sent at ten per second. Therefore, failover occurs after one tenth of a second.

Field	Description
Latency	Amount of latency in milliseconds. Thresholds for Latency, Loss, or Jitter are checked once every second. Receiving three successive measurements in a row that exceed the threshold puts the tunnel into a brownout situation and flows will attempt to fail over to another tunnel within the next 100 ms. Receiving three successive measurements in a row that drop below the threshold will drop the tunnel out of brownout.
Loss	Amount of data lost as a percentage.
Jitter	Amount of jitter in milliseconds.
Fastfail wait-time base offset	Fastfail basic timeout time in milliseconds.
Fastfail RTT multiplication factor	Amount of RTT (Round Trip Time) added to the basic timeout.

IKE tab (for manually created underlay)

Access the following fields by clicking the IKE tab on the Add Tunnel dialog box. This tab is displayed only if the Mode field on the General tab is set to IPSec.

IKE

Field	Description
Peer authentication	There are two options for IKE authentication, End entity certificate or Pre-shared key, choose one of the options. **End entity certificate** – If selected, select the certificate (./Mediate protection against future “harvest now, decrypt later” attacks by quantum computers. Select an encoding standard from the drop-down menu to use for the post quantum pre-shared key (**ascii**, **hex**, or **base 64**). You can enter your own values for the PPK ID (identifier) and PPK Value (secret), or allow Orchestrator to generate both. To have Orchestrator generate the PPK ID and PPK Value, you must select either **hex** or **base 64** encoding, and then click **Generate**. To enter your own PPK ID and PPK Value, select any encoding standard and enter the values in the respective fields. **NOTE:** For orchestrated tunnels, Orchestrator automatically generates the PPK ID and PPK Value for the PPK. **Post quantum pre-shared key mandatory** – If selected, RFC 8784 must be used.
Authentication Algorithm	Authentication algorithm used for IKE security association (SA). The default is SHA1. If the Encryption Algorithm field is set to AES-GCM-128 or AES-GCM-256, this field is not applicable. If the IPSec Suite B Preset field on the General tab is set to None, you can select **SHA1**, **SHA2-256**, **SHA2-384**, or **SHA2-512**. If the IPSec Suite B Preset field is set to any other setting, this field is automatically set to the appropriate algorithm. **NOTE:** With IKEv2 and the Encryption algorithm field set to auto, AES-GCM will probably be negotiated, which includes encryption and authentication. In this case, this field might show a SHA setting that is not actually used. If the Encryption algorithm field is set to AES-GCM-128 or AES-GCM-256, this field will show as NA because the authentication algorithm is already included.

Field	Description
Encryption Algorithm	Encryption algorithm used for IKE security association (SA). The recommendation is to select **AES-GCM-256**. This algorithm also includes authentication (Authentication Algorithm will show as “NA”). If the IPSec Suite B Preset field on the General tab is set to None, and the IKE Version field is set to IKE v1, you can select **AES-CBC-128**, **AES-CBC-256**, or **auto**. The default setting is auto. If the IPSec Suite B Preset field on the General tab is set to None, and the IKE Version field is set to IKE v2, you can select **AES-CBC-128**, **AES-CBC-256**, **AES-GCM-128**, **AES-GCM-256**, or **auto**. The default setting is auto. If the IPSec Suite B Preset field is set to any other setting, this field is automatically set to the appropriate algorithm.
Pseudo Random Function	This field is displayed only if the IKE Encryption Algorithm field is set to AES-GCM-128 or AES-GCM-256. For AES-GCM-128, you can select **SHA2-256**, **SHA2-384**, or **SHA2-512**. For AES-GCM-256, you can select **SHA-384** or **SHA-512**. The recommendation is to select **SHA-384**.
Diffie-Hellman Group	Diffie-Hellman Group used for IKE security association (SA) negotiation. The default setting is DH 14. If the IPSec Suite B Preset field on the General tab is set to None, you can select the appropriate group. Available groups are 14 through 21, 26, and 31. If the IPSec Suite B Preset field is set to any other setting, this field is automatically set to the appropriate group. If IPSec Suite B preset is not selected, then groups 19 or higher are recommended.
Rekey interval/lifetime	Rekey interval/lifetime of IKE security association (SA) in minutes. The default is 360 minutes.
Dead peer detection	**Delay time:** The interval (in seconds) to check the lifetime of the IKE peer. **Retry count:** Number of times to retry the connection before determining that the connection is dead. This field is not editable.

Field	Description
Local IKE identifier	Specify the local IKE identifier. If you are using an end entity certificate for authentication, the certificate Subject Name (SN) is auto populated in this field. The Subject Alternative Name (SAN) may also be used but refer to the following note for guidance. **NOTE:** A bi-directional IPSec tunnel is configured as two uni-directional tunnel objects, each with a local IKE identifier and a remote IKE identifier. In a tunnel from A-to-B, the local IKE identifier for A matches the end-entity certificate for A, while the remote IKE identifier matches the end-entity certificate for B. There is a dependency between the A-to-B tunnel configuration and the B-to-A tunnel configuration. When provisioning the A-to-B tunnel, the Local IKE identifier field is pre-populated with the Subject Name (SN) for A as specified on the end-entity certificate for A, e.g., Subject Name (O=HPE, OU=Aruba, CN=10.81.87.64). If the SN is used for the local IKE identifier for A, then the full SN must also be used as the remote IKE identifier when configuring the tunnel B-to-A. Either the SN or SAN can be used, but whichever is chosen, the local IKE identifier of one direction must match the remote IKE identifier of the other direction. Also, whenever SN is used, the full SN must be used as specified in the end-entity certificate, e.g., O=HPE, OU=Aruba, CN=10.81.87.64 (where O=organization, OU=organizational unit, CN=common name) and it must be entered in the exact order it appears on the certificate. As a best practice, if the SN is preferred, then SN must be used as identifiers for BOTH tunnels. If SAN is preferred, enter the value for SAN without any prefix letters such as “IP:”, and only enter the IP address e.g., 10.81.87.65. When SN/SAN are used interchangeably, the tunnel will not be established.
Remote IKE identifier	Specify the remote IKE identifier. If you are using an end entity certificate for authentication, enter the entire SN or SAN of the peer certificate, and refer to the note in the Local IKE identifier description above for guidance.
Phase 1 mode	Exchange mode for the IKE security association (SA) negotiation. If the IKE Version field is set to IKE v1, you can select **Main** or **Aggressive**. If the IKE Version field is set to IKE v2, this field is automatically set to Aggressive.

Field	Description
IKE version	If the IPSec Suite B Preset field on the General tab is set to None, you can select **IKE v1** or **IKE v2**. The recommendation is to select **IKE v2**. If the IPSec Suite B Preset field is set to any other setting, this field is automatically set to IKE v2.

IPSec tab (for manually created underlay)

Access the following fields by clicking the IPSec tab on the Add Tunnel dialog box. This tab is displayed only if the Mode field on the General tab is set to IPSec.

IPSec

Field	Description
Authentication algorithm	Authentication algorithm used for the IPSec security association (SA). The default is SHA1. If the Encryption Algorithm field is set to AES-GCM-128 or AES-GCM-256, this field is not applicable. If the IPSec Suite B Preset field on the General tab is set to None, you can select **SHA1**, **SHA2-256**, **SHA2-384**, **SHA2-512**, **AES-GMAC-128**, or **AES-GMAC-256**. If the IPSec Suite B Preset field is set to GMAC-128 or GMAC-256, this field is automatically set to the appropriate algorithm.
Encryption algorithm	Encryption algorithm used for the IPSec security association (SA). The recommendation is to select AES-GCM-256. This algorithm also includes authentication (Authentication Algorithm will show as “NA”). If the IPSec Suite B Preset field on the General tab is set to None, and the IPSec Authentication algorithm field is set to SHA1, SHA2-256, SHA2-384, or SHA2-512, you can select **AES-CBC-128**, **AEC-CBC-256**, **AES-GCM-128**, **AES-GCM-256**, **NULL**, or **Auto**. The default setting is auto. If the IPSec Suite B Preset field is set to None, and the IPSec Authentication algorithm field is set to AES-GMAC-128 or AES-GMAC-256, this field is automatically set to NULL.
IPSec anti-replay window	Select a size from the drop-down list or **Disable** to disable the IPSec anti-replay window. If a size is selected, protection is provided against an attacker duplicating encrypted packets by assigning a unique sequence number to each encrypted packet.
Rekey interval/lifetime	Rekey interval/lifetime of the IPSec security association (SA) in minutes. The default is 360 minutes.

Field	Description
Perfect forward secrecy group	Diffie-Hellman group used for IPSec security association (SA) negotiation. The recommendation is to select groups 19 or higher. Based on the setting of the IPSec Suite B Preset field on the the General tab, this field is set to the following Diffie-Hellman group: For None: 14 (by default) For GCM-128 or GMAC-128: 19 For GCM-256 or GMAC-256: 20

6. Click **Save**.

Add or Modify a Manually Created Passthrough Tunnel

To add a manually created passthrough tunnel, perform the following steps:

1. Navigate to **Configuration > Networking > Tunnels > Tunnels**.
2. Click the edit icon next to the appliance for which you want to add or modify a tunnel.

   The Tunnels dialog box opens.

   NOTE: To modify a tunnel, click the edit icon next to the tunnel. The Modify Tunnel dialog box opens. Change the fields as described below, and then click **Save**.
3. Click **Passthrough**.
4. Click **Add Tunnel**.

   The Add Passthrough Tunnel dialog box opens.
5. Complete the following fields as appropriate.

   __*Add Passthrough Tunnel dialog box__*

   The Add Passthrough Tunnel dialog box displays a General tab. If you set the Mode field on this tab to IPSec, the IKE and IPSec tabs are also displayed.

   General tab (for manually created passthrough)

   Access the following fields by clicking the General tab on the Add Passthrough Tunnel dialog box.

Field	Description
Alias	Alias name of the tunnel.

Field	Description
Mode	Indicates whether the tunnel protocol is **GRE**, **No Encap**, **IPSec**, or **IPSec OTO**. **NOTE:** If using IKE-based IPSec with IKEv2 you can leave this field set to Auto, otherwise it is recommended that you use the AES_256_GCM_16 algorithm, which performs both encryption and authentication, resulting in better performance.
IPSec Suite B Preset	This field is available only if the Mode field is set to IPSec. Select an IPSec Suite B preset if required by the security service (**GCM-128**, **GCM-256**, **GMAC-128**, or **GMAC-256**). The default setting is None. If IPSec Suite B Preset is set to None, no preset is selected, but GCM and GMAC algorithms are available to set independently. If an IPSec Suite B preset is selected, various settings on the IKE and IPSec tabs are configured automatically based on the selected preset.
Admin	Indicates whether the tunnel has been set to admin **up** or **down**.
Local IP / Interface	IP address or interface name of the local endpoint.
Remote IP / Hostname	IP address or hostname of the remote endpoint.
NAT	Whether Network Address Translation (NAT) has been applied.
Peer/Service	Enter the peer or service being used.
Auto max BW enabled	When enabled, allows the appliances to auto-negotiate the maximum tunnel bandwidth.
Max BW Kbps	Maximum amount of bandwidth in kilobits per second. This field is not available if the Auto max BW enabled check box is selected.

IKE Tab (for manually created passthrough)

Access the following fields by clicking the IKE tab on the Add Passthrough Tunnel dialog box. This tab is displayed only if the Mode field on the General tab is set to IPSec.

IKE

Field	Description
Pre-shared key	The pre-shared key used for IKE authentication. A default value of “silverpeak” is pre-populated in the Pre-shared key field. It is recommended to change the pre-shared key per the following requirements: The pre-shared key must contain at least 8 characters, and cannot contain [ ] { } ” # * characters. Max length is 64 characters. **NOTE:** If you change the pre-shared key, record the new pre-shared key you entered, as the pre-shared key configuration on both peers should match.
Authentication Algorithm	Authentication algorithm used for IKE security association (SA). The default is SHA1. If the Encryption Algorithm field is set to AES-GCM-128 or AES-GCM-256, this field is not applicable. If the IPSec Suite B Preset field on the General tab is set to None, you can select **SHA1**, **SHA2-256**, **SHA2-384**, or **SHA2-512**. If the IPSec Suite B Preset field is set to any other setting, this field is automatically set to the appropriate algorithm. **NOTE:** With IKEv2 and the Encryption algorithm field set to auto, AES-GCM will probably be negotiated, which includes encryption and authentication. In this case, this field might show a SHA setting that is not actually used.
Encryption Algorithm	Encryption algorithm used for IKE security association (SA). The recommendation is to select **AES-GCM-256**. This algorithm also includes authentication (Authentication Algorithm will show as “NA”). If the IPSec Suite B Preset field on the General tab is set to None, and the IKE Version field is set to IKE v1, you can select **AES-CBC-128**, **AES-CBC-256**, or **auto**. The default setting is auto. If the IPSec Suite B Preset field on the General tab is set to None, and the IKE Version field is set to IKE v2, you can select **AES-CBC-128**, **AES-CBC-256**, **AES-GCM-128**, **AES-GCM-256**, or **auto**. The default setting is auto. If the IPSec Suite B Preset field is set to any other setting, this field is automatically set to the appropriate algorithm.

Field	Description
Pseudo Random Function	This field is displayed only if the IKE Encryption Algorithm field is set to AES-GCM-128 or AES-GCM-256. For AES-GCM-128, you can select **SHA2-256**, **SHA2-384**, or **SHA2-512**. For AES-GCM-256, you can select **SHA-384** or **SHA-512**. The recommendation is to select **SHA-384**.
Diffie-Hellman Group	Diffie-Hellman Group used for IKE security association (SA) negotiation. The default setting is DH 14. If the IPSec Suite B Preset field on the General tab is set to None, you can select the appropriate group. Available groups are 14 through 21, 26, and 31. If the IPSec Suite B Preset field is set to any other setting, this field is automatically set to the appropriate group. If IPSec Suite B preset is not selected, then groups 19 or higher are recommended.
Rekey interval/lifetime	Rekey interval/lifetime of IKE security association (SA) in minutes. The default is 360 minutes.
Dead peer detection	**Delay time:** The interval (in seconds) to check the lifetime of the IKE peer. **Retry count:** Number of times to retry the connection before determining that the connection is dead. This field is not editable.
Local IKE identifier	Specify the local IKE identifier. This field is displayed only if the IKE Version field is set to IKE v2.
Remote IKE identifier	Specify the remote IKE identifier. This field is displayed only if the IKE Version field is set to IKE v2.
Phase 1 mode	Exchange mode for the IKE security association (SA) negotiation. If the IKE Version field is set to IKE v1, you can select **Main** or **Aggressive**. If the IKE Version field is set to IKE v2, this field is automatically set to Aggressive.
IKE version	If the IPSec Suite B Preset field on the General tab is set to None, you can select **IKE v1** or **IKE v2**. The recommendation is to select **IKE v2**. If the IPSec Suite B Preset field is set to any other setting, this field is automatically set to IKE v2.

IPSec tab (for manually created passthrough)

Access the following fields by clicking the IPSec tab on the Add Passthrough Tunnel dialog box. This tab is displayed only if the Mode field on the General tab is set to IPSec.

IPSec

Field	Description
Authentication algorithm	Authentication algorithm used for the IPSec security association (SA). The default is SHA1. If the Encryption Algorithm field is set to AES-GCM-128 or AES-GCM-256, this field is not applicable. If the IPSec Suite B Preset field on the General tab is set to None, you can select **SHA1**, **SHA2-256**, **SHA2-384**, **SHA2-512**, **AES-GMAC-128**, or **AES-GMAC-256**. If the IPSec Suite B Preset field is set to GMAC-128 or GMAC-256, this field is automatically set to the appropriate algorithm.
Encryption algorithm	Encryption algorithm used for the IPSec security association (SA). The recommendation is to select AES-GCM-256. This algorithm also includes authentication (Authentication Algorithm will show as “NA”). If the IPSec Suite B Preset field on the General tab is set to None, and the IPSec Authentication algorithm field is set to SHA1, SHA2-256, SHA2-384, or SHA2-512, you can select **AES-CBC-128**, **AEC-CBC-256**, **AES-GCM-128**, **AES-GCM-256**, **NULL**, or **Auto**. The default setting is auto. If the IPSec Suite B Preset field is set to None, and the IPSec Authentication algorithm field is set to AES-GMAC-128 or AES-GMAC-256, this field is automatically set to NULL.
IPSec anti-replay window	Select a size from the drop-down list or **Disable** to disable the IPSec anti-replay window. If a size is selected, protection is provided against an attacker duplicating encrypted packets by assigning a unique sequence number to each encrypted packet.
Rekey interval/lifetime	Rekey interval/lifetime of the IPSec security association (SA) in minutes. The default is 360 minutes.

Field	Description
Perfect forward secrecy group	Diffie-Hellman group used for IPSec security association (SA) negotiation. The recommendation is to select groups 19 or higher. Based on the setting of the IPSec Suite B Preset field on the the General tab, this field is set to the following Diffie-Hellman group: For None: 14 (by default) For GCM-128 or GMAC-128: 19 For GCM-256 or GMAC-256: 20

6. Click **Save**.

Delete a Tunnel

To delete a tunnel listed in the table on the Underlay or Passthrough subtab of the Tunnels dialog box, click the corresponding delete icon (**X**) in the last column.

Use Passthrough Tunnels

NOTE: On upgrade, Orchestrator will delete previously created Passthrough tunnels for alias interfaces.

Use passthrough tunnels in the following situations:

- For internet breakout to a trusted SaaS application, like Office 365
- For service chaining to a cloud security service, like Zscaler or Netskope
  - This requires building secure and compatible third-party IPSec tunnels from EdgeConnect devices to non-EdgeConnect devices in the data center or cloud.
  - When you create the tunnel, the Service Name in the Business Intent Overlay's Internet Traffic Policies must exactly match the Peer/Service specified in the Passthrough tunnel configuration.
  - To load balance, create two or more passthrough IPSec tunnels and, in the Business Intent Overlay, ensure that they all specify the same Service Name in the Internet Traffic Policies.

IPSec Suite B Presets

As of version 9.2, Orchestrator provides you with four IPSec Suite B presets, as follows:

- GCM-128
- GCM-256
- GMAC-128
- GMAC-256

Each preset includes a predetermined set of IKE and ESP (IPSec) cryptographic algorithms. By selecting an IPSec Suite B preset, you can streamline the algorithm aspect of your tunnel setup rather than selecting individual algorithms. However, you can select individual algorithms if you want to. To select a preset, use the **IPSec Suite B Preset** drop-down field on the Add Tunnel or Modify Tunnel dialog box.

The following tables show the IPSec Suite B presets in the header row and provide the associated algorithm setups for the IKEv2 and ESP (IPSec) stages.

IKEv2 Stage

	GCM-128	GCM-256	GMAC-128	GMAC-256
Encryption (Note)	AES-128-CBC	AES-256-CBC	AES-128-CBC	AES-256-CBC
Pseudo Random Function	HMAC-SHA-256	HMAC-SHA-384	HMAC-SHA-256	HMAC-SHA-384
Integrity (IKE Data Authentication)	HMAC-SHA-256-128	HMAC-SHA-384-192	HMAC-SHA-256-128	HMAC-SHA-384-192
Key Exchange (NIST Elliptic Curve Groups)	DH-19 256-bit Prime Size	DH-20 384-bit Prime Size	DH-19 256-bit Prime Size	DH-20 384-bit Prime Size

ESP (IPSec) Stage

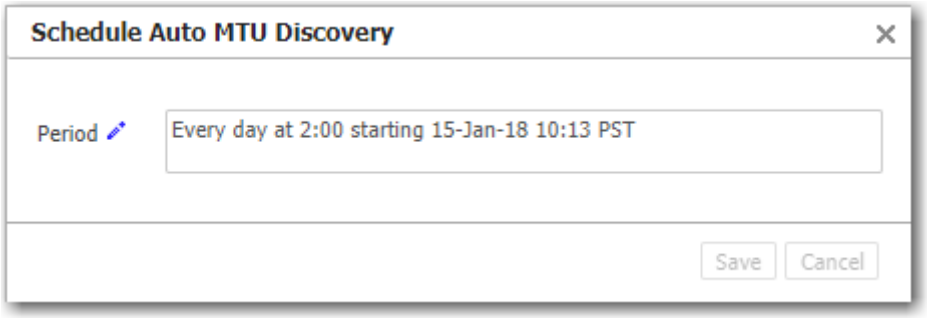
	GCM-128	GCM-256	GMAC-128	GMAC-256
Encryption	AES-128-GCM with 16 octet ICV	AES-256-GCM with 16 octet ICV	NULL	NULL
Integrity (Data Authentication)	NULL	NULL	AES-128-GMAC	AES-256-GMAC

Notice in the second table that the encryption and data authentication is done in one step for GCM. For GMAC, there is no encryption.

Schedule Auto MTU Discovery

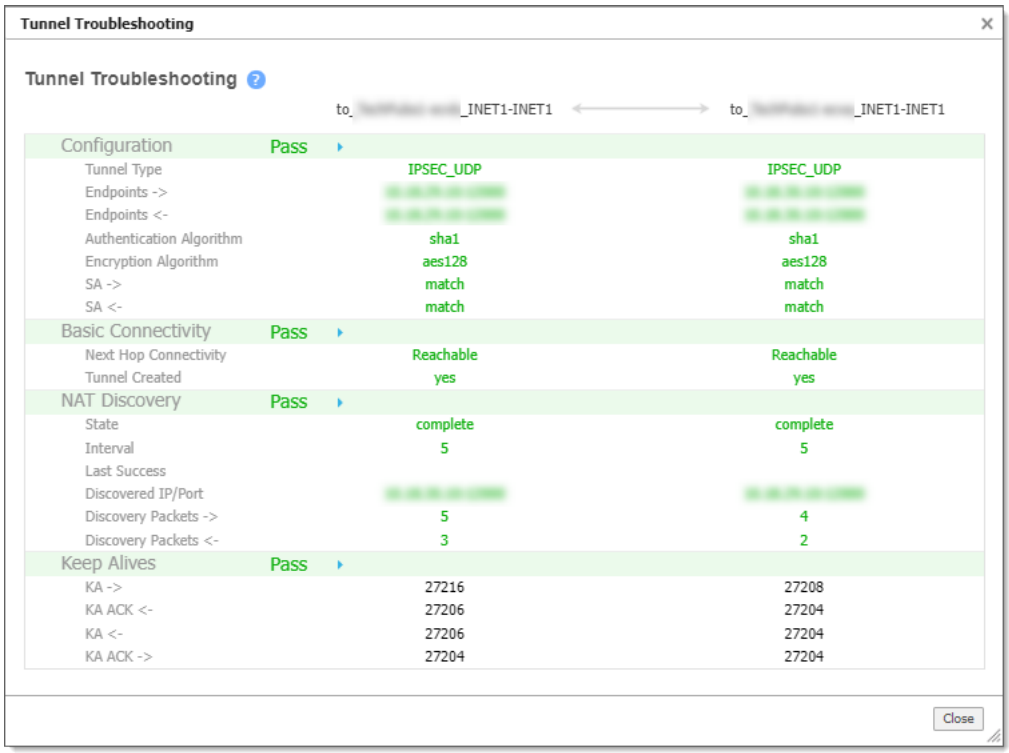
Configuration > Networking > Tunnels > Auto MTU Discovery

Use this dialog box to schedule when to discover Auto MTU.



Tunnel Troubleshooting

The Tunnel Troubleshooting dialog box provides some basic diagnostic results for SD-WAN fabric tunnels. HPE Networking Technical Assistance Center (TAC) can request that you capture and send this diagnostic information.



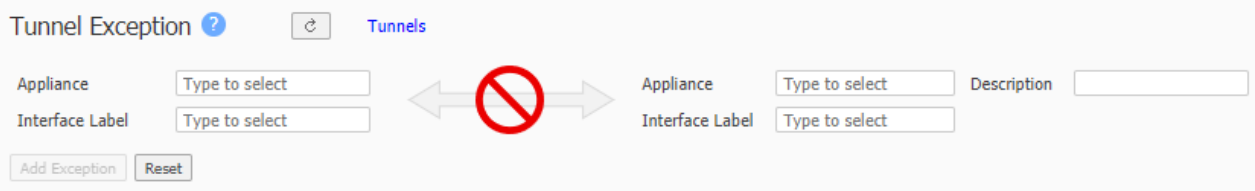
Tunnel Exception

Configuration > Networking > Tunnels > Tunnel Exception

Orchestrator includes a tunnel exception feature that enables you to specify tunnel transactions between overlays. There are two ways you can enable this feature in Orchestrator.

You can configure tunnel exceptions through the Tunnel Exception tab.

1. Select the two appliances that you do not want connected via a tunnel.
2. Enter the Interface Labels.



The interface label can be any type of connection, such as **any**, **MPLS**, **Internet**, or **LTE**. Specifying the label excludes appliances within a given network to communicate with that particular appliance.

NOTE: Use the description field to add a comment if you want to indicate why you are adding an exception.

Configuration > Policies

The options under **Configuration > Policies** focus on managing access lists and policies.

Route Policies Tab

Configuration > Templates & Policies > Policies > Route Policies

Use the Route Policies tab to view route policies that exist on the appliances selected in the appliance tree. This includes the appliance-based defaults, entries applied manually by using the Appliance Manager or CLI, and entries that result from applying an Orchestrator Route Policies template or Business Intent Overlay (if you are deploying an SD-WAN).

To directly manage route policies for a particular appliance, click the associated edit icon. To create and manage route policies with a template, click **Manage Route Policies with Templates**.

For details about fields on this tab, see Route Policies Dialog Box below.

Route Policies Dialog box

Use the Route Policies tab to directly manage route policies for a particular appliance.

The default behavior for each appliance is to auto-optimize all IP traffic, automatically directing flows to the appropriate tunnel. Auto-optimization strategies reduce the need to create explicit route map entries for optimization. The three strategies provided are:

- TCP-based auto-opt
- IP-based auto-opt

- Subnet sharing

By default, all three are enabled on the Route Policies template.

Route policies only require entries for flows that are to be:

- Sent passthrough (shaped or unshaped)
- Dropped
- Configured for a specific high-availability deployment
- Routed based on application, VLAN, DSCP, or ACL (Access Control List)

When multiple tunnels exist to a remote peer, you can create a route policy entry that enables the appliance to dynamically select the best path based on one of these criteria:

- Load balancing
- Lowest loss
- Lowest latency
- Specified tunnel

If you are deploying an SD-WAN network and setting up Internet breakout from the branch, you must create manual route policy entries for sanctioned SaaS applications or Guest WiFi.

Priority

- If you are using Orchestrator templates to add rules, Orchestrator will delete all entries from **1000 – 9999** before applying its policies.
- You can create rules with higher priority than Orchestrator rules (**1 – 999**) and rules with lower priority (**10000 – 19999** and **25000 – 65534**).

  NOTE: The priority range from **20000** to **24999** is reserved for Orchestrator.
- When adding a rule, the priority is incremented by 10 from the previous rule. The priority can be changed, but this default behavior helps to ensure you can insert new rules without having to change subsequent priorities.

Match Criteria

Match criteria are used universally across all policy maps for route, QoS, optimization, SaaS NAT, and firewall zone security policies.

To use the same match criteria in different maps, create an ACL (Access Control List), which is a named, reusable set of rules. For efficiency, navigate to **Configuration > Templates & Policies > ACLs > Access Lists** to create ACLs, and then apply them across your appliances.

Use the Match Criteria dialog box to select and configure a variety of match criteria options. Some are explained below:

- **Address Map:** Use this option to sort by country, IP address owner, or SaaS application. You can also select and configure Microsoft Instance, Microsoft Category, and Proxy attributes for an address map. These attributes are secondary parameters to the address map. They are evaluated for a policy match only when the configured address map matches the flow. To select and configure these attributes, click **+Attributes**.

- Match criteria options related to Secure Web Services include:
  - **URL:** Omit the protocol in URLs and include a slash character (/) or a slash followed by a query parameter. For example, *google.com/* or *google.com/maps* are valid; *https://google.com/* is not valid. You can also use the asterisk (*) wildcard character, such as in *google.com/**, to specify a domain, but it is more appropriate to use the Domain match criteria option to specify *google.com/*. Separate multiple URL addresses with the pipe character (|). For example, you can specify two URL addresses as *bing.com/*|google.com/**.
  - **Web Category:** If you select this option, click **Select**, and then click **Strict**, **Moderate**, or **Custom**. Strict and Moderate include predetermined web category selections, which you cannot change. However, if you attempt to select or clear a web category for these, the Custom group of web category selections opens automatically. You can customize web category selections in the Custom group. After you make your selections, click **Ok**.
  - **Web Reputation:** If you select this option, select one or more web reputation categories (High Risk, Suspicious, Moderate Risk, Low Risk, or Trustworthy) from the list. To select more than one, press and hold the **Shift** key, and then click the appropriate categories.
  - **Bad IP Reputation:** Selecting this option matches on flows for which the source or destination IP address (or both) has a High Risk reputation score.
- Match criteria options related to RADIUS user role include:
  - **User Role:** This is the user role as specified in the authentication exchange with the ClearPass RADIUS server.
  - **User Name**
  - **User Group**
  - **User Device**
  - **User MAC**
  - **User Vlan**

  Configuring these match criteria related to user role enables an EdgeConnect to automatically assign traffic steering and firewall zone policies.
- Use the Src:Dest check box associated with several match criteria options to specify separate criteria for inbound and outbound traffic. You can configure source and destination role-based policies when both source and destination users are in the same network.

Source or Destination

- An IP address can specify a subnet; for example, 10.10.10.0/24 (IPv4) or fe80::204:23ff:fed8:4ba2/64 (IPv6).
- To allow **any IP address**, use 0.0.0.0/0 (IPv4) or ::/0 (IPv6).
- Ports are available only for the protocols **tcp**, **udp**, and **tcp/udp**.
- To allow **any port**, use **0**.

Wildcard-based Prefix Matching Rules

- Even when using a range or a wildcard, the IPv4 address must be specified in the 4-octet format, separated by the dot notation. For example, **A.B.C.D**.
- Range is specified using a single dash. For example, **128-129**.
- Wildcard is specified as an asterisk (___*).
- Range and Wildcard can both be used in the same address, but an octet can only contain one or the other. For example, **10.136-137.*.64-95**.
- A wildcard can only be used to define an entire octet. For example, **10.13*.*.64-95** is not supported. Use **10.130-139.*.64-95** to specify this range.
- The same rules apply to IPv6 addressing.
- CIDR notation and (Range or Wildcard) are mutually exclusive in the same address. For example, **192.168.0.1-127/24** is not supported. Use either **192.168.0.0/24** or **192.168.0.1-127**.
- These prefix-matching rules apply to the following policies only: Route, QoS, Optimization, NAT, Security, and ACLs.

QoS Policies Tab

Configuration > *Templates & Policies* > *Policies* > *QoS Policies*

QoS policies determine how flows are queued and marked. Use the QoS Policies tab to view the QoS policies that exist on the appliances selected in the appliance tree. This includes the appliance-based defaults, entries applied manually by using the Appliance Manager or CLI, and entries that result from applying an Orchestrator QoS Policy template or Business Intent Overlay.

Use the Shaper to define, prioritize, and name traffic classes. The Shaper defines, and the QoS policy assigns. For information about the Shaper, see Shaper Tab.

To directly manage QoS policies for a particular appliance, click the associated edit icon. To create and manage QoS policies by using a template, click **Manage QoS Policies with Templates**.

For details about fields on this tab, see QoS Policies Dialog Box below.

QoS Policies Dialog Box

Use the QoS Policies dialog box to directly manage QoS policies for a particular appliance.

Set Actions for QoS policies determine two things:

- To what traffic class a shaped flow—optimized or passthrough—is assigned
- Whether to trust incoming DSCP markings for LAN QoS and WAN QoS, or to remark them as they leave for the WAN

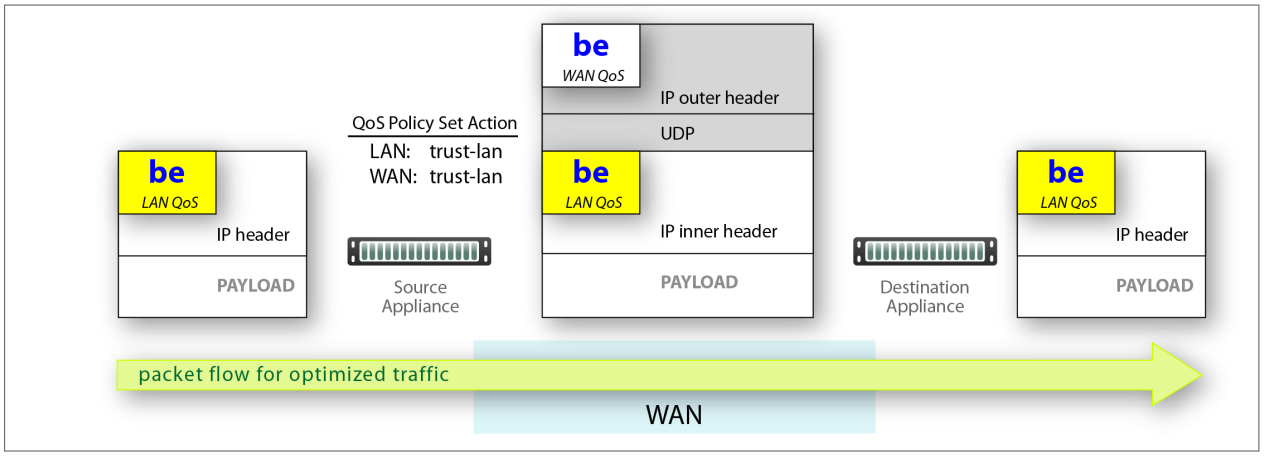
Handle and Mark DSCP Packets

- DSCP markings specify end-to-end QoS policies throughout a network.
- The default values for LAN QoS and WAN QoS are “trust-lan”.

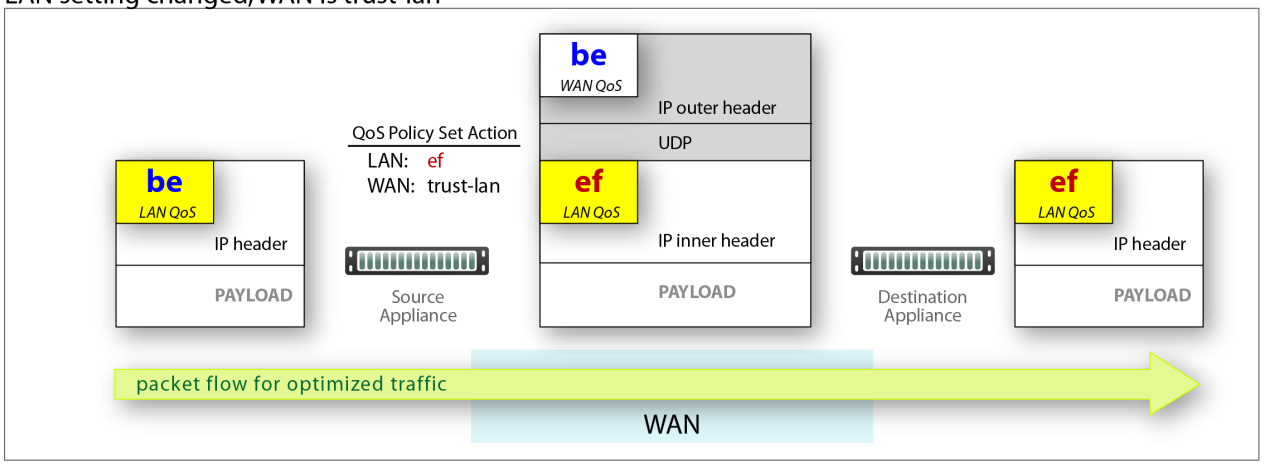
Apply DSCP Markings to Optimized (Tunnelized) Traffic

- The appliance encapsulates optimized traffic. This adds an IP outer header to packets for travel across the WAN. This outer header contains the WAN QoS DSCP marking.
- **LAN QoS:** The DSCP marking applied to the IP header before encapsulation.
- **WAN QoS:** The DSCP marking in the encapsulating outer IP header. The remote appliance removes the outer IP header.

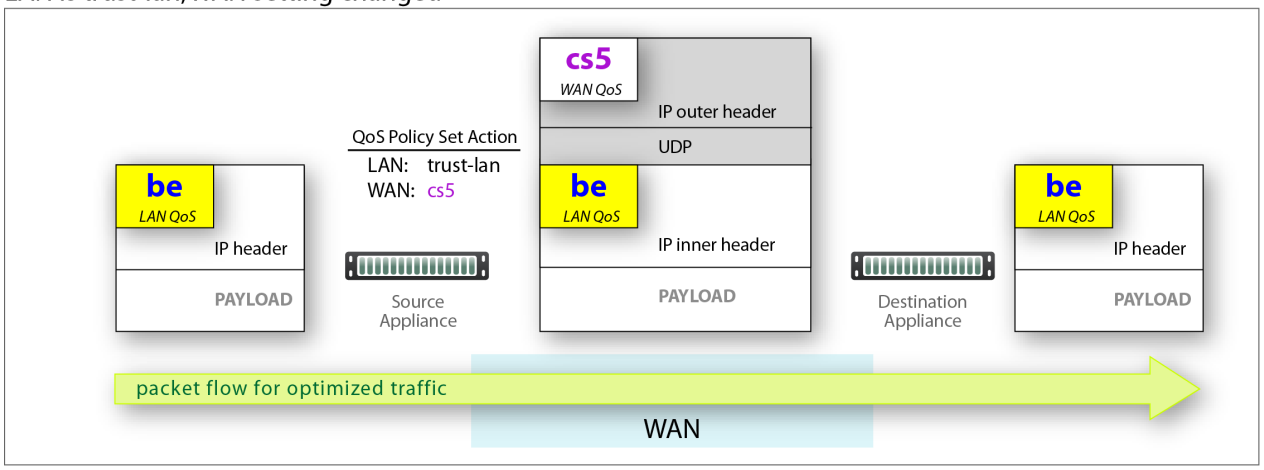
LAN and WAN set to trust-lan



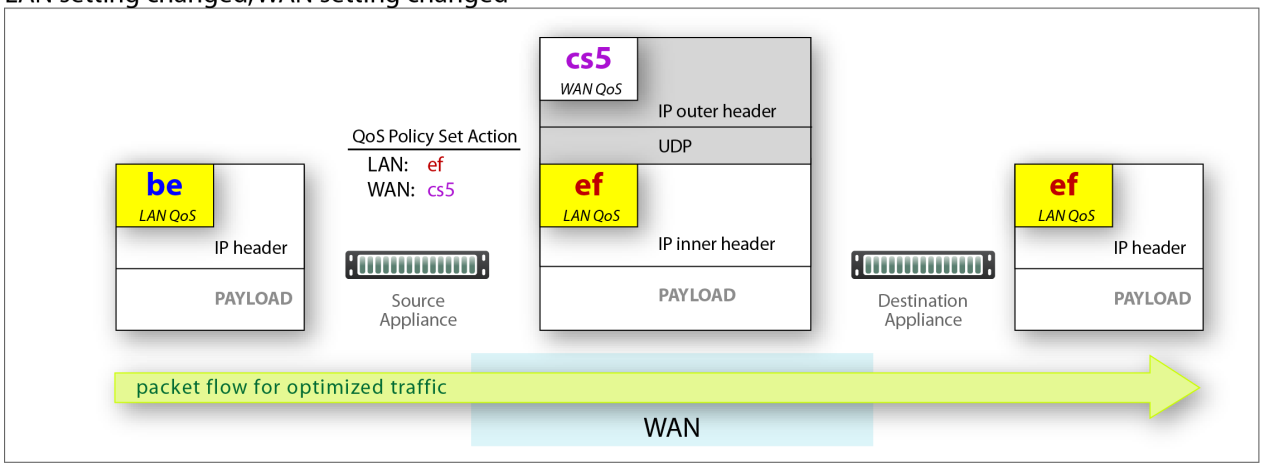
LAN setting changed, WAN is trust-lan



LAN is trust-lan, WAN setting changed



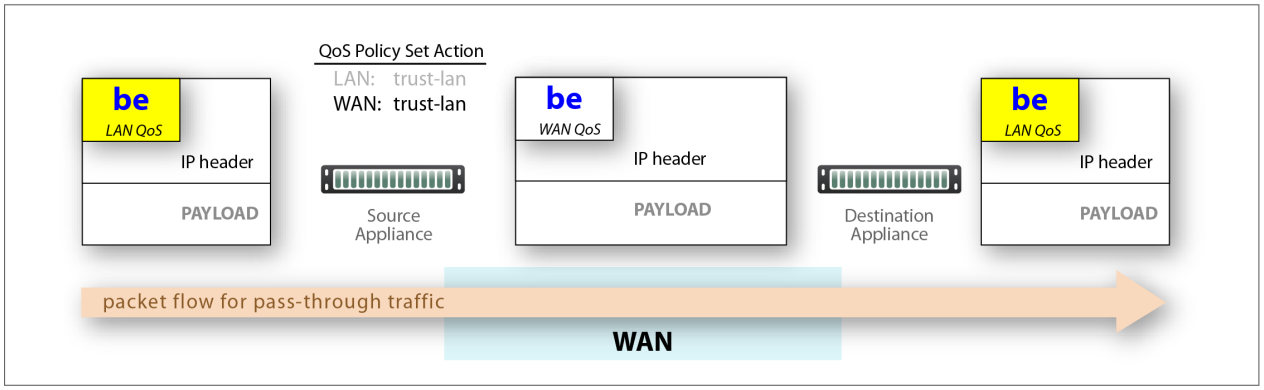
LAN setting changed, WAN setting changed



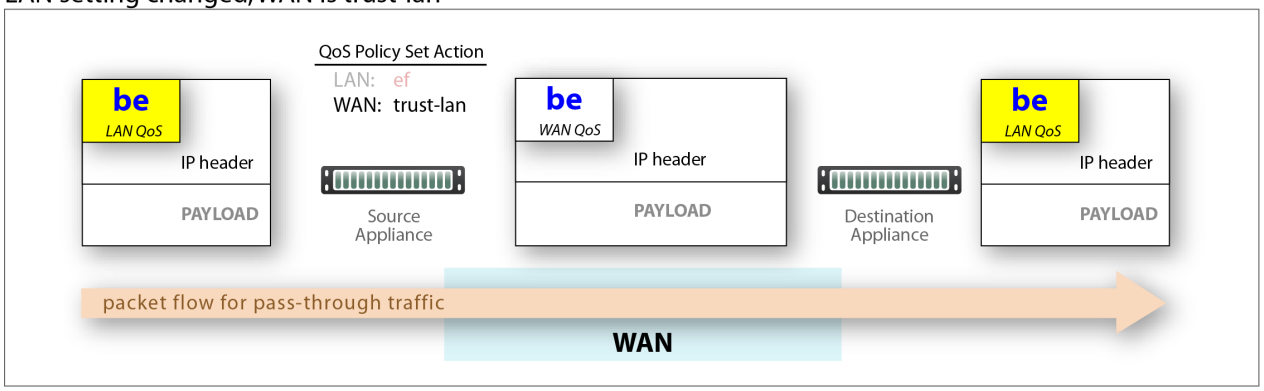
Apply DSCP Markings to Passthrough Traffic

- The appliance applies the QoS policy's DSCP markings to all passthrough flows—shaped and unshaped.
- Passthrough traffic does not receive an additional header, so it is handled differently:
  - The QoS policy's LAN QoS Set Action is ignored.
  - The specified WAN QoS marking replaces the packet's existing LAN QoS DSCP marking.
  - When the packet reaches the remote appliance, it retains the modified QoS setting as it travels to its destination.

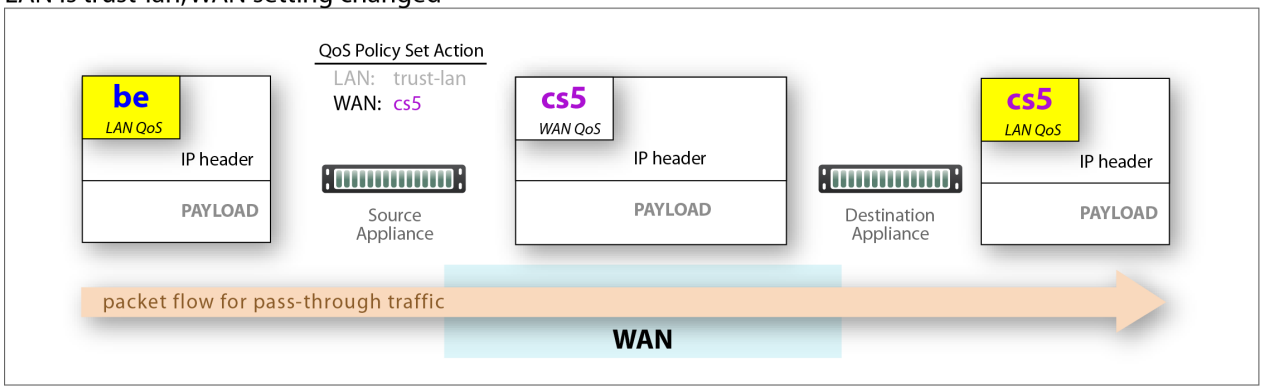
LAN and WAN set to trust-lan



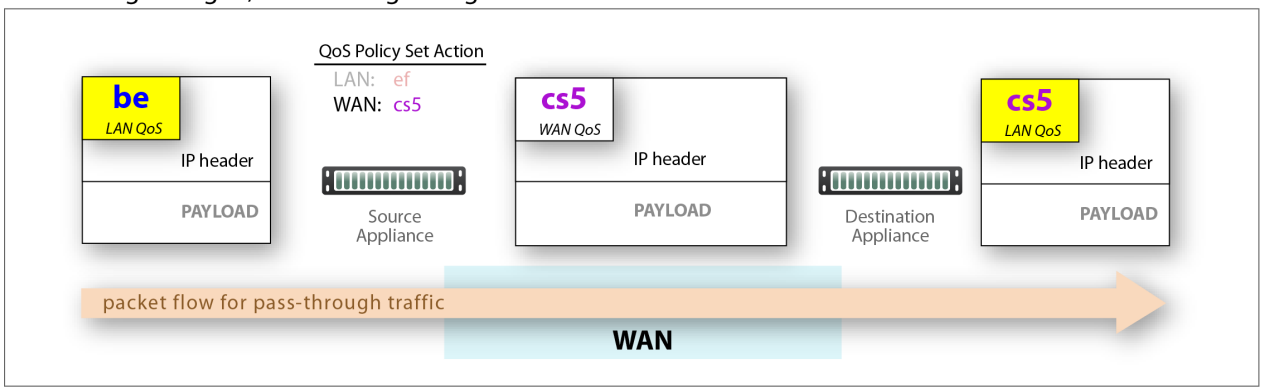
LAN setting changed, WAN is trust-lan



LAN is trust-lan, WAN setting changed



LAN setting changed,WAN setting changed



Priority

- If you are using Orchestrator templates to add rules, Orchestrator will delete all entries from **1000 – 9999** before applying its policies.
- You can create rules with higher priority than Orchestrator rules (**1 – 999**) and rules with lower priority (**10000 – 19999** and **25000 – 65534**).

  NOTE: The priority range from **20000** to **24999** is reserved for Orchestrator.
- When adding a rule, the priority is incremented by 10 from the previous rule. The priority can be changed, but this default behavior helps to ensure you can insert new rules without having to change subsequent priorities.

Match Criteria

Match criteria are used universally across all policy maps for route, QoS, optimization, SaaS NAT, and firewall zone security policies.

To use the same match criteria in different maps, create an ACL (Access Control List), which is a named, reusable set of rules. For efficiency, navigate to **Configuration > Templates & Policies > ACLs > Access Lists** to create ACLs, and then apply them across your appliances.

Use the Match Criteria dialog box to select and configure a variety of match criteria options. Some are explained below:

- **Address Map:** Use this option to sort by country, IP address owner, or SaaS application. You can also select and configure Microsoft Instance, Microsoft Category, and Proxy attributes for an address map. These attributes are secondary parameters to the address map. They are evaluated for a policy match only when the configured address map matches the flow. To select and configure these attributes, click **+Attributes**.
- Match criteria options related to Secure Web Services include:

  - **URL:** Omit the protocol in URLs and include a slash character (/) or a slash followed by a query parameter. For example, *google.com/* or *google.com/maps* are valid; *https://google.com/* is not valid. You can also use the asterisk (*) wildcard character, such as in *google.com/**, to specify a domain, but it is more appropriate to use the Domain match criteria option to specify *google.com/*. Separate multiple URL addresses with the pipe character (|). For example, you can specify two URL addresses as *bing.com/*|google.com/**.
  - **Web Category:** If you select this option, click **Select**, and then click **Strict**, **Moderate**, or **Custom**. Strict and Moderate include predetermined web category selections, which you cannot change. However, if you attempt to select or clear a web category for these, the Custom group of web category selections opens automatically. You can customize web category selections in the Custom group. After you make your selections, click **Ok**.
  - **Web Reputation:** If you select this option, select one or more web reputation categories (High Risk, Suspicious, Moderate Risk, Low Risk, or Trustworthy) from the list. To select more than one, press and hold the **Shift** key, and then click the appropriate categories.
  - **Bad IP Reputation:** Selecting this option matches on flows for which the source or destination IP address (or both) has a High Risk reputation score.

- Match criteria options related to RADIUS user role include:

  - **User Role:** This is the user role as specified in the authentication exchange with the ClearPass RADIUS server.
  - **User Name**
  - **User Group**
  - **User Device**
  - **User MAC**
  - **User Vlan**

  Configuring these match criteria related to user role enables an EdgeConnect to automatically assign traffic steering and firewall zone policies.

- Use the Src:Dest check box associated with several match criteria options to specify separate criteria for inbound and outbound traffic. You can configure source and destination role-based policies when both source and destination users are in the same network.

Source or Destination

- An IP address can specify a subnet; for example, 10.10.10.0/24 (IPv4) or fe80::204:23ff:fed8:4ba2/64 (IPv6).
- To allow **any IP address**, use 0.0.0.0/0 (IPv4) or ::/0 (IPv6).
- Ports are available only for the protocols **tcp**, **udp**, and **tcp/udp**.
- To allow **any port**, use **0**.

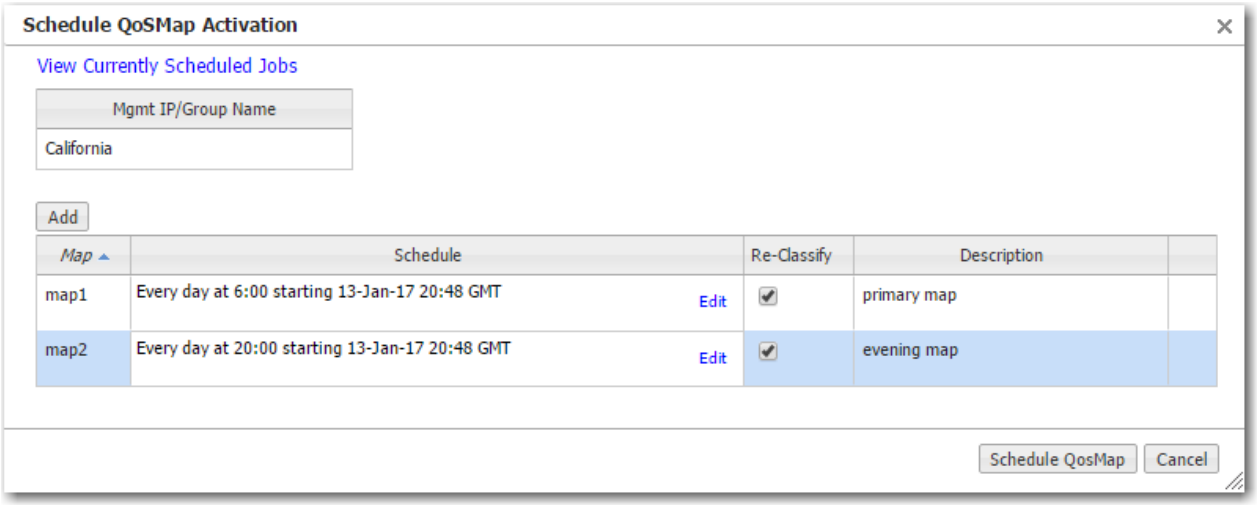
Wildcard-based Prefix Matching Rules

- Even when using a range or a wildcard, the IPv4 address must be specified in the 4-octet format, separated by the dot notation. For example, **A.B.C.D**.
- Range is specified using a single dash. For example, **128-129**.
- Wildcard is specified as an asterisk (____*).
- Range and Wildcard can both be used in the same address, but an octet can only contain one or the other. For example, **10.136-137.*.64-95**.
- A wildcard can only be used to define an entire octet. For example, **10.13*.*.64-95** is not supported. Use **10.130-139.*.64-95** to specify this range.
- The same rules apply to IPv6 addressing.
- CIDR notation and (Range or Wildcard) are mutually exclusive in the same address. For example, **192.168.0.1-127/24** is not supported. Use either **192.168.0.0/24** or **192.168.0.1-127**.
- These prefix-matching rules apply to the following policies only: Route, QoS, Optimization, NAT, Security, and ACLs.

Schedule QoS Map Activation

Configuration > Templates & Policies > Policies > Schedule QoSMap Activation

You can schedule appliances to apply different QoS maps at different times.



Before using this option, verify the following:

- The desired Template Group has the QoS maps you need.
- You have applied the Template Group to the appliances you want to schedule.

TIP: To specify the timezone for scheduled jobs and reports, use the Schedule Timezone window (**Orchestrator > Software & Setup > Setup > Timezone for Scheduled Jobs**).

Optimization Policies Tab

Configuration > Templates & Policies > Policies > Optimization Policies

Use the Optimization Policies tab to view the optimization policies that exist on the appliances selected in the appliance tree. This includes the appliance-based defaults, entries applied manually by using the Appliance Manager or CLI, and entries that result from applying an Orchestrator Optimization Policy template or Business Intent Overlay.

To directly manage optimization policies for a particular appliance, click the associated edit icon. To create and manage optimization policies by using a template, click **Manage Optimization Policies with Templates**.

For details about fields on this tab, see Optimization Policies Dialog Box below.

Optimization Policies Dialog Box

Use the Optimization Policies tab to directly manage optimization policies for a particular appliance.

Priority

- If you are using Orchestrator templates to add rules, Orchestrator deletes all entries from **1000 – 9999** before applying its policies.
- You can create rules with higher priority than Orchestrator rules (**1 – 999**) and rules with lower priority (**10000 – 19999** and **25000 – 65534**).

  NOTE: The priority range from **20000** to **24999** is reserved for Orchestrator.
- When adding a rule, the priority is incremented by 10 from the previous rule. The priority can be changed, but this default behavior helps ensure that you can insert new rules without having to change subsequent priorities.

Match Criteria

Match criteria are used universally across all policy maps for route, QoS, optimization, SaaS NAT, and firewall zone security policies.

To use the same match criteria in different maps, create an ACL (Access Control List), which is a named, reusable set of rules. For efficiency, navigate to **Configuration > Templates & Policies > ACLs > Access Lists** to create ACLs, and then apply them across your appliances.

Use the Match Criteria dialog box to select and configure a variety of match criteria options. Some are explained below:

- **Address Map:** Use this option to sort by country, IP address owner, or SaaS application. You can also select and configure Microsoft Instance, Microsoft Category, and Proxy attributes for an address map. These attributes are secondary parameters to the address map. They are evaluated for a policy match only when the configured address map matches the flow. To select and configure these attributes, click **+Attributes**.
- Match criteria options related to Secure Web Services include:

  - **URL:** Omit the protocol in URLs and include a slash character (/) or a slash followed by a query parameter. For example, *google.com/* or *google.com/maps* are valid; *https://google.com/* is not valid. You can also use the asterisk (*) wildcard character, such as in *google.com/**, to specify a domain, but it is more appropriate to use the Domain match criteria option to specify *google.com/*. Separate multiple URL addresses with the pipe character (|). For example, you can specify two URL addresses as *bing.com/*|google.com/**.
  - **Web Category:** If you select this option, click **Select**, and then click **Strict**, **Moderate**, or **Custom**. Strict and Moderate include predetermined web category selections, which you cannot change. However, if you attempt to select or clear a web category for these, the Custom group of web category selections opens automatically. You can customize web category selections in the Custom group. After you make your selections, click **Ok**.
  - **Web Reputation:** If you select this option, select one or more web reputation categories (High Risk, Suspicious, Moderate Risk, Low Risk, or Trustworthy) from the list. To select more than one, press and hold the **Shift** key, and then click the appropriate categories.
  - **Bad IP Reputation:** Selecting this option matches on flows for which the source or destination IP address (or both) has a High Risk reputation score.

- Match criteria options related to RADIUS user role include:

  - **User Role:** This is the user role as specified in the authentication exchange with the ClearPass RADIUS server.
  - **User Name**
  - **User Group**
  - **User Device**
  - **User MAC**
  - **User Vlan**

  Configuring these match criteria related to user role enables an EdgeConnect to automatically assign traffic steering and firewall zone policies.

- Use the Src:Dest check box associated with several match criteria options to specify separate criteria for inbound and outbound traffic. You can configure source and destination role-based policies when both source and destination users are in the same network.

Source or Destination

- An IP address can specify a subnet; for example, 10.10.10.0/24 (IPv4) or fe80::204:23ff:fed8:4ba2/64 (IPv6).
- To allow **any IP address**, use 0.0.0.0/0 (IPv4) or ::/0 (IPv6).
- Ports are available only for the protocols **tcp**, **udp**, and **tcp/udp**.
- To allow **any port**, use **0**.

Wildcard-based Prefix Matching Rules

- Even when using a range or a wildcard, the IPv4 address must be specified in the 4-octet format, separated by the dot notation. For example, **A.B.C.D**.
- Range is specified using a single dash. For example, **128-129**.
- Wildcard is specified as an asterisk (____*).
- Range and Wildcard can both be used in the same address, but an octet can only contain one or the other. For example, **10.136-137.*.64-95**.
- A wildcard can only be used to define an entire octet. For example, **10.13*.*.64-95** is not supported. Use **10.130-139.*.64-95** to specify this range.
- The same rules apply to IPv6 addressing.
- CIDR notation and (Range or Wildcard) are mutually exclusive in the same address. For example, **192.168.0.1-127/24** is not supported. Use either **192.168.0.0/24** or **192.168.0.1-127**.
- These prefix-matching rules apply to the following policies only: Route, QoS, Optimization, NAT, Security, and ACLs.

Set Actions

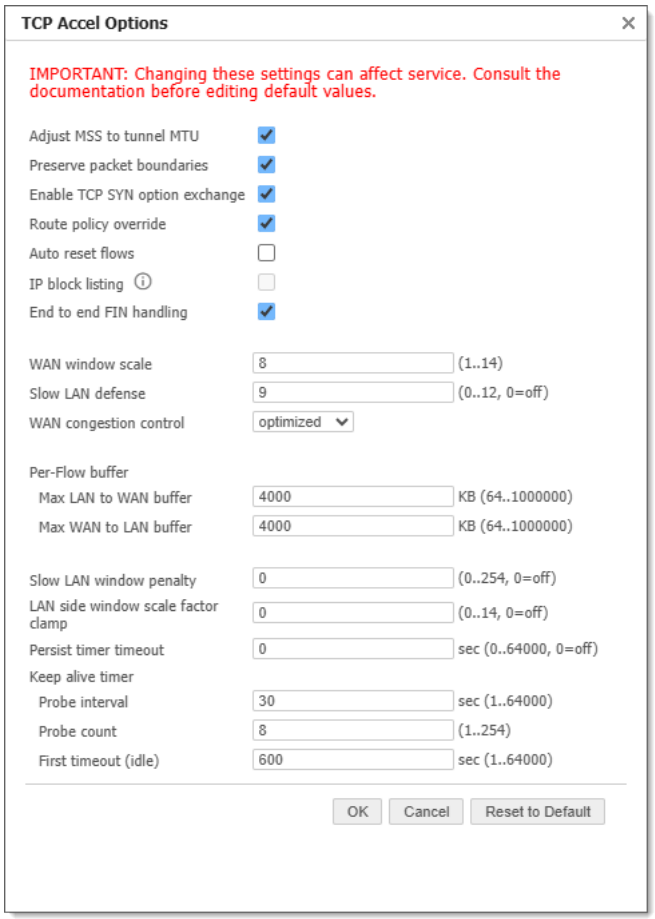
Set Action	Description
Network Memory	Addresses limited bandwidth. This technology uses advanced fingerprinting algorithms to examine all incoming and outgoing WAN traffic. Network Memory localizes information and transmits only modifications between locations. **maximize reduction:** Optimizes for maximum data reduction at the potential cost of slightly lower throughput and/or some increase in latency. It is appropriate for bulk data transfers, such as file transfers and FTP, when bandwidth savings are the primary concern. **minimize latency:** Ensures that Network Memory processing adds no latency. This might come at the cost of lower data reduction. It is appropriate for extremely latency-sensitive interactive or transactional traffic. It is also appropriate when the primary objective is to fully utilize the WAN pipe to increase the LAN-side throughput, as opposed to conserving WAN bandwidth. **balanced:** This is the default setting. It dynamically balances latency and data reduction objectives. It is the best choice for most traffic types. **disabled:** Turns off Network Memory.

Set Action	Description
IP Header Compression	Process of compressing excess protocol headers before transmitting them on a link and uncompressing them to their original state at the other end. It is possible to compress the protocol headers due to the redundancy in header fields of the same packet as well as in consecutive packets of a packet stream.
Payload Compression	Uses algorithms to identify relatively short byte sequences that are repeated frequently. These are then replaced with shorter segments of code to reduce the size of transmitted data. Simple algorithms can find repeated bytes within a single packet; more sophisticated algorithms can find duplication across packets and even across flows.
TCP Accel	TCP Acceleration. Uses techniques such as selective acknowledgments, window scaling, and maximum segment size adjustment to mitigate poor performance on high-latency links. **NOTE:** The slow LAN alert is issued when the loss has fallen below 80% of the specified value configured in the TCP Accel Options dialog box. To open the TCP Accel Options dialog box and view the available advanced options, click the info icon. For more information, see TCP Acceleration below.
Protocol Accel	Protocol Accelertion. Provides explicit configuration for optimizing CIFS, SSL, SRDF, Citrix, and iSCSI protocols. In a network environment, it is possible that not every appliance has the same optimization configurations enabled. Therefore, the site that initiates the flow (the client) determines the state of the protocol-specific optimization.

TCP Acceleration

TCP acceleration uses techniques such as selective acknowledgment, window scaling, and message segment size adjustment to compensate for poor performance on high latency links.

The TCP Accel Options dialog box includes a set of advanced options, provided with default settings.



CAUTION: Because changing the default settings can affect service, it is highly recommended that you contact Support for guidance before modifying them.

Option	Description
Adjust MSS to tunnel MTU	Limits the TCP MSS (Maximum Segment Size) advertised by the end hosts in the SYN segment to a value derived from the Tunnel MTU (Maximum Transmission Unit). This is *TCP MSS = Tunnel MTU – Tunnel Packet Overhead*. This feature is enabled by default so that the maximum value of the end host MSS is always coupled to the Tunnel MSS. If the end host MSS is smaller than the tunnel MSS, the end host MSS is used instead. A use case for disabling this feature is when the end host uses Jumbo frames.

Option	Description
Preserve packet boundaries	Preserves the packet boundaries end to end. If this feature is disabled, the appliances in the path can coalesce consecutive packets of a flow to use bandwidth more efficiently. This feature is enabled by default so that applications that require packet boundaries to match do not fail.
Enable TCP SYN option exchange	Controls whether the proprietary TCP SYN option is forwarded on the LAN side. Enabled by default, this feature detects whether more than two EdgeConnect appliances are in the flow's data path and optimizes accordingly. Disable this feature if there is a LAN-side firewall or a third-party appliance that would drop a SYN packet when it encounters an unfamiliar TCP option.
Route policy override	Attempts to override asymmetric route policy settings. It emulates auto-opt behavior by using the same tunnel for the returning SYN+ACK as it did for the original SYN packet. Disable this feature if the asymmetric route policy setting is necessary to correctly route packets. In this case, you might need to configure flow redirection to ensure optimization of TCP flows.
Auto reset flows	If enabled, this feature resets all TCP flows that are not accelerated but should be (based on policy and internal criteria like a Tunnel Up event). **NOTE:** Regardless of whether this feature is enabled, the default behavior when a tunnel goes down is to automatically reset the flows. The internal criteria can also include: Resetting all TCP accelerated flows on a Tunnel Down event. TCP acceleration is enabled. SYN packet was not seen, so this flow was either part of WCCP redirection or it already existed when the appliance was inserted in the data path.

Option	Description
IP block listing	If selected, and if the appliance does not receive a TCP SYN-ACK from the remote end within five seconds, the flow proceeds without acceleration and the destination IP address is blocked for one minute.
End to end FIN handling	Assists in fine tuning TCP behavior during a connection's graceful shutdown event. When this feature is enabled (default), TCP on the local appliance synchronizes this graceful shutdown of the local LAN side with the LAN side of the remote appliance. When not enabled (default TCP), synchronization does not occur and the two LAN segments at the ends gracefully and independently shut down.
WAN window scale	WAN-side TCP Window scale factor that is used internally for WAN-side traffic. This works independently of the WAN-side factor advertised by the end hosts.
Slow LAN defense	Resets all flows that consume a disproportionate amount of buffer and have very slow throughput on the LAN side. Because of a few slower end hosts or a lossy LAN, these flows affect the performance of all other flows so that no flows see the customary throughput improvement gained through TCP acceleration. This feature is enabled by default. The number relates indirectly to the amount of time the system waits before resetting slow flows.
WAN congestion control	This selects the internal Congestion Control parameter: **optimized:** This is the default setting. This mode offers optimized performance in almost all scenarios. **standard:** In some unique cases, it might be necessary to downgrade to standard performance to better inter-operate with other flows on the WAN link. **aggressive:** Provides aggressive performance. Use with caution. Recommended mostly for Data Replication scenarios.

Option	Description
Per-Flow buffer	Provides settings for **Max LAN to WAN buffer** and **Max WAN to LAN buffer**. These settings clamp the maximum buffer space that can be allocated to a flow in each direction.
Slow LAN window penalty	Penalizes flows that are slow to send data on the LAN side by artificially reducing their TCP receive window. This causes less data to be received and helps to reach a balance with the data sending rate on the LAN side. This setting is not selected by default.
LAN side window scale factor clamp	Allows the appliance to present an artificially lowered Window Scale Factor (WSF) to the end host. This reduces the need for memory in scenarios in which many out-of-order packets are being received from the LAN side. These out-of-order packets cause significant buffer utilization and maintenance.
Persist timer timeout	Allows TCP to terminate connections that are in the Persist timeout stage after the configured number of seconds.
Keep alive timer	Allows the keep alive timer for the TCP connections to change. **Probe interval:** Time interval in seconds between two consecutive keep alive probes. **Probe count:** Maximum number of keep alive probes to send. **First timeout (idle):** Time interval until the first keep alive timeout.

Access Lists Tab

Configuration > *Templates & Policies* > *Policies* > *ACLs* > *Access Lists*

This tab lists the configured **Access Control List** (ACL) rules. An **ACL** is a reusable MATCH criteria for filtering flows. It is associated with an action: **permit** or **deny**. An ACL can be a MATCH condition in more than one policy—Route, QoS, or Optimization.

Field	Description
Appliance	Name of the appliance selected.
ACLs	Access Control Lists. A list of one or more ordered access control rules. **NOTE:** An ACL only becomes active when it is used in a policy.
Priority	For ACL rules, you can set the priority to a value within the range 1 to 65535. When adding a rule, the priority is incremented by ten from the previous rule. You can change the priority, but this default behavior helps ensure that you can insert new rules without having to change subsequent priorities.
Match Criteria	Configured ACL match criteria associated to the appliance. See below for more information about Match Criteria.
Permit	Whether the ACL is set to **Permit** or **Deny**. **Permit** allows the matching traffic flow to proceed to the policy entry's associated SET actions. **Deny** prevents further processing of the flow by **that ACL, specifically**. The appliance continues to the next entry in the policy.
Comment	Any additional information about the ACL.

Click the edit icon to make add, delete, or modify rules to your ACLs.

Match Criteria

Match criteria are used universally across all policy maps for route, QoS, optimization, SaaS NAT, and firewall zone security policies.

To use the same match criteria in different maps, create an ACL (Access Control List), which is a named, reusable set of rules. For efficiency, navigate to **Configuration > Templates & Policies > ACLs > Access Lists** to create ACLs, and then apply them across your appliances.

Use the Match Criteria dialog box to select and configure a variety of match criteria options. Some are explained below:

- **Address Map:** Use this option to sort by country, IP address owner, or SaaS application. You can also select and configure Microsoft Instance, Microsoft Category, and Proxy attributes for an address map. These attributes are secondary parameters to the address map. They are evaluated for a policy match only when the configured address map matches the flow. To select and configure these attributes, click **+Attributes**.
- Match criteria options related to Secure Web Services include:

  - **URL:** Omit the protocol in URLs and include a slash character (/) or a slash followed by a query parameter. For example, *google.com/* or *google.com/maps* are valid; *https://google.com/* is not valid. You can also use the asterisk (*) wildcard character, such as in *google.com/**, to specify a domain, but it is more appropriate to use the Domain match criteria option to specify *google.com/*. Separate multiple URL addresses with the pipe character (|). For example, you can specify two URL addresses as *bing.com/*|google.com/**.
  - **Web Category:** If you select this option, click **Select**, and then click **Strict**, **Moderate**, or **Custom**. Strict and Moderate include predetermined web category selections, which you cannot change. However, if you attempt to select or clear a web category for these, the Custom group of web category selections opens automatically. You can customize web category selections in the Custom group. After you make your selections, click **Ok**.
  - **Web Reputation:** If you select this option, select one or more web reputation categories (High Risk, Suspicious, Moderate Risk, Low Risk, or Trustworthy) from the list. To select more than one, press and hold the **Shift** key, and then click the appropriate categories.
  - **Bad IP Reputation:** Selecting this option matches on flows for which the source or destination IP address (or both) has a High Risk reputation score.

- Match criteria options related to RADIUS user role include:

  - **User Role:** This is the user role as specified in the authentication exchange with the ClearPass RADIUS server.
  - **User Name**
  - **User Group**
  - **User Device**
  - **User MAC**
  - **User Vlan**

  Configuring these match criteria related to user role enables an EdgeConnect to automatically assign traffic steering and firewall zone policies.

- Use the Src:Dest check box associated with several match criteria options to specify separate criteria for inbound and outbound traffic. You can configure source and destination role-based policies when both source and destination users are in the same network.

Wildcard-based Prefix Matching Rules

- Even when using a range or a wildcard, the IPv4 address must be specified in the 4-octet format, separated by the dot notation. For example, **A.B.C.D**.
- Range is specified using a single dash. For example, **128-129**.
- Wildcard is specified as an asterisk (___*).
- Range and Wildcard can both be used in the same address, but an octet can only contain one or the other. For example, **10.136-137.*.64-95**.

- A wildcard can only be used to define an entire octet. For example, **10.13*.*.64-95** is not supported. Use **10.130-139.*.64-95** to specify this range.
- The same rules apply to IPv6 addressing.
- CIDR notation and (Range or Wildcard) are mutually exclusive in the same address. For example, **192.168.0.1-127/24** is not supported. Use either **192.168.0.0/24** or **192.168.0.1-127**.
- These prefix-matching rules apply to the following policies only: Route, QoS, Optimization, NAT, Security, and ACLs.

Access Lists Edit Row

The Access Lists dialog box lists the configured **Access Control List** (ACL) rules.

You can add, delete, or rename an ACL by clicking the buttons at the top of this dialog box. You can also add rules to an ACL.

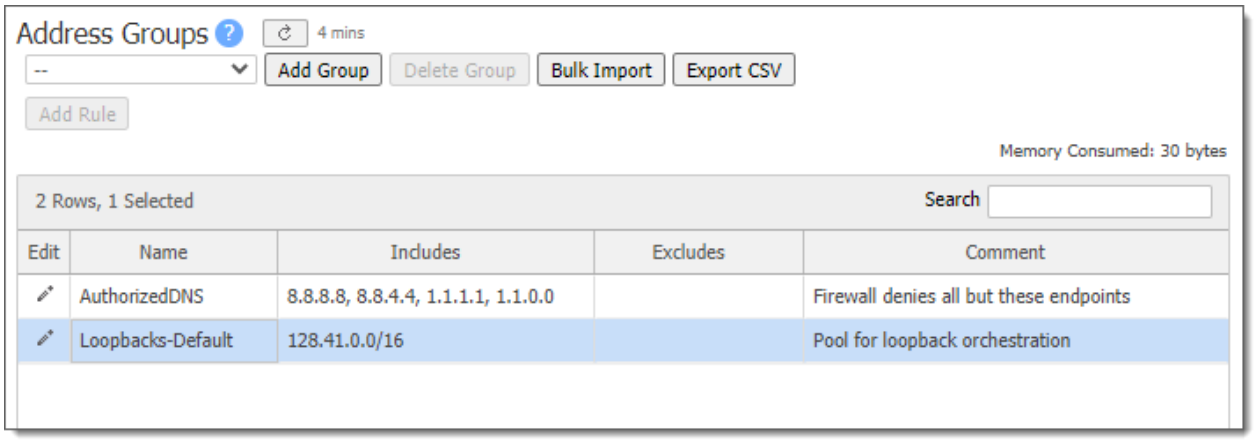
1. Click **Add Rule**.
2. Enter a priority value.
3. Click the edit icon to configure the match criteria. The **Match Criteria** dialog box opens and you can specify the match criteria. Click **More Options** to apply more rules.
4. Select if you want to **Permit** or **Deny** traffic in the ACL.
5. Enter any comments if you decide to do so.

Address Groups

Configuration > Templates & Policies > ACLs > Address Groups

Use the Address Groups tab to view and manage address groups in your SD-WAN network. An address group is a logical collection of IP hosts or subnets that can be referenced in source or destination matching criteria in the zone based firewall and security policies (route, QOS, optimization, and so forth). You can specify both IPs to include and IPs to exclude within a group. Groups may also be nested, meaning one address group can reference another.

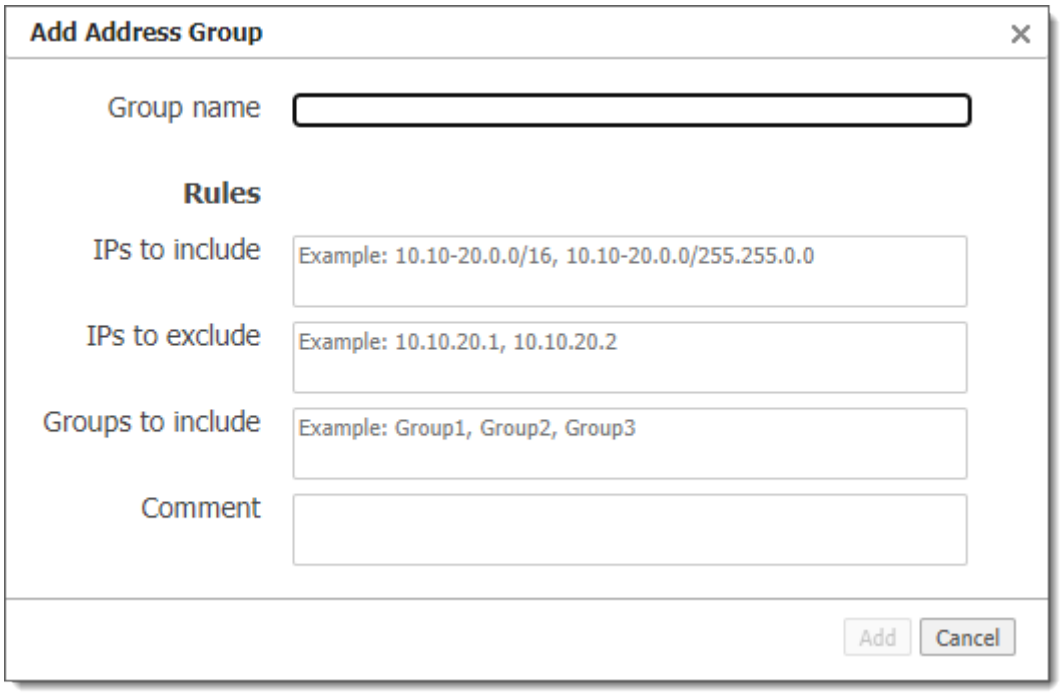
NOTE: Orchestrator supports up to 8MB of address group definitions. For current usage, check the Address Groups UI.



Add an Address Group

Follow the steps below to create a new address group:

1. Click **Add Group** to open the Add Address Group dialog box.



2. Provide the following details in the fields provided:

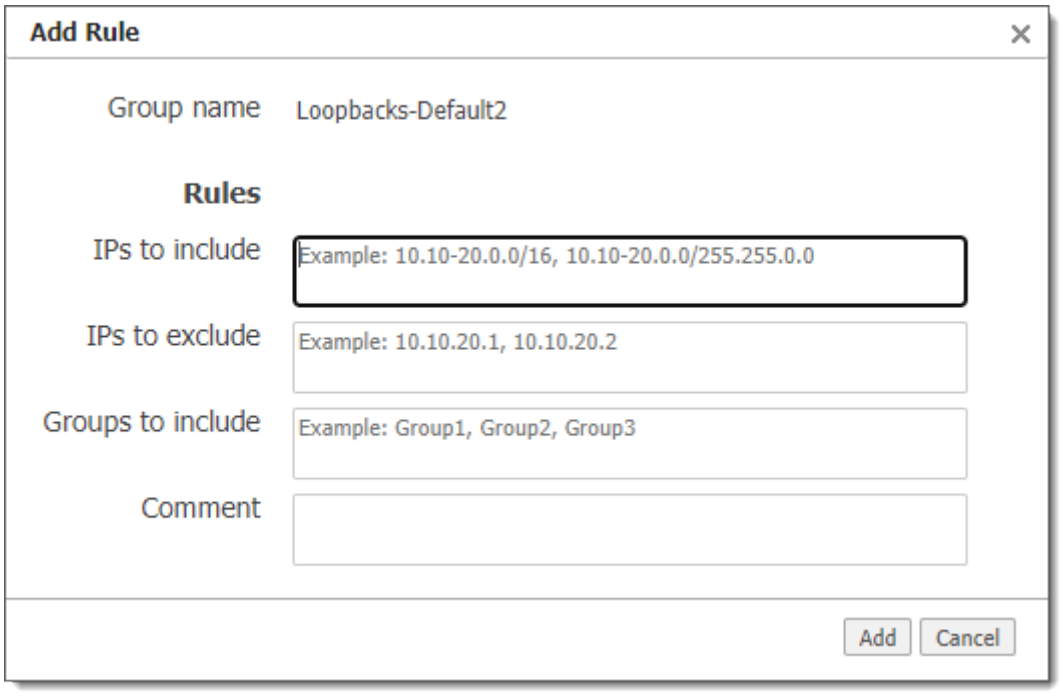
Field	Description
Group name	Enter a unique name for the group, up to 64 characters long. **NOTE:** Group names can only contain uppercase and lowercase letters, numbers, dots, underscores, and hyphens.
IPs to include	Enter one or more IP addresses or subnets to include in the group (see **Address Group Formats** below).
IPs to exclude	Enter one or more IP addresses to exclude, in the case where you are including an IP range.
Groups to include	Enter the name of one or more address groups to include. **NOTE:** Group inclusion only supports two levels of nesting. For example, if Group1 includes Group2 and Group2 includes Group3, you could not include Group1 anywhere because it already contains two levels of nested groups.
Comment	Enter an optional comment that describes the address group and how it might be used.

3. Click **Add** to create the address group, or click **Cancel** to close the dialog box without making any changes.

Add a Rule to an Address Group

Follow the steps below to add a rule to an existing address group:

1. Select the address group to which you want to add a rule from the drop-down list above the table.
2. Click **Add Rule** to open the Add Rule dialog box.



3. Provide the details for the new rule in the fields provided (see field descriptions in **Add an Address Group**).
4. Click **Add** to create the rule or click **Cancel** to close the dialog box without making any changes.

Delete an Address Group

Follow the steps below to delete an address group:

1. Select the address group you want to delete from the drop-down list above the table.
2. Click **Delete Group**.

   A confirmation dialog box opens.
3. Click **Delete** to confirm your choice and permanently remove the selected group and all of its rules. Otherwise, click **Cancel** to return to the list without deleting the group.

Export Address Groups

You can export the current address groups to a CSV file as a backup to make bulk modifications outside of the Orchestrator UI.

To export address groups:

1. Click **Export CSV**.
2. In the save dialog box, browse to the location where you want to save the file, provide a name for the file, and then click **Save**.

3. Open the saved file in Excel or another program to view or modify its contents.

	A	B	C	D	E
1	Name	IncludedIPs	ExcludedIPs	IncludedGroups	Comment
2	AuthorizedDNS	8.8.8.8,8.8.4.4,1.1.1.1,1.1.0.0			Firewall denies all but these endpoints
3	Loopbacks-Default	128.41.0.0/16			Pool for loopback orchestration
4					

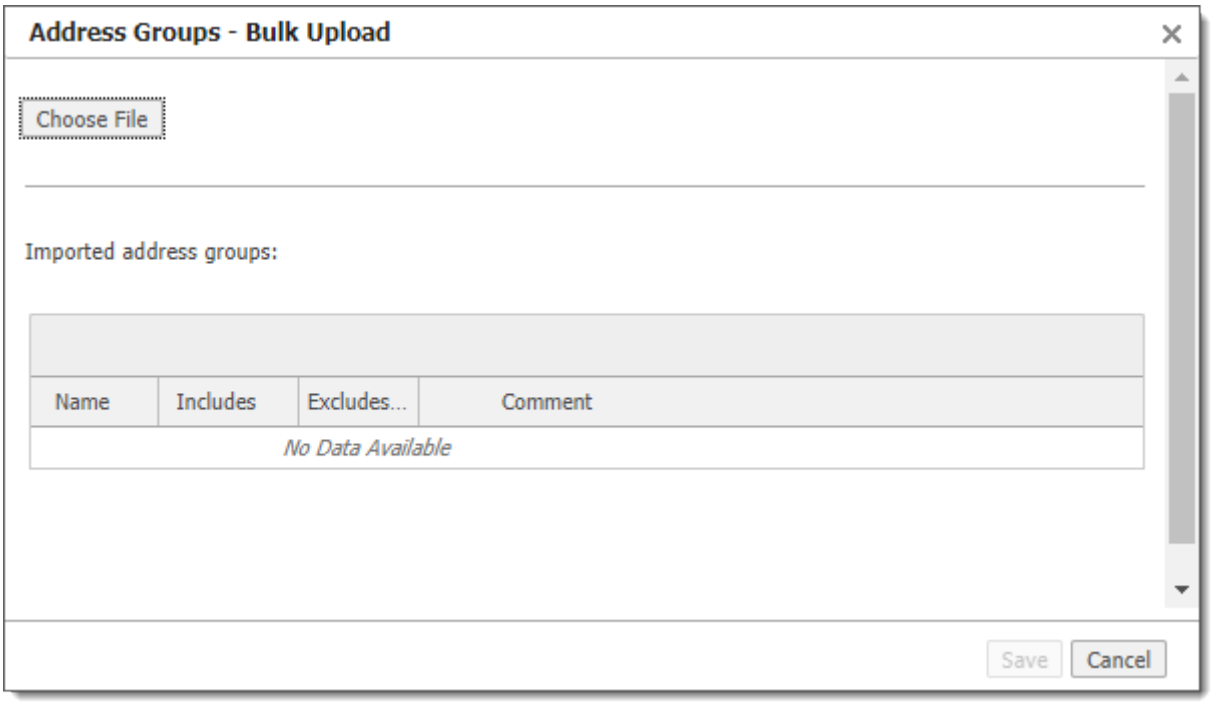
NOTE: When editing exported rules and address groups, you can modify the included or excluded IPs, included groups, or comments to overwrite the same rule when imported. If you modify the group name on a rule, however, it will create a new rule when imported.

Import Address Groups

To import address groups from a CSV file:

NOTE: You can import a file that was exported and modified, or a new file that contains data in the same rows and columns as the exported file. Columns are ordered as Name, Included IPs, Excluded IPs, Included Groups, and Comment. The first row of the import file will be ignored.

1. Click **Bulk Import** to open the Address Groups - Bulk Upload dialog box.



2. Click **Choose File**, locate and select the CSV file to be imported, and then click **Open**.
3. Review the groups and rules to be imported.
4. Click **Save** to import the file and merge with or replace the existing address groups, or click **Cancel** to close the dialog box without making any changes.

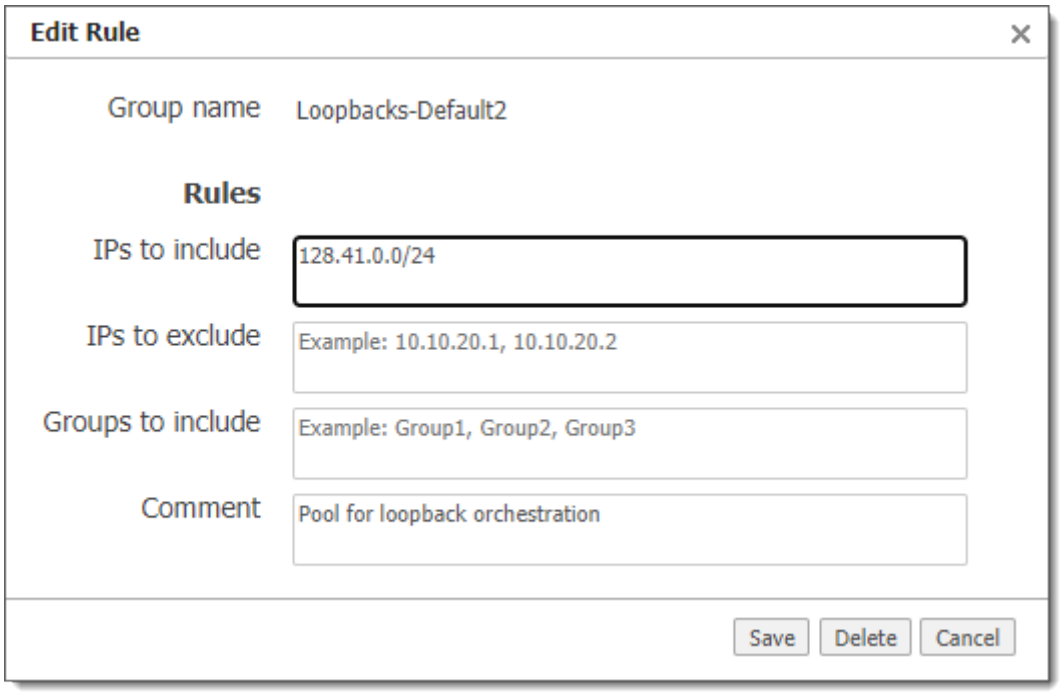
View a Single Address Group

By default, all address groups are displayed in the table on the Address Groups tab. To filter the table to a single address group, select the group from the drop-down list above the table.

NOTE: You can only add rules to an existing group when viewing a single address group. You cannot add a group with the same name as an existing group.

Edit or Delete a Rule

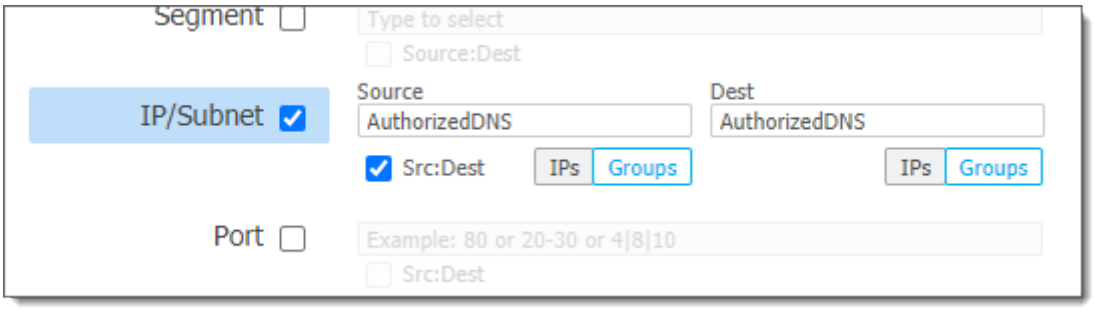
To edit or delete an existing rule, click the edit icon to the right of the rule. The Edit Rule dialog box opens.



- To edit the rule, modify the available fields, and then click **Save**.
- To delete the rule, click **Delete**.

Using Address Groups in Match Criteria

When specifying match criteria for IP/Subnet, you can use an address group by enabling the **Src:Dest** and **Groups** options.



Requirement and Behavior of Excluded IPs in Nested Groups

When you nest address groups, only the included IPs from the referenced (nested) group are considered for policy matching. Any IPs specified in the “exclude” field of the nested group are not evaluated or inherited by the parent group.

- Excluded IPs are local to the group in which they are defined. If you create Group A and exclude certain IPs, and then nest Group A inside Group B, Group B will only inherit the included IPs from Group A. The excluded IPs from Group A do not propagate to Group B and are not considered when Group B is used in policy match criteria.
- This behavior prevents unintended exclusions and maintains predictable, explicit group definitions. Each group’s exclusions apply only when that group is referenced directly in a policy or configuration.
- Observe the following best practice. If you need to exclude specific IPs across multiple nested groups, you must explicitly exclude those IPs in each parent group where exclusion is required.

Example

Suppose you have:

- **Group A:**

Include: 10.1.1.0/24

Exclude: 10.1.1.10

- **Group B:**

Include: Group A, 10.2.2.0/24

Exclude: 10.2.2.20

When Group B is used in a policy, the excluded IP 10.1.1.10 from Group A is not automatically excluded by Group B. To exclude 10.1.1.10 when using Group B, you must add it to Group B’s exclude list.

Troubleshooting and Validation - If you observe that excluded IPs from nested groups are not being excluded in policy matches, this is expected behavior.

- To enforce exclusions at all levels, review and update the exclude lists in each group as needed.

Address Group Formats

An address group can include IP addresses, subnets, address groups, or any combination thereof. For IPs and subnets, the following formats are allowed:

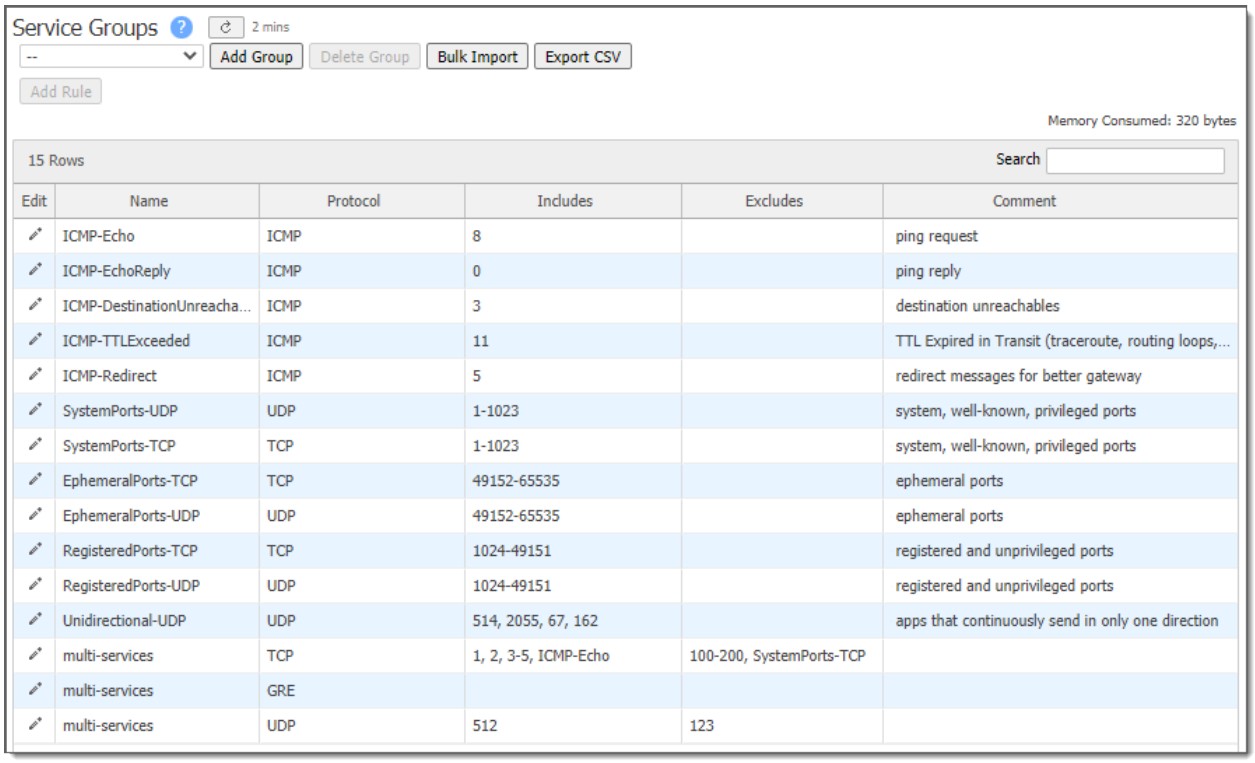
- One or more IP addresses: 10.10.10.1 or 10.10.10.2, 10.10.10.2, 10.10.10.3
- IP subnet: 10.10.0.0/16 or 10.10.0.0/255.255.0.0
- IP range: 10.10.10.10-20
- IP range and subnet: 10.10-20.0.0/16, 10.10-20.0.0/255.255.0.0
- IP wildcard: 10.10.10.* (you can use the wildcard in any octet)
- Wildcard and mask: 10.*.0.0/16, 10.*.0.0/255.255.0.0

Service Groups

Configuration > Templates & Policies > ACLs > Service Groups

Use the Service Groups tab to view and manage service groups in your SD-WAN network. A service group is a logical collection of protocols and ports that can be referenced in source or destination matching criteria in the zone based firewall and security policies (route, QOS, optimization, and so forth).

NOTE: Orchestrator supports up to 4MB of service group definitions. For current usage, check the Service Groups UI.

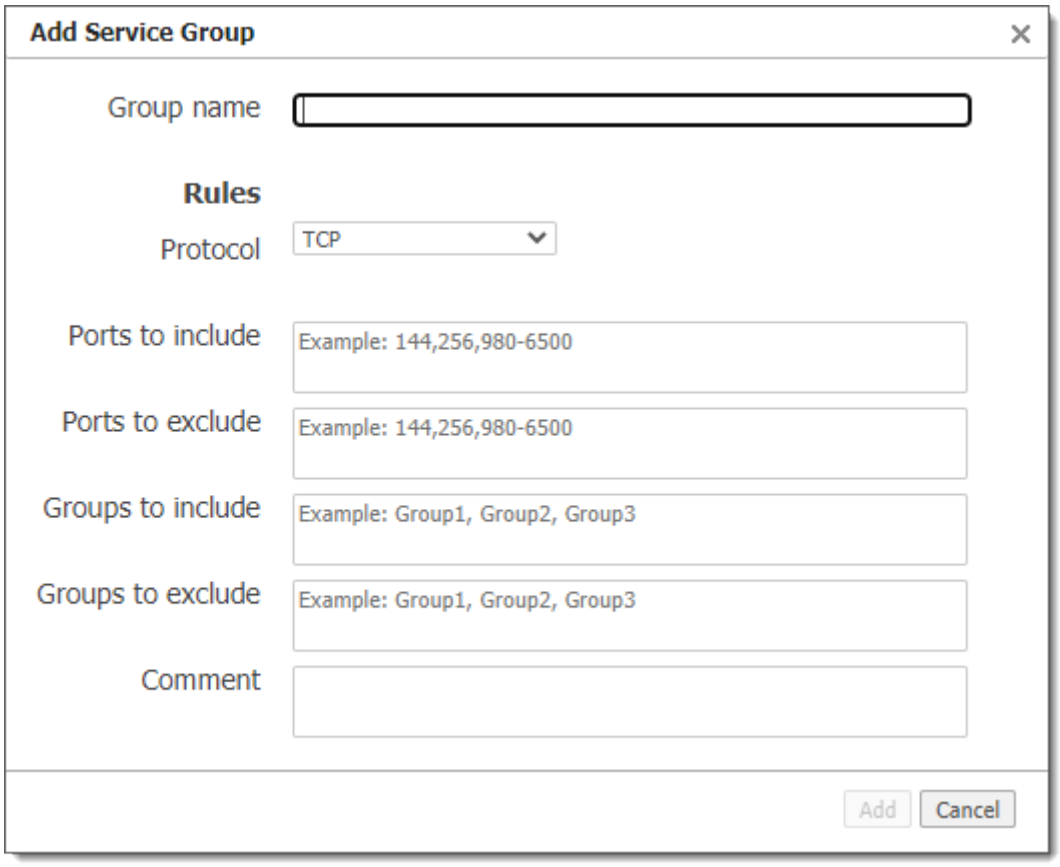


Edit	Name	Protocol	Includes	Excludes	Comment
	ICMP-Echo	ICMP	8		ping request
	ICMP-EchoReply	ICMP	0		ping reply
	ICMP-DestinationUnreacha...	ICMP	3		destination unreachables
	ICMP-TTLExceeded	ICMP	11		TTL Expired in Transit (traceroute, routing loops,...
	ICMP-Redirect	ICMP	5		redirect messages for better gateway
	SystemPorts-UDP	UDP	1-1023		system, well-known, privileged ports
	SystemPorts-TCP	TCP	1-1023		system, well-known, privileged ports
	EphemeralPorts-TCP	TCP	49152-65535		ephemeral ports
	EphemeralPorts-UDP	UDP	49152-65535		ephemeral ports
	RegisteredPorts-TCP	TCP	1024-49151		registered and unprivileged ports
	RegisteredPorts-UDP	UDP	1024-49151		registered and unprivileged ports
	Unidirectional-UDP	UDP	514, 2055, 67, 162		apps that continuously send in only one direction
	multi-services	TCP	1, 2, 3-5, ICMP-Echo	100-200, SystemPorts-TCP	
	multi-services	GRE			
	multi-services	UDP	512	123	

Add a Service Group

Follow the steps below to create a new service group:

1. Click **Add Group**. The Add Service Group dialog box opens.



2. Provide the following details in the fields provided:

Field	Used in	Description
Group name	All	Enter a unique name for the group, up to 64 characters long. **NOTE:** Group names can only contain uppercase and lowercase letters, numbers, dots, underscores, and hyphens.
Protocol	All	Select a protocol from the list of those available.
Ports to include	TCP, UDP	Enter one or more ports to include in the group. A single port, multiple comma-separated ports, and a range of ports are supported (e.g., 20, 22, 24-30).

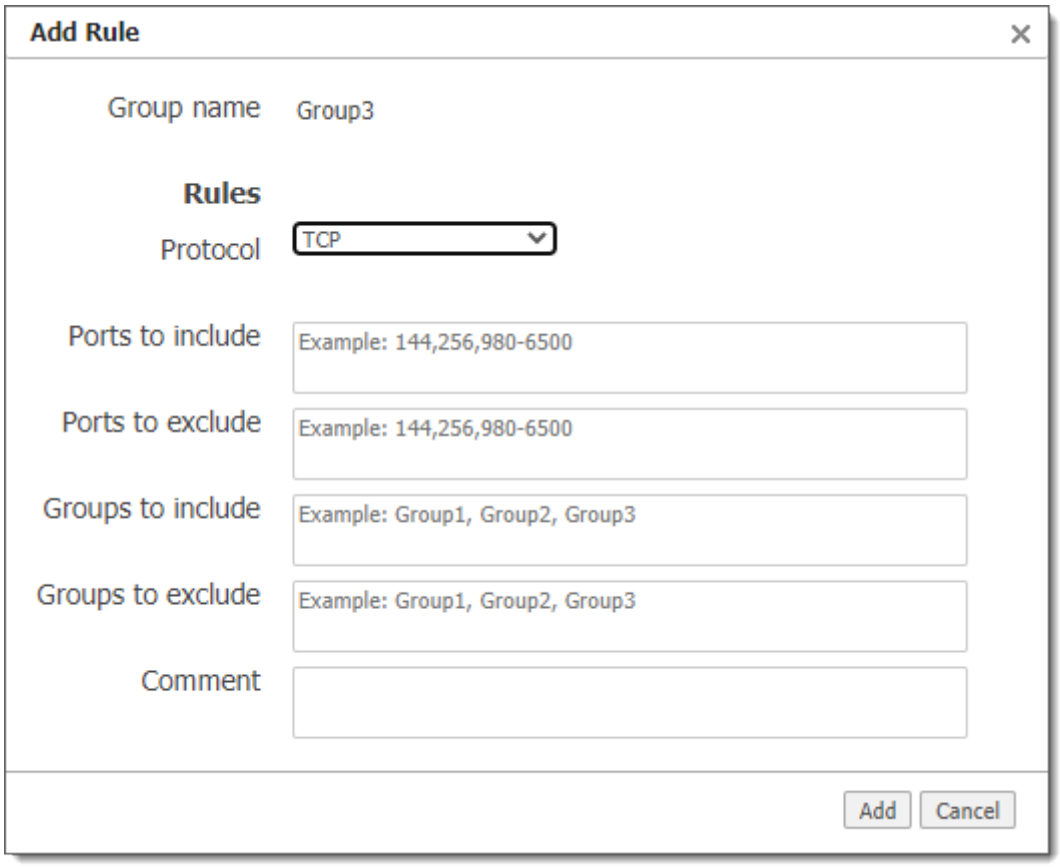
Field	Used in	Description
Ports to exclude	TCP, UDP	Enter one or more ports to exclude from the group, in the case where you are including a range of ports. A single port, multiple comma-separated ports, and a range of ports are supported (e.g., 20, 22, 24-30).
Groups to include	TCP, UDP	Enter the name of one or more service groups to include. **NOTE:** Group inclusion only supports two levels of nesting. For example, if Group1 includes Group2 and Group2 includes Group3, you could not include Group1 anywhere because it already contains two levels of nested groups.
Groups to exclude	TCP, UDP	Enter the name of one or more service groups to exclude, in the case where you are already including a group that includes multiple groups.
ICMP types	ICMP	For ICMP, add one or more message types to include. Multiple types and ranges are supported (e.g., 1, 2, 4-8).
Comment	All	Enter an optional comment that describes the service group and how it might be used.

3. Click **Add** to create the service group or click **Cancel** to close the dialog box without making any changes.

Add a Rule to a Service Group

Follow the steps below to add a rule to an existing service group:

1. Select the service group to which you want to add a rule from the drop-down list above the table.
2. Click **Add Rule**. The Add Rule dialog box opens.



3. Provide the details for the new rule in the fields provided (see field descriptions in **Add a Service Group**).
4. Click **Add** to create the rule or click **Cancel** to close the dialog box without making any changes.

Delete a Service Group

Follow the steps below to delete a service group:

1. Select the service group you want to delete from the drop-down list above the table.
2. Click **Delete Group**.

   A confirmation dialog box opens.
3. Click **Delete** to confirm your choice and permanently remove the selected group and all of its rules. Otherwise, click **Cancel** to return to the list without deleting the group.

Export Service Groups

You can export the current service groups to a CSV file as a backup to make bulk modifications outside of the Orchestrator UI. Follow the steps below to export service groups.

1. Click **Export CSV**.

2. In the save dialog box, browse to the location where you want to save the file, provide a name for the file, and then click **Save**.
3. Open the saved file in Excel or another program to view or modify its contents.

	A	B	C	D	E	F	G	H
1	Name	Protocol	IncludedPorts	ExcludedPorts	IncludedGroups	ExcludedGroups	IcmpTypes	Comment
2	ICMP-Echo	ICMP					8	ping request
3	ICMP-EchoReply	ICMP					0	ping reply
4	SystemPorts-UDP	UDP	1-1023					system, well-known, privileged ports
5	SystemPorts-TCP	TCP	1-1023					system, well-known, privileged ports
6	EphemeralPorts-TCP	TCP	49152-65535					ephemeral ports
7	EphemeralPorts-UDP	UDP	49152-65535					ephemeral ports
8	RegisteredPorts-TCP	TCP	1024-49151					registered and unprivileged ports
9	RegisteredPorts-UDP	UDP	1024-49151					registered and unprivileged ports
10	Unidirectional-UDP	UDP	514,2055,67,162					apps that continuously send in only one direction

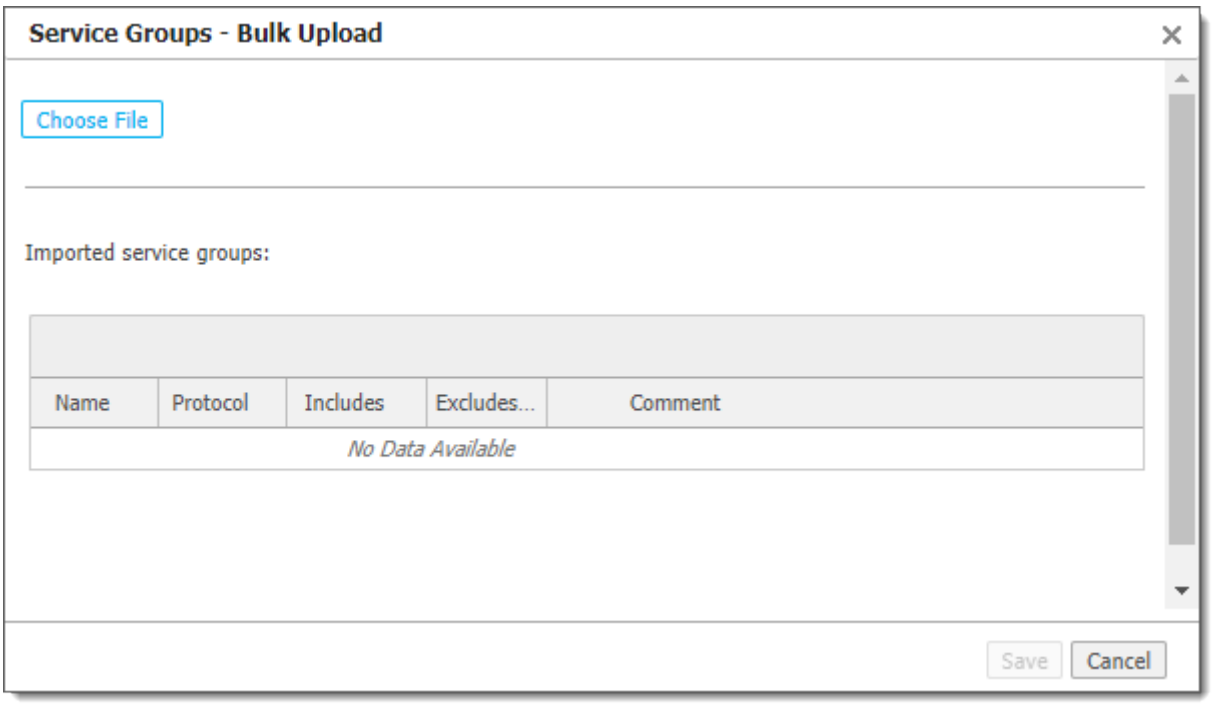
NOTE: When editing exported rules and service groups, you can modify the protocol, inclusions, exclusions, ICMP types, or comments to overwrite the same rule when imported. If you modify the group name on a rule, however, it will create a new rule when imported.

Import Service Groups

Follow the steps below to import service groups from a CSV file:

NOTE: You can import a file that was exported and modified, or a new file that contains data in the same rows and columns as the exported file. Columns are ordered as Name, Protocol, Included Ports, Excluded Ports, Included Groups, Excluded Groups, ICMP types, and Comment. The first row of the import file will be ignored.

1. Click **Bulk Import**. The Service Groups - Bulk Upload dialog box opens.



2. Click **Choose File**, locate and select the CSV file to be imported, and then click **Open**.
3. Review the groups and rules to be imported.
4. Click **Save** to import the file and merge with or replace the existing service groups, or click **Cancel** to close the dialog box without making any changes.

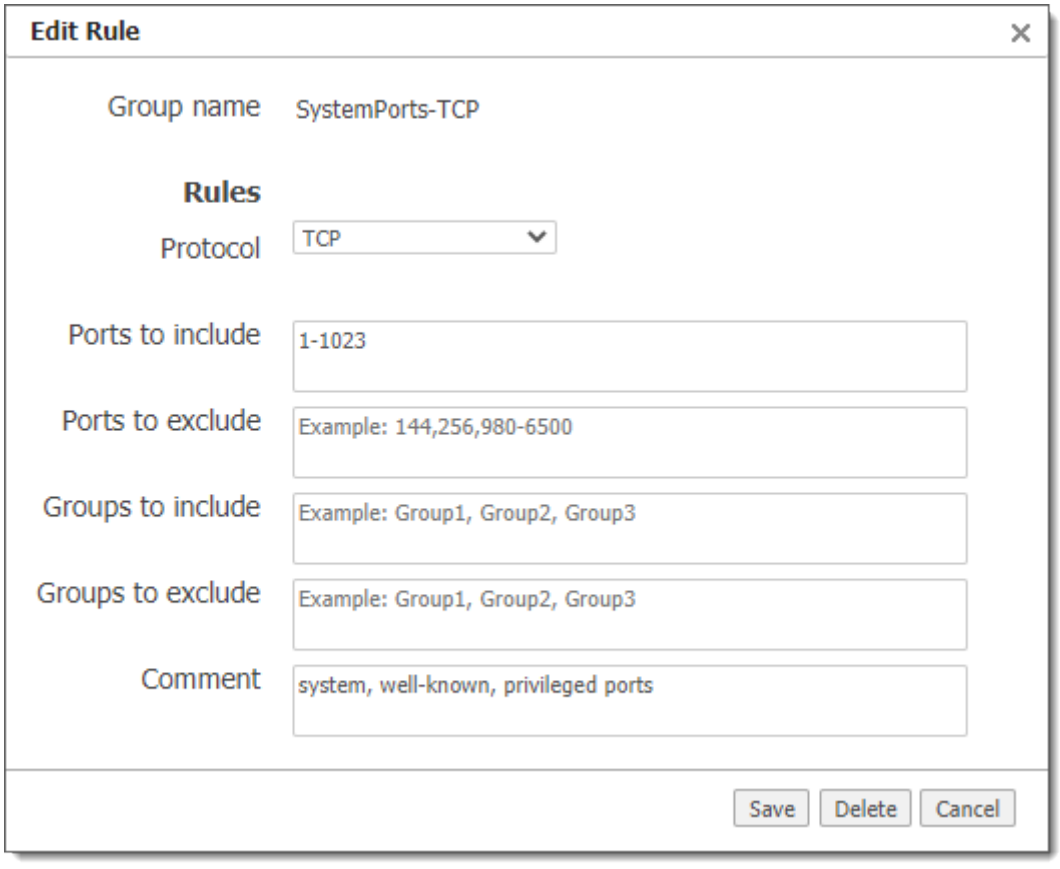
View a Single Service Group

By default, all service groups are displayed in the table on the Service Groups tab. To filter the table to a single service group, select the group from the drop-down list above the table.

NOTE: You can only add rules to an existing group when viewing a single service group. You cannot add a group with the same name as an existing group.

Edit or Delete a Rule

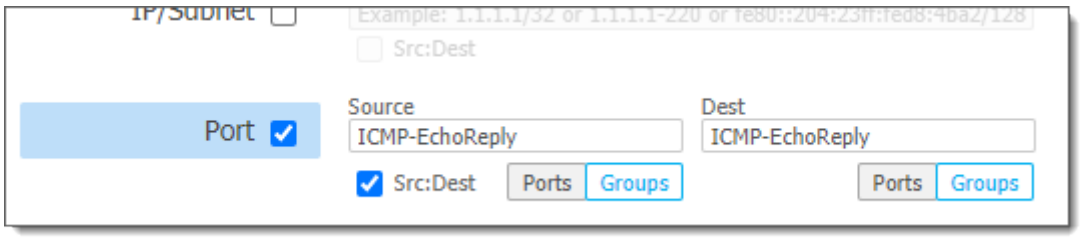
To edit or delete an existing rule, click the edit icon to the right of the rule and the Edit Rule dialog box opens.



- To edit the rule, modify the available fields, and then click **Save**.
- To delete the rule, click **Delete**.

Using Service Groups in Match Criteria

When specifying match criteria for Port, you can use a service group by enabling the **Src:Dest** and **Groups** options.



Shaper Tab

Configuration > Templates & Policies > Shaping > Shaper

The Shaper tab provides a simplified way to globally configure QoS (Quality of Service) on the appliances.

- Outbound Shaping is always enabled.
- Inbound Shaping is disabled by default and can be enabled by clicking the edit icon for a specific interface.

Traffic shaping allocates bandwidth as a percentage of system bandwidth. Shaper parameters are organized into ten traffic classes. Four traffic classes are preconfigured (RealTime, CriticalApps, BulkApps, and Default). After compressing (deduplicating) all the outbound tunnelized and passthrough-shaped traffic, the system either applies policy settings globally or upon each interface, shaping traffic as it exits the interface.

To manage Shaper settings for an appliance's system-level WAN Shaper, access the Shaper template. For minimum and maximum bandwidth, you can configure traffic class values as a percentage of total available system bandwidth and as an absolute value. The appliance always provides the larger of the minimum values and limits bandwidth to the lower of the maximum values.

If you set Min Bandwidth to a value greater than Max Bandwidth, Max overrides Min.

Shaper Tab Settings

Field	Description
Appliance	Name of the appliance.

Field	Description
Interface Shaper	Enables a separate shaper for a specific WAN interface. For WAN optimization, the interface shaper can be used, but it is not recommended. For SD-WAN, it should never be used because overlay traffic is not directed to an interface shaper; traffic is always shaped by the default WAN shaper.
Max WAN BW	This limits the maximum bandwidth that a traffic class can use for WAN traffic.
Recalc on IF State Changes	When an interface state changes to UP or DOWN, selecting this recalculates the total bandwidth based on the configured bandwidth of all UP interfaces. For example, when wan0 goes down, wan0 bandwidth is removed from the total bandwidth when recalculating.
Traffic ID	The number assigned to the traffic class.
Traffic Name	The name assigned to a traffic class, either prescriptively or by the user.
Priority	Determines the order in which to allocate each class's minimum bandwidth. 1 is first, 10 is last.
Min BW %	Refers to the percentage of bandwidth guaranteed to each traffic class, allocated by priority. However, if the sum of the percentages is greater than 100%, lower-priority traffic classes might not receive their guaranteed bandwidth if it is all consumed by higher-priority traffic. If you set Min Bandwidth to a value greater than Max Bandwidth, Max overrides Min.
Min BW Absolute (kbps)	This guarantees a specific level of service when total system bandwidth declines. This is useful for maintaining the quality of VoIP, for example.
Min BW Actual (kbps)	This specifies the actual minimum level of service when total system bandwidth declines. Max BW is equal to the largest of the following values: Licensed BW, system BW, or the sum of the BW of the configured interfaces.
Excess Weighting	If there is bandwidth left over after satisfying the minimum bandwidth percentages, the excess is distributed among the traffic classes in proportion to the weightings specified in the Excess Weighting column. Values range from 1 to 10,000.
Max BW %	This limits the maximum bandwidth that a traffic class can use to a percentage of total available system bandwidth.

Field	Description
Max BW Absolute (kbps)	This limits the maximum bandwidth that a traffic class can use to an absolute value (kbps). You can specify a maximum absolute value to cap the bandwidth for downloads and streaming.
Max BW Actual (kbps)	This specifies the actual maximum level of service when total system bandwidth declines. Max BW is equal to the smallest of the following values: Licensed BW, system BW, or the sum of the BW of the configured interfaces.
Max Wait Time (ms)	Any packets waiting longer than the specified Max Wait Time are dropped.
Rate Limit (kbps)	You can set per-flow rate limit that a traffic class uses by specifying a number in the Rate Limit column. For no limit, use 0 (zero).
Enable	Specifies the shaping enabled status. Shaping is always enabled for outbound traffic and by configuration for inbound traffic.
Per Interface	Indicates if shaping is for each interface or for the gateway as a unit.

Application Definitions Tab

Configuration > Templates & Policies > Applications & SaaS > Application Definitions

This tab provides application visibility and control.

- You can search to see if a definition exists for an application and, if so, how it is defined.
- You can set AppExpress performance monitoring for up to 50 specific applications.
- To filter the application definition list to show only monitored applications, click the **Monitored** button in the tab header. Click **All** to remove the filter.

Orchestrator uses the following eight "application pipelines" to assign an application name to each flow within the SD-WAN fabric:

- **IP Protocol** – Matches flows based on IP protocol number. Matches on the first packet.
- **UDP Port** – Matches flows based on layer 4 UDP port. Matches on the first packet.
- **TCP Port** – Matches flows based on layer 4 TCP port. Matches on the first packet.
- **Domain Name** – Uses DNS, HTTP Host Header, and HTTPS/SNI snooping to derive a domain name for each flow. It is expressed as *example.com* or **.example.com*. The DNS snooping mechanism matches on the first packet. The snooped DNS queries must be cleartext and must transit the EdgeConnect in both directions.

NOTE: Multiple domain names can be associated with the same IP address. When this happens, the snooping mechanism needs up to six packets to determine the domain name and cannot match on the first packet. However, when a policy specifies a domain name, the system gives precedence to it when associating the flow with a policy. To more deterministically ensure that a domain name you specify is given precedence, use the Application Definition confidence level to enable the system to match the domain name to a flow according to your intention on the first packet received back from the DNS server.

- **Address Map** – Formerly known as *IP Intelligence*. Given a range of IP addresses, the Address Map reveals the organization that owns the segment, along with the country of origin. Matches on the first packet.
- **DPI** – Deep Packet Inspection. These applications are derived by looking into the packet payload. Examples include RTP, FTP, and HTTP. The DPI pipeline requires multiple packets to read the required payload elements and cannot match on the first packet.
- **Compound** – Created using multiple application match criteria. First-packet matching varies based on the configuration of the Compound application.
- **SaaS** – Deprecated. For use with legacy SaaS Optimization system. AppExpress should be used for SaaS optimization use cases.

You can use any of these application pipelines to define a new application, and you can modify or disable an existing application. Multiple application definitions can match at the same time. When this occurs the application with the highest confidence configured (1-100) is used.

Orchestrator automatically checks the Cloud Portal for updated application definitions every 24 hours by default (Auto update set to ON). Application definition data on the Cloud Portal is updated generally once per month. If new definitions are discovered, Orchestrator downloads the data, merges it with the applications, and pushes the changes to appliances in the network. You can also force an update at any time by clicking **Update Now**.

Application Definition Dialog Box

From this dialog box you can add, edit, disable, or delete an application definition and enable AppExpress for an application.

Add an Application Definition

Complete the following steps to add an application definition.

1. Navigate to **Configuration > Templates & Policies > Applications & SaaS > Application Definitions**.
2. Click **Show Advanced App Definitions**.
3. Click **+Add New Application**.

   The Application Definition dialog box opens.
4. From the Type drop-down menu, select an application pipeline.

5. Complete the fields that apply to the application pipeline you selected.

   For a description of these fields, refer to the table at the end of this procedure.

6. To enable AppExpress for the application, continue to Enable AppExpress for an Application.
7. Click **Apply**.

   The definition appears in the Advanced App Definitions section.

   NOTE: To find a user-created definition in the Advanced App Definitions section, click the appropriate tab for the type of definition, such as IP Protocol or UDP Port, and then click **Modified**.

   Application Definition Fields

Field	Description
Name	Enter a name for the application. This application name is used throughout the EdgeConnect system to match and apply various policies. Application names are not case-sensitive. **NOTE:** When you change an application name, you must also change it in any associated policies, such as Overlay ACLs or Firewall policies.
Protocol Number	Applies to IP Protocol. Enter the protocol number for the application.
Port Number	Applies to UDP Port and TCP Port. Enter the port number for the application.
Domain	Applies to UDP Port, TCP Port, and Compound. Enter the domain for the application.
IP range(IPV4 only)	Applies to Address Map. Enter the range of IP addresses that are included.
Organization	Applies to Address Map. Enter the name of the organization that owns the range of IP addresses.
Country	Applies to Address Map. Select the country where the organization that owns the range of IP addresses resides.
Protocol	Applies to Compound. Select the type of protocol used for the application.
Port	Applies to Compound. Enter the port number for the application.
IP/Subnet	Applies to Compound. Enter the IP address or subnet for the application.
Geo Location	Applies to Compound. Matches flows with IPs associated with a specific country. Select the country from the drop-down list.
Address Map	Applies to Compound. Matches flows with IPs contained within an Address Map. Select the address map from the drop-down list.

Field	Description
Interface	Applies to Compound. Matches flows that are inbound to the EdgeConnect through the specified interface or label. Select the interface or label from the drop-down list.
DSCP	Applies to Compound. Click the check box to match the first DSCP value observed for the flow.
Domains	Deprecated. Applies to SaaS Optimization.
Addresses	Deprecated. Applies to SaaS Optimization.
Notes	Applies to all except SaaS. This is a text-entry field where you can enter any notes or information about the application definition.
Confidence	Used when two or more application definitions match the same flow. The application with the highest Confidence value is assigned to the flow. Enter a value of 1 to 100. The higher the number, the higher the confidence.
Microsoft Instance	Applies to Address Map. Allows filtering by Microsoft Instance type, such as “WorldWide”, “USGovDoD”, and “China”. Select the instance from the drop-down list.
Microsoft Category	Applies to Address Map. Matches the Microsoft-assigned endpoint category, which includes the following (select one): **Optimize** – High-priority traffic that should get priority QoS treatment, take the most optimal path, and bypass security inspection mechanisms. **Allow** – Lower-priority traffic that should bypass security inspection mechanisms. **Default** – Traffic that should be treated as “regular internet traffic”.
Proxy	Applies to Address Map. If this attribute is set, the EdgeConnect does not learn domain names for the given IP(s). Select **No** or **Yes** from the drop-down menu.
Disabled	Applies to all except Address Map and SaaS. Click this check box to disable the application definition. This action does not delete the application definition.

Enable AppExpress for an Application

You can enable AppExpress performance monitoring for any application. If you choose to have AppExpress steer traffic for an application, you must enable steering and add the application to an AppExpress group. For more information, see AppExpress Groups Tab.

1. Navigate to **Configuration > Templates & Policies > Applications & SaaS > Application Definitions**.

2. Click **Show Advanced App Definitions**.
3. Locate the application in the advanced definitions list, and then click the edit icon.

   The Application Definition dialog box opens.
4. Enter the following information based on the level of monitoring you want to apply to the application:

Field	Description
AppExpress	**Off** – Click to disable AppExpress for the application. **Monitor only** – Click to enable only AppExpress monitoring for the application. Data is collected about the performance of the application, which is reflected on the AppExpress Summary tab and in reports, but the data is not used to steer the application traffic. **Monitor and Steer** – Click to enable AppExpress monitoring and steering for the application. Data is collected about the performance of the application, which the system uses to steer the application traffic from one transport to another. The application must be added to an AppExpress group for traffic steering to occur. For more information, see AppExpress Groups Tab.
Use Cloud Portal Config	Select this check box to use the HPE Aruba Networking-provided AppExpress settings for this application. **NOTE:** Not all applications have HPE Aruba Networking-provided AppExpress settings. You can override these settings at any time.
Ping Type	Select the method used to send probes to the application across the loopback interface. Options include **ICMP echo-request/response**, **TCP connect**, **HTTP**, or **HTTPS**. **NOTE:** TCP is the default selection and recommended for most AppExpress applications.
Ping Hostname	Enter the hostname or IP address of the server from which the probes originate.

Field	Description
User Experience Thresholds	Enter the two threshold values (in milliseconds) for the application performance. These values are measures of latency and are what the Target QoE is derived from. The defaults are 100 ms for the threshold between Satisfied and Tolerable, and 300 ms for the threshold between Tolerable and Frustrated. When determining the threshold values, keep the following definitions in mind: Satisfied – If latency for the application is at or below this threshold, users will have the best experience with the application. Tolerable – If latency for the application falls in this range, users will have a tolerable experience with the application, but it could be better. Frustrated – If latency for the application falls at or above this threshold, users could have a negative experience with the application.

5. Click **Apply** to save your changes.

   The Performance Monitor column shows “Yes” for the application.

Edit an Application Definition

1. Click the edit icon for an application definition.

   The Application Definition dialog box opens.
2. Edit the settings as needed.
3. Click **Apply**.

Disable an Application Definition

1. Click the edit icon for an application definition.

   The Application Definition dialog box opens.
2. Click the **Disabled** check box.
3. Click **Apply**.

Delete a User-created Application Definition

You can only delete user-created application definitions.

1. Click the edit icon for an application definition.

   The Application Definition dialog box opens.
2. Click **Delete**.

   The Delete Record dialog box opens.

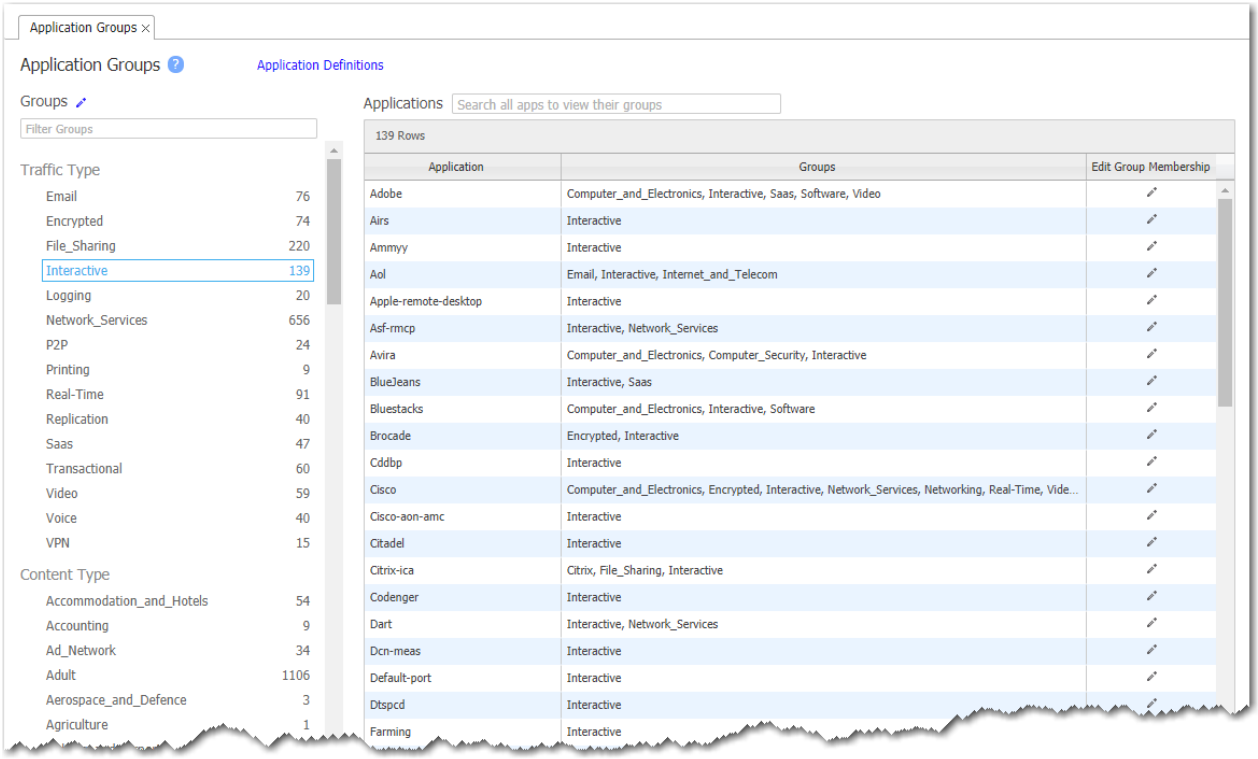
3. Click **Delete**.

   The definition is deleted and all dialog boxes close.

Application Groups Tab

Configuration > Templates & Policies > Applications & SaaS > Application Groups

Application groups associate applications into a common group you can use as a MATCH criteria. The applications can be built-in, user-defined, or a combination of both.



- The **Group Name** cannot be blank.
- Group names are case-insensitive.
- An application group cannot contain another application group.
- A group name followed by * indicates a group defined by a user.
- You cannot change the name of a group provided by Orchestrator, but you can modify the applications those groups contain.

NOTE: To avoid performance issues, it is strongly recommended that you assign an application to no more than three groups.

Apply AppExpress Groups Tab

Configuration > Templates & Policies > Applications & SaaS > Apply AppExpress Groups

From this tab you can add and remove AppExpress groups from appliances. If you choose to have AppExpress steer traffic for an application, it must be added to an AppExpress group. The AppExpress group determines the available transport paths for the application and the frequency that the Ping QoE and User QoE are updated.

NOTE: If you plan to enable DNS proxy and deploy AppExpress Only mode, you must push an empty DNS template to the appliances before you apply any AppExpress groups to the appliances. See Templates Overview.

To create an AppExpress group, click **Create AppExpress Group**. The AppExpress Groups tab opens. Click **+Add Group**. For more information about adding AppExpress Groups, see AppExpress Groups Tab.

AppExpress Groups Tab

Configuration > Templates & Policies > Applications & SaaS > AppExpress Groups

On this tab, you can view all AppExpress groups, edit or remove groups, and add new groups. To edit a group, click the edit icon next to the group you want to edit. To add a new group click **+Add Group**.

The information in the following table is displayed for each end entity certificate on this tab.

Field	Description
Group Name	The name of the AppExpress group.
Overlay	The overlay that the flows for the AppExpress group are put into.
Eligible transports	The transports that AppExpress polls across for this group. **NOTE:** The Overlay configuration determines which transports are available.
Target QoE	The quality of experience (QoE) target measure for this AppExpress group.
Application(s)	The applications that are assigned to the AppExpress group.

Add or Edit AppExpress Group Dialog Box

From this dialog box you can add a new AppExpress group or edit an existing group. Orchestrator includes a few default preconfigured groups, or you can create your own groups.

To add a new AppExpress group:

1. Click **+Add Group**.

   The Add AppExpress Group dialog box opens. Enter information in the following fields:

Field	Description
Group Name	Enter a name for the AppExpress group.
Overlay	Select an overlay for the group. The flows for the group are put into the selected overlay. The overlay you select determines the available transports for the applications included in the AppExpress group. **NOTE:** Applications that are in the group are automatically added to the BIO that is chosen and you do not need to add them to the BIO manually because AppExpress overrides Overlay ACLs and sits logically in front of Overlays.
Eligible transports	Click the check box for each transport that that you want AppExpress to poll across for this group. These selections determine which paths are available for AppExpress to optimize for this group for transport. The transports listed are determined by the overlay selected in the Overlay field. See Transport Types for more information about Transport Types that AppExpress supports.
Target QoE	This is the desired quality of experience (QoE) for the applications in the group and determines how traffic is routed to the transport paths. Select from these options: **Excellent**, **Good**, **Fair**, **Best Effort**. Most applications will be set to Good or Excellent.

2. To add applications to the group, drag and drop applications from the AppExpress Applications box to the [Realtime] Group box.
3. (optional) Click **Show Advanced** to see additional settings. If you want to change from the default options for these settings, enter information in the following fields:

Field	Description
Ping interval	Determines how frequently synthetic probes are sent to the applications in the group. Measured in seconds. The default is 10 seconds.
Source loopback	The loopback interface for the group. It is recommended to use LOOPBACK for sourcing synthetic probes and for proxied DNS queries. It automatically sets itself to use Orchestrated Loopbacks.

Field	Description
DNS for ping and user traffic	Enter the DNS that will be used to send probes (pings) and monitor user traffic. The default is 8.8.8.8. **NOTE:** Leave this field blank to disable DNS proxy for applications in this group. **NOTE:** If you plan to enable DNS proxy and deploy AppExpress Only mode, you must push an empty DNS template to the appliances before you apply any AppExpress groups to the appliances. See Templates Overview.
User QoE interval	Determines how frequently the Apdex score is calculated for user flows and how often a new path can be chosen.
Ping QoE interval	Determines how frequently the Ping QoE is calculated and how often the path hunting is reset.

4. Click **Save**.

To edit a group:

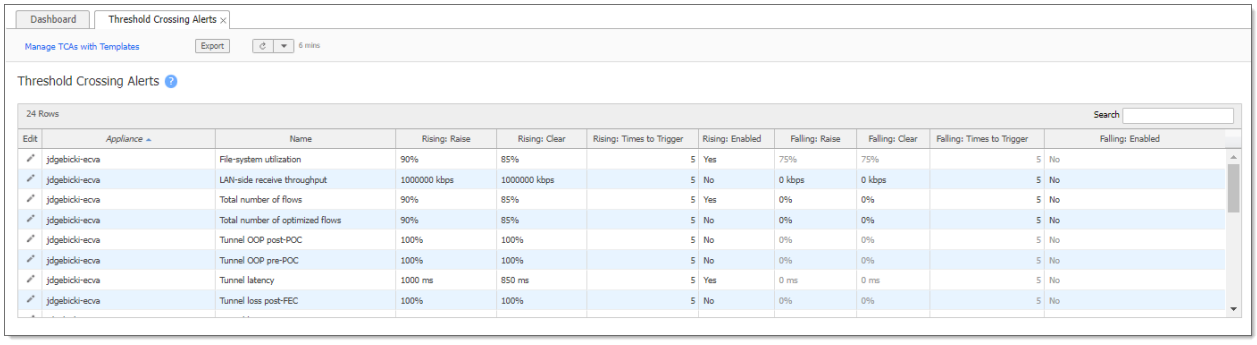
1. Click the edit icon next to the group you want to edit.

   The Edit AppExpress group dialog box opens.
2. Edit the settings as needed.
3. Click **Save**.

Threshold Crossing Alerts Tab

Configuration > Templates & Policies > TCAs > Threshold Crossing Alerts

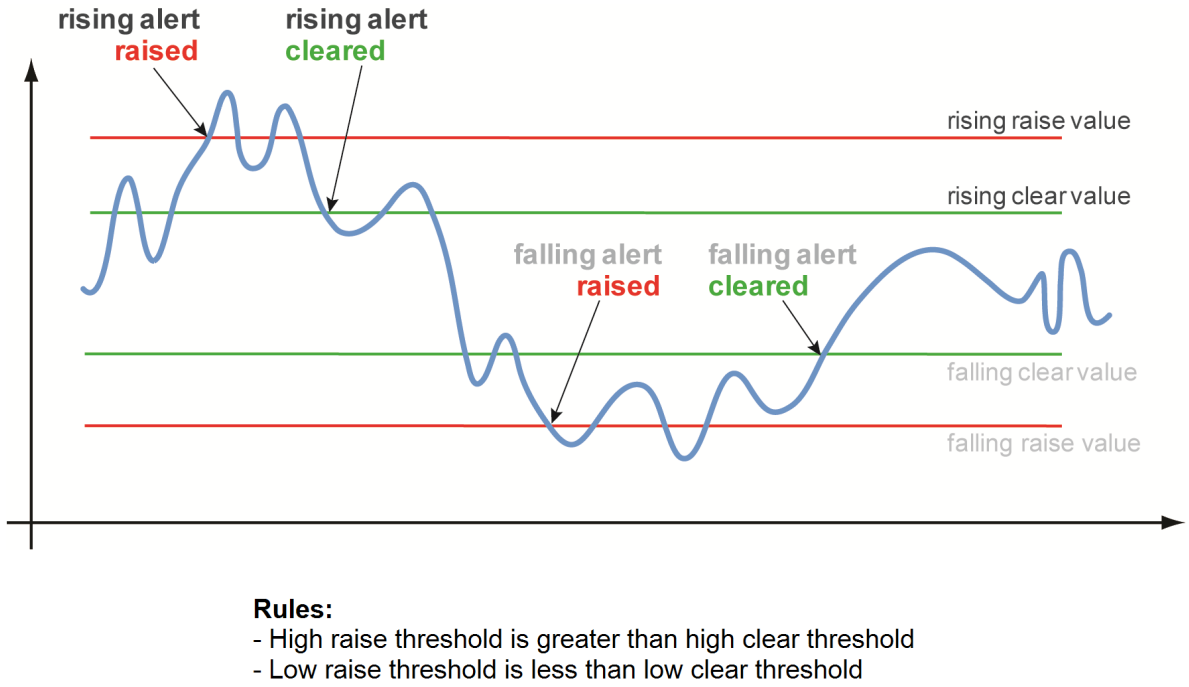
Threshold Crossing Alerts (TCAs) are pre-emptive, configurable alarms triggered when specific thresholds are crossed.



The alerts are triggered with rising and falling threshold crossing events (that is, floor and ceiling levels). For both levels, one value raises the alarm while another value clears it.

When you configure appliance and tunnel TCAs with an Orchestrator template, all alerts apply globally, so all of an appliance's tunnels have the same alerts.

Times to Trigger – A value of 1 triggers an alarm on the first threshold crossing instance.



ON by Default

- **Appliance Capacity** – Triggers when an appliance reaches 95% of its total flow capacity. It is not configurable and can be cleared only by an operator.
- **File-system utilization** – Percent of non-Network Memory disk space filled by the appliance. This TCA cannot be disabled.
- **Tunnel latency** – Measured in milliseconds, the maximum latency of a one-second sample within a 60-second span.

OFF by Default

- **LAN-side receive throughput** – Based on a one-minute average, the LAN-side receive **TOTAL** for all interfaces.
- **WAN-side transmit throughput** – Based on a one-minute average, the WAN-side transmit **TOTAL** for all interfaces.
- **TCAs based on an end-of-minute count:**
  - Total number of flows

  - Total number of optimized flows

- **TCAs based on a one-minute average:**
  - Tunnel loss post-FEC
  - Tunnel loss post-FEC
  - Tunnel OOP post-POC
  - Tunnel OOP post-POC
  - Tunnel reduction
  - Tunnel utilization (based on percent of configured maximum [system] bandwidth)

Threshold Crossing Alerts Edit Row

Click any cell in the table to edit and configure the Threshold Crossing Alerts.

This table lists the **defaults** of each type of threshold crossing alert:

TCA Name	Default [ON or OFF]	Default Values [Rising Raise; Rising Clear; Falling Raise; Falling Clear]	allow rising	allow falling
Appliance Level				
WAN-side transmit throughput	OFF	1 Gbps; 1 Gbps; 0; 0	4	4
LAN-side receive throughput	OFF	1 Gbps; 1 Gbps; 0; 0	4	4
Total number of optimized flows	OFF	256,000, 256,000; 0; 0	4	4
Total number of flows	OFF	256,000, 256,000; 0; 0	4	4
File-system-utilization	**ON** (cannot be disabled)	95%; 85%; 0%; 0%	4	–
Tunnel Level				
Tunnel latency	**ON**	1000; 850; 0; 0	4	–
Tunnel loss pre-FEC	OFF	100%; 100%; 0%; 0%	4	–
Tunnel loss post-FEC	OFF	100%; 100%; 0%; 0%	4	–
Tunnel OOP pre-POC	OFF	100%; 100%; 0%; 0%	4	–
Tunnel OOP post-POC	OFF	100%; 100%; 0%; 0%	4	–
Tunnel utilization	OFF	95%; 90%; 0%; 0%	4	4

TCA Name	Default [ON or OFF]	Default Values [Rising Raise; Rising Clear; Falling Raise; Falling Clear]	allow rising	allow falling
Tunnel reduction	OFF	100%; 100%; 0%; 0%	–	4

IP SLA

Configuration > Templates & Policies > TCAs > IP SLA

Use the IP SLA tab to create and manage IP SLA rules.

Using a polling process, IP SLA (Internet Protocol Service Level Agreement) tracking generates specific network actions that depend on the state of an IP interface or tunnel. The goal is to prevent dropped traffic. For example, associated IP subnets can be removed from the subnet table and from subnet sharing if the LAN-side interfaces on an appliance go down.

The IP SLA tab lists every IP SLA rule configured on the selected appliances. From this tab you can:

- Click **Add** to create a rule.
- Click the icon in the Charts column to view real-time and historical latency and loss trends for an interface or tunnel.
- Click the **IP SLA Summary** link at the top of the tab to open the read-only, aggregated performance view across all appliances.

NOTE: For 9.6.0 or later deployments, your IP SLA Rule Destination addresses must be compatible with IPv6 if you want to use IPv6 local breakout addressing. The following IP SLA internet endpoints are recommended:

```
192.151.28.254,2600:9000:a60f:6c50:ca04:acb1:b183:9524,2600:9000:a70d:506e:85a8:bd96
:7482:16,sp-ipsla.silverpeak.cloud
```

Replace the contents of the **Address** field on the IP SLA Rule Destination dialog box with these endpoints as one continuous string. For details, see Business Intent Overlays.

IP SLA Tab

Each row on the IP SLA tab represents one IP SLA rule on one appliance.

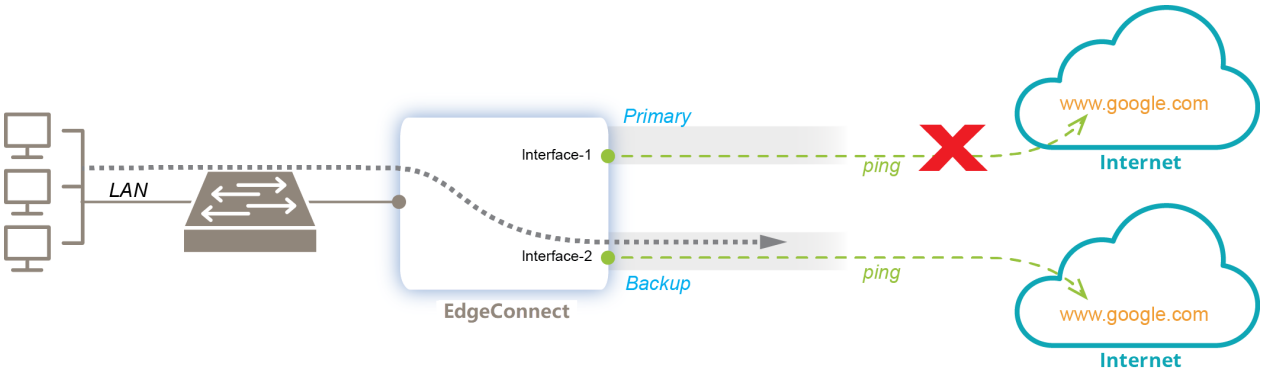
Column	Description
Appliance	The EdgeConnect appliance on which the rule is configured.
Active	Whether the rule is currently enabled (Yes) or disabled (No). Set this state with the ON/OFF toggle in the IP SLA Rule dialog box.

Column	Description
State	The current operational state of the rule, such as Up or Down. Down states are highlighted.
Monitor	The monitor type and its target, such as Ping with an address or Interface with an interface name.
Segment	The routing segment (VRF) in which the rule operates. Appears when segmentation is enabled.
Down action	The action performed when the monitored target transitions to the Down state.
Up action	The action performed when the monitored target returns to the Up state.
Route label	The route label, if any, associated with the rule.
Charts	Opens real-time and historical latency and loss trend charts for the rule.
Comment	An optional description of the rule. The first 64 characters appear in any alarms generated for the rule.
Up stats	The duration of the most recent up period.
Down stats	The number of down events and the duration of the most recent down period.
Actions	Edit or delete the rule.

IP SLA Monitor Use Cases

The following examples describe six basic use cases for IP SLA monitoring. The diagrams illustrate the network scenario and the intent of each rule; for the current field layout of the IP SLA Rule dialog box, see IP SLA Rule.

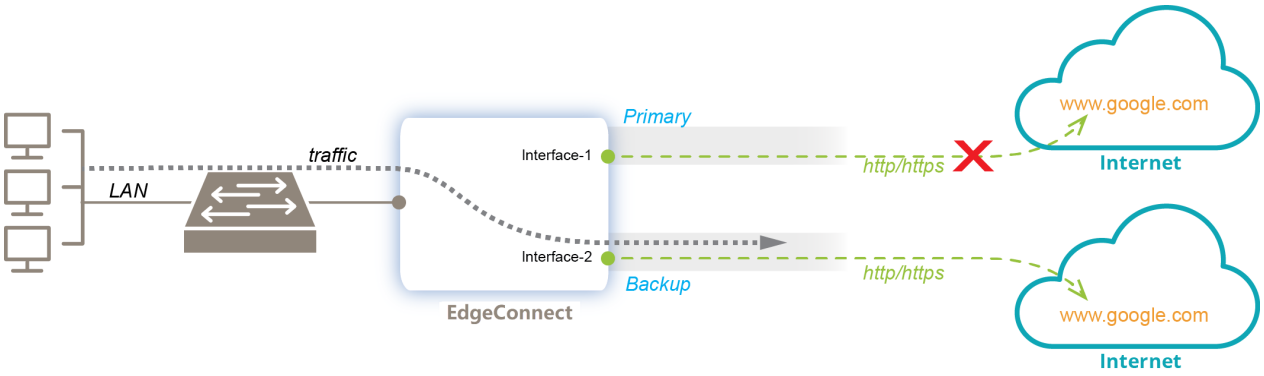
Example #1 – Ping via Interface



- Two passthrough tunnels configured for Internet breakout and High Availability.
- If the Primary passthrough tunnel goes down, traffic goes to the Backup tunnel.

- Configure the IP SLA Rule with a Ping monitor and specify the same tunnel for both the Down Action and the Up Action. See IP SLA Rule for the field-by-field layout.

Example #2 – HTTP/HTTPS via Interface

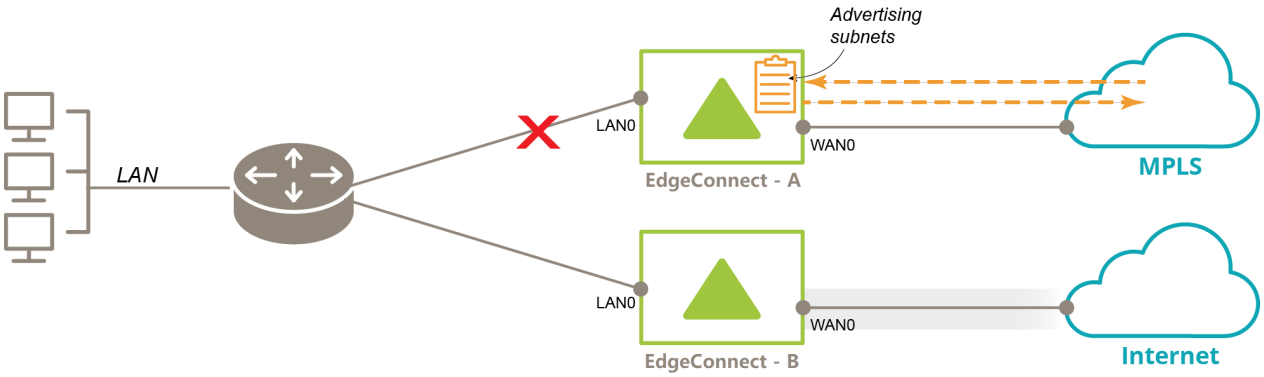


- Two passthrough tunnels configured for Internet breakout and High Availability.
- If the Primary passthrough tunnel goes down, traffic goes to the Backup tunnel.
- Configure the IP SLA Rule with an HTTP/HTTPS monitor and specify the same tunnel for both the Down Action and the Up Action. See IP SLA Rule for the field-by-field layout.

NOTE: Do not select a loopback interface for HTTP/HTTPS IP SLA monitoring. A loopback interface is invalid because packets cannot be sent to that interface. If the source interface selection is a tunnel, select a second source interface. This second source interface is used as the source IP address of the packet to be sent through the tunnel.

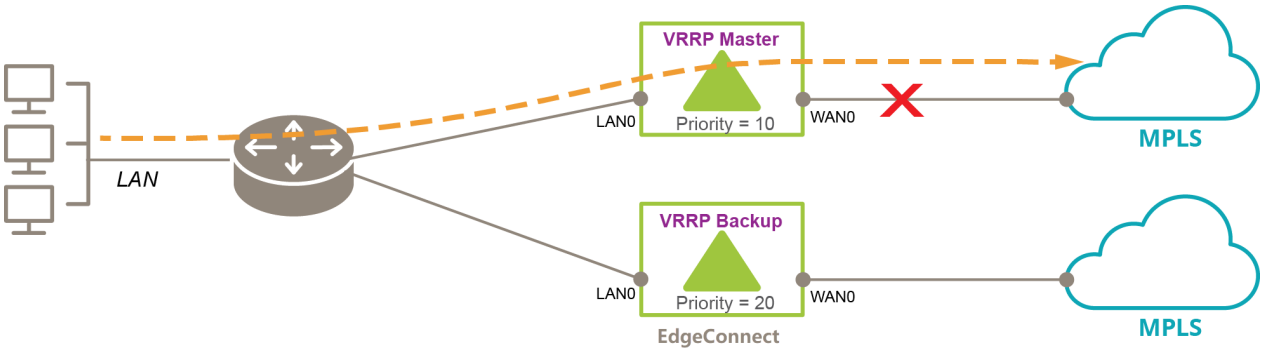
- In the URL(s) field, the protocol identifier is required only when specifying HTTPS, as in `https://www.google.com`.

Example #3 – Monitor Interface



- On *EdgeConnect - A*, we want subnet advertising conditional on LAN0 being up.
- Configure the IP SLA Rule with an Interface monitor on LAN0 and set the Default Subnet Action to resume advertising subnets.

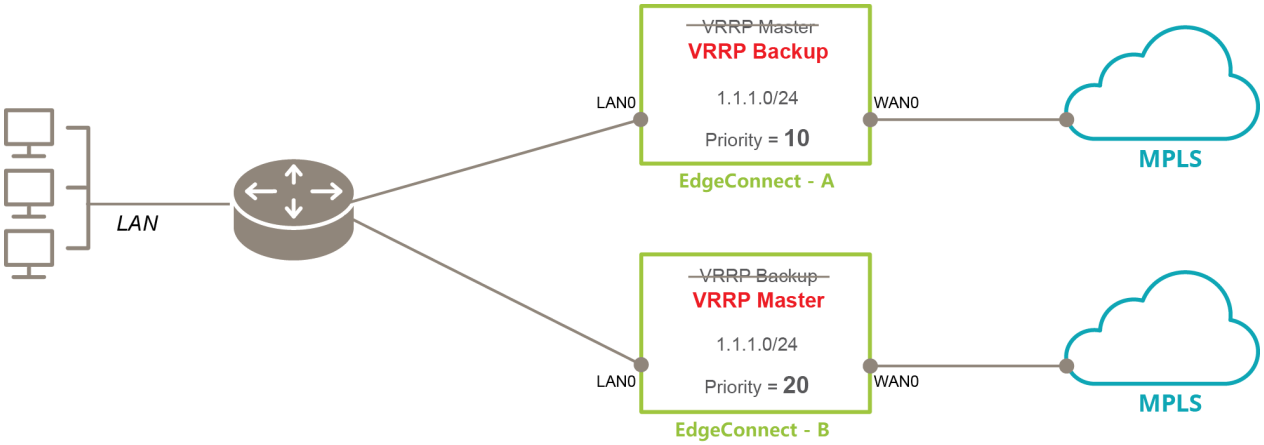
Example #4 – Monitor Interface (WAN0) to Ensure High Availability



- If WAN0 goes down on the VRRP Master, we want to decrease its Priority so traffic goes to the VRRP Backup.
- Configure the IP SLA Rule with an Interface monitor on WAN0 and set the Default Subnet Action to revert to the original priority.

  NOTE: In this instance, the WAN0 interface was given the label MPLS to match the service to which it connected.

Example #5 – Monitor VRRP



- To monitor the VRRP router state, use VRRP Monitor and specify the interface on which the VRRP instance is configured.

  In this example, it is LAN0.
- Here we are looking at an instance where the VRRP role changes, but priority does not, for whatever reason.
- Configure the IP SLA Rule with a VRRP monitor and set the Default Subnet Action to revert to the original priority.

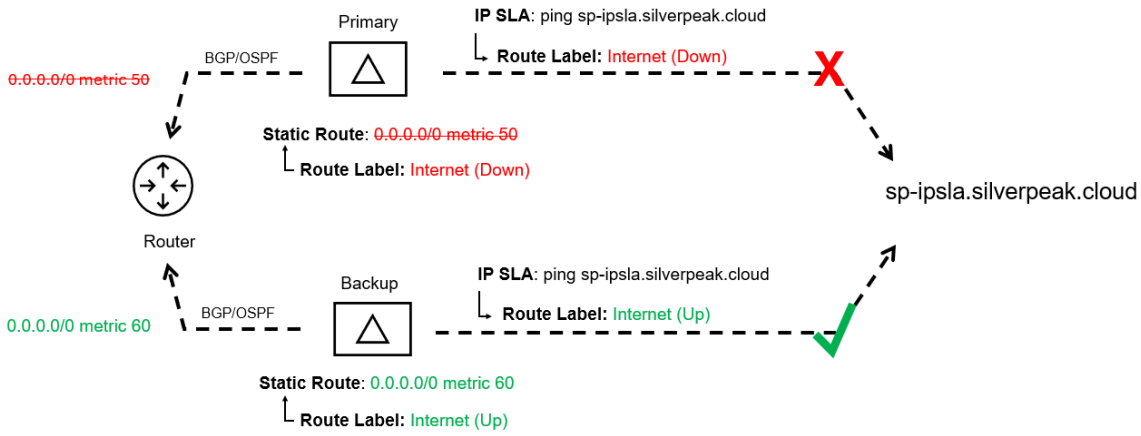
NOTE: In this instance, the WAN0 interface was given the label MPLS to match the service to which it connected.

- Another option is to leave the VRRP priority unchanged and instead influence path selection through subnet metrics. Set Down Action = Modify Subnet Metric, then enter a positive value in the field that Orchestrator adds to increase the current subnet metric. Set Up Action = Default Subnet Action to restore the original metric when the monitored state recovers.

  Use this approach when you want traffic to shift based on overlay route preference rather than by changing which appliance is the VRRP master. Raising the subnet metric makes the affected subnets less preferred to EdgeConnect peers, so traffic favors an alternate path across the SD-WAN fabric while the local VRRP master election stays the same. This produces a different result from adjusting VRRP priority: modifying the subnet metric steers routing among EdgeConnect peers, whereas changing VRRP priority steers first-hop gateway selection for locally attached clients.

  In most scenarios, adjust VRRP priority when the goal is local gateway (first-hop) redundancy, and modify the subnet metric when the goal is to influence path selection across the fabric. If a single monitored event must affect both first-hop and fabric behavior, you can use both mechanisms together.

Example #6 – Revoke Static Route Using IP SLA



- In example 6, the route label Internet is associated with an IP SLA rule to ping sp-ipsla.silverpeak.cloud. The route label is assigned to two static routes and each route is assigned a metric. When the static route with a metric of 50 goes down, traffic is routed to the backup route (with a metric of 60). See Route Labels for more information.
- Configure the IP SLA Rule with the Down Action set to Raise Alarm Only and the Up Action set to Clear Alarm Only.

IP SLA Rule

Use the IP SLA Rule dialog box to create or modify an IP SLA rule. On the IP SLA tab, click **Add** to create a rule, or use the edit control in the Actions column to modify an existing rule.

1. Set the rule state with the **ON**/**OFF** toggle at the top of the dialog box. ON enables the rule; OFF disables the rule without deleting it. This state is reflected in the Active column on the IP SLA tab.
2. Under **Monitor**, select the type of target to monitor.

Option	Description
Interface	Monitors the operational status of a specific local interface.
Ping	Monitors the reachability of one or more IPv4 or IPv6 addresses.
HTTP/HTTPS	Monitors the reachability of an HTTP/HTTPS endpoint. **NOTE:** Both HTTP and HTTPS require a response of 200. Redirects are not supported. **NOTE:** Using HTTPS as a monitor for IP SLA with multiple targets can cause potential problems. HTTPS does not provide any additional benefit about the path check.
VRRP Monitor	Monitors the VRRP router state (TRUE if Master; FALSE if Backup) for a VRRP instance on an interface.

3. Complete the fields that appear for the monitor you selected. The dialog box displays a different set of fields for each monitor type.

   Interface and VRRP Monitor

Field	Description
Interface	The local interface to monitor. For VRRP Monitor, this is the interface on which the VRRP instance is configured.
Check IP SLA status every	The polling interval, in seconds, at which the appliance checks the monitored interface or VRRP state. The default is 30 seconds.

   Ping

Field	Description
Address	One or more target IPv4 or IPv6 addresses to ping. Separate multiple addresses with commas on a single line.
Source	Whether probes are sourced from an Interface or a Tunnel.
Tunnel	(Source = Tunnel) The passthrough tunnel used to send probes.
Source Interface / Interface	The interface used to source probes. When the source is a tunnel, this second interface supplies the source IP address of the probe packet.
Ping interval	The interval, in seconds, between successive pings.
Rolling average window for loss and latency	The averaging window, in seconds (1–900), used to smooth the loss and latency measurements.
Force DNS through probe path / Do not force DNS through probe path	Whether DNS resolution for the target is forced through the probe path.
Mark up after X sequential successes / Mark down after X sequential failures	The reachability thresholds that transition the rule up or down.
Mark up after loss below X % / Mark down after loss above X %	The packet-loss thresholds, as a percentage (0–100), that transition the rule up or down.
Mark up after average latency below X / Mark down after average latency above X	The average-latency thresholds, in milliseconds (0–9999), that transition the rule up or down.

HTTP/HTTPS

Field	Description
URL(s)	One or more HTTP or HTTPS endpoints to monitor. Separate multiple URLs with commas. The protocol identifier is required only when specifying HTTPS, as in `https://www.example.com`.
Proxy address	*(Optional)* The IPv4 or IPv6 address of a proxy server used to reach the endpoint.
Proxy port	(Required when a proxy address is set) The proxy server port (1–65535).
User agent	*(Optional)* A custom User-Agent string to send with the request.

Field	Description
HTTP request timeout	The time, in seconds, to wait for a response before the request is considered failed.
Source	Whether probes are sourced from an Interface or a Tunnel.
Tunnel	(Source = Tunnel) The passthrough tunnel used to send probes.
Source Interface / Interface	The interface used to source probes. When the source is a tunnel, this second interface supplies the source IP address of the probe packet.
Polling frequency	The interval, in seconds, between successive requests. The default is 60 seconds.
Rolling average window for loss and latency	The averaging window, in seconds (1–900), used to smooth the loss and latency measurements.
Force DNS through probe path / Do not force DNS through probe path	Whether DNS resolution for the target is forced through the probe path.
Mark up after X sequential successes / Mark down after X sequential failures	The reachability thresholds that transition the rule up or down.
Mark up after loss below X % / Mark down after loss above X %	The packet-loss thresholds, as a percentage (0–100), that transition the rule up or down.
Mark up after average latency below X / Mark down after average latency above X	The average-latency thresholds, in milliseconds (0–9999), that transition the rule up or down.

NOTE: Do not select a loopback interface for HTTP/HTTPS IP SLA monitoring. A loopback interface is invalid because packets cannot be sent to that interface. If the source is a tunnel, select a second source interface; this interface is used as the source IP address of the packet sent through the tunnel.

NOTE: For Ping and HTTP/HTTPS monitors, the IP version of the target must match the IP version of the source interface or path used to send the probes. If the target is an IPv6 address (or resolves only to IPv6) but the source interface has no IPv6 address configured—or the reverse—the monitor is marked Down and its Down Action is applied. Make sure the source interface is configured for the same IP version (IPv4 or IPv6) as the monitored target.

4. Select an action from the **Down Action** list. This action runs when the monitored target transitions to the Down state.

Down Action	Description
Enable Tunnel	Enable a passthrough (internet breakout) tunnel for IP SLA tracking purposes. Select the tunnel in the **Tunnel** field that appears.
Disable Tunnel	Disable a passthrough (internet breakout) tunnel for IP SLA tracking purposes. The tunnel can no longer be used for load balancing (when load balancing traffic between multiple passthrough tunnels), although it can still be used as a last resort for traffic forwarding. Select the tunnel in the **Tunnel** field that appears.
Advertise Subnets	Advertise subnets to EdgeConnect peers. Select the interface in the **Interface** field that appears.
Decrease VRRP Priority	Decrease the configured VRRP router priority by a delta amount. Select the **Interface** and enter the **Priority** delta in the fields that appear.
Disable Subnet Sharing	Disable the sharing of subnets with other EdgeConnect peers on the appliance.
Increase VRRP Priority	Increase the configured VRRP router priority by a delta amount. Select the **Interface** and enter the **Priority** delta in the fields that appear.
Modify Subnet Metric	Add a metric delta to the metric of all subnets shared with EdgeConnect peers. Enter the amount in the **Delta** field that appears.
Remove Auto Subnets	Remove from the subnet table an auto subnet for a port (including all VLAN and subinterface subnets). Select the interface in the **Interface** field that appears.
Raise Alarm Only	Raise an alarm when the IP SLA rule goes down.

NOTE: If an IP SLA rule is misconfigured, the appliance marks the rule Down and applies its Down Action. Verify the monitor and source settings if a rule remains Down unexpectedly.

5. Select an action from the **Up Action** list. This action runs when the monitored target returns to the Up state.

Up Action	Description
Enable Tunnel	Enable a passthrough (internet breakout) tunnel for IP SLA tracking purposes. Select the tunnel in the **Tunnel** field that appears.

Up Action	Description
Disable Tunnel	Disable a passthrough (internet breakout) tunnel for IP SLA tracking purposes. The tunnel can no longer be used for load balancing (when load balancing traffic between multiple passthrough tunnels), although it can still be used as a last resort for traffic forwarding. Select the tunnel in the **Tunnel** field that appears.
Decrease VRRP Priority	Decrease the configured VRRP router priority by a delta amount.
Increase VRRP Priority	Increase the configured VRRP router priority by a delta amount.
Default Subnet Action	Reverts whatever the Down Action was back to the normal state. Examples: If Down Action = Disable Subnet Sharing, the Up Action re-enables subnet sharing. If Down Action = Remove Auto Subnets, the Up Action re-adds the auto subnet. If Down Action = Modify Subnet Metric, the Up Action restores subnet metrics to their original values.
VRRP Default	Reverts the VRRP priority back to the configured value.
Clear Alarm Only	Clears any alarms created as a result of an IP SLA rule going down.

NOTE: If a default Up Action (Default Subnet Action, VRRP Default, or Clear Alarm Only) is used, it must match the corresponding Down Action.

6. *(Optional)* Select a **Route Label** from the list of configured route labels to assign to the rule.

   The Route Label field is in addition to the Down Action and Up Action fields. Route labels are used with IP SLA rules to mark associated, locally configured static or announce-only routes Down when the route label state is Down.

   - Route labels are configured globally at **Configuration > Overlays & Security > Route Labels**.
   - Multiple IP SLA rules can reference the same route label. The route label remains Up as long as at least one associated IP SLA rule is Up. The route label is marked Down only when all IP SLA rules referencing it are Down. The exact behavior also depends on the route label's IP SLA Condition (ALL or ANY).
   - Route labels apply only to locally configured static or announce-only routes. They do not affect routes received from BGP, OSPF, or fabric peers.
   - The Route Label action is independent of the Down Action field; both can be configured on the same rule.

   For more information, see Route Labels.

7. *(Optional)* In the **Comment** field, enter a description of the rule. The first 64 characters are included in any alarms generated for the rule.
8. Click **Save**.

Configuration > Templates

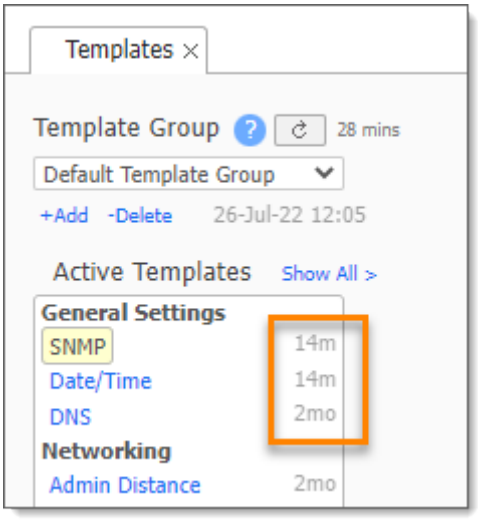
The options under **Configuration > Templates** focus on setting up a variety of templates that you can apply to various aspects of Orchestrator and applying template groups.

Templates Overview

Use templates to manage and assign common configuration parameters to appliances.

CAUTION: After saving, templates are applied automatically and replace all settings on an appliance with those configured in the template. Some templates support a MERGE option. Refer to the Help for more information.

- Each template that appears under Active Templates includes a timestamp that indicates the amount of time that has passed since it was last edited, and the most recently edited templates appear at the top of each template section in the list.

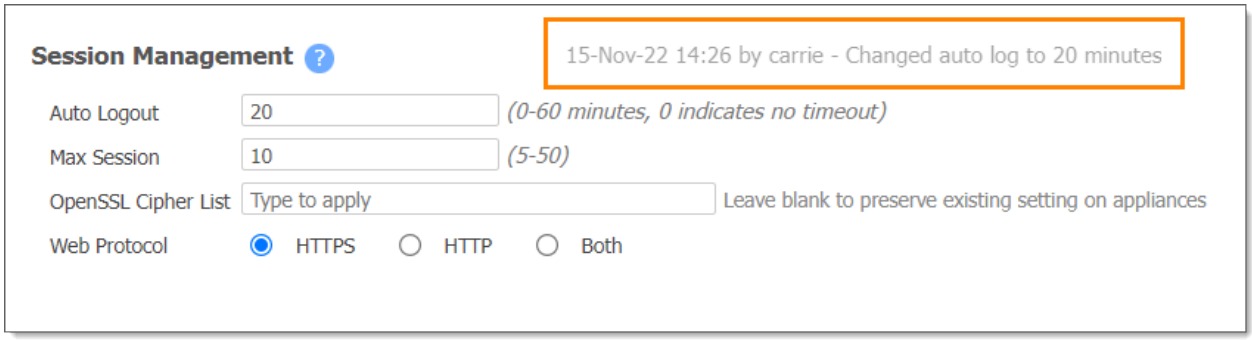


- You can edit only a template that appears under Active Templates.
- Click **Show All >** to view available templates that are not part of the selected template group.
- To add a template to Active Templates, double-click it or drag it from Available Templates.
- To copy and save the current Active Templates as a new template group, click **Save As**.

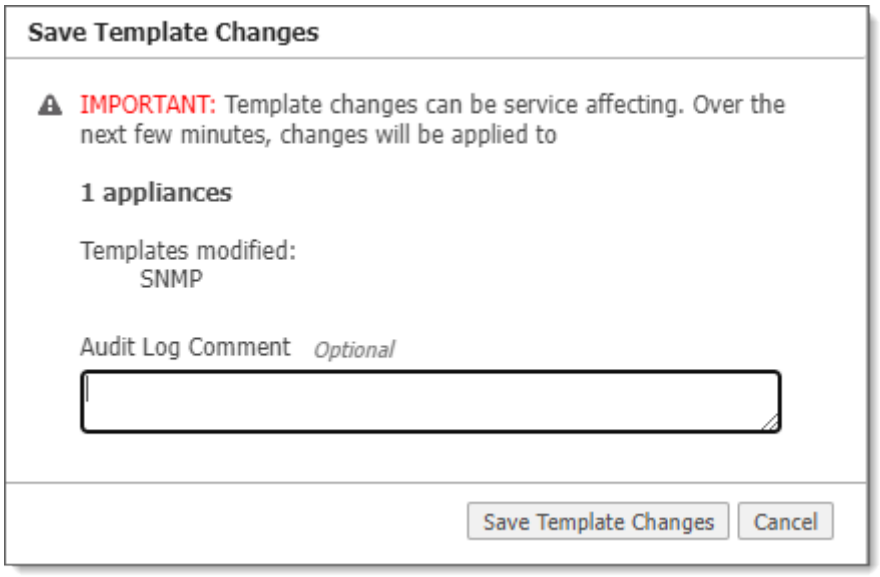
Modifying a Template

1. Click a template under Active Templates to modify it.

   The template has a timestamp that indicates when it was last modified, the user who made the changes, and any comment (optional) that was entered by the user who made the changes. The timestamp appears in the format "DD-MMM-YY HH:MM by [user] - [optional comment text]" and the time is expressed in a 24-hour format.



2. To save the changes you made, click **Save**. The Save Template Changes dialog box opens.



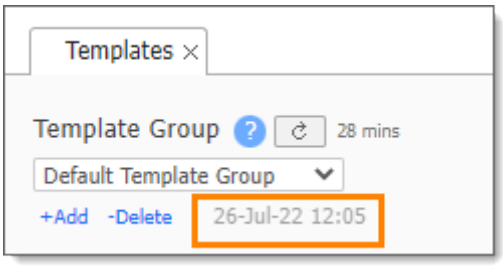
3. Enter a comment (optional) in the **Audit Log Comment** field, and then click **Save Template Changes**. Any text entered in the **Audit Log Comment** field appears on the template timestamp and the Audit Logs tab.

NOTE: If the text of a timestamp comment is truncated, hover your cursor over the timestamp to display the full text of the comment.

Template Groups

A Template Group contains one or more templates you can assign to some or all of the appliances in your network.

- A timestamp for the selected template group appears below the template group drop-down list and it indicates when one of the templates in the template group was last modified. The timestamp appears in the format “DD-MMM-YY HH:MM” and the time is expressed in a 24-hour format.



- To create a template group, click **+Add** below the template group drop-down list.
  - To save changes you made to the new template group, click **Save**. The Save Template Changes dialog box opens. Enter a comment (optional) in the **Audit Log Comment** field, and then click **Save Template Changes**. Any text entered in the **Audit Log Comment** field appears on the template timestamp and the Audit Logs tab.
- To delete the selected template group, click **-Delete** below the template group drop-down list.
- When you apply a template group to an appliance, Orchestrator automatically keeps the templates in the group in sync with the appliance.
- To apply template groups, click **Apply Template Groups** at the bottom of the page. This will bring you to the Apply Templates tab where you can permanently associate appliances with specific template groups.
- When returning to the Templates page, Orchestrator displays the last template group viewed.

System Template

Use this template to configure system-level features.

Optimization

Field	Description
IP ID auto optimization	Enables any IP flow to automatically identify the outbound tunnel and gain optimization benefits. Enabling this option reduces the number of required static routing rules (route map policies).
TCP auto optimization	Enables any TCP flow to automatically identify the outbound tunnel and gain optimization benefits. Enabling this option reduces the number of required static routing rules (route map policies).
Flows and tunnel failure	If there are parallel tunnels and one fails, __*Dynamic Path Control__* determines where to send the flows. There are three options: **fail-stick** – When the failed tunnel comes back up, the flows do not return to the original tunnel. They stay where they are. **fail-back** – When the failed tunnel comes back up, the flows return to the original tunnel. **disable** – When the original tunnel fails, the flows are not routed to another tunnel.

Network Memory

Field	Description
Encrypt data on disk	Enables encryption of all the cached data on the disks. Disabling this option is not recommended.

Excess Flow Handling

Field	Description
Excess flow policy	Specifies what happens to flows when the appliance reaches its maximum capacity for optimizing flows. The default is to **bypass** flows. Or, you can choose to **drop** the packets.

Miscellaneous

Field	Description
SSL optimization for non-IPSec tunnels	Specifies whether the appliance should perform SSL optimization when the outbound tunnel for SSL packets is not encrypted (for example, a GRE or UDP tunnel). To enable Network Memory for encrypted SSL-based applications, you must provision server certificates by using the Orchestrator. This activity can apply to the entire distributed network of EdgeConnect appliances or just to a specified group of appliances.
Bridge loop test	Only valid for virtual appliances. When enabled, the appliance can detect bridge loops. If it detects a loop, the appliance stops forwarding traffic and raises an alarm. Appliance alarms include recommended actions.
Always send pass-through traffic to original sender	If the tunnel goes down when using WCCP and PBR, traffic that was intended for the tunnel is sent back the way it came.
Enable default DNS lookup	Enables the default DNS server to be included with other configured DNS servers for associating cloud portal domain names to network IP addresses.
Enable HTTP/HTTPS snooping	Enables a more granular application classification of HTTP/HTTPS traffic by inspection of the HTTP/HTTPS header, Host. This is enabled by default.
Quiescent tunnel keep alive time	Specifies the rate at which to send keep alive packets after a tunnel has become idle (quiescent mode). The default is 60 seconds.
UDP flow timeout	Specifies how long to keep the UDP session open after traffic stops flowing. The default is 120 seconds (2 minutes).
Non-accelerated TCP flow timeout	Specifies how long to keep the TCP session open after traffic stops flowing. The default is 1800 seconds (30 minutes).

Field	Description
Maximum TCP MSS	Maximum Segment Size. The default value is 1328 bytes. This setting ensures that packets larger than the actual maximum transmission unit (MTU) are not dropped if fragmentation is not possible. Some services such as (but not limited to) Zscaler or PPPoE require this setting to be 1328 for successful packet transmission. You can set the value from 500 to 9000, but 1328 works successfully across all known link, tunnel, and traffic types without any performance degradation. **NOTE:** This setting applies only to passthrough or third-party IPSec tunnelled flows. Flows routed through IPSec UDP tunnels or other underlay tunnels do not use this setting.
NAT-T keep alive time	If a device is behind a NAT, this specifies the rate at which to send keep alive packets between hosts to keep the mappings in the NAT device intact.
Tunnel alarm aggregation Threshold	Specifies the number of alarms to allow before alerting the tunnel alarm.
Maintain end-to-end overlay mapping	Enforces the same overlay to be used end-to-end when traffic is forwarded on multiple nodes.
IP directed broadcast	Allows an entire network to receive data that only the target subnet initially receives.
Allow WAN to WAN routing	Redirects inbound WAN traffic back to the WAN.
Allow Unknown Destination Role	Indicates whether to allow unknown destination roles.
Stateful-SNAT exceptions	Name of the address group configured for Stateful-SNAT exceptions (for example, Stateful-SNAT-Exceptions). To set up this address group, see Disable Stateful+SNAT Processing for Selected LAN-side Subnets below.

Disable Stateful+SNAT Processing for Selected LAN-side Subnets

Most internet providers require that flows originate from the WAN-side IP address assigned to the appliance. When Stateful+SNAT is configured on a WAN-side interface, all traffic that leaves the interface will be Source NATed to the IP address of the WAN-side interface.

In certain situations, you want the original LAN-side IP address to be seen by the upstream network. You can use the Stateful-SNAT exceptions feature to avoid Source NATing for specific IP addresses or subnets.

Considerations:

- Stateful-SNAT exceptions apply only to appliances with firewall mode set to “Stateful+SNAT”.
- Exceptions apply only to outbound flows destined to external addresses.
- Inbound flows initiated from the WAN side toward IP addresses within the address group rely on existing inbound port-forwarding functionality.
- SNAT exceptions apply to the default segment only, not VRF SNAT.
- This feature does not support IPv6 because the address groups feature does not support IPv6.

You can use the System template to set up Stateful-SNAT exceptions for all appliances or the System Information dialog box for individual appliances. To set up exceptions for individual appliances, see System Information.

To set up Stateful-SNAT exceptions for all appliances:

1. Create an address group for all public IP space (subnets) used by your network across all branches, as follows:
    1. Navigate to **Configuration > Templates & Policies > ACLs > Address Groups**.

        The Address Groups tab opens.
    2. Click **Add Group**.

        The Add Address Group dialog box opens.
    3. In the **Group name** field, enter an appropriate name for the Stateful-SNAT exceptions (for example, Stateful-SNAT-Exceptions).
    4. In the **IPs to include** and **IPs to exclude** fields, enter IP addresses/masks to include/exclude individually or IP prefixes to include/exclude multiple addresses at once, as appropriate. Use commas to separate entries.
    5. If desired, use the **Comment** field to state the purpose of this address group.
    6. Click **Add**.
2. In the System template’s **Stateful-SNAT Exceptions** field, enter the name of the address group you created for Stateful-SNAT exceptions.

Auth/Radius/TACACS+ Template

EdgeConnect appliances support user **authentication** and **authorization** as a condition of providing access rights.

- **Authentication** is the process of validating that the end user, or a device, is who they claim to be.
- **Authorization** is the action of determining what a user is allowed to do. Generally, authentication precedes authorization.
- **Map order** refers to the order in which the authorization servers are queried.
- The configuration specified for authentication and authorization **applies globally** to all users accessing that appliance.
- If a logged-in user is inactive for an interval that exceeds the inactivity time-out, the appliance logs them out and returns them to the login page. You can change that value, as well as the maximum number of sessions, in the **Session Management template**.

Authentication and Authorization

To provide authentication and authorization services, EdgeConnect appliances:

- Support a built-in, **local database**.
- Can be linked to a **RADIUS** (Remote Authentication Dial-In User Service) server.
- Can be linked to a **TACACS+** (Terminal Access Controller Access Control System) server.

Both RADIUS and TACACS+ are client-server protocols.

Appliance-based User Database

- The local, built-in user database supports user names, groups, and passwords.
- The two user groups are **admin** and **monitor**. You must associate each user name with one or the other. Neither group can be modified or deleted.
- The **monitor** group supports reading and monitoring of all data, in addition to performing all actions. This is equivalent to the Command Line Interface's (CLI) __*enable__* mode privileges.
- The **admin** group supports full privileges, along with permission to add, modify, and delete. This is equivalent to the Command Line Interface's (CLI) __*configuration__* mode privileges.

RADIUS

- RADIUS uses UDP as its transport.
- With RADIUS, the authentication and authorization functions are coupled together.
- RADIUS authentication requests must be accompanied by a shared secret. The shared secret must be the same as defined in the RADIUS setup. Refer to your RADIUS documentation for details.

- **IMPORTANT:** Configure your RADIUS server's __*priv levels__* within the following ranges:
  - **admin** = 7 - 15
  - **monitor** = 1 - 6

TACACS+

- TACACS+ uses TCP as its transport.
- TACACS+ provides separated authentication, authorization, and accounting services.
- Transactions between the TACACS+ client and TACACS+ servers are also authenticated through the use of a shared secret. Refer to your TACACS+ documentation for details.
- **IMPORTANT:** Configure your TACACS+ server's roles to be **admin** and **monitor**.

What Is Recommended

- Use either RADIUS or TACACS+, but not both.
- For **Authentication Order**, configure the following:
  - **First** – Remote first.
  - **Second** – Local. If not using either, then None.
  - **Third** – None.
- When using RADIUS or TACACS+ to authenticate users, configure **Authorization Information** as follows:
  - **Map Order** – Remote First
  - **Default Role** – admin

Network Access Control (NAC) Template

Configuration > Templates & Policies > Templates > Active Templates > General Settings > Network Access Control (NAC)

Use the Network Access Control (NAC) template to apply NAC configurations to multiple appliances.

The Network Access Control (NAC) tab displays the configuration settings for NAC security using 802.1x and MAC authentication. When Network Access Control (NAC) is enabled on an appliance, the appliance authenticates traffic that accesses the network over untrusted interfaces.

NAC security supports EAP-TLS, EAP-TTLS, and EAP-PEAP methods for 802.1x authentication.

By default, authentication for all interfaces is set to “trusted.” When authentication is set to “trusted”, no authentication is required to access the network. NAC security allows you to authenticate supplicants on a specific interface. You must enable NAC security, configure the NAC settings, and assign a policy to the interface. When NAC security is enabled, the appliance authenticates clients (supplicants) that are trying to access the network using the policy you assign to the interface.

NOTE: Some devices cannot act as an 802.1x client. You must enable the MAC address authentication on the interface connected to the client. The interface connected to the client uses the client’s MAC address as the username and password and uses the MAC address for authentication.

Click **Enable NAC**, and then complete the following four steps.

1. Create an 802.1x/MAC authentication profile. See 802.1x/MAC Authentication Profiles.
2. Define the servers and optional server groups used for authenticating supplicants on the selected interface. See Server.
3. Add or edit the AAA profiles used for authentication. See AAA Profile.
4. Apply the Network Access Control (NAC) policies to the interface labels. See Apply Policies.

802.1x/MAC Authentication Profiles

Use the 802.1x/MAC tab to add or edit authentication profiles. You should create both 802.1x authentication and MAC authentication profiles. If the supplicant is 802.1x compliant, the appliance will use the 802.1x profile to authenticate the supplicant. If the supplicant is not 802.1x compliant, the appliance will use the MAC profile to authenticate the supplicant.

802.1x Authentication Profile Fields

1. Click **Add** to add a new 802.1x authentication profile or click the pencil icon to edit an existing 802.1x profile.

   The Add 802.1x Authentication Profile dialog box opens.

   NOTE: To delete a profile listed in the table on the Network Access Control (NAC) dialog box, click the corresponding delete icon (X) in the last column.

2. Complete the following fields:

Field	Description
Profile	The name for the 802.1x profile.
Max Auth Failure	The maximum number of authentication failures allowed before the supplicant is denied access.
Max Request	The maximum number of authentication requests that the appliance will send to the server.
Identity Requests Interval	The interval in seconds between identity request retries.

Field	Description
Quiet Period	The interval in seconds to wait between attempting to reauthenticate after a failed authentication.
Server Retry Count	The maximum number of retries that can be made on each server in a server group.
Server Group Retry Period	The timeout duration. If the appliance cannot reach the server in the specified duration, the session times out.
Reauthentication	Select this option to force the appliance to do a reauthentication with the configured reauthentication interval.
Max Reauthentication	The maximum number of reauthentication attempts.
Reauthentication Interval	The interval in seconds, between reauthentication attempts. The configured interval will be overridden if the RADIUS server provided the reauthentication period.
Ignore EAPOL-START After Authentication	Select whether the appliance should ignore the EAPOL-START messages after authentication.
Handle EAPOL-Logoff	Select whether to handle the EAPOL-LOGOFF messages sent by the supplicants.

3. Click **Add**.

MAC Authentication Profile Fields

1. Click **Add** to add a new MAC authentication profile or click the pencil icon to edit an existing MAC authentication profile.

   NOTE: To delete a profile listed in the table on the Network Access Control (NAC) dialog box, click the corresponding delete icon (X) in the last column.

2. Complete the following fields:

Field	Description
Profile	Enter a name for the MAC authentication profile.
Max Auth Failure	The maximum number of authentication failures allowed before the supplicant is denied access.
Quiet Period	The interval in seconds to wait before attempting the retry after the failed authentication.
Server Retry Count	The maximum number of retires that can be made to each server in a server group. If a server is not available, after the specified number of retries, Orchestrator attempts to access the next server in the server group.

Field	Description
Server Group Retry Period	Set the timeout duration in seconds. If the appliance cannot reach the server in the specified duration, the session will time out.
Reauthentication	Select this option to force the appliance to do a reauthentication with the configured reauthentication interval.
Max Reauthentication	The maximum number of reauthentication attempts.
Reauthentication Interval	The interval in seconds between reauthentication attempts.

3. Click **Add**.

Navigate to the Server tab to configure the servers and server groups you want to use to authenticate the supplicants.

Server

Use the Server tab to add or edit the servers and server groups you want to use to authenticate the supplicants attempting to log in to the network.

Servers Fields

1. Click **Add** to add a new server.

   NOTE: Click in any cell to modify an existing data. To delete a server listed in the table, click the corresponding delete icon (X) in the last column.

2. Complete the following fields:

Field	Description
ID	The unique identifier of the server.
Server Name	Enter a name for the server.
IP Address	The IPv4 or IPv6 address of the RADIUS server.
Key	The pre-shared key of the authentication server. This key is shared between the Mobility Conductor and the server. The maximum length is 128 characters.
Auth Port	The server port on the sever.
Interface for Source IP Address	The IP address of the RADIUS server. This option allows the user to configure the interface to reach the RADIUS server.

Navigate to the AAA Profile tab to add or edit AAA profiles. AAA profiles define the authentication profile and server and server groups you want to use to authenticate supplicants.

Server Group Fields

You can create groups of servers. If one server is not reachable based on the server retry count configured on the 802.1x/MAC tab, the appliance will switch to another server.

1. Click **Add** to add a new server group.

NOTE: To modify an existing server group, modify the existing data and click Save. To delete a server group, click the corresponding delete icon (X) in the last column.

2. Complete the following fields:

Field	Description
ID	The unique identifier of the server group.
Server Group Name	Enter a name for the server group.
Servers	The servers in the server group. Click the cell to select servers from the list.

Navigate to the AAA Profile tab to add or edit AAA profiles used to authenticate supplicants. AAA profiles define the authentication profile and server and server groups you want to use to authenticate supplicants.

AAA Profile

Use the AAA Profile tab to create profiles to map the 802.1x and MAC authentication profile to a server group you want to use to authenticate supplicants. This profile is used for dynamic authorization. For example, when a supplicant needs to be reauthenticated or the when the existing session is disconnected. After you create a AAA profile, you will assign that profile to an interface label.

You can edit an existing AAA profile or add a new AAA profile.

1. Click **Add** to add a AAA profile or click the pencil icon to edit an existing AAA profile.

   The Add or Edit AAA Profile dialog box opens.

2. Complete the following fields:

Field	Description
Profile	The name of the AAA profile.
DA Enable	Select this option to enable Dynamic Authorization functionality.
DA Server	Select the server to be used for Dynamic Authorization.
802.1x Auth Profile	Select the name of the 802.1x authentication profile.
802.1x Default Role	Select the default role assigned to the 802.1x clients.
802.1x Auth Server Group	Select the server group used for 802.1x authentication.

Field	Description
MAC Auth Profile	Select the name of the MAC authentication profile.
MAC Default Role	Select the role assigned to the client for MAC clients.
MAC Auth Server Group	Select the name of the server group used for MAC authentication.

3. Click **Add**.

Navigate to the Apply Policies tab to assign policies to interface labels.

Apply Policies

Use the Apply Policies tab to modify the policies that are assigned to each interface label. Supplicants plugged into the LAN port with the assigned interface label will be authenticated using the policy you select.

Each LAN interface label defined in your Orchestrator deployment is assigned the default authentication policy. The default authentication policy is set to "trusted." When authentication is set to "trusted", no authentication is required to access the network.

1. Click **Add** to add a policy.
2. Click in the cell to modify a LAN Label or AAA profile.

   NOTE: If AAA Profile is set to "none", the authentication type is automatically set to "trusted".
3. Click in the cell to modify the Auth Type.
   - **trusted**: Select **trusted** if no authentication is required.
   - **both**: Select **both** to first attempt 802.1x authentication and then fall back to MAC authentication.
   - **802.1x**: Select **802.1x** if the port only supports 802.1x authentication.
   - **mac**: Select **mac** if the port only supports MAC authentication.

Delete a Policy

To delete a policy from an LAN interface, click the corresponding delete icon (**X**) in the last column. The NAC security settings for this LAN interface will return to the default values.

Merge / Replace

At the top of the page, choose:

Merge to use the values in the template, but keep any values set on the appliance as is (producing a mix of template and appliance rules),

-OR-

Replace (recommended) to replace all values with those in the template.

Appliance IP Allow List Template

Configuration > Overlays & Security > Security > Appliance IP Allow List Template

The Appliance IP Allow List Template enables administrators to define and enforce IP-based access controls for management protocols such as SSH, HTTPS, and SNMP on EdgeConnect appliances. In Orchestrator version 9.6.0 and later, this feature enhances control plane security by restricting access to trusted IP addresses only.

Appliance IP Allow List Template Settings

For a target appliance selected in the appliance tree, the configuration options below are available.

Field	Description
Enable	Enable/Disable the IP Allow List feature. The status indicator displays whether the IP Allow List feature is enabled or disabled for the selected appliance.
IP/Mask	Specify a valid IPv4 Remote Address by setting the Mask which determines the range of IP addresses allowed. Possible values are 24 or 32. If the last octet of the IP address is zero, use a mask of 24. Use a mask of 32 for absolute IP addresses.
Authentication Type	Choose one or more protocols (SSH, HTTPS, SNMP) for which the IP entry applies.
Comments	Optionally annotate entries to clarify purpose and ownership for audit or documentation purposes.

Upon saving changes, the configuration tab auto-closes to confirm successful application.

Note the following conditions and guidelines.

- Ensure the IP Allow List feature is enabled and protocols are correctly assigned before relying on its enforcement. Entries may appear even when the feature is disabled.
- Adding a new entry via the CLI with different protocols will override existing protocol settings for that IP.
- Review entries periodically to maintain alignment with security policies.

Flow Export Template

You can configure your appliance to export statistical data to NetFlow and IPFIX collectors.

- The appliance exports flows against two virtual interfaces—**sp_lan** and **sp_wan**—that accumulate the total of LAN–side and WAN–side traffic, regardless of physical interface.
- These interfaces appear in SNMP and are, therefore, “discoverable” by NetFlow and IPFIX collectors.

- For more information about Custom Information Elements, see Flow Export Tab.
- **Enable Flow Exporting** allows the appliance to export the data to collectors (and makes the configuration fields accessible).
- The Collector's **IP Address** is the IP address of the device to which you are exporting the NetFlow/IPFIX statistics. The default Collector Port is **2055**.
- In **Traffic Type**, you can select as many of the traffic types as you want. The default is **WAN TX**.
- Click the **Anonymize IPs** check box to enable anonymizing IP addresses in log messages. If enabled, select an option from the **Bit Masking** drop-down menu to indicate how IP addresses have bit masking applied in log messages.

Firewall Protection Profiles Template

Use this template to enable baseline learning for appliances and to add or modify a protection profile on any appliance with a firewall and map it to a segment and zone of your firewall. For more information about firewall protection profile settings and baseline learning, see Firewall Protection Profiles.

NOTE: If you apply a Firewall Protection Profile template to an appliance that is running an older version of ECOS, certain DoS threshold settings are not applied if the software version does not support them.

Enable Baseline Learning

The following instructions describe how to enable baseline learning for appliances using a template.

NOTE: Baseline learning, Auto rate limit, and Smart burst all require either an AS (Advanced Security) license or an AAS-DTD (Dynamic Threat Defense) license.

1. Select the **Baseline learning** check box.
2. To customize the baseline learning settings, click **Baseline settings**.

   The Baseline Settings dialog box opens.
3. Enter the following information based on your network or click **Cancel** to use the default settings.

Field	Description
Data aggregation method	Technique used for data aggregation. The default is percentile and there are currently no other options.
Data aggregation limit	Indicates what percentage of the sample data is used to determine baseline values. The default setting is 95%, which means the top 5% of the sample is discarded and the other 95% is considered when computing the baselines. You can enter a value between 75-100%.

Field	Description
Computation interval	Time that passes before the system computes new baselines. The default is 8 hours. For example, when using the default, the baselines are computed every 8 hours using the latest sample data collected during the Model training interval. This can be configured in 4-hour units (e.g., 4, 8, 12, and so on) up to 240 hours.
Model training interval	During this period, data is collected for various metrics every five minutes and is aggregated into a data file. This data is used to compute the baselines. The default is 14 days, the minimum is 7 days, and the maximum is 56 days. **NOTE:** This period should include a diverse set of data that covers various types of legitimate traffic and captures the characteristics that distinguish normal traffic from malicious traffic during an attack.
Baseline upper limit	Upper limit for the minimum baseline. An alarm is raised when this value is reached. This setting is useful if Auto rate limit is configured without Smart burst. The setting is a percentage of the maximum baseline value, which is set manually. The default is 90%. You can enter a value between 50-100%.
TCP inactivity timeout	Inactivity timeout used for TCP flows created using burst support levels. Inactive flow gets deleted after this timeout. The default is 300 seconds. You can enter a value between 30-1800 seconds.
Headroom for baseline plus	Percentage of headroom that is added to the baseline. The default is 20%. You can enter a value between 5-100%.
Per-source limit for committed burst	Committed burst for a zone is available to all sources in the zone. This determines the percentage of committed burst in a zone that one source can use. The default is 50%. You can enter a value between 1-50%.
Reserve flow capacity distribution	Spare flow capacity is distributed among all zones by Smart burst using different methods (**Proportional** or **Equal**). The default method is Proportional.
Excess burst credit interval	On a per second basis, the zone is supposed to use a portion of committed burst capacity. Unused committed burst capacity of zones is made available as excess burst capacity every second. After this interval of time, unused excess burst capacity goes back to the respective committed burst. The default is 30 seconds. Enter a value between 30-100 seconds.
Minimum reserve capacity limit	Minimum amount of reserve flow capacity that should be available before Smart burst redistributes new reserve capacity after a baseline computation interval. Smart burst continues with previously distributed capacities if the minimum reserve capacity limit is not available. The default is 20%. You can enter a value between 10-50%.

4. Click **Done**.

Enable Internal IP Reputation

This section describes how to enable Internal IP Reputation for appliances using a template.

NOTE: Smart SYN cookie and dynamic IP reputation require either an AS (Advanced Security) license or an AAS-DTD (Dynamic Threat Defense) license.

1. Select the **Internal IP reputation** check box.
2. To customize the reputation settings, click **Reputation settings**.

   The Internal IP Reputation dialog box opens.
3. Enter the following information based on your network or click **Cancel** to use the default settings.

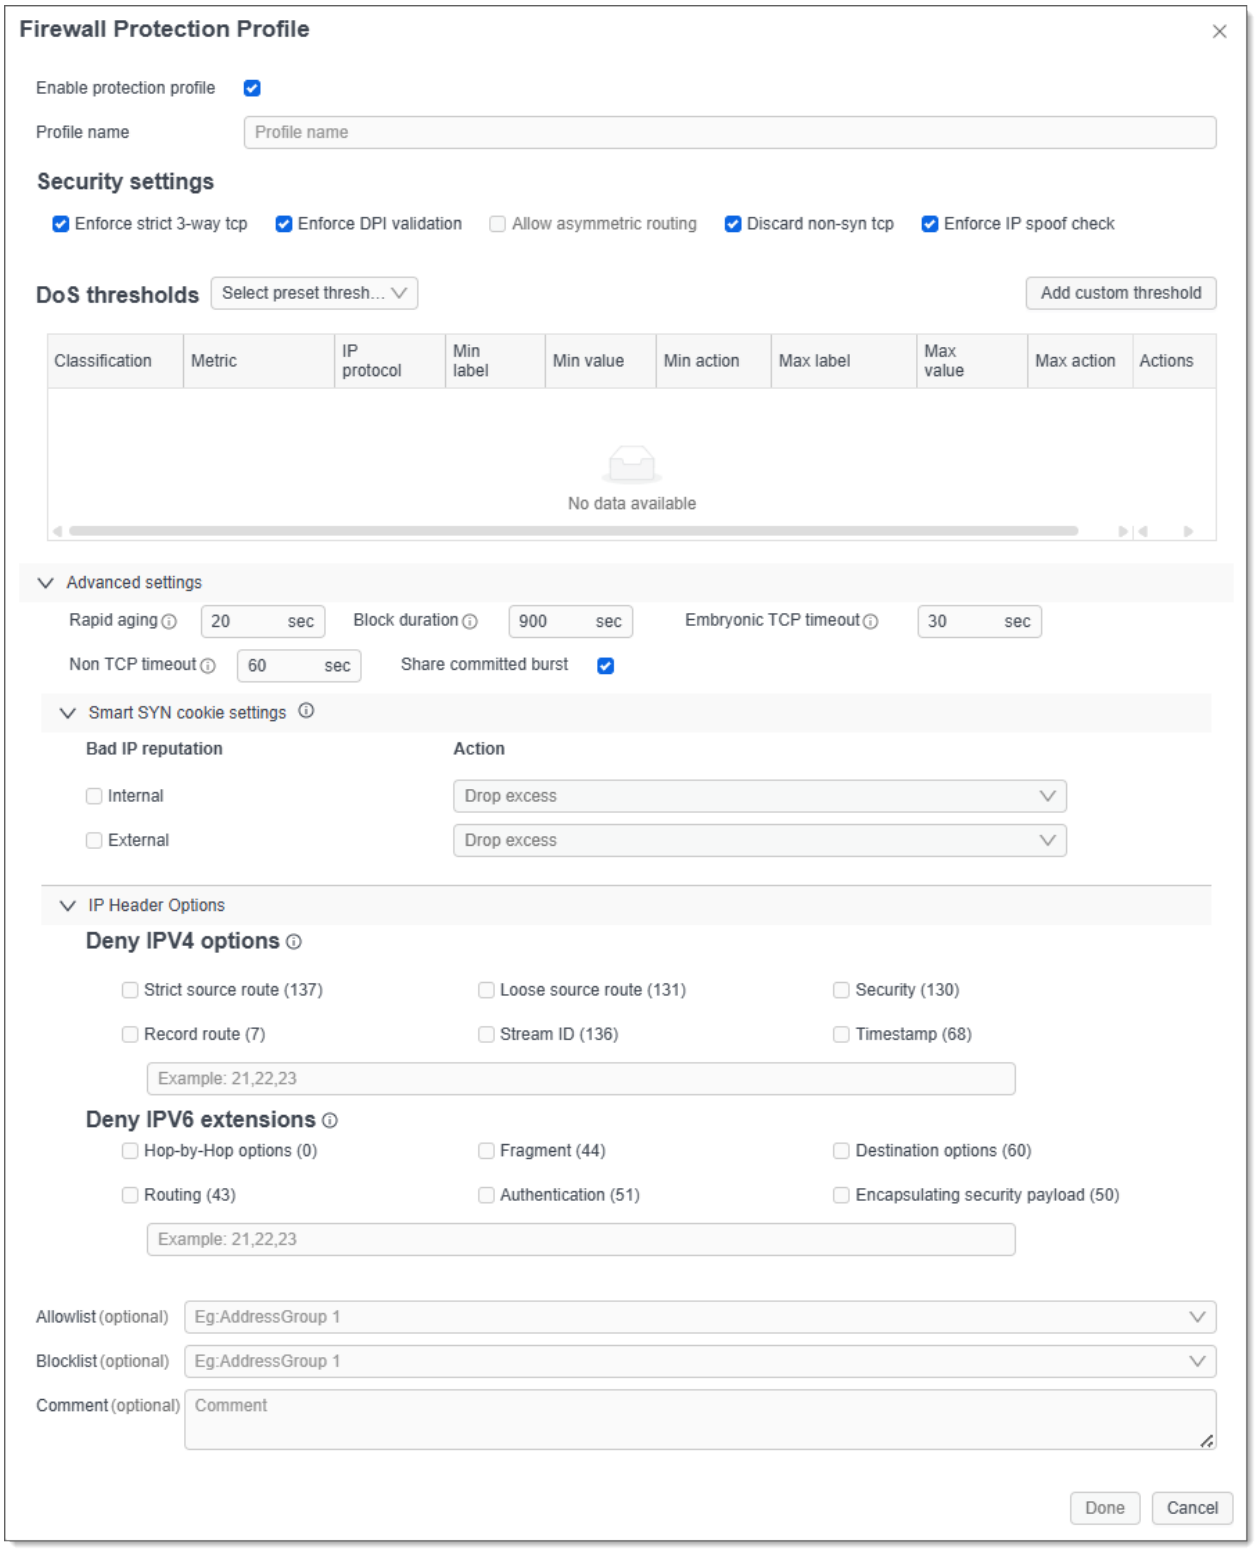
Field	Description
Successive Good Transactions (Trusted)	Number of consecutive good (trusted) transactions that an IP address must incur to be considered trusted. The default is 20 transactions. You can enter a value between 5-1024.
Successive Bad Transactions (Untrusted)	Number of consecutive bad (untrusted) transactions that an IP address can incur before it is considered untrusted. The default is 20 transactions. You can enter a value between 5-1024.

4. Click **Save** to close the Internal IP reputation dialog box.

Add New Profiles

1. In the Profiles section, click **Add**.

   The Firewall Protection Profile dialog box opens.



2. Enter a name for the profile.
3. Select or clear any of the Security settings check boxes.

NOTE: When asymmetric routing is configured, strict three-way TCP enforcement and deep packet inspection (DPI) validation cannot be performed. To enable these settings, turn off asymmetric routing.

4. From the DoS thresholds drop-down list, select a preset threshold (**Lenient**, **Moderate**, **Strict**, **Auto rate limit**, or **Smart burst**).

   NOTE: To clear your selection and reset the DoS thresholds table, choose **Select preset threshold** from the drop-down list.

   NOTE: To use Auto rate limit or Smart burst, you must enable baseline learning first. These options only appear in the drop-down list after baseline learning is enabled.

5. *(Optional)* Add exceptions to the following fields:

Field	Description
Allowlist	Enter an existing Address Group. Any IP address contained within the Address Group will be exempt from DoS threshold analysis. The Allowlist does not exempt flows from the options shown in the Security Settings section.
Blocklist	Enter an existing Address Group to explicitly block any IP address contained within the configured Address Group.

6. *(Optional)* Click **Advanced settings** and set the following fields:

Field	Description
Rapid aging	Set a threshold value (in seconds) to enforce the tearing down of TCP connections when the period of inactivity matches the configured value (for example, 30s).
Block duration	Enforce dynamic blocking of flows originating from a source for a specified duration (for example, 300s).
Embryonic TCP timeout	Set this value so that the firewall can tear down half-open TCP connections when the timeout value is reached (for example, 30s). While a TCP connection goes through the three-way handshake (SYN, ACK, SYN-ACK), an embryonic connection is a half-open connection that produces, for example, a SYN without the other two parts of the handshake. This is a popular form of denial of service (DoS) attack.
Non TCP timeout	Set this value so that the firewall can tear down unidirectional/half-open non-TCP connections when the timeout value is reached. Default is 60 seconds.
Share committed burst	Select this check box to enable unused committed burst to be shared with other zones. This check box is enabled by default. For critical zones, you can disable this option, which retains the committed burst capacity for the zone itself.

Field	Description
Smart SYN cookie settings	Bad IP Reputation **Internal**: Select the check box to enable, and then select the action to take if the internal IP reputation for an IP address that is requesting a TCP connection is Untrusted (Drop excess or Block source). Default is Drop excess. **External**: Select the check box to enable, and then select the action to take if the external IP reputation for an IP address that is requesting a TCP connection is Untrusted (Drop excess or Block source). Default is Drop excess.
Deny IPV4 options	For appliances running ECOS version 9.4.0.0 or later, select the check box for any of these items to deny their inclusion in the Options field of IPv4 packet headers: **Strict source route**, **Loose source route**, **Security**, **Record route**, **Stream ID**, and **Timestamp**. In the text box, enter any options using comma-separated format.
Deny IPV6 extensions	For appliances running ECOS version 9.4.0.0 or later, select the check box for any of these optional extension headers of IPv6 packets to deny their inclusion: **Hop-by-Hop options**, **Fragment**, **Destination options**, **Routing**, **Authentication**, and **Encapsulating security payload**. In the text box, enter any extension options using comma-separated format.

7. Click **Done** to close the Firewall Protection Profile dialog box.
8. Click **Save** to save your changes to the Firewall Protection Profile template.

Set Firewall Protection Profile Thresholds

NOTE: You can use both preset and custom thresholds.

To change the threshold settings:

1. Hover over the Actions column for the profile you want to edit and click **Edit**.

   The Firewall Protection Profile dialog box opens.
2. Either select a preset threshold from the DoS thresholds drop-down list, or click **Add Custom Threshold**.

   The DoS Threshold dialog box opens.

   NOTE: To edit the settings for an existing preset or custom threshold, hover over the Actions column for the threshold and click **Edit**.
3. Set the following parameters.

   NOTE: To return preset threshold settings to their defaults, click **Reset to default**, and then click **Reset** to confirm the change. This does not impact custom thresholds.

Field	Description
Classification	Classify flows in two ways: **Zone-level**: Flows originating from multiple endpoints that are part of a single firewall zone. **Source-level**: All flows originating from a single endpoint or source device.
Metric	DoS thresholds can be configured with any or all of the four metrics available in a firewall protection profile: **Flows per second**: Rate of flow (fps). A single flow is a unidirectional set of packets containing common attributes (source and destination IP, ports, protocols). **Concurrent Flows**: Number of flows that are active at a given point in time. **Embryonic Flows**: An incomplete TCP connection where only the first step in the three-way handshake has been completed; the EdgeConnect has received the SYN packet from the requesting client. For other protocols, such as UDP or ICMP or other, this indicates uni-directional traffic and the flow is called Embryonic. **Embryonic Two-Way flows**: A half-open TCP connection occurs when only the SYN and SYN-ACK packets of the three-way handshake are completed. In this state, EdgeConnect has received the SYN-ACK from the server, but it has not received the final ACK from the Client.
IP Protocol	Select an IP protocol (**TCP**, **UDP**, **ICMP**, **Others**, or **All**) for use in threshold settings.
Min Label	Select the method used to determine the min value: **Baseline**: If selected, the min value is determined by the system using baseline learning, and the corresponding Value field shows “Dynamic”.This option is available only if Baseline Learning is enabled. **Custom**: If selected, you configure the min value by entering a percentage in the corresponding Value field.
Value	Minimum threshold value as a percentage of target appliance flow capacity. When this value is breached, the protection profile takes a corresponding minimum action. If Baseline is selected as the Min Label, the system determines this value, and it cannot be configured.
Action	Action to take when the min value is breached (**Log**, **Rapid aging**, **Drop excess**, **Block source**, or **SYN cookie**). Because this corresponds to the min value, a less intensive action can be configured.

Field	Description
Max Label	Select the method used to determine the max value: **Custom**: If selected, you configure the max value by entering a percentage in the corresponding Value field. **Baseline plus**: A buffer of 20% is added to the computed baseline when determining flow capacity. If selected, the max value is determined by the system using baseline learning and the corresponding Value field shows “Dynamic”. **Committed burst**: Reserve flow capacity is divided equally or proportionally among all zones configured for Smart burst. If selected, the max value is determined by the system using baseline learning and the corresponding Value field shows “Dynamic”. **Excess burst**: Continuously, on a per second basis, unused committed burst (distributed reserve flow capacity) is collected from all zones and shared as a second level of support for all zones. If selected, the max value is determined by the system using baseline learning and the corresponding Value field shows “Dynamic”.
Value	Maximum threshold value as a percentage of target appliance flow capacity. When this value is breached, the protection profile takes a corresponding maximum action. If Baseline plus, Committed burst, or Excess burst are selected as the Max Label, the system determines this value, and it cannot be configured.
Action	Action to take when the max value is breached (**Log**, **Rapid aging**, **Drop excess**, **Block source**, **SYN cookie**, or **Smart SYN cookie**). Because this corresponds to the max value, a more intensive action can be configured.

4. Click **Done** to close the DoS Threshold dialog box.
5. Click **Save** to save your changes to the Firewall Protection Profile template.

Add Profile Mappings

After you create a profile, you can map it to a segment and zone of your firewall to achieve the expected behavior.

To map a profile to a segment:

1. In the Profile Mappings section, click **Add**.
2. From the Segment drop-down list, select the segment you want to map to your profile. To search for a segment in the list, type the segment name.
3. From the Zone drop-down list, select the zone you want to assign to your profile. To search for a zone in the list, type the zone name.

4. From the Profile drop-down list, select the profile you created earlier. To search for the profile in the list, type the profile name.
5. Click **Done** to close the Profile Mapping dialog box.
6. Click **Save** to save your changes to the Firewall Protection Profile template.

Logging Template

Use this template to configure local and remote logging parameters. Each requires that you specify the minimum severity level of event to log.

WARNING: Appliance logging levels should only be set to "Notice" unless TAC asks you to set it differently. This applies to both the Minimum severity level field in the Log Configuration area of this template and the Minimum Severity field in the Remote Log Receivers area. Be aware that setting this level to "Debug" will generate logs for all modules that are turned on, which causes the packet processing engine to spend excessive time logging instead of forwarding packets.

- Set up local logging in the **Log Configuration** and **Log Facilities Configuration** sections.
  - Click the **Anonymize IPs** check box to enable anonymizing IP addresses in log messages. If enabled, select an option from the **Bit Masking** drop-down menu to indicate how IP addresses have bit masking applied in log messages.
  - Click the **Jsonify** check box to convert firewall policy and IDS/IPS log messages to JSON when exported.
  - Uniquely assign log facilities for System, Audit, Firewall, and IDS/IPS Events; they cannot overlap. For example, System can be assigned to local2 and Audit to local3, but both cannot be assigned to local2.
- Set up remote logging by using the **Remote Log Receivers** section.

For detailed information on logging, see Logging Tab

Minimum Severity Levels

In decreasing order of severity, the levels are as follows. See the WARNING note above.

Severity Level	Description
EMERGENCY	System is unusable.
ALERT	Includes all alarms the appliance generates: **CRITICAL**, **MAJOR**, **MINOR**, and **WARNING**.
CRITICAL	Critical event.
ERROR	An error. This is a non-urgent failure.
WARNING	A warning condition. Indicates an error will occur if action is not taken.
NOTICE	A normal, but significant, condition. No immediate action required.

Severity Level	Description
INFO	Informational. Used by Support for debugging.
DEBUG	Used by Support for debugging.
NONE	If you select **NONE**, no events are logged.

- If **NOTICE** is selected (the default setting), the log records any event with a severity level of NOTICE, WARNING, ERROR, CRITICAL, ALERT, and EMERGENCY.
- These are related to event logging levels, not alarm severities, even though some naming conventions overlap. Events and alarms have different sources. Alarms, after they clear, list as the ALERT level in the Event Log.
- In the Log Facilities Configuration section, assign each message/event type (System / Audit / Firewall / IDS/IPS) to a syslog facility level (**local0** to **local7**).

Configure Remote Logging

You can configure the appliance to forward all events, at and above a specified severity, to a remote syslog server.

A syslog server is independently configured for the minimum severity level that it will accept. Without reconfiguring, it might not accept as low a severity level as you are forwarding to it.

To configure remote logging:

1. Under Remote Log Receivers, click **Add**.
2. For each remote syslog server that you add to receive the events, complete the following fields with the appropriate information.

Field	Description
Host Name	Host's IP address.
Port	Port number of the remote syslog server. Valid values range from 2 through 65535.
Protocol	Select the protocol you want to apply: **UDP**, **TCP**, or **TCP SSL**.
Minimum Severity	Select the minimum severity level of messages you want to log (see the WARNING message above): **None**, **Emergency**, **Alert**, **Critical**, **Error**, **Warning**, **Notice**, **Info**, or **Debug**.
Facility	Select **all**, **local1**, **local2**, **local3**, **local4**, **local5**, **local6**, or **local7**.

Field	Description
Client Certificate	If you selected TCP SSL protocol, do one of the following: Complete the instructions below for adding a client certificate using an orchestrated appliance end entity profile. This is the recommended method. Complete the instructions below for adding a client certificate using the legacy method of uploading a certificate and key files. Click **View** to view the client certificate. Click **Don't Apply** if you do not want to apply a client certificate.
Verify	Click this cell to display a check box, and then select it to verify the server certificate.

Add a Client Certificate

You can add a client certificate in two ways. The recommended method is to use an orchestrated appliance end entity profile. The legacy method is to upload a certificate and key files.

To add a client certificate using an appliance end entity profile:

1. In the Client Certificate column, click **Add**.

   The Add Remote Receiver SSL Certificate dialog box opens.
2. Click **Use End Entity Certificate**.
3. Select the end entity certificate profile from the menu.
4. Click **Add**.

To add a client certificate using the legacy method of uploading a certificate and key files:

1. In the Client Certificate column, click **Add**.

   The Add Remote Receiver SSL Certificate dialog box opens.
2. Click **Upload Certificate**.
3. Complete the following fields.

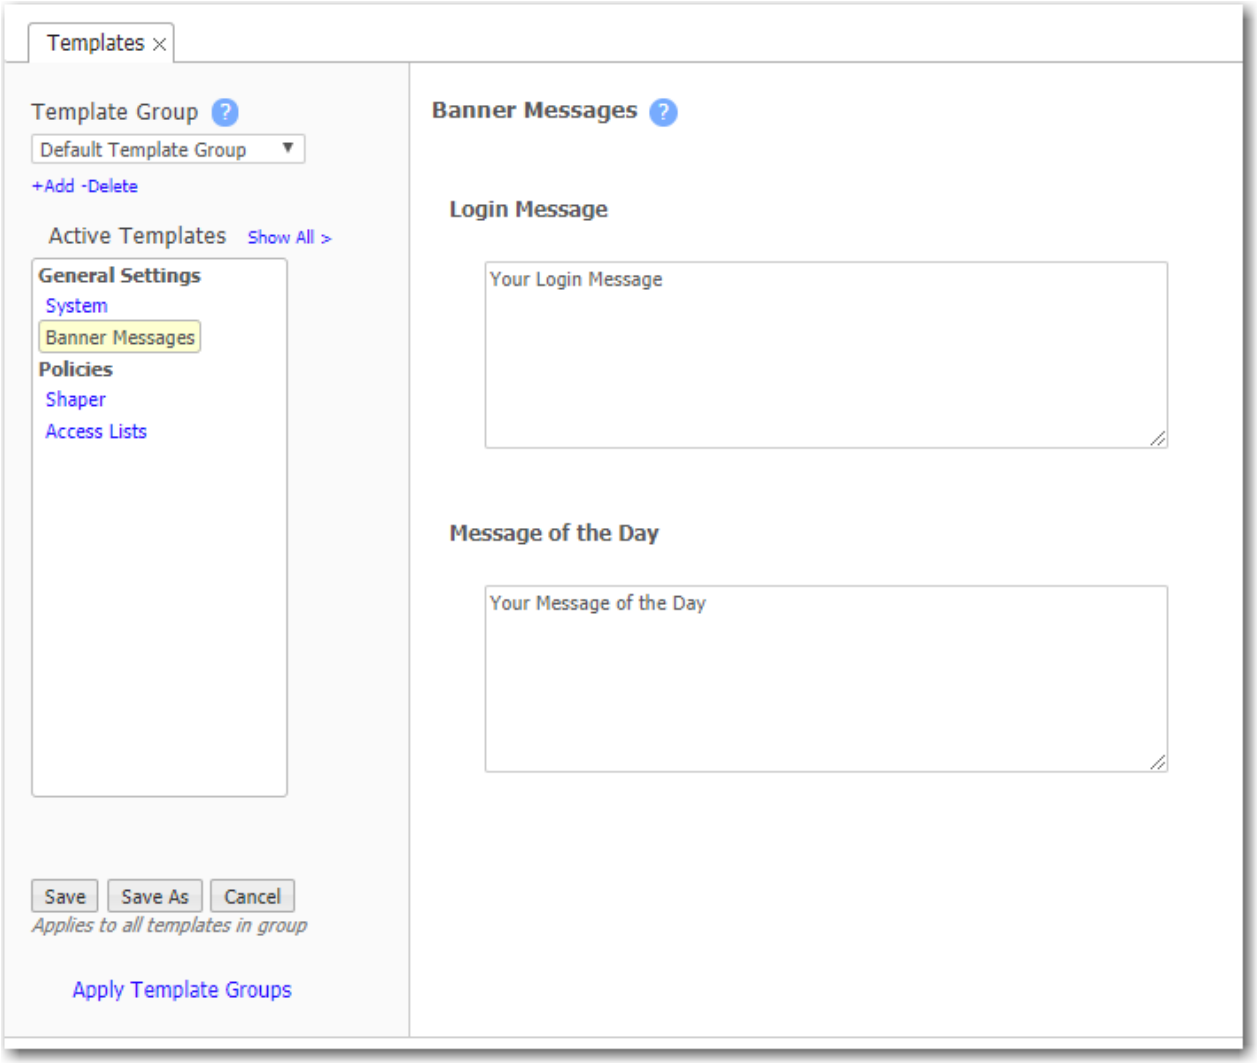
Field	Description
PFX Certificate File	To use a PFX certificate file, select this check box.
Certificate File	Click **Choose File**. Locate and select the certificate file, and then click **Open**.

Field	Description
Private Key File	Click **Choose File**. Locate and select the private key file, and then click **Open**. If you selected PFX Certificate File, this field is disabled.
Import Password	Enter the import password for the certificate.
Passphrase	Enter the passphrase for the certificate.

4. Click **Add**.

Banner Messages Template

- The **Login Message** appears before the login prompt.
- The **Message of the Day** appears after a successful login.



HTTPS Certificate Template

On this template you select the type of certificate to use for EdgeConnect appliances within a template group. There are three options to set up an HTTPS certificate for appliances using templates.

- End Entity Certificate: This is the recommended option. It automates certificate enrollment using an EST server and globally orchestrated end entity profiles if the profile Purpose is set to "TLS Server".
- Custom Certificate: This is a legacy option. You install your own custom certificate from a CA certificate authority.
- Self Signed Certificate: This is the default option. Browsers will not show this as secure, and most IT departments will now allow this. If your enterprise intends to use the EdgeConnect web UI directly, you need to use one of the other options.

To use an end entity certificate (recommended):

Before configuring the HTTPS template to use an end entity certificate, you must first add an EST server profile and create an appliance end entity profile with a Purpose of "TLS Server". If you have not created an appliance end entity profile, see End Entity Certificates Tab. After you have completed the profile, proceed with the following steps.

1. Navigate to **Configuration > Templates & Policies> Setup > Templates**.
2. Select the template group.
3. If the HTTPS Certificate template is not active, click **Show All** and drag the template to the Active Templates column.
4. Click **End Entity Certificate**.
5. Select an appliance end entity profile from the drop-down menu.
6. Click **Save** to apply the template changes to the template group.

To use a custom certificate (legacy method):

1. Consult with your IT security team to generate a certificate signing request (CSR) and submit it to your organization's chosen SSL Certificate Authority (CA).
   - Examples of Certificate Authorities include GoDaddy, Verisign, Comodo, Symantec, Microsoft Entrust, GeoTrust, and so forth.
   - All certificate and key files must be in **PEM** format.
2. After the Certificate Authority provides a CA-verified certificate, navigate to **Configuration > Templates & Policies> Setup > Templates**.
3. Select the template group.
4. If the HTTPS Certificate template is not active, click **Show All** and drag the template to the Active Templates column.
5. Click **Custom Certificate**, and then click **Upload and Replace**.

   The Add HTTPS Certificate dialog box appears.

6. If your IT security team advises the use of an Intermediate CA, upload an Intermediate Certificate File. Otherwise, skip this file.
7. Upload the Certificate File from the CA.
8. Upload the Private Key File that was generated as part of the CSR.
9. Click **Add** to close the Add HTTPS Certificate dialog box.
10. Click **Save** to apply the template changes to the template group.

Password Settings Template

Use the Password Settings template to manage password criteria and enforcement for all locally configured appliance users. When new users are created, password validation ensures that enterprise security standards are enforced.

NOTE: You can only apply the Password Settings template to appliances running version 9.6.0 and later.

NOTE: If the enforcePasswordCriteria field on the Orchestrator Advanced Properties is set to true, all properties in the following table are enabled. If the enforcePasswordCriteria field on the Orchestrator Advanced Properties is set to false, this template will not be applied.

WARNING: The Password Settings template enables password enforcement and user inactivity controls on appliances. You cannot disable password enforcement on the appliance after this template is applied.

Field	Description
Password Minimum Length	Minimum number of characters for the password. Allowable range: 8-64. The default is 8.
Use Lowercase Letter in Password	Select this check box to require a lowercase letter in the password.
Use Uppercase Letter in Password	Select this check box to require an uppercase letter in the password.
Use Digit in Password	Select this check box to require a number in the password.
Use Special Symbol in Password	Select this check box to require a special symbol in the password.
Password Expire Days	Number of days before the password expires. Allowable range: 30-360. The default is 90.
Notify Password Expiry Days	Number of days before the password expires to notify the user of the upcoming expiration. The user is notified via a popup. Allowable range: 5-30. The default is 7.

Field	Description
Password History Count	Number of previous passwords to store for the user to ensure passwords are not repeated. Allowable range: 4-8. The default is 8.
User Inactive Days	Automatic deletion of the user account occurs based on the configured number of inactive days that have transpired. Enter the number of inactive days (30 - 360) that will trigger the user account to be deleted. Enter 0 to disable automatic user account deletion. The default is 0. The only exception to this policy is that the user “admin” will never be deleted.

Use the following CLI commands to set the threshold on the number of failed password attempts and the lockout duration if the number of failed password attempts is exceeded for appliance users:

`aaa auth-failure deny n` where `n` is the threshold number of failed attempts

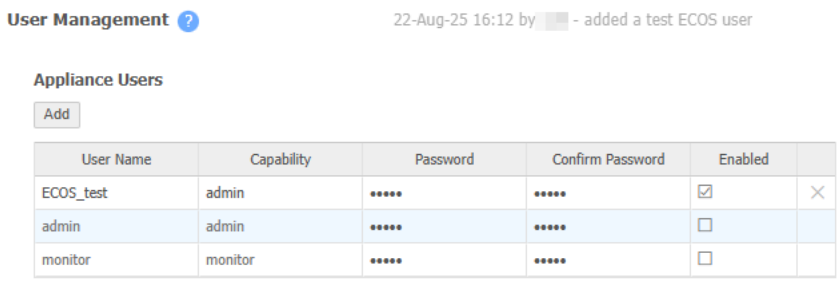
`aaa auth-failure unlock-time n` where `n` is the number of seconds that the account will be locked

User Management Template

Use this tab to manage EdgeConnect local users and, if desired, require a password with the highest user privilege level when using the Command Line Interface.

ECOS Local User Accounts

- Each appliance has two default user accounts, **admin** and **monitor**.
  - Admin users have read/write privileges.
  - Monitor users have read-only privileges.
- Passwords for these users are controlled from this template.
- You can enable or disable users by selecting or clearing the check box in the Enabled column.
- Click **Add** to create new users (for example, the ECOS_test user in the following screen capture).



Command Line Interface Privileges

- The Command Line Interface (CLI) for physical EdgeConnect appliances has three command modes. In order of increasing permissions, they are User EXEC Mode, Privileged EXEC Mode, and Global Configuration Mode.
- When you first log in to an EdgeConnect appliance via a console port, you are in User EXEC Mode. This provides access to commands for many non-configuration tasks, such as checking the appliance status.
- To access the next level, Privileged EXEC Mode, enter the `enable` command. With this template, you can choose to associate and enforce a password with the `enable` command.



DNS Template

A **Domain Name Server** (DNS) stores the IP addresses with their associated domain names. It enables you to reference locations by domain name, such as *mycompany.com*, instead of using the routable IP address.

- You can configure up to three name servers.
- Under **Domain Names**, add the network domains to which your appliances belong.

Date/Time Setting

Configure an appliance's **date and time** manually, or complete the following steps to configure it to use an NTP (Network Time Protocol) server.

1. From the Time Zone list, select the appliance's geographical location.

2. If you select Manual, the appliance is matched to your web client system when the template is applied. This eliminates the delay between configuring time manually and applying the template.
3. To use an NTP server, select **NTP Time Synchronization** and complete the following steps.
    1. Click **Add**.
    2. Enter the IP address or host name of the server.
    3. Select the version of NTP protocol to use.

NOTE: The server is selected in the order listed when you list more than one NTP server.

Data Collection

- Orchestrator collects and puts all statistics in its own database in Coordinated Universal Time (UTC).
- When a user views statistics, the appliance (or Orchestrator server) returning the statistics always presents the information relative to the browser time zone.

SNMP Template

EdgeConnect appliances support Management Information Base (MIB-II) as described in RFC 1213 for cold start traps, warm start traps, and EdgeConnect private MIBs. Appliances issue an SNMP trap during reset when loading a new image, recovering from a crash, or rebooting.

An appliance sends a trap whenever an alarm is raised or cleared. Traps contain additional information about alarms, including severity, sequence number, a text-based description of the alarm, and the time the alarm was created. To download a .zip file of standard and proprietary MIB files you can load into the MIB browser you use, see MIB Files.

Use the SNMP template to configure the SNMP agent and trap receivers.

1. To enable SNMP on the appliance, select the **Enable SNMP** check box. Additional areas on the template appear (SNMP V1/V2, SNMP V3, and Trap Receivers).
2. To enable the SNMP agent on the appliance to send traps to configured receivers, select the **Enable SNMP Traps** check box.
3. In the **Default Trap Community** field, specify the string the trap receiver should use to accept traps sent to it. The default value is “public”. You can modify this value.
4. Configure SNMP v1/v2, SNMP v3, and trap receivers, as described below.

SNMP V1/V2

Configure the following fields for SNMP v1 and v2c.

Field	Description
Enable SNMP	Enables the SNMP agent on the appliance to send traps to configured receivers.
Read-Only Community	The SNMP application needs to present this text string (secret) to poll the appliance's SNMP agent. The default value is "public". You can modify this value.

SNMP v3

For additional security, configure SNMP v3 if you want to authenticate without using clear text. To add an SNMP v3 user, click **Add** above the SNMP V3 table, and then complete the following fields.

Field	Description
Enabled	Select this check box to enable the selected user. Clear it to disable the user and maintain the configuration.
Username	Enter the username to identify the SNMP v3 user.
Authentication Type	Select the authentication type to use for SNMP requests from the user. **NOTE:** Authentication type is required. SHA-1 is the only supported algorithm.
Authentication Password	Enter a password that the SNMP agent can use to authenticate requests sent by the user. **NOTE:** The password must have at least 20 characters.
Privacy Type	Select the encryption type to use for encrypting requests from the SNMP user. **NOTE:** Encryption is required. AES-128 is the only supported algorithm.
Privacy Password	Enter a password (key) to use for encrypting requests sent by the user. **NOTE:** The password must have at least 20 characters.

To remove an SNMP v3 user, click the corresponding delete icon (**X**) in the last column.

Trap Receivers

To configure a trap receiver, click **Add** above the Trap Receivers table, and then complete the following fields.

NOTE: You can configure up to three trap receivers per appliance.

Field	Description
Host	IP address of the host to which to send traps.
Version	Select the SNMP version of the trap receiver.
Community/Username	For v1 and v2c, enter the community string the receiver should use to accept traps. If left blank, the default community string (public) is used. If a different community string is configured on the trap receiver, enter it in this field. For v3, specify the SNMP v3 user sending traps to the receiver.
Enabled	Select this check box to enable the receiver. Clear it to disable the receiver and maintain the configuration.

To remove a receiver, click the corresponding delete icon (**X**) in the last column.

SSL WAN Optimization Certificates Template

Use this template to provision SSL WAN Optimization certificates when the server is __*part of your enterprise network__* and has its own enterprise SSL certificates and key pairs.

EdgeConnect provides deduplication for Secure Sockets Layer (SSL) encrypted WAN traffic by supporting the use of SSL WAN Optimization certificates and other keys:

- EdgeConnect decrypts SSL data using the configured certificates and keys, optimizes the data, and transmits data over an IPSec tunnel. The peer EdgeConnect appliance uses configured SSL WAN Optimization certificates to re-encrypt data before transmitting.
- Peers that exchange and optimize SSL traffic must use the same certificate and key.
- Use this template to provision a certificate and its associated key across multiple appliances.
  - You can add either a PFX certificate (generally, for Microsoft servers) or a PEM certificate.
  - The default is PEM when PFX Certificate File is deselected.
  - If the key file has an encrypted key, enter the passphrase needed to decrypt it.
- Before installing the certificates, you must do the following:
  - Configure the tunnels bilaterally for IPSec (or IPSec UDP) mode. To do so, navigate to **Configuration > Networking > Tunnels > Tunnels**, select the tunnel, and set the Mode field to **IPSec**.
  - Verify that TCP acceleration and SSL acceleration are enabled. To do so, navigate to **Configuration > Templates & Policies > Policies > Optimization Policies**, and then review the Set Actions.

- If you choose to be able to decrypt the flow, optimize it, and send it in the clear between appliances, access the **System** template and select **SSL optimization for non-IPSec tunnels**.

TIP: For a historical matrix of EdgeConnect and Orchestrator security algorithms, view the Security Algorithms PDF.

SSL WAN Optimization CA Certificates Template

If the enterprise certificate you used for signing substitute certificates is subordinate to higher-level Certificate Authorities (CA), you must add those CA certificates here. If the browser cannot validate up the chain to the root CA, it will warn you that it cannot trust the certificate.

TIP: For a historical matrix of EdgeConnect and Orchestrator security algorithms, view the Security Algorithms PDF.

Multicast Template

Use the Multicast template to configure multicast settings across multiple appliances in the SD-WAN fabric.

NOTE: The Multicast template only configures segments. You cannot configure Hello interval, Join/prune interval, or DR priority using the Multicast template. You must configure these settings on an appliance level using the Multicast feature. See Multicast.

To add a segment to the template:

1. Click **Add**.
2. Select a segment from the drop-down list. Only segments not already in the grid are available.
3. From the **Enable** drop-down list, select one of the following:
   - **Do not apply**—(Default) Does not configure multicast for this segment. Rendezvous point IP and Allowed group fields are disabled.
   - **Yes**—Enables multicast on the segment. Rendezvous point IP becomes a required field.
   - **No**—Explicitly disables multicast on the segment.
4. If you selected Yes, enter a valid Rendezvous point IP and optionally an Allowed group.

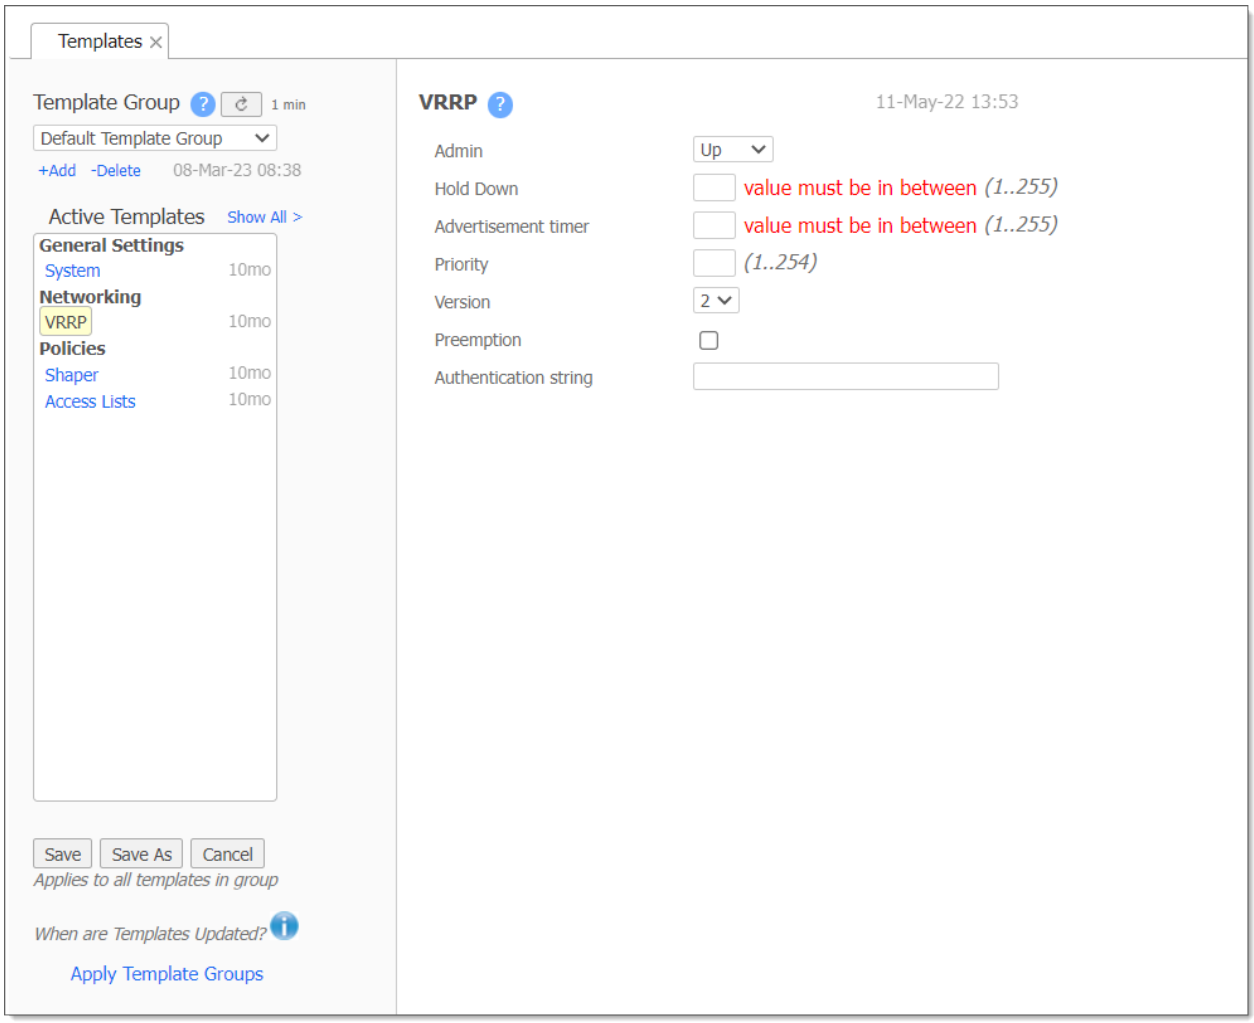
   NOTE: The Rendezvous point IP must be a valid IPv4 address. Allowed group must be alphanumeric characters, dots, hyphens, and underscores (up to 64 characters). No more than eight segments can be enabled.
5. Click **Done** to add the segment to the grid.

To edit or delete a segment, use the **Actions** menu on the corresponding row.

VRRP Template

Use this template to distribute common parameters for appliances deployed with Virtual Router Redundancy Protocol (VRRP).

VRRP specifies an election protocol that dynamically assigns responsibility for a virtual router to one of the VRRP routers on a LAN. The VRRP router controlling the IP address(es) associated with a virtual router is called the Master and forwards packets sent to these IP addresses. The election process provides dynamic failover in the forwarding responsibility should the Master become unavailable. This allows end hosts to use any virtual router IP addresses on the LAN as the default first-hop router. The advantage gained from using VRRP is a higher availability default path without configuring dynamic routing protocols such as BGP or OSPF.



Field	Description
Admin	Select **up** (enable) or **down** (disable).

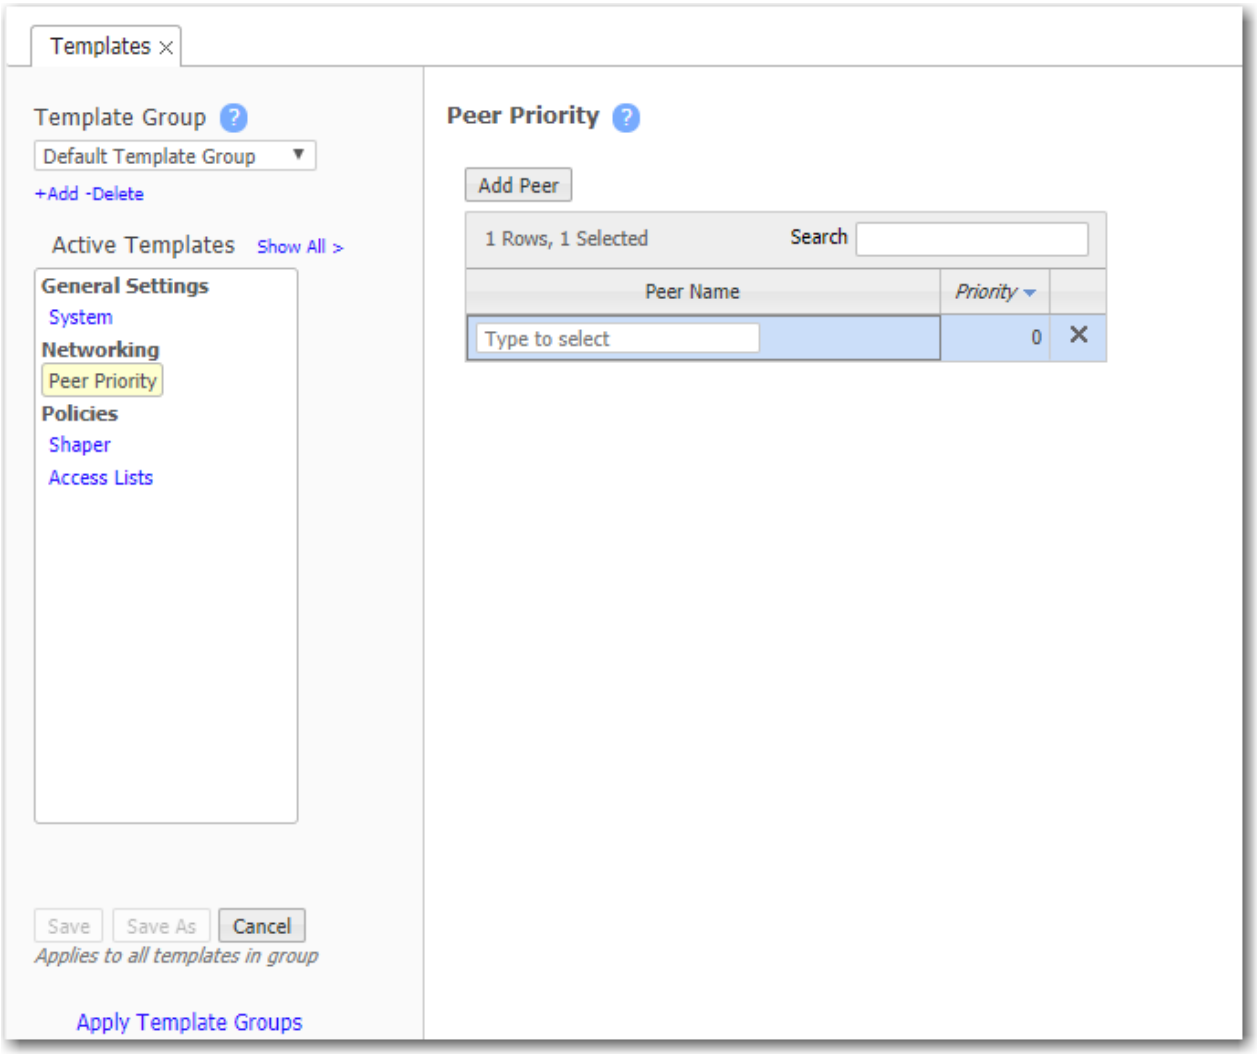
Field	Description
Hold Down	The number of seconds a higher-priority backup router that has just started up waits before preempting the primary router. It is best practice to configure a hold time so that routing protocols converge and tunnels come back up before preemption occurs. The default value is **60 sec** to account for the default *quiescent tunnel keep alive time*. The minimum value is 1 second.
Advertisement Timer	The time interval between sent advertisements. For version 2, the time is measured in seconds, and the default is **1 sec**. For version 3, the time is measured in centi-seconds, and the default is **1 centi-sec**. 1 centi-sec = 10ms.
Priority	The greater the number, the higher the priority. The appliance with the higher priority is the VRRP Master.
Version	Select the VRRP version that applies for your system: **2** – Supports only IPv4. **3** – Supports IPv4 and IPv6; does not support authentication strings.
Preemption	Leave this selected/enabled so that after a failure, the appliance with the highest priority comes back online and again assumes primary responsibility.
Authentication String	Clear text password for authenticating VRRP version 2 group members. You cannot use an authentication string if you are using VRRP version 3.

Peer Priority Template

When an appliance receives a **Subnet** with the same **Metric** from multiple remote/peer appliances, it uses the Peer Priority list as a tie-breaker.

- If a **Peer Priority** is not configured, the appliance randomly distributes flows among multiple peers.
- The lower the number, the higher the peer's priority.

NOTE: This feature requires appliance software 8.3.3.0 or higher for version 8 releases, and requires 9.0.2.0 or higher for version 9 releases.



Route Redistribution Maps Template

Before you use the Route Redistribution Maps template, ensure that your route maps are configured for either SD-WAN, BGP, or OSPF. For information about configuring route maps and defining rules for them, see Routes Tab.

To set up the template:

1. From the **Redistribute Routes To** drop-down list, select the appropriate target (direction of traffic) to which to redistribute your routes (**SD-WAN Fabric**, **BGP Inbound**, **BGP Outbound**, or **OSPF**).
2. Select one of the following:
   - **Merge** – Adds new maps to your existing maps. If the map already exists, the new map matches appliance rules in the Orchestrator range. If the configured rules do not match, the new map's rules are appended to the existing rules.

- **Replace** – Replaces all existing maps with the new maps. It does not include rules that match outside of the configured range.

The range for Orchestrator rules is 1000–9999.

WARNING: If you select **Merge**, only rules with matching priorities are overwritten. If you select **Replace**, all maps on the appliance are deleted and replaced.

3. After making your selection, click **Add Map**.

   NOTE: You can use the other buttons to add the default route redistribution map to the Map Name list or to delete, rename, or clone a map selected in the Map Name field.

   The Create Route Redistribution Map dialog box opens.

4. Enter a map name in **Map Name**, and then click **Add**.

   The name of your map now appears in the Map Name drop-down list.

   NOTE: You can create up to 14 customized route redistribution maps for each type listed in the Redistribute Routes To drop-down list.

5. Ensure that the appropriate map is selected in the Map Name field, and then click **Add Rule**.

   The Add Rule dialog box opens.

   Use this dialog box to define the rules applied to your route map, including Match Criteria and Set Actions. Each route map has a **match** command and a **set** command. A route map rule is used to check whether all attribute values specified in the match criteria match the corresponding attribute values in the route and declares “match” or “no-match”.

   NOTE: You can apply 128 rules per map.

6. Click **Add**.

Routes Template

Use the following settings to apply subnet sharing configuration to appliances associated with this template group. *Subnet sharing* is the protocol used to exchange routes between EdgeConnect appliances across the SD-WAN fabric.

- **Automatically advertise local LAN subnets**: The appliance will advertise LAN and virtual interface subnets to SD-WAN fabric peers.
- **Automatically advertised local WAN subnets**: The appliance will advertise WAN interface subnets to SD-WAN fabric peers.

Enter specific values for the following:

Field	Description
Use SD WAN Fabric Learned Routes	When selected, this enables appliances to receive routes via subnet sharing. When this is disabled, the subnet sharing protocol is disabled, and appliances no longer receive routes via subnet sharing.
Enable Equal Cost Multi Path (ECMP)	When selected, this enables ECMP routing support for all segments.
Metric for automatically added routes	50 (default value).
Route map name to redistribute routes to SD-WAN fabric	Name of the route map being redistributed to the SD-WAN.
Include BGP local ASN to routes sent to SD-WAN fabric	Select **Don't apply**, **Yes**, or **No**.
Filter routes from SD-WAN fabric with matching local ASN	Select **Don't apply**, **Yes**, or **No**.
Tag BGP communities to routes	Send the specified communities with routes that are advertised to both SD-WAN fabric peers and BGP peers, if the routes are learned from any of the following source protocols: Local/Static SD-WAN (Local/Static) SD-WAN (BGP) SD-WAN (OSPF) Select **Don't apply**, **Yes**, or **No**. If you select **Yes**, enter the BGP communities you want to be tagged in the field. **NOTE:** A community must be a combination of two numbers (0 to 65535) separated by a colon. For multiple communities, use a comma to separate them.

NOTE: If you select **Don't apply**, Orchestrator ignores this field when applying this template to appliances.

BGP Template

Use the BGP template to apply BGP configurations per segment to all appliances in the SD-WAN fabric.

1. Under the Common settings for all segments header, complete the following fields:

Field	Description
Max route updates per peer	The maximum number of route updates for each peer running route loop detection. The default value is 10, and the range is 5–100.
Detection interval	The interval, in minutes, at which route advertisement loop detection runs. A route advertisement loop occurs when the same route is being advertised, removed, and re-advertised repeatedly within a short time period. The default value is 15, and the range is 1–60.

2. Click the edit icon next to the segment for which you want to modify the configuration.
3. Configure the following elements as needed:

Field	Description
AS Path Propagate	Select **Yes** to enable this appliance to send the full AS path associated with a prefix to other routers and appliances, avoiding routing loops. This will provide the learned path from an external prepend between a remote BGP site to local BGP peers.
Graceful Restart	Select **Yes** to enable receiver-side graceful restart capability. EdgeConnect retains routes learned from the peer and continues to use it for forwarding (if possible) if/when a BGP peer goes down. Retained routes are considered stale routes. They will be deleted and replaced when new routes are received.
Max Restart Time	If Graceful Restart is enabled, specifies the maximum time in seconds to wait for a capable peer to come back after a restart or peer session failure.
Stale Path Time	If Graceful Restart is enabled, specifies the maximum time in seconds following a peer restart before removing stale routes associated with a peer.
Log BGP update messages	Enable logging of BGP peer messages on the segment.
Next-Hop-Self	Advertised route connected to a CE router that an EdgeConnect appliance learns from a PE router.

Field	Description
Override ASN	Indicates whether routes are advertised to the BGP peer where the BGP peer's own ASN is in the AS-Path.
Keep Alive Timer	This is the interval, in seconds, between keep alive signals to a peer.
Hold Timer	When availability to a peer is lost, this value specifies how long to wait before dropping the session.
Soft Reconfiguration	Select **Yes** to prevent the appliance from sending a route-refresh message to the BGP peer when a policy is changed. If you select **Don't apply**, Orchestrator ignores this field when applying this template to appliances. When enabled, the appliance applies policy changes against BGP peer learned routes stored in memory. **NOTE:** To request a route update from the peer, click the **Soft Reset** button for the peer on the BGP tab. Before you perform a soft reset, ensure that Soft Reconfiguration is disabled for the BGP peer.
Enable MD5 Password	If applied, adds a password to authenticate TCP sessions with peers.
Password / Confirm Password	If the MD5 password is enabled, use these fields to specify the password.

4. Click **Update**.

OSPF Template

Use the OSPF template to apply OSPF configurations per segment to all appliances in the SD-WAN fabric.

1. Click the edit icon next to the segment for which you want to modify the configuration.
2. Configure the following elements as needed:

Field	Description
Enable	Indicates whether the segment can access OSPF protocol. If you select **Don't apply**, Orchestrator ignores this field when applying this template to appliances.

Field	Description
Route Map name to Redistribute routes to OSPF	Name of the route map being redistributed to the SD-WAN. The OSPF template is used in conjunction with the Route Redistribution Maps template. OSPF route maps are configured in the Route Redistribution Maps template, and then applied in the OSPF template. The default OSPF route map name is “default_rtmap_to_ospf”. **NOTE:** Leave this field blank to preserve the current setting on the appliance.
Opaque LSA support	Enable for acknowledgment of opaque LSAs through OSPF protocol.
Admin Status	Indicates whether the interface admin status is up or down. If you select **Don’t apply**, Orchestrator ignores this field when applying this template to appliances.
Hello Interval	Length of time (in seconds) that must transpire between hello packets that a router sends on an OSPF interface.
Dead Interval	Length of time (in seconds) that must transpire before neighbors that have not detected a router’s hello packets can declare the OSPF router down.
Transmit Delay	Length of time (in seconds) that must transpire before transmitting a link state update packet. Specify a value from 1 to 65535.
Retransmit Interval	Length of time (in seconds) that a router that has received no acknowledgment must wait before resending transmissions.

Field	Description
Authentication Type	Type of authentication to use for requests. Select one of the following drop-down list options: **Don't apply** – Orchestrator ignores this field when applying this template to appliances. **None** – Authentication not performed. **Text** – Simple password authentication, which allows a key (password) to be configured per area. **MD5** – Message Digest cryptographic authentication. A key ID and key (password) are configured on each router. The router uses an algorithm based on the OSPF packet, the key ID, and the key to generate a message digest that gets appended to the packet.
Authentication Key	Key (password) to use for authentication of requests. This field is available only if Authentication Type is set to Text.
MD5 Key	Key ID to use for MD5 authentication of requests. This field is available only if Authentication Type is set to MD5.
MD5 Password / MD5 Confirm Password	Password for the MD5 key. These fields are available only if Authentication Type is set to MD5. Specify and confirm the password.

3. Click **Update**.

BFD Template

Use the BFD template to apply BFD configurations per segment to all appliances in the SD-WAN fabric, as follows:

1. Click the edit icon next to the segment for which you want to modify the configuration.
2. Configure the following fields:

Field	Description
Min Tx Interval	Minimum transmit interval in milliseconds (ms). Specify a value from 300 to 5000. The default setting is 300.

Field	Description
Min Rx Interval	Minimum receive interval in milliseconds (ms). Specify a value from 300 to 5000. The default setting is 300.
Detection Multiplier	Detection time multiplier. In BFD, the detection time is the transmit interval multiplied by the detection multiplier. If BFD data is not received within the detection time, a failure occurs. Specify a value from 3 to 10. The default setting is 3.

3. Click **Update**.

VXLAN Template

Use the VXLAN template to to efficiently deploy Virtual Network Identifier (VNI) instances for Virtual Extensible Local Area Network (VXLAN) segments. A VNI maps a routing segment to a firewall zone and a fallback role. Each segment is identified by a 24-bit VNI that can be configured for up to 16 million virtual networks. For additional information, see the VXLAN tab

Prerequisites

Before you can assign a VNI to a VXLAN segment, you must configure the following settings:

- Segmentation must be enabled to support VXLAN. See the Routing Segmentation (VRF) tab
- The IP routing on the BGP Layer 3 network that connects the EdgeConnect appliance VTEPs must already be configured. This is necessary to enable VXLAN traffic to traverse the network. Therefore, only in-line router mode is supported.
- Currently, the EdgeConnect EVPN address family is only supported for BGP EVPN peers in the Default segment (VRF ID = 0).
- One or more loopback interfaces must already be available.
- VXLAN is only supported on LAN interfaces. Route-Targets must be defined, and BGP enabled for all segments, even if no BGP peers are configured in non-default segments.

Common Settings for all VNIs

Use this section of the VXLAN Tab to configure these common settings for all VNIs:

- **Destination UDP Port**: You can configure a custom destination UDP port for VXLAN. If not selected, the appliance uses the default port of 4789.
- **VTEP Source Interface**: Select a loopback interface from the list.
  NOTE: Only loopback interfaces are valid. The loopback interface you choose will automatically be configured in the local interface field of the BGP Peer configuration if EVPN Peer is enabled.

VNI Mappings

For this dialog box, use the steps belwo to map a VNI to a routing segment, a firewall zone, and a fallback role.

Add

1. Click **Add** to create a new VNI for a segment.
2. Enter a value for the VNI segment. Valid values are 1-16777215.
3. Select the **Segment**, **Firewall Zone**, and **Fallback Role** (Don't Apply, Guest IOT, Untrusted).
4. Click **OK**.

Edit

1. Select an existing VNI from the list.
2. Click the **Edit** icon to modify an existing VNI.

Note: In the Flows tab, enable the VNI Tx and VNI Rx columns to display the number of the VNI that received or sent the VXLAN traffic. Both values should match for every flow. If not, there might be a misconfiguration downstream from the EdgeConnect.

Role to GPID Mapping

Use the Roles dialog box to map a policy enforcement role to a VXLAN Group Policy Identifier (GPID). Mapping policy enforcement roles to a VXLAN GPID is optional. Policy enforcement role mapping to a GPID propagates globally across the SD-WAN Fabric. Enabling the identity-based policy enforcement capability of the HPE Aruba Networking SD-WAN solution in VXLAN segments provides a highly automated extensible way of enabling a zero-trust security architecture.

Admin Distance Template

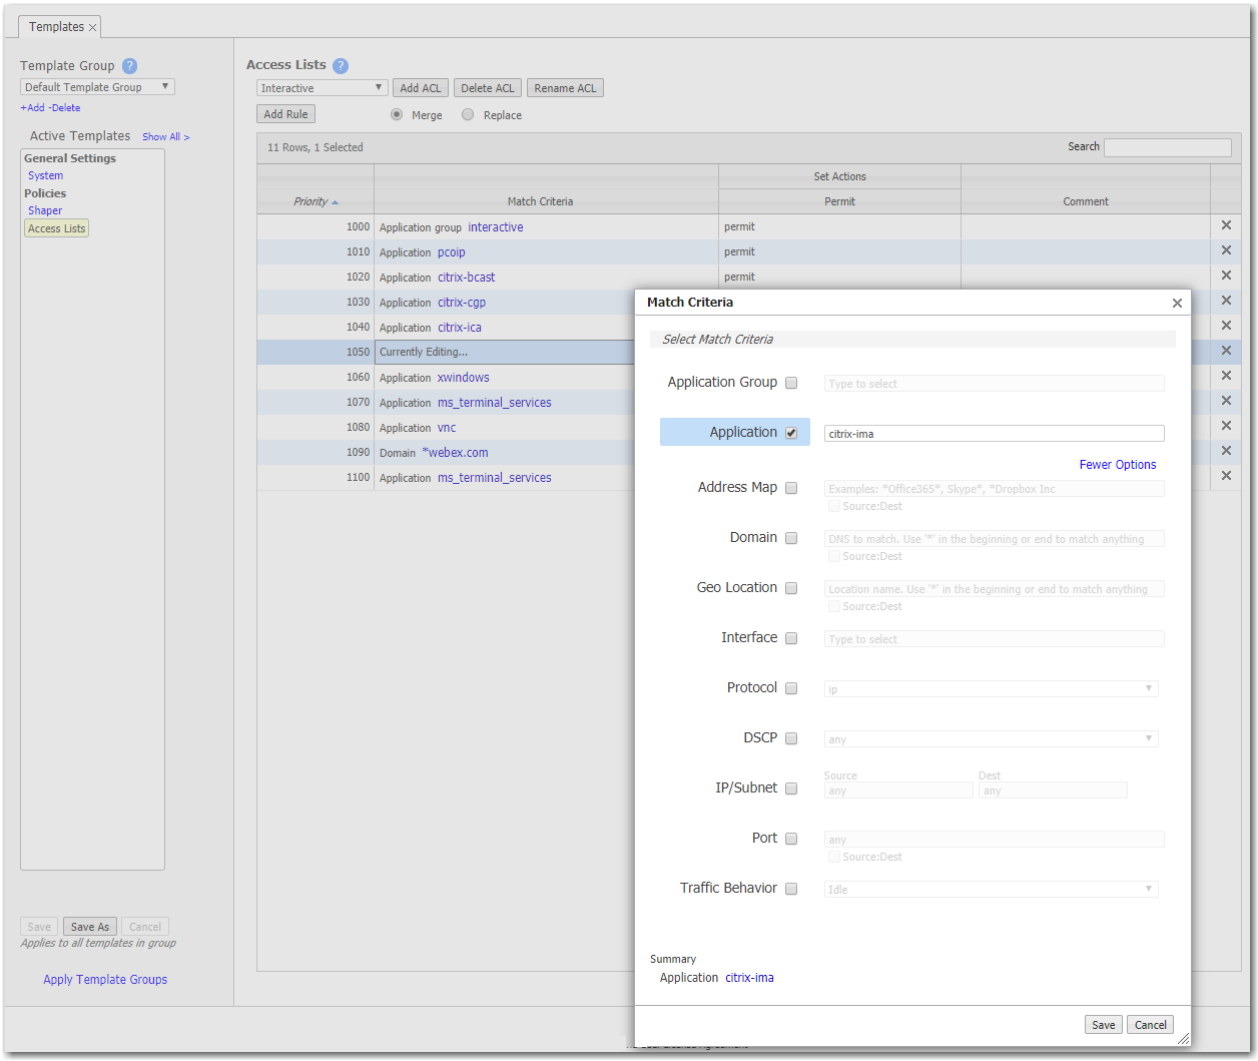
This table shows values associated with various types of **Admin Distance**. Admin Distance (AD) is the route preference value assigned to dynamic routes, static routes, and directly connected routes. When the appliance's Routes table has multiple routes to the same destination, the appliance uses the route with the lowest administrative distance.

Field	Description
Local	Manually configured route, or one learned from locally-connected subnets.
Subnet Shared - Static Routes	Route learned from an EdgeConnect peer.

Field	Description
Subnet Shared - BGP Remote	Route shared from an EdgeConnect peer in an external network.
Subnet Shared - OSPF Remote	Route shared from an EdgeConnect peer within the same network.
BGP Branch (pre-8.1.9.4)	Type of dynamic route learned from a local BGP branch peer before version 8.1.9.4.
BGP Transit (pre-8.1.9.4)	Type of dynamic route learned from a local BGP branch-transit peer before version 8.1.9.4.
EBGP (post-8.1.9.4)	External BGP: exchanging routing information with a router outside the company-wide network after version 8.1.9.4.
Subnet Shared CFGSET (9.5.0.0+)	ORO (Overlay Route Orchestrator) BGW (branch gateway) route learned from the SD-WAN fabric.
Subnet Shared IAPVPN (9.5.0.0+)	ORO Instant Access Point route learned from the SD-WAN fabric.
Subnet Shared Overlay (9.5.0.0+)	ORO modified/added route learned from the SD-WAN fabric.
Subnet Shared RIP (9.5.0.0+)	ORO learned Routing Information Protocol route learned from the SD-WAN fabric.
OAP BGP (9.5.0.0+)	Route learned from an OAP (Overlay Route Orchestrator) BGP peer in an external network.
OAP CFGSET (9.5.0.0+)	BGW (branch gateway) route learned from ORO.
OAP IAPVPN (9.5.0.0+)	Instant Access Point route learned from ORO.
OAP OSPF (9.5.0.0+)	Route learned from an OAP OSPF neighbor.
OAP Overlay (9.5.0.0+)	ORO modified/added route.
OAP RIP (9.5.0.0+)	Routing Information Protocol route learned from ORO.
OAP Static (9.5.0.0+)	Static route learned from ORO.
OAP Direct (9.5.0.0+)	Direct (connected) route learned from ORO.
BGP PE (pre-8.1.9.4)	Type of dynamic route learned from a local BGP PE (Provider Edge) router before version 8.1.9.4.
OSPF	Route learned from an OSPF (Open Shortest Path First) neighbor.
IBGP (post-8.1.9.4)	Internal BGP: exchanging routing information with a router inside the company-wide network after version 8.1.9.4.

Access Lists Template

Use this page to create, modify, delete, and rename Access Control Lists (ACLs).



An ACL is a reusable MATCH criteria for filtering flows.

NOTE: All selected match criteria options are ANDed to form the final match specification. For example, if you select Application Group and Application, both specifications are used to determine how the policy is applied. Review the summary of your match criteria at the bottom of this dialog.

Click **More Options** and **More User Profile Options** to display more match criteria selections. The More User Profile Options displays match criteria for identity-based traffic monitoring (User MAC, User Name, User Device, User Group, and User Vlan). For more information on these options, see HPE Aruba Networking SD-WAN Identity-Based Traffic Management User Guide.

It is associated with an action, **permit** or **deny**. You can use the same ACL as the MATCH condition in more than one policy: Route, QoS, Optimization, or NAT.

- An ACL consists of one or more ordered access control rules.
- An ACL only becomes active when it is used in a policy.

- **Deny** prevents further processing of the flow by **that ACL, specifically**. The appliance continues to the next entry in the policy.
- **Permit** allows the matching traffic flow to proceed on to the policy entry's associated SET actions. The default is **permit**.
- When creating ACL rules, list **deny** statements first, and prioritize less restrictive rules ahead of more restrictive rules.

Priority

- For ACL rules, you can set the priority to a value within the range 1 to 65535. When adding a rule, the priority is incremented by ten from the previous rule. You can change the priority, but this default behavior helps ensure that you can insert new rules without having to change subsequent priorities.

Match Criteria

- To specify different criteria for inbound versus outbound traffic, select the **Source:Dest** check box.

Source or Destination

- An IP address can specify a subnet - for example: 10.10.10.0/24.
- To allow **any IP address**, use 0.0.0.0/0.
- Ports are available only for the protocols **tcp**, **udp**, and **tcp/udp**.
- To allow **any port**, use **0**.

Wildcard-based Prefix Matching Rules

- Even when using a range or a wildcard, the IPv4 address must be specified in the 4-octet format, separated by the dot notation. For example, **A.B.C.D**.
- Range is specified using a single dash. For example, **128-129**.
- Wildcard is specified as an asterisk (____*).
- Range and Wildcard can both be used in the same address, but an octet can only contain one or the other. For example, **10.136-137.*.64-95**.
- A wildcard can only be used to define an entire octet. For example, **10.13*.*.64-95** is not supported. Use **10.130-139.*.64-95** to specify this range.
- The same rules apply to IPv6 addressing.
- CIDR notation and (Range or Wildcard) are mutually exclusive in the same address. For example, **192.168.0.1-127/24** is not supported. Use either **192.168.0.0/24** or **192.168.0.1-127**.
- These prefix-matching rules apply to the following policies only: Route, QoS, Optimization, NAT, Security, and ACLs.

Route Policies Template

NOTE: If you have deployed an SD-WAN network by using Business Intent Overlays (BIO), Orchestrator uses BIOs to automatically create the necessary Route Policies.

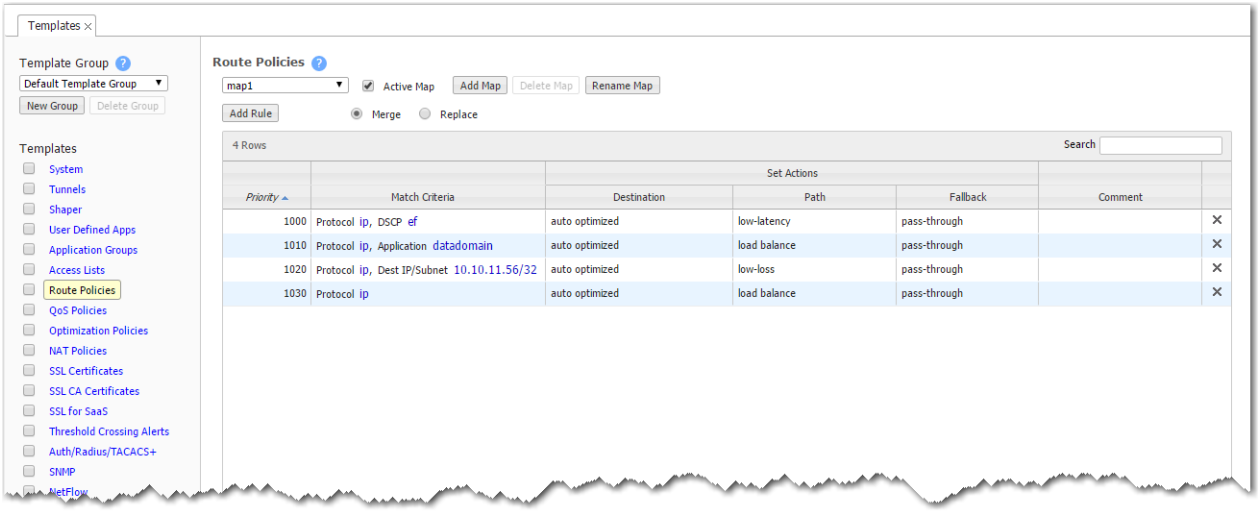
If you are creating a conventional WAN optimization network, there might be occasions when you need to directly configure Route Policies. Then, the following applies.

Only use the Route Policy template to create (and apply) rules for flows that are to be:

- Sent passthrough (shaped or unshaped)
- Dropped
- Configured for a specific high-availability deployment
- Routed based on application, ports, VLAN, DSCP, or ACL (Access Control List)

You might also want to create a Route Policy entry when multiple tunnels exist to the remote **peer**, and you want the appliance to dynamically select the best path based on one of these criteria:

- Load balancing
- Lowest loss
- Lowest latency
- A preferred interface
- A specific tunnel



Why?

Each appliance's default routing behavior is to auto-optimize all IP traffic, automatically directing flows to the appropriate tunnel. **Auto-optimization** strategies reduce the need to create explicit route map entries for optimization. The three strategies provided are **TCP-based** auto-opt, **IP-based** auto-opt, and **subnet sharing**. By default, all three are enabled on the **System** template.

Priority

- With this template, you can create rules with a priority from **1000 – 9999**. When the template is applied to an appliance, Orchestrator will delete all rules having a priority in that range before applying its policies.
- If you access an appliance directly, you can create rules with higher priority than Orchestrator rules (**1 – 999**) and rules with lower priority (**10000 – 19999** and **25000 – 65534**).

  NOTE: The priority range from **20000** to **24999** is reserved for Orchestrator.
- When adding a rule, the priority is incremented by ten from the previous rule. The priority can be changed, but this default behavior helps to ensure you can insert new rules without having to change subsequent priorities.

Match Criteria

Match criteria are used universally across all policy maps for route, QoS, optimization, SaaS NAT, and firewall zone security policies.

To use the same match criteria in different maps, create an ACL (Access Control List), which is a named, reusable set of rules. For efficiency, navigate to **Configuration > Templates & Policies > ACLs > Access Lists** to create ACLs, and then apply them across your appliances.

Use the Match Criteria dialog box to select and configure a variety of match criteria options. Some are explained below:

- **Address Map:** Use this option to sort by country, IP address owner, or SaaS application. You can also select and configure Microsoft Instance, Microsoft Category, and Proxy attributes for an address map. These attributes are secondary parameters to the address map. They are evaluated for a policy match only when the configured address map matches the flow. To select and configure these attributes, click **+Attributes**.
- Match criteria options related to Secure Web Services include:
  - **URL:** Omit the protocol in URLs and include a slash character (/) or a slash followed by a query parameter. For example, *google.com/* or *google.com/maps* are valid; *https://google.com/* is not valid. You can also use the asterisk (*) wildcard character, such as in *google.com/**, to specify a domain, but it is more appropriate to use the Domain match criteria option to specify *google.com/*. Separate multiple URL addresses with the pipe character (|). For example, you can specify two URL addresses as *bing.com/*|google.com/**.
  - **Web Category:** If you select this option, click **Select**, and then click **Strict**, **Moderate**, or **Custom**. Strict and Moderate include predetermined web category selections, which you cannot change. However, if you attempt to select or clear a web category for these, the Custom group of web category selections opens automatically. You can customize web category selections in the Custom group. After you make your selections, click **Ok**.
  - **Web Reputation:** If you select this option, select one or more web reputation categories (High Risk, Suspicious, Moderate Risk, Low Risk, or Trustworthy) from the list. To select more than one, press and hold the **Shift** key, and then click the appropriate categories.

  - **Bad IP Reputation:** Selecting this option matches on flows for which the source or destination IP address (or both) has a High Risk reputation score.

- Match criteria options related to RADIUS user role include:

  - **User Role:** This is the user role as specified in the authentication exchange with the ClearPass RADIUS server.
  - **User Name**
  - **User Group**
  - **User Device**
  - **User MAC**
  - **User Vlan**

  Configuring these match criteria related to user role enables an EdgeConnect to automatically assign traffic steering and firewall zone policies.

- Use the Src:Dest check box associated with several match criteria options to specify separate criteria for inbound and outbound traffic. You can configure source and destination role-based policies when both source and destination users are in the same network.

Source or Destination

- An IP address can specify a subnet; for example, 10.10.10.0/24 (IPv4) or fe80::204:23ff:fed8:4ba2/64 (IPv6).
- To allow **any IP address**, use 0.0.0.0/0 (IPv4) or ::/0 (IPv6).
- Ports are available only for the protocols **tcp**, **udp**, and **tcp/udp**.
- To allow **any port**, use **0**.

Wildcard-based Prefix Matching Rules

- Even when using a range or a wildcard, the IPv4 address must be specified in the 4-octet format, separated by the dot notation. For example, **A.B.C.D**.
- Range is specified using a single dash. For example, **128-129**.
- Wildcard is specified as an asterisk (___*).
- Range and Wildcard can both be used in the same address, but an octet can only contain one or the other. For example, **10.136-137.*.64-95**.
- A wildcard can only be used to define an entire octet. For example, **10.13*.*.64-95** is not supported. Use **10.130-139.*.64-95** to specify this range.
- The same rules apply to IPv6 addressing.
- CIDR notation and (Range or Wildcard) are mutually exclusive in the same address. For example, **192.168.0.1-127/24** is not supported. Use either **192.168.0.0/24** or **192.168.0.1-127**.
- These prefix-matching rules apply to the following policies only: Route, QoS, Optimization, NAT, Security, and ACLs.

Set Actions Fields

The Route Policy template's SET actions determine where to direct traffic and what the fallback is when a tunnel is down.

Where the Appliance Directs Traffic

- In the **Destination** field, you specify how to characterize the flow. The options are a specific overlay, **auto-optimized**, **pass-through** [shaped], **pass-through-unshaped**, or __drop__ped.
- When **auto-optimized**, a flow is directed to the appropriate tunnel. If you choose, you can specify that the appliance use metrics to dynamically select the best path based on one of these criteria:
  - Load balancing
  - Lowest loss
  - Lowest latency
- When configuring the Route Policy for an **individual** appliance when multiple tunnels exist to the remote **peer**, you can also select the path based on a preferred interface or a specific tunnel.

How Traffic Is Managed If a Tunnel Is Down

- The **Fallback** can be **pass-through** [shaped], **pass-through-unshaped**, or __drop__ped.
- When configuring the Route Policy for an **individual** appliance, the **continue** option is available if a specific tunnel is named in the **Destination** column. That option enables the appliance to read subsequent entries in the individual Route Policy in the event that the tunnel used in a previous entry goes down.

QoS Policies Template

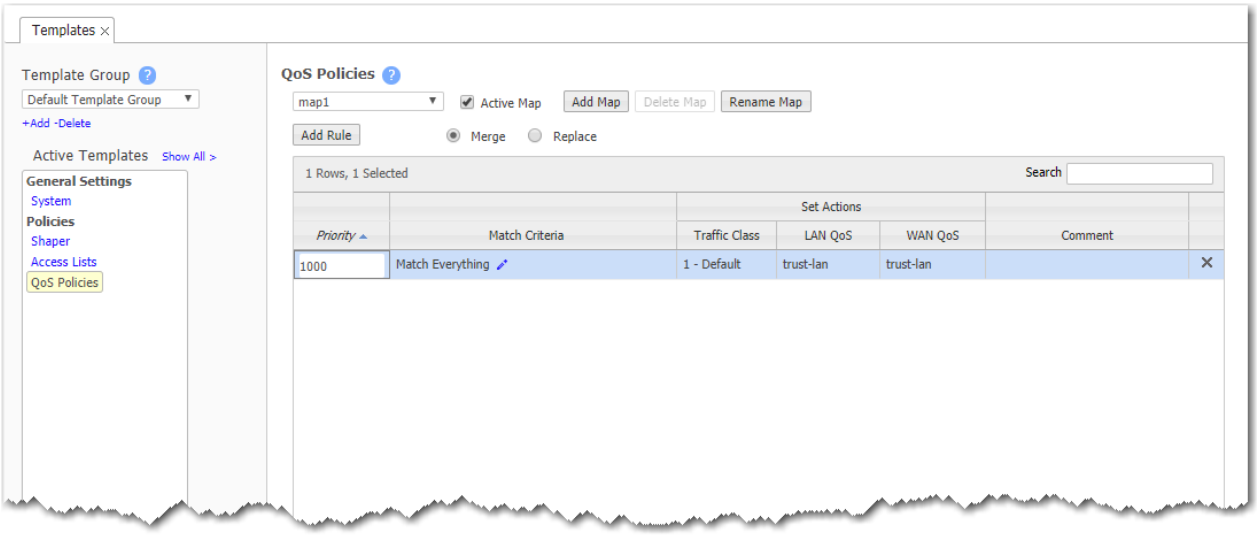
QoS Policy determines how flows are queued and marked.

The QoS Policy's SET actions determine two things:

- What traffic class a shaped flow—whether optimized or pass-through—is assigned
- Whether to trust incoming DSCP markings for LAN QoS and WAN QoS, or to remark them as they leave for the WAN

Use the **Shaper** to define, prioritize, and name traffic classes.

Think of it as the Shaper **defines** and the QoS Policy **assigns**.



Priority

- With this template, you can create rules with a priority from **1000 – 9999**. When the template is applied to an appliance, Orchestrator will delete all rules having a priority in that range before applying its policies.
- If you access an appliance directly, you can create rules with higher priority than Orchestrator rules (**1 – 999**) and rules with lower priority (**10000 – 19999** and **25000 – 65534**).

  NOTE: The priority range from **20000** to **24999** is reserved for Orchestrator.
- When adding a rule, the priority is incremented by ten from the previous rule. The priority can be changed, but this default behavior helps to ensure you can insert new rules without having to change subsequent priorities.

Match Criteria

Match criteria are used universally across all policy maps for route, QoS, optimization, SaaS NAT, and firewall zone security policies.

To use the same match criteria in different maps, create an ACL (Access Control List), which is a named, reusable set of rules. For efficiency, navigate to **Configuration > Templates & Policies > ACLs > Access Lists** to create ACLs, and then apply them across your appliances.

Use the Match Criteria dialog box to select and configure a variety of match criteria options. Some are explained below:

- **Address Map:** Use this option to sort by country, IP address owner, or SaaS application. You can also select and configure Microsoft Instance, Microsoft Category, and Proxy attributes for an address map. These attributes are secondary parameters to the address map. They are evaluated for a policy match only when the configured address map matches the flow. To select and configure these attributes, click **+Attributes**.
- Match criteria options related to Secure Web Services include:

  - **URL:** Omit the protocol in URLs and include a slash character (/) or a slash followed by a query parameter. For example, *google.com/* or *google.com/maps* are valid; *https://google.com/* is not valid. You can also use the asterisk (*) wildcard character, such as in *google.com/**, to specify a domain, but it is more appropriate to use the Domain match criteria option to specify *google.com/*. Separate multiple URL addresses with the pipe character (|). For example, you can specify two URL addresses as *bing.com/*|google.com/**.
  - **Web Category:** If you select this option, click **Select**, and then click **Strict**, **Moderate**, or **Custom**. Strict and Moderate include predetermined web category selections, which you cannot change. However, if you attempt to select or clear a web category for these, the Custom group of web category selections opens automatically. You can customize web category selections in the Custom group. After you make your selections, click **Ok**.
  - **Web Reputation:** If you select this option, select one or more web reputation categories (High Risk, Suspicious, Moderate Risk, Low Risk, or Trustworthy) from the list. To select more than one, press and hold the **Shift** key, and then click the appropriate categories.
  - **Bad IP Reputation:** Selecting this option matches on flows for which the source or destination IP address (or both) has a High Risk reputation score.

- Match criteria options related to RADIUS user role include:

  - **User Role:** This is the user role as specified in the authentication exchange with the ClearPass RADIUS server.
  - **User Name**
  - **User Group**
  - **User Device**
  - **User MAC**
  - **User Vlan**

  Configuring these match criteria related to user role enables an EdgeConnect to automatically assign traffic steering and firewall zone policies.

- Use the Src:Dest check box associated with several match criteria options to specify separate criteria for inbound and outbound traffic. You can configure source and destination role-based policies when both source and destination users are in the same network.

Source or Destination

- An IP address can specify a subnet; for example, 10.10.10.0/24 (IPv4) or fe80::204:23ff:fed8:4ba2/64 (IPv6).
- To allow **any IP address**, use 0.0.0.0/0 (IPv4) or ::/0 (IPv6).
- Ports are available only for the protocols **tcp**, **udp**, and **tcp/udp**.
- To allow **any port**, use **0**.

Wildcard-based Prefix Matching Rules

- Even when using a range or a wildcard, the IPv4 address must be specified in the 4-octet format, separated by the dot notation. For example, **A.B.C.D**.
- Range is specified using a single dash. For example, **128-129**.
- Wildcard is specified as an asterisk (___*).
- Range and Wildcard can both be used in the same address, but an octet can only contain one or the other. For example, **10.136-137.*.64-95**.
- A wildcard can only be used to define an entire octet. For example, **10.13*.*.64-95** is not supported. Use **10.130-139.*.64-95** to specify this range.
- The same rules apply to IPv6 addressing.
- CIDR notation and (Range or Wildcard) are mutually exclusive in the same address. For example, **192.168.0.1-127/24** is not supported. Use either **192.168.0.0/24** or **192.168.0.1-127**.
- These prefix-matching rules apply to the following policies only: Route, QoS, Optimization, NAT, Security, and ACLs.

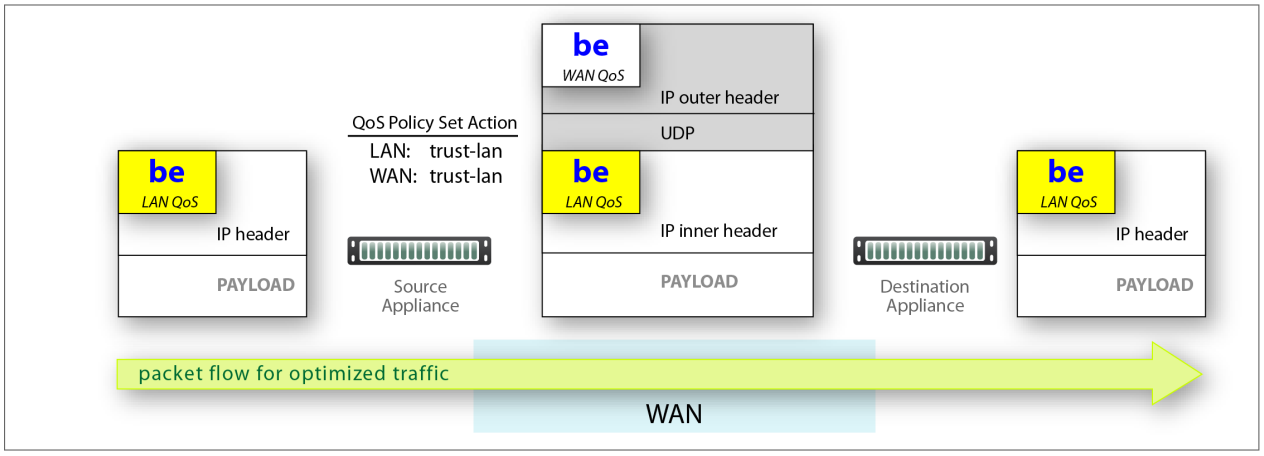
Handle and Mark DSCP Packets

- DSCP markings specify end-to-end QoS policies throughout a network.
- The default values for **LAN QoS** and **WAN QoS** are **trust-lan**.

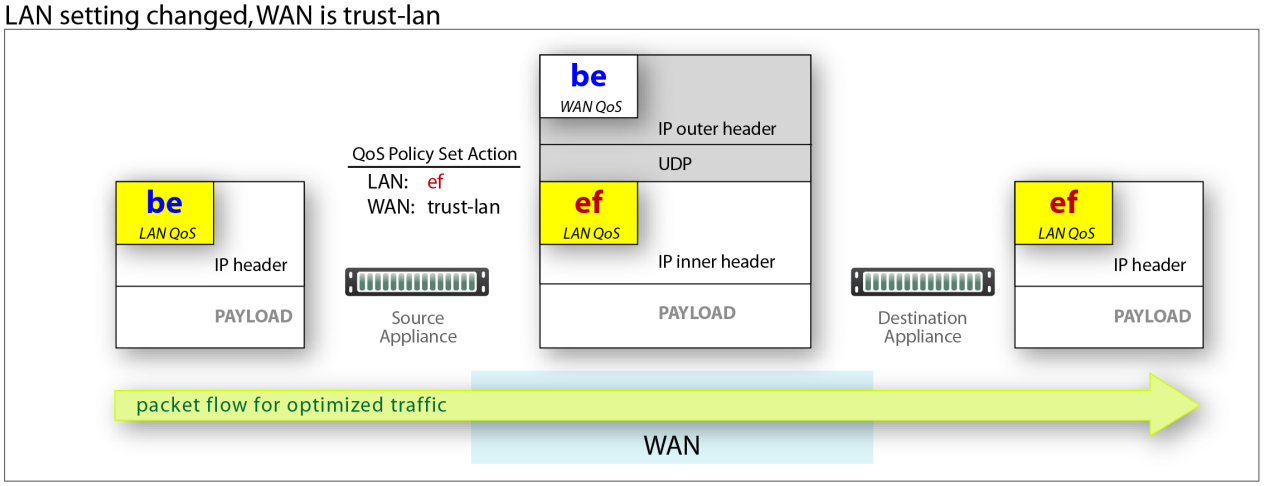
Apply DSCP Markings to Optimized (Tunnelized) Traffic

- The appliance encapsulates optimized traffic. This adds an IP outer header to packets for travel across the WAN. This outer header contains the **WAN QoS** DSCP marking.
- **LAN QoS:** The DSCP marking applied to the IP header before encapsulation.
- **WAN QoS:** The DSCP marking in the encapsulating outer IP header. The remote appliance removes the outer IP header.

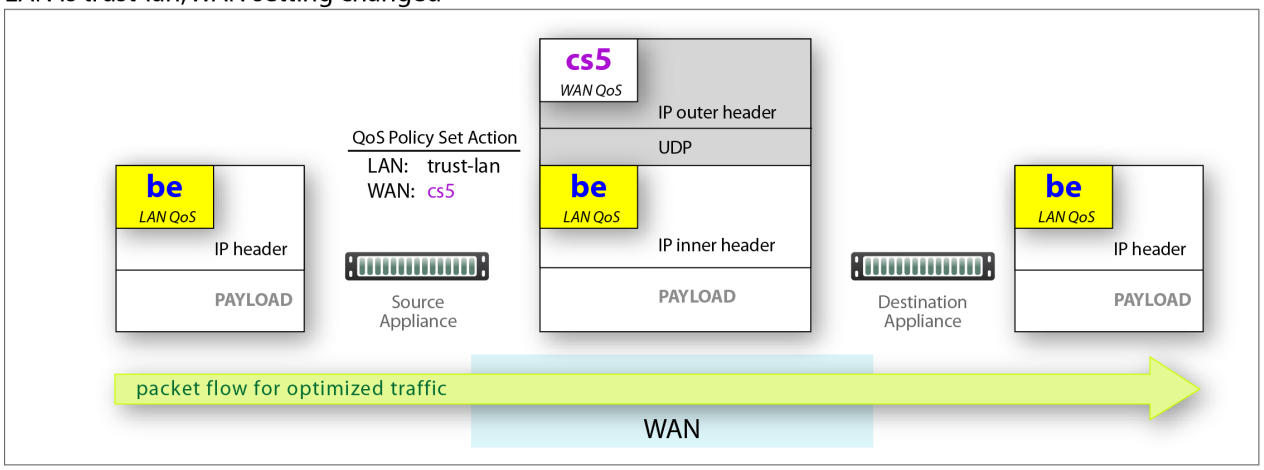
LAN and WAN set to trust-lan



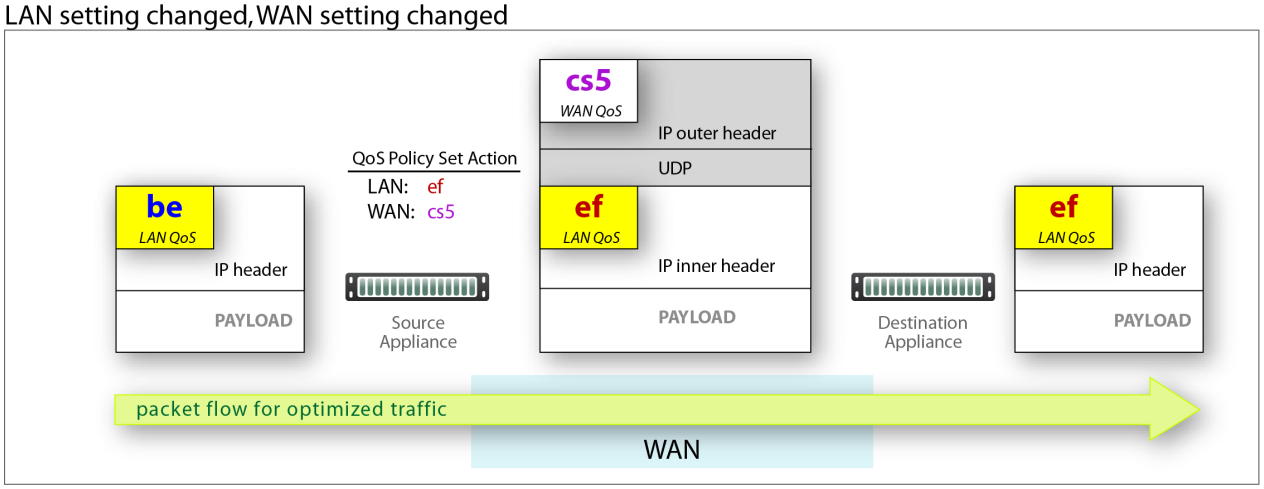
LAN setting changed, WAN is trust-lan



LAN is trust-lan, WAN setting changed



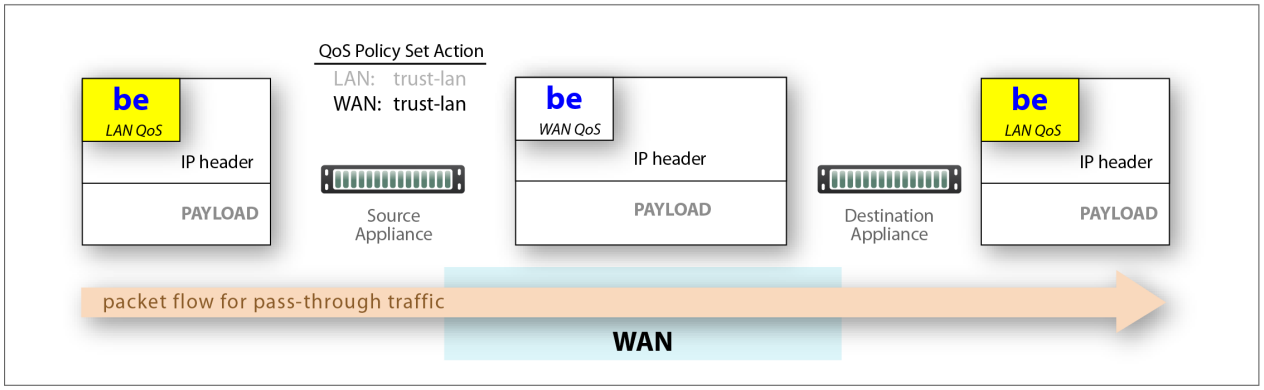
LAN setting changed, WAN setting changed



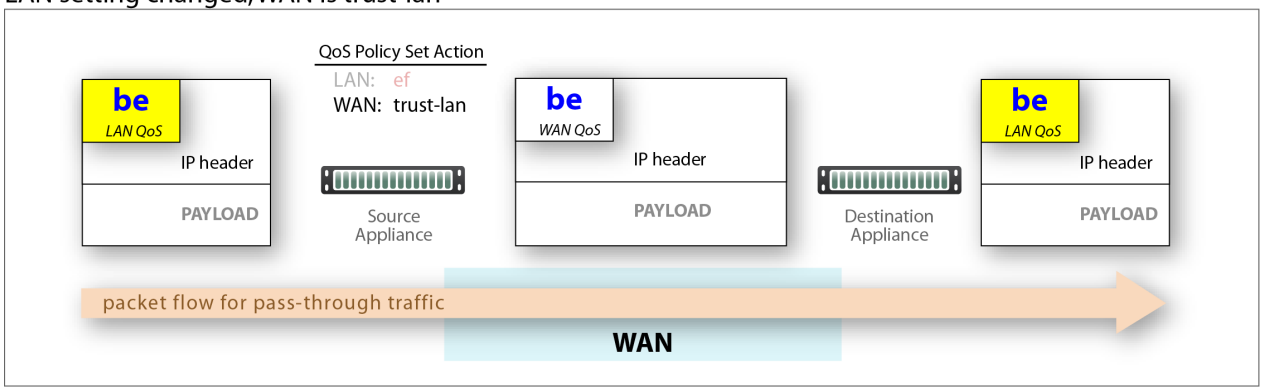
Apply DSCP Markings to Pass-through Traffic

- The appliance applies the QoS Policy's DSCP markings to all pass-through flows—shaped and unshaped.
- Pass-through traffic does not receive an additional header, so it is handled differently:
  - The Optimization Policy's **LAN QoS** Set Action is ignored.
  - The specified **WAN QoS** marking replaces the packet's existing **LAN QoS** DSCP marking.
  - When the packet reaches the remote appliance, it retains the modified QoS setting as it travels to its destination.

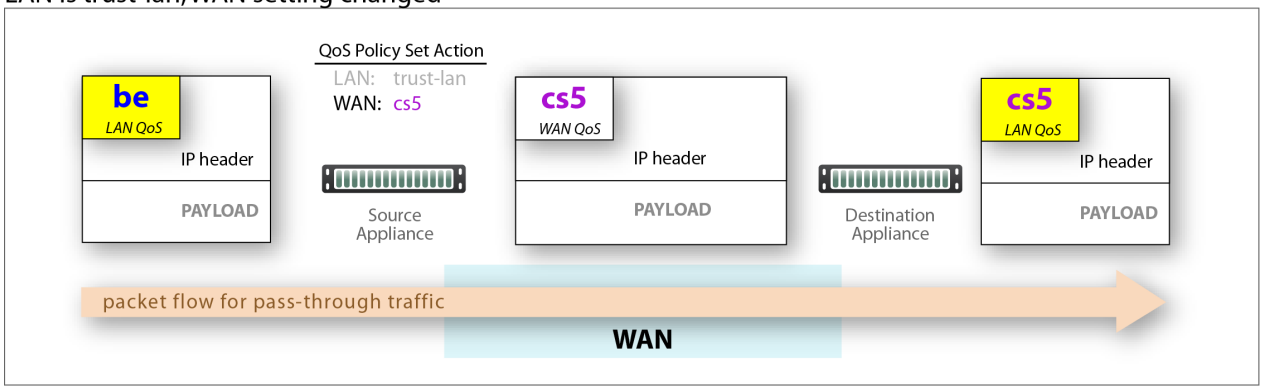
LAN and WAN set to trust-lan



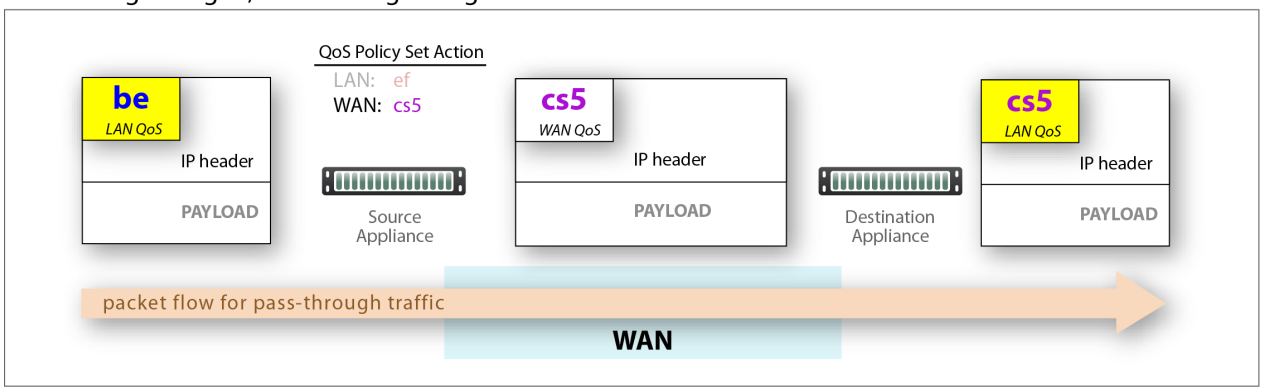
LAN setting changed, WAN is trust-lan



LAN is trust-lan, WAN setting changed

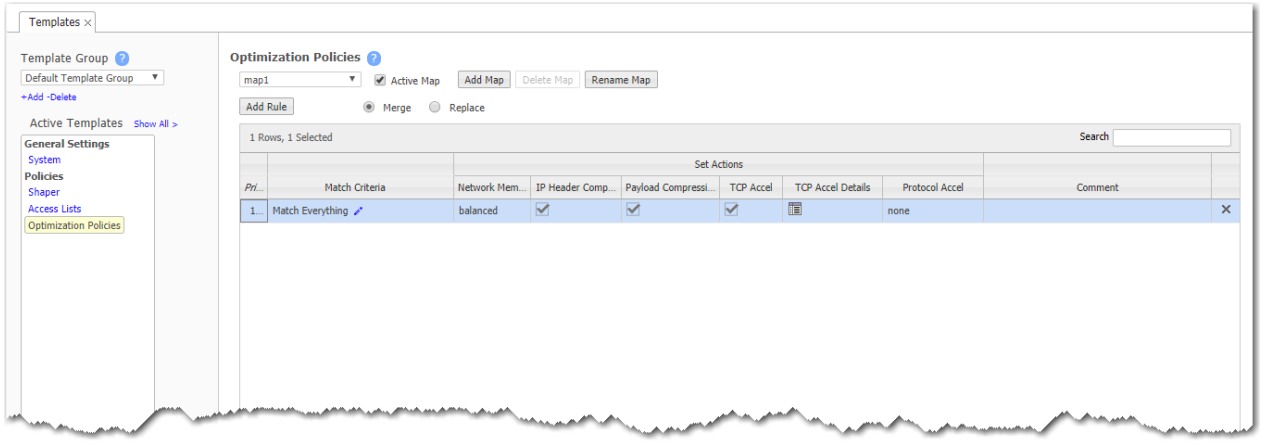


LAN setting changed, WAN setting changed



Optimization Policies Template

Optimization templates apply Optimization policies to appliances.



Priority

- With this template, you can create rules with a priority from **1000 – 9999**. When the template is applied to an appliance, Orchestrator will delete all rules having a priority in that range before applying its policies.
- If you access an appliance directly, you can create rules with higher priority than Orchestrator rules (**1 – 999**) and rules with lower priority (**10000 – 19999** and **25000 – 65534**).

  NOTE: The priority range from **20000** to **24999** is reserved for Orchestrator.
- When adding a rule, the priority is incremented by ten from the previous rule. The priority can be changed, but this default behavior helps to ensure you can insert new rules without having to change subsequent priorities.

Match Criteria

Match criteria are used universally across all policy maps for route, QoS, optimization, SaaS NAT, and firewall zone security policies.

To use the same match criteria in different maps, create an ACL (Access Control List), which is a named, reusable set of rules. For efficiency, navigate to **Configuration > Templates & Policies > ACLs > Access Lists** to create ACLs, and then apply them across your appliances.

Use the Match Criteria dialog box to select and configure a variety of match criteria options. Some are explained below:

- **Address Map:** Use this option to sort by country, IP address owner, or SaaS application. You can also select and configure Microsoft Instance, Microsoft Category, and Proxy attributes for an address map. These attributes are secondary parameters to the address map. They are evaluated for a policy match only when the configured address map matches the flow. To select and configure these attributes, click **+Attributes**.
- Match criteria options related to Secure Web Services include:
  - **URL:** Omit the protocol in URLs and include a slash character (/) or a slash followed by a query parameter. For example, *google.com/* or *google.com/maps* are valid; *https://google.com/* is not valid. You can also use the asterisk (*) wildcard character, such as in *google.com/**, to specify a domain, but it is more appropriate to use the Domain match criteria option to specify *google.com/*. Separate multiple URL addresses with the pipe character (|). For example, you can specify two URL addresses as *bing.com/*|google.com/**.
  - **Web Category:** If you select this option, click **Select**, and then click **Strict**, **Moderate**, or **Custom**. Strict and Moderate include predetermined web category selections, which you cannot change. However, if you attempt to select or clear a web category for these, the Custom group of web category selections opens automatically. You can customize web category selections in the Custom group. After you make your selections, click **Ok**.
  - **Web Reputation:** If you select this option, select one or more web reputation categories (High Risk, Suspicious, Moderate Risk, Low Risk, or Trustworthy) from the list. To select more than one, press and hold the **Shift** key, and then click the appropriate categories.

  - **Bad IP Reputation:** Selecting this option matches on flows for which the source or destination IP address (or both) has a High Risk reputation score.

- Match criteria options related to RADIUS user role include:

  - **User Role:** This is the user role as specified in the authentication exchange with the ClearPass RADIUS server.
  - **User Name**
  - **User Group**
  - **User Device**
  - **User MAC**
  - **User Vlan**

  Configuring these match criteria related to user role enables an EdgeConnect to automatically assign traffic steering and firewall zone policies.

- Use the Src:Dest check box associated with several match criteria options to specify separate criteria for inbound and outbound traffic. You can configure source and destination role-based policies when both source and destination users are in the same network.

Source or Destination

- An IP address can specify a subnet; for example, 10.10.10.0/24 (IPv4) or fe80::204:23ff:fed8:4ba2/64 (IPv6).
- To allow **any IP address**, use 0.0.0.0/0 (IPv4) or ::/0 (IPv6).
- Ports are available only for the protocols **tcp**, **udp**, and **tcp/udp**.
- To allow **any port**, use **0**.

Wildcard-based Prefix Matching Rules

- Even when using a range or a wildcard, the IPv4 address must be specified in the 4-octet format, separated by the dot notation. For example, **A.B.C.D**.
- Range is specified using a single dash. For example, **128-129**.
- Wildcard is specified as an asterisk (____*).
- Range and Wildcard can both be used in the same address, but an octet can only contain one or the other. For example, **10.136-137.*.64-95**.
- A wildcard can only be used to define an entire octet. For example, **10.13*.*.64-95** is not supported. Use **10.130-139.*.64-95** to specify this range.
- The same rules apply to IPv6 addressing.
- CIDR notation and (Range or Wildcard) are mutually exclusive in the same address. For example, **192.168.0.1-127/24** is not supported. Use either **192.168.0.0/24** or **192.168.0.1-127**.
- These prefix-matching rules apply to the following policies only: Route, QoS, Optimization, NAT, Security, and ACLs.

Set Actions Fields

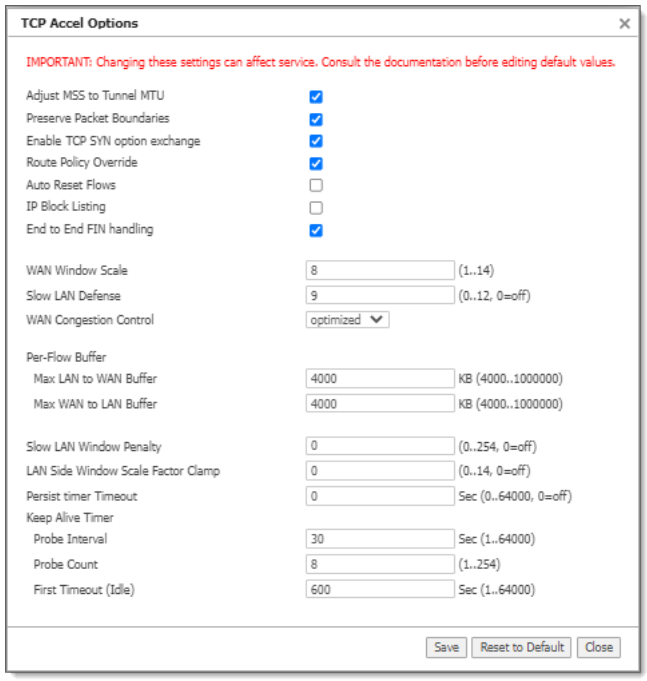
Set Action	Description
Network Memory	Addresses limited bandwidth. This technology uses advanced fingerprinting algorithms to examine all incoming and outgoing WAN traffic. Network Memory localizes information and transmits only modifications between locations. **Maximize Reduction:** Optimizes for maximum data reduction at the potential cost of slightly lower throughput and/or some increase in latency. It is appropriate for bulk data transfers such as file transfers and FTP, where bandwidth savings are the primary concern. **Minimize Latency:** Ensures that Network Memory processing adds no latency. This might come at the cost of lower data reduction. It is appropriate for extremely latency-sensitive interactive or transactional traffic. It is also appropriate when the primary objective is to fully utilize the WAN pipe to increase the LAN-side throughput, as opposed to conserving WAN bandwidth. **Balanced:** Is the default setting. It dynamically balances latency and data reduction objectives and is the best choice for most traffic types. **Disabled:** Turns off Network Memory.
IP Header Compression	Process of compressing excess protocol headers before transmitting them on a link and uncompressing them to their original state at the other end. It is possible to compress the protocol headers due to the redundancy in header fields of the same packet, as well as in consecutive packets of a packet stream.
Payload Compression	Uses algorithms to identify relatively short byte sequences that are repeated frequently. These are then replaced with shorter segments of code to reduce the size of transmitted data. Simple algorithms can find repeated bytes within a single packet; more sophisticated algorithms can find duplication across packets and even across flows.
TCP Acceleration	Uses techniques such as selective acknowledgments, window scaling, and maximum segment size adjustment to mitigate poor performance on high-latency links. **NOTE:** Slow LAN alert goes off when the loss has fallen below 80% of the specified value configured in the **TCP Accel Options** window. For more information, see TCP Acceleration Options.

Set Action	Description
Protocol Acceleration	Provides explicit configuration for optimizing SSL, SRDF, Citrix, and iSCSI protocols. In a network environment, it is possible that not every appliance has the same optimization configurations enabled. Therefore, the site that initiates the flow (the **client**) determines the state of the protocol-specific optimization.

TCP Acceleration Options

TCP acceleration uses techniques such as selective acknowledgment, window scaling, and message segment size adjustment to compensate for poor performance on high latency links.

This feature has a set of advanced options with default values.



CAUTION: Because changing these settings can affect service, it is recommended that you **do not modify** these without direction from Support.

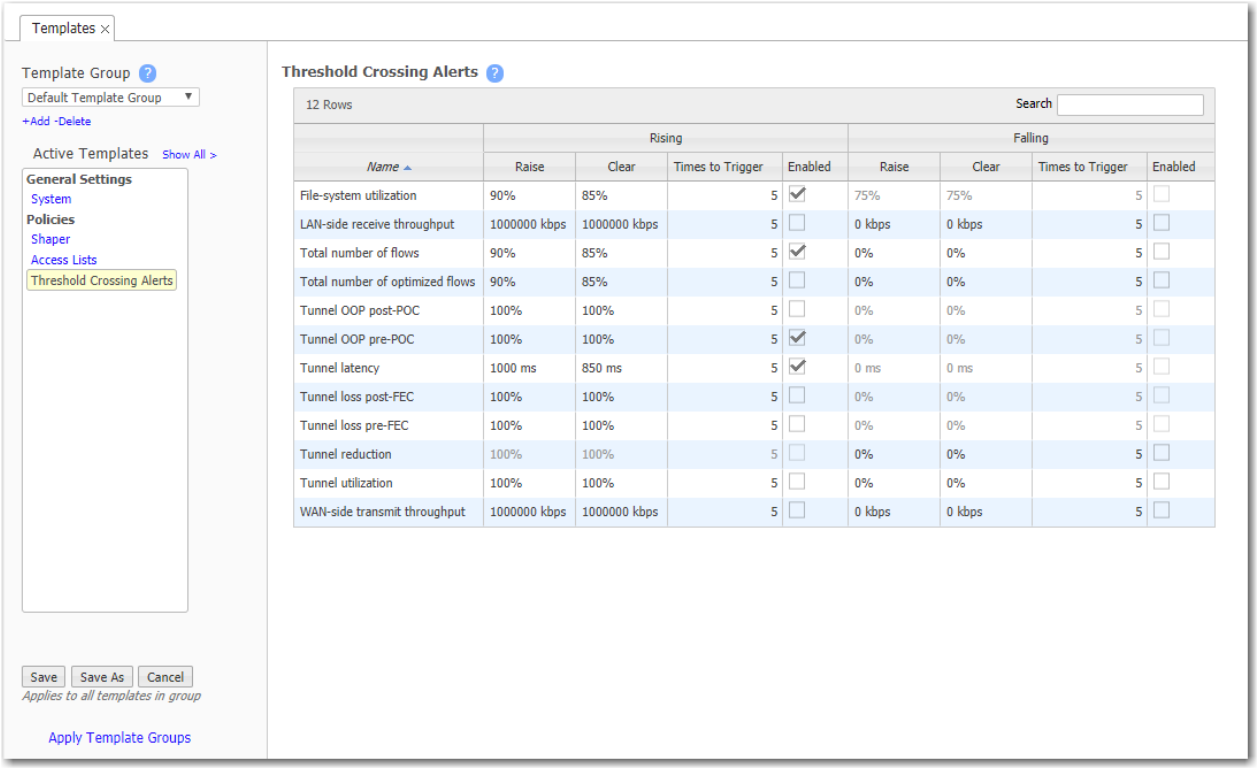
Option	Description
Adjust MSS to Tunnel MTU	Limits the TCP MSS (Maximum Segment Size) advertised by the end hosts in the SYN segment to a value derived from the Tunnel MTU (Maximum Transmission Unit). This is *TCP MSS = Tunnel MTU – Tunnel Packet Overhead*. This feature is enabled by default so that the **maximum value** of the end host MSS is always coupled to the Tunnel MSS. If the end host MSS is smaller than the tunnel MSS, the end host MSS is used instead. A use case for disabling this feature is when the end host uses Jumbo frames.
Auto Reset Flows	**NOTE:** Whether this feature is enabled or not, the default behavior when a tunnel goes Down is to automatically reset the flows. If enabled, it resets all TCP flows that are not accelerated, but should be (based on policy and on internal criteria like a Tunnel Up event). The internal criteria can also include: Resetting all TCP accelerated flows on a Tunnel Down event. Resetting TCP acceleration is enabled. SYN packet was not seen (so this flow was either part of WCCP redirection or it already existed when the appliance was inserted in the data path).
Enable TCP SYN option exchange	Controls whether or not the proprietary TCP SYN option is forwarded on the LAN side. Enabled by default, this feature detects if there are more than two EdgeConnect appliances in the flow's data path, and optimizes accordingly. Disable this feature if there is a LAN-side firewall or a third-party appliance that would drop a SYN packet when it encounters an unfamiliar TCP option.

Option	Description
End to End FIN Handling	This feature helps to fine tune TCP behavior during a connection's graceful shutdown event. When this feature is **ON** (Default), TCP on the local appliance synchronizes this graceful shutdown of the local LAN side with the LAN side of the remote appliance. When this feature is **OFF** (Default TCP), no such synchronization happens and the two LAN segments at the ends gracefully shut down, independently.
IP Block Listing	If selected, and if the appliance does not receive a TCP SYN-ACK from the remote end within five seconds, the flow proceeds without acceleration and the destination IP address is blocked for one minute.
Keep Alive Timer	Allows changing the Keep Alive timer for the TCP connections. **Probe Interval:** Time interval in seconds between two consecutive Keep Alive probes. **Probe Count:** Maximum number of Keep Alive probes to send. **First Timeout (Idle):** Time interval until the first Keep Alive timeout.
LAN Side Window Scale Factor Clamp	This setting allows the appliance to present an artificially lowered Window Scale Factor (WSF) to the end host. This reduces the need for memory in scenarios in which there are many out-of-order packets being received from the LAN side. These out-of-order packets cause much buffer utilization and maintenance.
Per-Flow Buffer	(**Max LAN to WAN Buffer** and **Max WAN to LAN Buffer**) This setting clamps the maximum buffer space that can be allocated to a flow, in each direction.
Persist timer Timeout	Allows the TCP to terminate connections that are in Persist timeout stage after the configured number of seconds.
Preserve Packet Boundaries	Preserves the packet boundaries end-to-end. If this feature is disabled, the appliances in the path can coalesce consecutive packets of a flow to use bandwidth more efficiently. It is enabled by default so that applications requiring packet boundaries to match do not fail.

Option	Description
Route Policy Override	Tries to override asymmetric route policy settings. It emulates auto-opt behavior by using the same tunnel for the returning SYN+ACK as it did for the original SYN packet. Disable this feature if the asymmetric route policy setting is necessary to correctly route packets. In this case, you might need to configure flow redirection to ensure optimization of TCP flows.
Slow LAN Defense	Resets all flows that consume a disproportionate amount of buffer and have a very slow throughput on the LAN side. Owing to a few slower end hosts or a lossy LAN, these flows affect the performance of all other flows so that no flows see the customary throughput improvement gained through TCP acceleration. This feature is enabled by default. The number relates indirectly to the amount of time the system waits before resetting such slow flows.
Slow LAN Window Penalty	This setting (**OFF** by default) penalizes flows that are slow to send data on the LAN side by artificially reducing their TCP receive window. This causes less data to be received and helps to reach a balance with the data sending rate on the LAN side.
WAN Congestion Control	Selects the internal Congestion Control parameter: **Optimized:** This is the default setting. This mode offers optimized performance in almost all scenarios. **Standard:** In some unique cases, it might be necessary to downgrade to Standard performance to better interoperate with other flows on the WAN link. **Aggressive:** Provides aggressive performance and should be used with caution. Recommended mostly for Data Replication scenarios.
WAN Window Scale	This is the WAN-side TCP Window scale factor that is used internally for WAN-side traffic. This is independent of the WAN-side factor advertised by the end hosts.

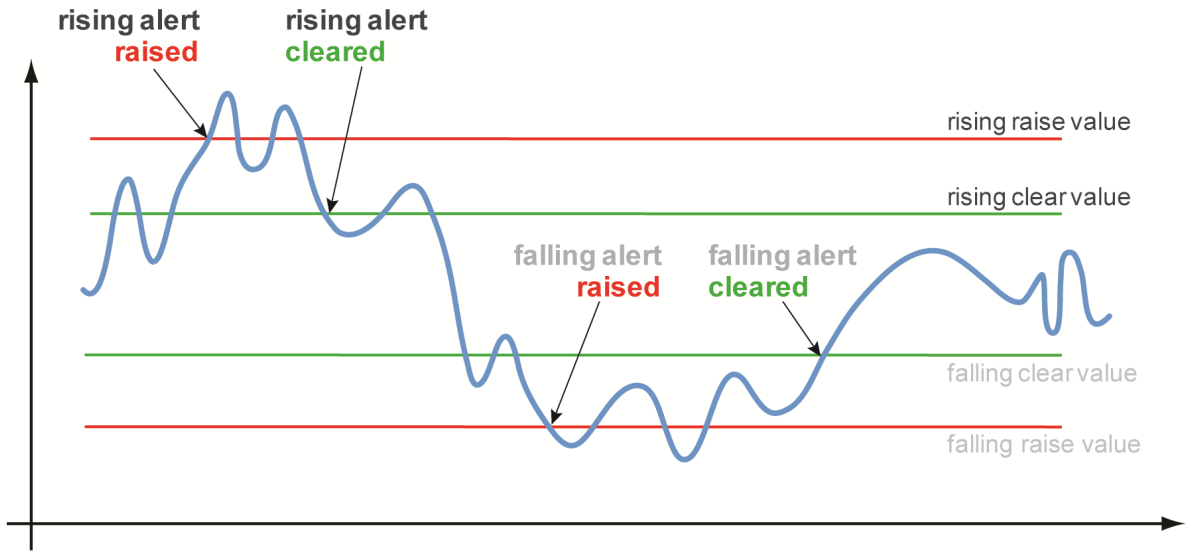
Threshold Crossing Alerts Template

Threshold Crossing Alerts (TCAs) are preemptive, user-configurable alarms that are triggered when the specific thresholds are crossed.



Name	Rising Raise	Rising Clear	Rising Times to Trigger	Rising Enabled	Falling Raise	Falling Clear	Falling Times to Trigger	Falling Enabled
File-system utilization	90%	85%	5	☑	75%	75%	5	☐
LAN-side receive throughput	1000000 kbps	1000000 kbps	5	☐	0 kbps	0 kbps	5	☐
Total number of flows	90%	85%	5	☑	0%	0%	5	☐
Total number of optimized flows	90%	85%	5	☐	0%	0%	5	☐
Tunnel OOP post-POC	100%	100%	5	☐	0%	0%	5	☐
Tunnel OOP pre-POC	100%	100%	5	☑	0%	0%	5	☐
Tunnel latency	1000 ms	850 ms	5	☑	0 ms	0 ms	5	☐
Tunnel loss post-FEC	100%	100%	5	☐	0%	0%	5	☐
Tunnel loss pre-FEC	100%	100%	5	☐	0%	0%	5	☐
Tunnel reduction	100%	100%	5	☐	0%	0%	5	☐
Tunnel utilization	100%	100%	5	☐	0%	0%	5	☐
WAN-side transmit throughput	1000000 kbps	1000000 kbps	5	☐	0 kbps	0 kbps	5	☐

They alarm on both rising and falling threshold crossing events (that is, floor and ceiling levels). For both levels, one value raises the alarm while another value clears it.



Rules:
- High raise threshold is greater than high clear threshold
- Low raise threshold is less than low clear threshold

ON by Default

- **Appliance Capacity** – Triggers when an appliance reaches 95% of its total flow capacity. It is not configurable and can be cleared only by an operator.
- **File-system utilization** – Percent of non-Network Memory disk space filled by the appliance. This TCA cannot be disabled.
- **Tunnel latency** – Measured in milliseconds, the maximum latency of a one-second sample within a 60-second span.

OFF by Default

- **LAN-side receive throughput** – Based on a one-minute average, the LAN-side receive **TOTAL** for all interfaces.
- **WAN-side transmit throughput** – Based on a one-minute average, the WAN-side transmit **TOTAL** for all interfaces.
- **TCAs based on an end-of-minute count**:
  - Total number of flows
  - Total number of optimized flows
- **TCAs based on a one-minute average**:
  - Tunnel loss post-FEC
  - Tunnel loss post-FEC
  - Tunnel OOP post-POC
  - Tunnel OOP post-POC
  - Tunnel reduction
  - Tunnel utilization (based on percent of configured maximum [system] bandwidth)

TCA Metrics

Times to Trigger – A value of **1** triggers an alarm on the first threshold crossing instance. The default sampling granularity (or *rate* or *interval*) is one minute.

This table lists the **metrics** of each type of threshold crossing alert:

Metrics for Threshold Crossing Alerts

TCA Name	Unit	Metric
Appliance Level		
WAN-side transmit throughput	kbps	Minute averageWAN–side transmit TOTAL for all interfaces
LAN-side receive throughput	kbps	Minute averageLAN–side receive TOTAL for all interfaces

TCA Name	Unit	Metric
Total number of optimized flows	flows	End of minute count
Total number of flows	flows	End of minute count
File-system-utilization	% (non–Network Memory)	End of minute count
Tunnel Level		
Tunnel latency	msec	Second-sampled maximum latency during the minute
Tunnel loss pre-FEC	1/10th %	Minute average
Tunnel loss post-FEC	1/10th %	Minute average
Tunnel OOP pre-POC	1/10th %	Minute average
Tunnel OOP post-POC	1/10th %	Minute average
Tunnel utilization	% of configured bandwidth	Minute average
Tunnel reduction	%	Minute average

Security Policies Template

Use this page to set up security policies, also known as __*zone-based firewalls__*.

CAUTION: If segmentation is enabled, do not use the Security Policies Template. Instead, configure Security Policies from the Routing Segmentation (VRF) tab.

- Zones are created on the Orchestrator and applied to an **Interface**.
- By default, traffic is allowed between interfaces labeled with the same zone. Any traffic between interfaces with different zones is dropped. Users can create exception rules (Security Policies) to allow traffic between interfaces with different zones. For information on troubleshooting flows that were denied by the firewall with the reason “outbound pkt new dst zone” or “Zone change detected on outbound packet,” see this troubleshooting video.
- When you create an interface, it is assigned **Default** zone.
- If you create a new zone and assign that to an interface, all traffic between that interface and rest of the interfaces (which are still in the **Default** zone) are dropped. This implies that zone creation and assignment to interfaces should be performed during a planned network maintenance.
- You can also assign a zone label to an **Overlay**. On a new system, all overlays are assigned the **Default** zone.

- Traffic between an Interface and an Overlay follows the same rules as traffic between Interfaces or two Overlays; traffic is allowed between zones with the same label and any traffic between different zones is dropped. Users can create Security Policies to allow traffic between different zones.

Implicit Drop Logging

Implicit Drop Logging enables you to configure implicit zone-based firewall drop logging levels. Implicit zone-based firewall drop is for inter-zone traffic by default. For example, if all the zone_x to zone_y traffic is the default **Deny All** (all the red cells from matrix), the traffic will be dropped by the zone-based firewall engine.

Select one of the following levels for the Implicit Drop Logging from the list: **None**, **Emergency**, **Alert**, **Critical**, **Error**, **Warning**, **Notice**, **Info**, or **Debug**.

NOTE: The default logging level is **Alert**.

Template

Complete the following steps to create a Security Policies Template:

1. Create zone names in **Configuration > Overlays & Security > Security > Firewall Zones**.
2. Create security policies to define exceptions.

   To edit or add a rule, select the desired square in the matrix, and when the Edit Rules pop-up appears, make the desired changes.
3. Select the edit icon in the Match Criteria column and the Match Criteria pop-up appears. Make the desired changes.
4. You can select **More Options** to customize your rules. Select the check box next to the specific match criteria and select your desired changes from the list.
5. Click **Save**.

Wildcard-based Prefix Matching Rules

- Even when using a range or a wildcard, the IPv4 address must be specified in the 4-octet format, separated by the dot notation. For example, **A.B.C.D**.
- Range is specified using a single dash. For example, **128-129**.
- Wildcard is specified as an asterisk (___*).
- Range and Wildcard can both be used in the same address, but an octet can only contain one or the other. For example, **10.136-137.*.64-95**.
- A wildcard can only be used to define an entire octet. For example, **10.13*.*.64-95** is not supported. Use **10.130-139.*.64-95** to specify this range.
- The same rules apply to IPv6 addressing.
- CIDR notation and (Range or Wildcard) are mutually exclusive in the same address. For example, **192.168.0.1-127/24** is not supported. Use either **192.168.0.0/24** or **192.168.0.1-127**.

- These prefix-matching rules apply to the following policies only: Route, QoS, Optimization, NAT, Security, and ACLs.

DNS Proxy Policies

Configuration > Templates & Policies > Templates

If you select ON, complete the following steps to configure and define your DNS Proxy policies.

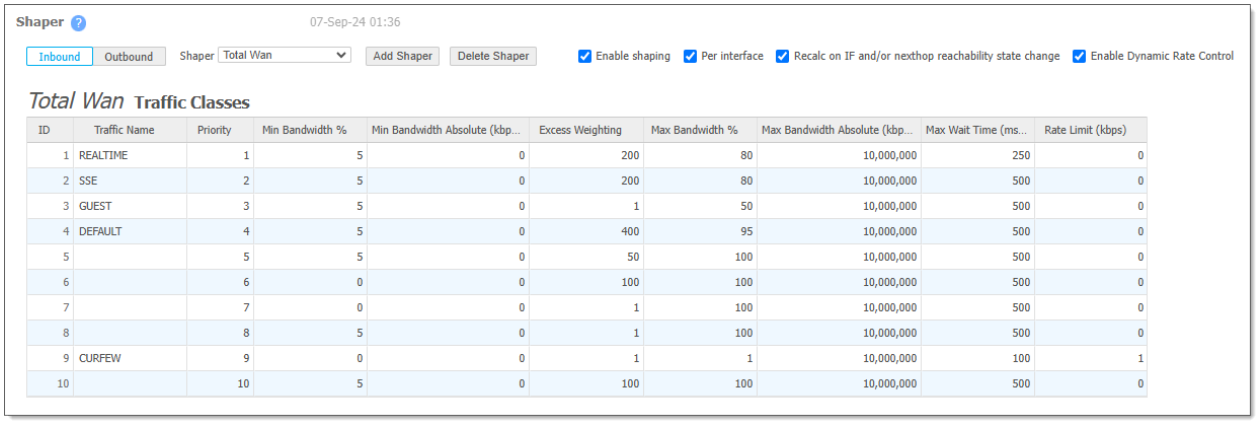
NOTE: This feature is configurable only if you have loopback interfaces configured.

1. Choose whether you want the DNS Proxy enabled by selecting **ON** or **OFF**.
2. Select the name of the loopback interface or LAN-side label associated with your DNS proxy.
3. Enter the IP addresses for Server A in the **Server A Addresses** field.
4. Choose whether you want Caching to be **ON** or **OFF**. If selected, the domain name to the IP address mapping is cached. By default, caching is **ON**.
5. Enter the domain names of the Server A for the above IP addresses.
6. Enter Server B IP addresses in the **Server B Addresses field**. Server B will be used if there are no matches to the Server A domains.

NOTE: You can **Clear DNS Cache**. This will erase the domain name to the IP address mapping you had cached for both Server A and B.

Shaper Template

The Shaper template provides a simplified way of globally configuring Quality of Service (QoS) on your appliances.



ID	Traffic Name	Priority	Min Bandwidth %	Min Bandwidth Absolute (kbp...	Excess Weighting	Max Bandwidth %	Max Bandwidth Absolute (kbp...	Max Wait Time (ms...	Rate Limit (kbps)
1	REALTIME	1	5	0	200	80	10,000,000	250	0
2	SSE	2	5	0	200	80	10,000,000	500	0
3	GUEST	3	5	0	1	50	10,000,000	500	0
4	DEFAULT	4	5	0	400	95	10,000,000	500	0
5		5	5	0	50	100	10,000,000	500	0
6		6	0	0	100	100	10,000,000	500	0
7		7	0	0	1	100	10,000,000	500	0
8		8	5	0	1	100	10,000,000	500	0
9	CURFEW	9	0	0	1	1	10,000,000	100	1
10		10	5	0	100	100	10,000,000	500	0

A shaper is a set of policies that control access and traffic flow on appliances by allocating bandwidth as a percentage of system bandwidth. Shaper parameters are organized into ten traffic classes. Four traffic classes are preconfigured (RealTime, CriticalApps, BulkApps, and Default). After compressing (deduplicating) all outbound tunnelized and passthrough-shaped traffic, Orchestrator either applies policy settings globally or upon each interface, shaping the traffic as it leaves the interface.

Applying the Shaper template to an appliance updates its system-level WAN shaper. Orchestrator preserves any interface-specific shapers that have been added. For minimum and maximum bandwidth, you can configure traffic class values as a percentage of total available system bandwidth and as an absolute value. The appliance always provides the larger of the minimum values and limits bandwidth to the lower of the maximum values. You can rename or edit any traffic class.

You can view applied configurations by navigating to **Configuration > Templates & Policies > Shaping > Shaper**. See Shaper Tab.

Add / Delete Shaper

To create an interface shaper, click **Add Shaper**. A shaper is defined only for a specified traffic direction (inbound or outbound).

To remove an interface shaper, click **Delete Shaper**. You cannot delete the Total WAN shapers (inbound or outbound).

Enable Interface Shaper

An outbound interface shaper is always enabled. You cannot clear the Enable shaping and Per interface check boxes for outbound interfaces.

Inbound interface shaping is disabled by default. To enable inbound traffic shaping, select the **Enable shaping** check box. To enable inbound interface shaping, select both the **Enable shaping** and **Per Interface** check boxes. For more information, see Per-interface Shaping below.

Shaper Options

Option	Description
Inbound / Outbound	Enables you to select the direction to which to apply shaper configurations. **Inbound** (for WAN to LAN) **Outbound** (for LAN to WAN)
Shaper	Select the shaper to use, both for inbound and outbound traffic. By default, the selected shaper is Total WAN, which provides shaping against the total WAN bandwidth on the appliance. Because this shaper is generally sufficient, you should not need to use any other shapers.

Option	Description
Add Shaper	If desired, you can add a shaper. However, because the Total WAN shaper provided by Orchestrator is generally sufficient, you should not need to add other shapers.
Delete Shaper	Enables you to delete shapers you have explicitly added.
Enable shaping	Indicates whether to enable shaping of traffic. Shaping is always enabled for outbound traffic. Shaping for inbound traffic is recommended and best practice, but optional.
Per interface	Indicates whether to shape traffic on an individual interface level. It is recommended and best practice to select this option. Shaping is based on per-interface bandwidth and total system bandwidth. For more information, see Per-interface Shaping below.
Recalc on IF and/or nexthop reachability state change	Indicates whether to recalculate shaper bandwidth based on the loss of any WAN-side interface or next-hop reachability. It is recommended and best practice to select this option.
Enable Dynamic Rate Control	Indicates whether to enable Dynamic Rate Control (DRC). This feature prevents many-to-one bandwidth oversubscription. This option is available only for inbound traffic. For more information, see Dynamic Rate Control below.

Shaper Configuration

Field	Description
Traffic Name	Name assigned to the traffic class (by Orchestrator or the user).
Priority	Order in which to allocate each class' minimum bandwidth. Valid values are 1 to 10 with 1 as first priority and 10 as last.
Min Bandwidth %	Minimum percentage of bandwidth guaranteed to the traffic class, allocated by priority. However, if the sum of the percentages is greater than 100%, lower-priority traffic classes might not receive their guaranteed bandwidth if it is all consumed by higher-priority traffic. Max overrides Min if you set Min Bandwidth % to a value greater than Max Bandwidth %.
Min Bandwidth Absolute (kbps)	Minimum bandwidth (in kbps) for the traffic class, which guarantees a minimum level of service when total system bandwidth declines. This is useful for maintaining, for example, the quality of VoIP.

Field	Description
Excess Weighting	If bandwidth remains after satisfying the minimum bandwidth percentages, the excess is distributed among the traffic classes in proportion to the weightings specified in this column. Valid values are 1 to 10,000.
Max Bandwidth %	Maximum percentage of bandwidth that a traffic class can use (as a percentage of total available system bandwidth).
Max Bandwidth Absolute (kbps)	Maximum bandwidth (in kbps) for the traffic class, which provides an absolute upper limit for bandwidth. This is useful for capping, for example, the bandwidth of downloads and streaming services.
Max Wait Time (ms)	Any packets waiting longer than this specified waiting time (in ms) are dropped.
Rate Limit (kbps)	Per-flow rate limit (in kbps) for the traffic class. For no limit, specify 0 (zero).

Per-interface Shaping

When you select the Per interface check box, the appliance applies shaping per WAN interface for inbound traffic in addition to system-level shaping.

NOTE: For outbound traffic, shaping is inherently per interface because packets leave on a specific WAN link.

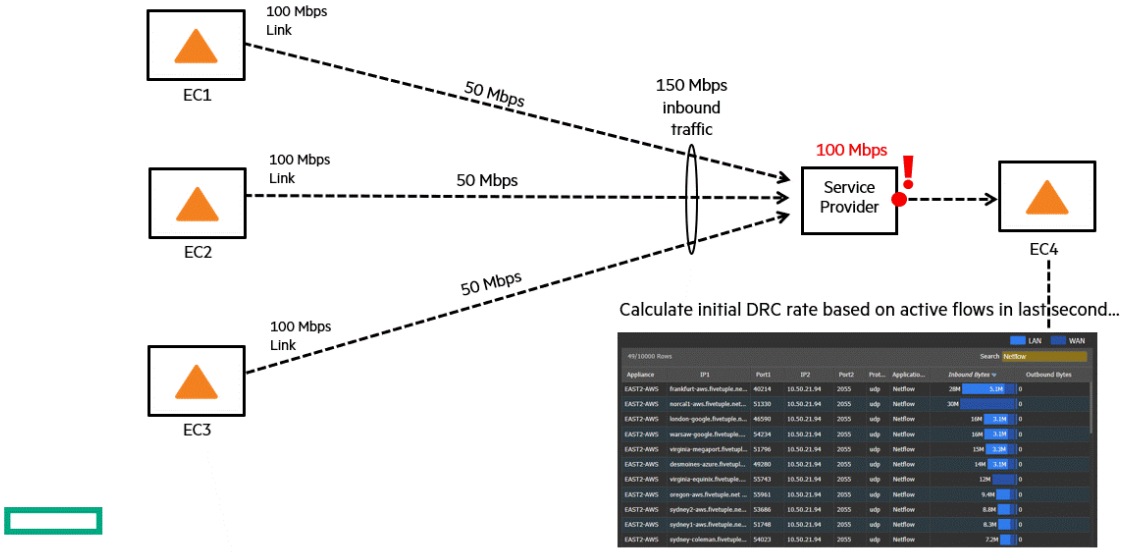
Shaping is handled in two layers: the system (appliance-wide) shaper enforces an aggregate bandwidth cap across all WAN interfaces while the per-interface shaper applies a cap separately on each WAN interface.

Each interface is shaped using that interface's configured bandwidth. Traffic arriving on one link cannot exceed that link's capacity or starve other links. If you do not select the Per interface check box, inbound shaping uses only the system-level aggregate rate. As a result, all interfaces share a single bandwidth cap, which can cause high-speed links to carry most of the traffic while overloads and increased loss can occur on lower-speed links. It is highly recommended that you always enable this feature for appliances with multiple WAN interfaces.

Dynamic Rate Control

Dynamic Rate Control (DRC) allows the EdgeConnect to prevent many-to-one bandwidth oversubscription by automatically adjusting per-flow bandwidth. If the EdgeConnect experiences congestion (drops or wait time), the EdgeConnect automatically regulates traffic by lowering each remote appliance's per-flow rate. The following animation illustrates this process.

Step 1: Oversubscribed EC4 inbound traffic



To enable this feature, select the Inbound filter, and then select the **Enable Dynamic Rate Control** check box on the Shaper template or the Shaper dialog box.

IMPORTANT: DRC is driven by QoS drops and shaper wait time. If you configure the EdgeConnect with an Inbound Shaper value greater than or equal to the service provider, the EdgeConnect will never see drops or wait time and your traffic will not receive the benefits of DRC.

Management Services Template

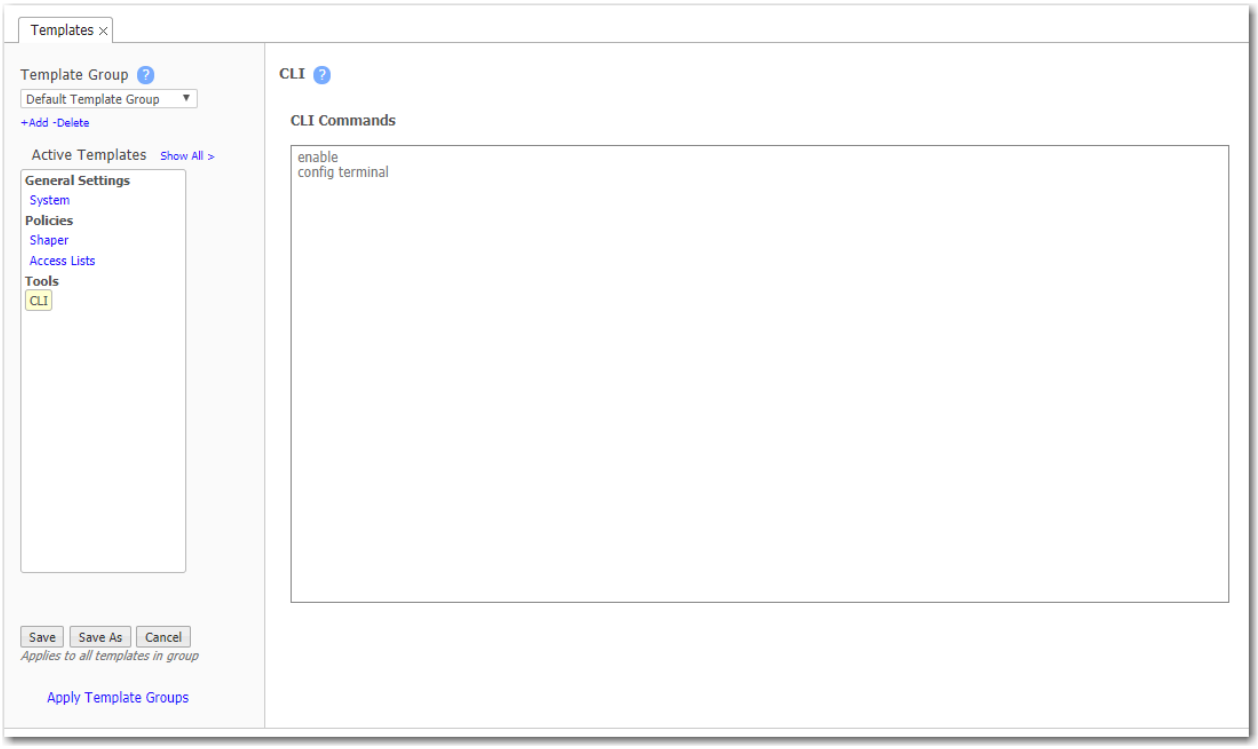
Use this template to globally apply the modifications made to your Management Services if segmentation is enabled or disabled. “any” is used as the default Interface for the Source IP address; however, you can change the interface with any interfaces you have previously configured on the Management Services tab. To modify the interface, click **any** in the table. For more information, see Management Services.

WARNING: Changing the source interface setting for “HTTP(S), Cloud Portal, Orchestrator, Secure Web Services” impacts robust appliance reachability to the Cloud Portal, Orchestrator, and Secure Web Services. If you change it, you could lose management plane connectivity. It is highly recommended that you leave this setting set to “any”.

CLI Template

Use this template to enter any sequence of **Command Line Interface (CLI)** commands.

Enter each CLI command on a new line.



Orchestrator sends all commands entered in the CLI template with a single batch command to EdgeConnect gateway. This batch command is sent through a single REST call that has a 60 second timeout. Placing enough commands to cause the REST call to timeout can result in unpredictable behavior.

CLI commands can be entered directly in the CLI to verify the time required to execute them. We recommend that the command set entered in the CLI Template page requires a maximum of 30 seconds when entered manually in the CLI.

Session Management Template

Use this page to configure settings that control access to the appliance web UI and enable Web Proxy.

Field	Description
Auto Logout	Specifies the amount of time in minutes after which an inactive session will be automatically logged out. The valid range is 0-60. Use 0 to disable automatic logout.
Max Sessions	Maximum number of active sessions on the appliance. If the maximum number of sessions is reached, users who try to log in to the appliance web UI will receive a message that the browser cannot access the appliance. On non-EdgeConnect appliance models, Orchestrator might not be able to access the appliance.

Field	Description
OpenSSL Cipher List	List of cipher suites to enable or disable on the appliance. For details about formatting this string, see this page. The string can only contain the following characters: a-z, A-Z, 0-9, and +-:.!_@ **WARNING:** Cipher format and availability are not validated. Ciphers should be thoroughly tested in a lab environment before being applied. When ciphers are applied from a template, an improperly formatted string or unavailable ciphers can cause an appliance crash.
Web Protocol	Select the web protocol to use for appliance UI sessions. HTTPS is recommended for maximum security.

Web Proxy

Web proxies are designed to intercept, process, and forward HTTP/HTTPS traffic at scale. They are built to handle resource-intensive tasks such as SSL decryption, content filtering, and malware scanning without degrading performance.

When you enable the **Web Proxy** option, provide the Host, Port, Username, and Password.

Secure Web Services Template

Use this template to globally enable and configure Secure Web Services on appliances. Your appliances must be licensed to use Secure Web Services. For detailed information about Secure Web Services and licensing, see Secure Web Services.

1. To enable Secure Web Services configuration, move the **Secure Web Services Config** toggle to the right.
2. Complete the following configuration fields as appropriate:

Field	Description
Enable URL Classification	Move the toggle to the right to enable the URL classification feature on appliances.
Incremental Update *(for URL Classification)*	Select this check box to enable real-time updates to appliance URL classification databases.

Field	Description
Query on miss	Select this check box to enable query on miss of URL classification. When categorization of a queried URL is not available in the current URL-related threat intelligence database, the query is sent over the network to the cloud-based OpenText™ Threat Intelligence (BrightCloud) platform. The asynchronous response is then internally cached so that future queries for the URL category and reputation will return information successfully.
Unclassified URL Reputation	From the drop-down list, select a risk category to assign to flows that do not have URL reputation score information in the current URL-related threat intelligence database. Risk categories are: High Risk (1 - 20) Suspicious (21 - 40) Moderate Risk (41 - 60) Low Risk (61 - 80) Trustworthy (81 - 100) The default risk category is High Risk.
Enable IP Reputation	Move the toggle to the right to enable the IP reputation feature on appliances.
Incremental Update *(for IP Reputation)*	Select this check box to enable real-time updates to appliance IP reputation databases.

3. Click **Save**.

Apply Template Groups

Configuration > Templates & Policies > Apply Template Groups

Use this tab to add or remove template groups from appliances.

- To add or remove template groups from appliances, select the appliances, select the **Add** check box for each template to add or clear the **Remove** check box for each template to remove, and then click **Apply** to save and apply the changes.
- Drag template groups up or down to reorder the list. Orchestrator automatically applies any changes to template groups to their associated appliances.

  NOTE: If multiple template groups are applied to an appliance, the order in which they are applied determines which template is used. Template groups applied later (lower on the apply order list) overwrite any conflicting templates applied earlier.
- Template groups that are already associated with appliances will be updated on those appliances when:
  - Any template in the template group is added, edited or removed.

- Orchestrator restarts.
- An appliance reboots.

The following table describes the information displayed on this tab for each appliance.

Field	Description
Appliance	The name of the appliance.
Present	The template groups applied to the appliance.
Changes	Any recent changes made to the template groups that are applied to the appliance.

Best Practices When Applying Template Groups

Do: Use templates for all fabric configurations

Avoid unique configurations in your network by using template groups and templates to apply configurations consistently to every appliance in your SD-WAN fabric.

Do: Use Default Template Group for global configuration settings

The Default Template Group cannot be removed. It is a recommended best practice to put all global appliance configuration items into the Default Template Group, including Security Policies, Access Lists, Shaper, and QoS Policies, as these are typically standard across the network.

Do: Use additional template groups for different configurations across appliance groups

Multiple template groups can be created, each containing different templates and different attributes of those templates. Additionally, each appliance or group of appliances can have multiple template groups applied to them. For more information, see Multiple Template Groups Example.

Do: Pay attention to the Template Apply Order

Template groups are applied to the appliances based on the Template Apply Order.

NOTE: The last template group in the list will be the final configuration applied to the appliance.



This order can be changed by simply dragging and dropping the template groups into the desired order. It is recommended to use the Default Template Group and keep it at the top of the list.

For more information, see Apply Template Groups Behavior and Multiple Template Groups Example.

Do: Make template changes in a limited scope for testing/validation of changes

Global changes applied via templates can be service impacting. It is recommended to make changes that are globally significant to a small portion of the network to ensure there is no service disruption related to the changes.

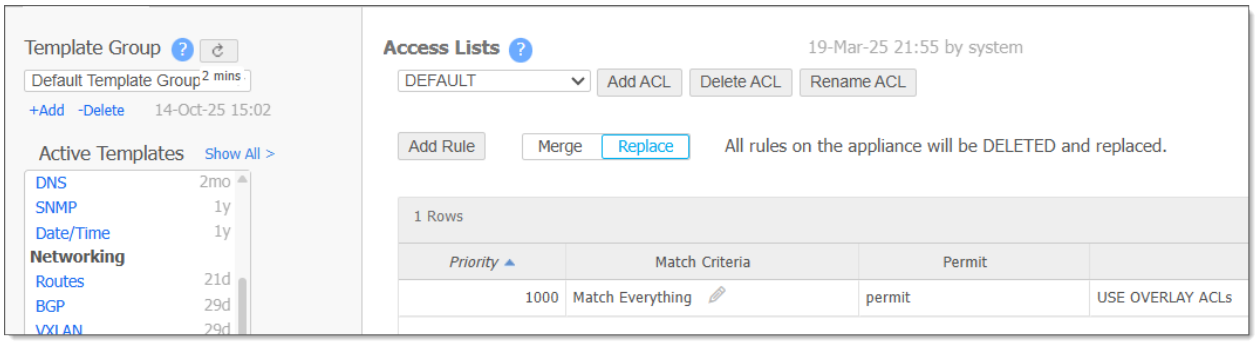
Do: Use the System template in template groups

The System template should be included inside any template group because it contains the MSS (Maximum Segment Size) value.

Do: Use Replace if Merge/Replace are available

Some templates have Merge and Replace options, which appear as Merge and Replace buttons on the template, as shown in the following figure. If a template does not have Merge and Replace options, the template is always replaced. In most cases, Replace should be used instead of Merge.

The Replace option replaces the configuration contained in the template on the appliance when any changes are made to the template. This ensures that the configuration is always consistent with the template.



Do Not: Use Merge

The Merge option allows changes in the template to be merged with what is currently in place on the appliance, which can lead to configuration inconsistencies. Merge should be avoided in most cases. One exception where Merge should always be used is the Network Access Control (NAC) template.

The following table is an example of Merge behavior. The newly applied template settings override all old settings (aclMap 1, 2 and 4). Where the new template does not have specific settings, the previous settings remain (aclMap3).

Template Settings Applied via Merge	Current Appliance Settings	Appliance Settings After Merge
aclMap1	aclMap1	aclMap1
100 abc	100 currabc	100 abc
200 def	100 currdef	200 def
aclMap2	aclMap2	aclMap2
100 mng	100 currmng	100 mng
200 twn	300 ttp	200 twn
——-	aclMap3	aclMap3
——-	200 efg	200 efg
——-	300 dng	300 dng
aclMap4	——-	aclMap4
300 hello	——-	300 hello
301 world	——-	301 world

Do Not: Make changes locally on an appliance

Templates ensure consistency of configurations across your entire fabric. Local changes to an appliance are unnecessary, except for the following:

1. Deployment: Each appliance has a unique deployment configuration that consists of IP addressing, Labels, Bandwidth, and WAN Optimization.
2. Routing and VRRP: BGP, OSPF, VRRP, and Multicast configurations are typically unique to each appliance.
3. VTI and Passthrough Tunnels: Third Party IPSec tunnels and VTIs are unique to each appliance.

All items noted above can be configured manually or through Preconfiguration, which can be used to model each of these unique configurations inside a YAML file stored on the Orchestrator. These configurations can be used to model everything on the appliance, or to add/update parts of the configuration.

For more information on preconfiguration, see Preconfigure Appliances.

Apply Template Groups Behavior

Using templates in Orchestrator helps maintain consistency across the network footprint. It is important to understand the behavior of how templates are applied and what actions trigger an update or synchronization of configuration in a template.

Templates are always pushed in a top-down manner from the Orchestrator to the appliances, and in the order of priority based on the Template Apply Order.

The following table details the actions that trigger an application or re-application of the template configuration.

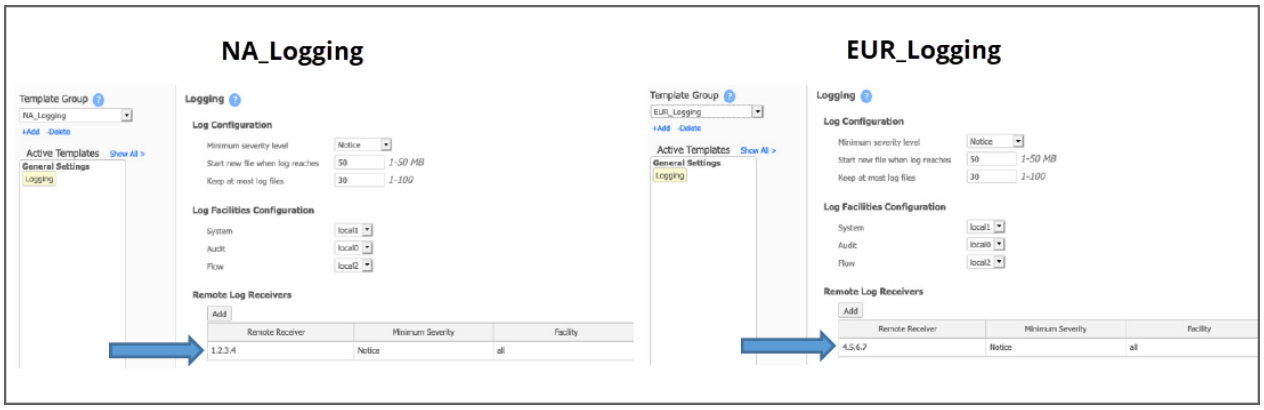
Event	Action	Result
Appliance Reboot	All associated template groups are re-applied to appliance, based on the Template Apply Order.	Any local changes to the appliance that are part of a template group are corrected by the Orchestrator.
Orchestrator Reboot	All template groups are re-applied to all associated appliances.	Any local changes to the appliance that are part of a template group are corrected by the Orchestrator.
Add, change, or delete new template group that is not yet associated with any appliance(s).	All template groups are re-applied to associated appliances, based on the Template Apply Order.	None.
Add, change, or delete template group that is associated with appliance(s).	All template groups are re-applied to associated appliances, based on the Template Apply Order.	Any local changes to the appliance that are part of a template group are corrected by the Orchestrator.
Apply template group to a previously non-associated appliance.	All associated template groups are re-applied to newly associated appliance(s), based on the Template Apply Order.	Any local changes to the appliance that are part of a template group are corrected by the Orchestrator.
Remove template group from associated appliance.	Template association is removed within Orchestrator.	Any configuration that was part of the template remains in place on the appliance, but any future updates to the template will not be applied.

Multiple Template Groups Example

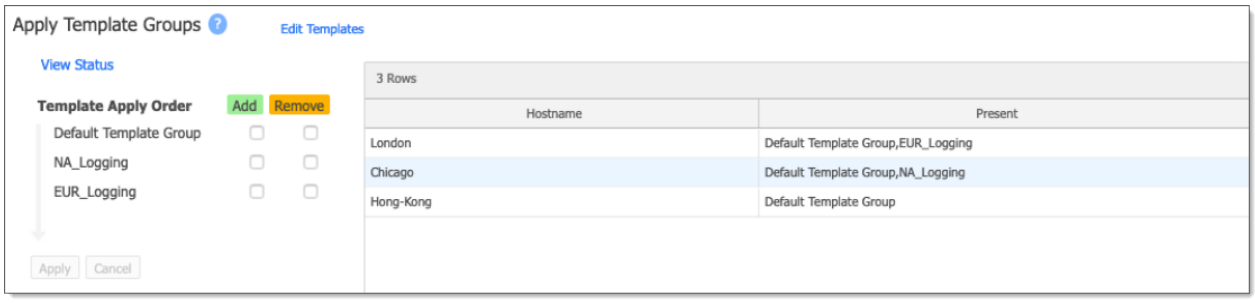
This section describes an example network that consists of appliances deployed in North America, Europe, and Asia.

- Appliances in North America should point to a regional syslog collector (IP = 1.2.3.4).
- Appliances in Europe should point to a different collector (IP = 4.5.6.7).
- Appliances in Asia should point to the global collector (IP = 8.9.10.11), which is configured in the Default Template Group.

To accomplish this, you would add two new template groups.

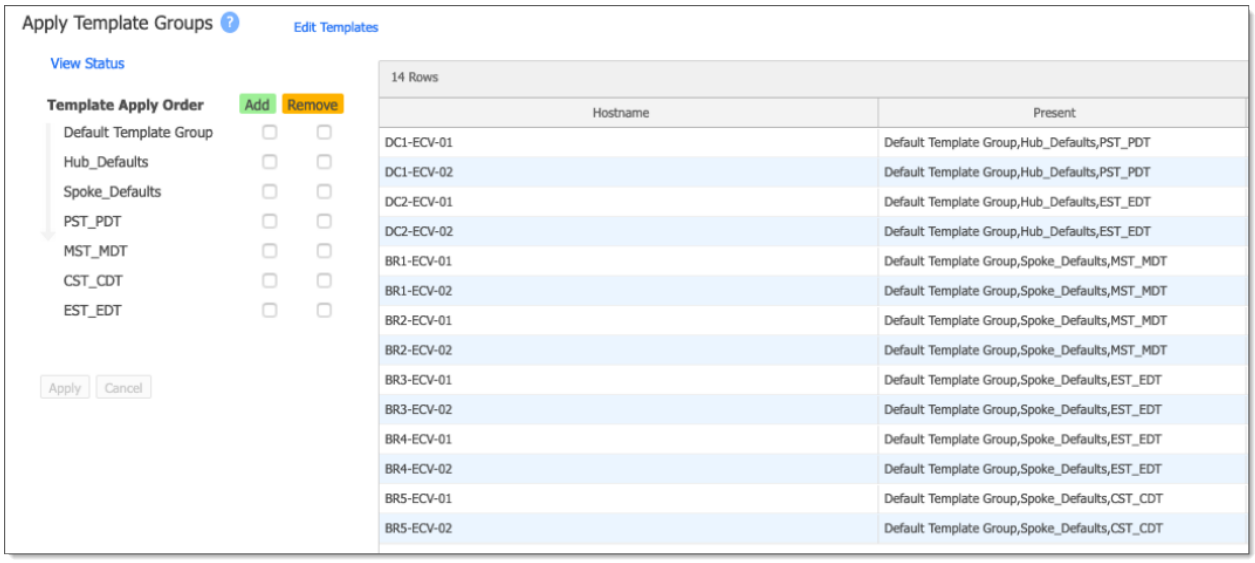


After configuration, apply the EUR_Logging template group to the Europe appliances and the NA_Logging template to the North America appliances.



This points the North America appliances to the logging server at 1.2.3.4 and the Europe appliances to 4.5.6.7. Any other appliances continue to use the logging server configured in the Default Template group, which is 8.9.10.11.

The following image shows the Apply Template Groups tab from an example Orchestrator that has a configuration with seven template groups. Each template group contains templates that are configured based on the appliance role, administrative requirements, geographic location, and other factors.

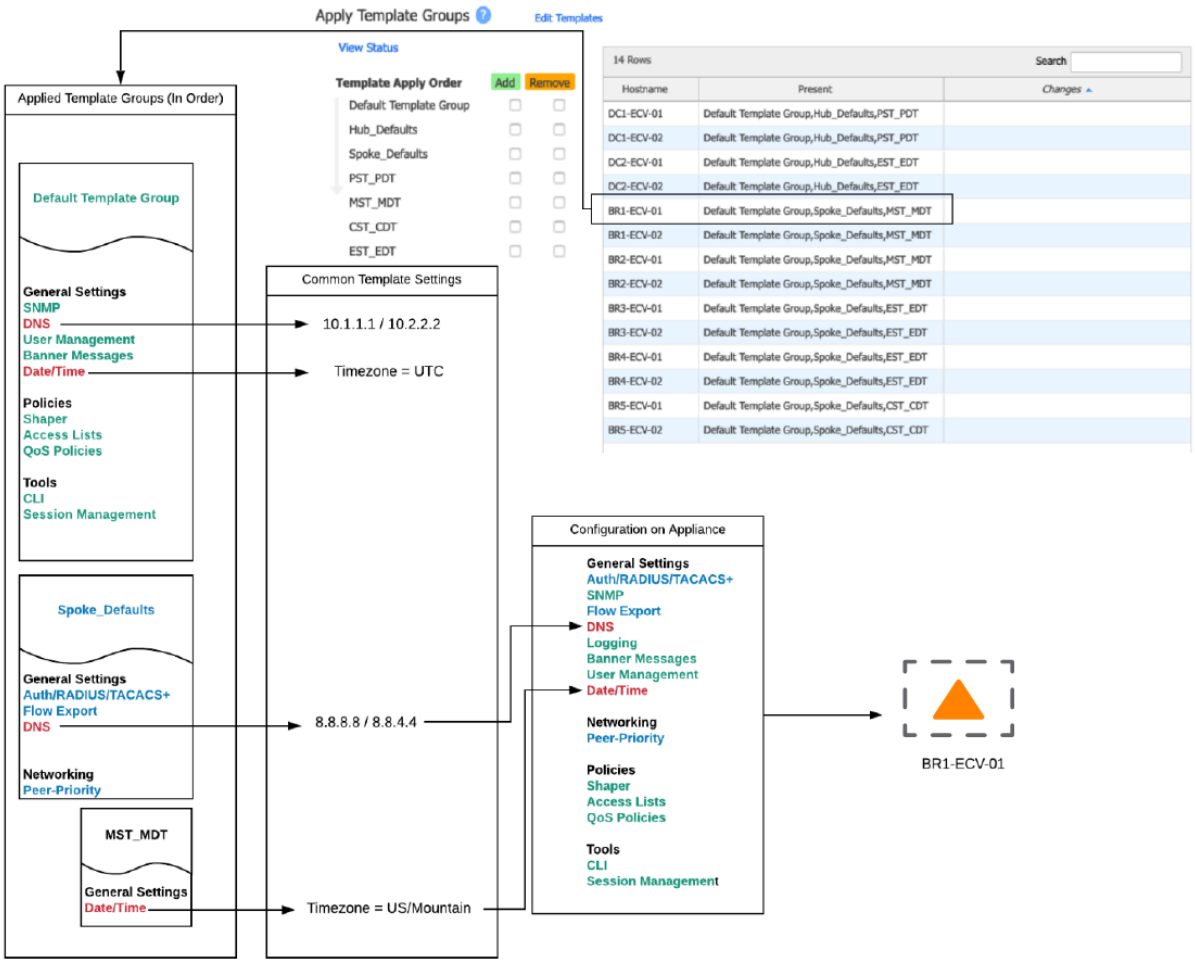


The following table outlines how the seven template groups are applied to the appliances.

Template Group Name	Appliances the Template Group is Applied To
Default Template Group	Applied to ALL appliances in the SD-WAN fabric.
Hub_Defaults	Applied to All Hub appliances in the SD-WAN fabric, and NOT applied to Spoke appliances.
Spoke_Defaults	Applied to All Spoke appliances in the SD-WAN fabric, and NOT applied to Hub appliances.
PST_PDT	Applied to all appliances in the Pacific Time Zone.
MST_MDT	Applied to all appliances in the Mountain Time Zone.
CST_CDT	Applied to all appliances in the Central Time Zone.
EST_EDT	Applied to all appliances in the Eastern Time Zone.

The following image shows a breakdown of the final appliance configuration for this example, when the same templates are part of multiple template groups. In this example the following is true:

- The Default Template Group contains DNS and Date/Time.
- Spoke_Defaults contains DNS with different values than the Default Template Group.
- MST_DST contains Date/Time with different values than the Default Template Group.



All the unique templates from each template group are consolidated and applied to the appliance for common template configuration between template groups. The last template group in the Apply Order is the final configuration on the appliance.

Template Options

The following table has general guidelines about whether templates are recommended in the default group or in a new or separate group. Consult with Technical Support for specific details on template configuration.

Template Name	Default Template Group	Separate Template Group	Notes
System	Yes	Case-by-case	Should be included in the Default Template Group because it contains the MSS (Maximum Segment Size) value.

Template Name	Default Template Group	Separate Template Group	Notes
Auth/Radius/TACAC	Yes		For local appliance authentication via HTTPS/SSH.
SNMP	Yes	Case-by-case	For appliance SNMP, which is disabled by default.
Network Access Control (NAC)	No	Case-by-case	Used to configure MAC Auth and 802.1x.
Appliance IP Allow List	Yes		Used to configure access to SNMP, HTTP/HTTPS, and SSH on the local appliance.
Flow Export	Yes		Typically global, but this could be configured in a region-specific template.
DNS	Yes		Typically global, but this could be configured in a region-specific template.
Logging	Yes		Typically global, but this could be configured in a region-specific template.
Banner Messages	Yes		Typically global, but this could be configured in a region-specific template.
HTTPS Certificate	No	Case-by-case	For the local appliance WebUI certificate.
User Management	Yes		Global configuration of local accounts for “password of last resort” access is recommended. This allows for easy lifecycle management of local passwords.
Date/Time	Case-by-case		Some deployments use a global time standard; others require appliance time to be based on physical location.
Password Settings	Yes		Used to configure password complexity requirements and expiration for local accounts.
SSL Certificates	Case-by-case		Needed for SSL Decryption/WAN Op only.
SSL CA Certificates	Case-by-case		Needed for SSL Decryption/WAN Op only.

Template Name	Default Template Group	Separate Template Group	Notes
VRRP	No	Case-by-case	Typically configured per site, but you could use this to set a global standard for “VRRP_Primary” and “VRRP_Secondary.”
Peer Priority	No	Yes	Typically set on branches for selecting a preferred peer when multiple duplicate routes are present.
Route Redistribution Maps	Yes		Typically global, but this could be configured in a region-specific template.
BGP	Yes		Typically global, but this could be configured in a region-specific template.
OSPF	Yes		Typically global, but this could be configured in a region-specific template.
BFD	Yes		Typically global, but this could be configured in a region-specific template.
VXLAN	Yes		Typically global, but this could be configured in a region-specific template.
Admin Distance	No	Yes	Admin distance is often configured differently for hub versus spoke appliances, and this varies based on environment.
Routes	No	Yes	Desired routes are often configured differently for hub versus spoke appliances, and this varies based on environment.
Shaper	Yes		Typically global, should remain global.
User Defined Apps	No	No	Deprecated. Use Application Definitions in Orchestrator.
Access Lists	Yes		Typically global, should remain global.

Template Name	Default Template Group	Separate Template Group	Notes
Route Policies	No	Case-by-case	Templatized or manual route policies should be avoided as the overlay should take care of this. Route policies are typically used only for troubleshooting purposes.
QoS Policies	Yes		Typically global, should remain global.
Optimization Policies	Yes		Typically global, should remain global.
SaaS NAT Policies	No	Case-by-case	If SaaS Optimization is used, hubs may have a different configuration than spokes.
Threshold Crossing Alerts	No	Yes	TCAs can be customized by region; including these in a region-specific template may be desired.
SaaS Optimization	No	Case-by-case	If SaaS Optimization is used, hubs may have a different configuration than spokes.
Security Policies	Yes		Typically global, should remain global.
DNS Proxy Policies	Yes		Typically global, should remain global.
Management Services	Yes		Typically global, should remain global.
Firewall Protection Profiles	Yes		Typically global, but this could be configured in a region-specific template.
CLI	Yes		Typically global, should remain global.
Session Management	Yes		Typically global, should remain global.
Secure Web Services	Yes		Typically global, should remain global.

Configuration > Cloud Services

The options under **Configuration > Cloud Services** focus on the various cloud services that are offered.

NOTE: Be aware that design changes that occur in third-party applications (especially the user interfaces) could affect instructions provided in topics in this section. Therefore, instructions are provided as guidelines rather than precise steps.

AWS Network Manager

Configuration > Cloud Services > AWS Network Manager

NOTE: In Orchestrator 9.5.2, this feature was updated to incorporate new functionality for Transit Gateways and Core Network Edge. This help topic has not yet been updated to incorporate these changes, so the information below may be out of date if you are running Orchestrator 9.5.2 or later. To view the most current version of the help for this topic, visit the AWS Network Manager page of the HPE Aruba Networking EdgeConnect SD-WAN Documentation site; this page will be updated with the most current help information when it is available.

Orchestrator supports association with Amazon Web Services and their Transit Gateway Network Manager. Orchestrator builds AWS Site-to-Site VPN tunnels, enabling you to securely connect your on-premises network to one or more Transit Gateways (TGWs).

Before you begin configuring AWS Transit Gateway Network Manager in Orchestrator, create an AWS account to authenticate and authorize Orchestrator with your AWS account. Then complete the prerequisite tasks in the following section.

Prerequisites for AWS Transit Gateway Network Manager

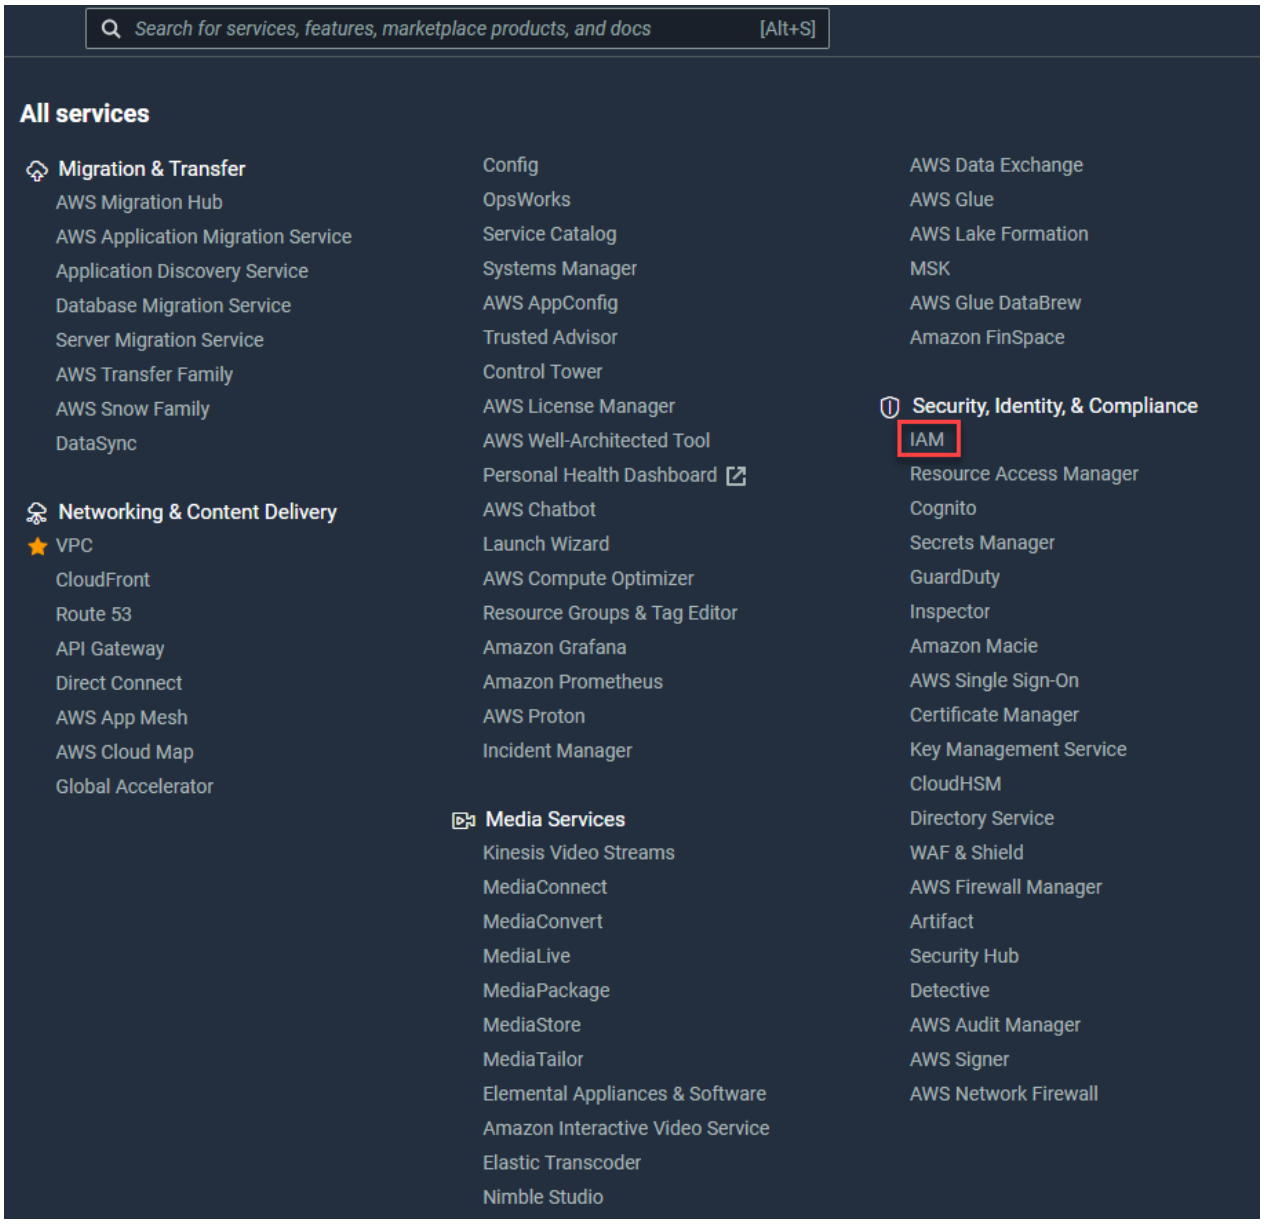
Make sure you complete the following tasks in AWS console before configuring Orchestrator:

- Navigate to the **Identity and Access Management** (IAM) under **Services** to create a user profile with permissions for Orchestrator.
- Navigate to the **Virtual Private Cloud (VPC) Dashboard** and configure your Transit Gateways for the regions you want.
- Navigate to **Network Manager** from the **VPC Dashboard** under **Transit Gateways** to create a Global Network.
- Associate your Transit Gateways to the Global Network.

Create a User Profile in AWS

To create a user profile in AWS, complete the following steps:

1. Sign in to AWS and navigate to the **Identity and Access Management (IAM)** service from the main AWS Management Console (**Services > Security, Identity, & Compliance > IAM**).



2. Click **User** in the left menu under **Access Management**.
3. Click **Add User**.
4. Enter a username in the **User name** field.
5. Choose the **Access Type**: **Programmatic Access** and **AWS Management Console Access**.
6. Click **Next: Permissions**.
7. Set the Permissions for your user on this page. You can do this in one of three ways:
    - **Adding a user to your group** – The user will inherit the permissions assigned to the group.

- **Copying permissions from an existing user** – Copy permissions from an existing user in AWS and assign them to the user you want.
- **Attaching existing policies directly** – Attach a file containing the permissions and assign it to the user.

8. Assign optional tags for your user. If you choose to add a tag, complete these steps:
    1. Enter a **key** – This represents the name of your tag.
    2. Enter a **value** – Enter text that you want the key/tag to represent.

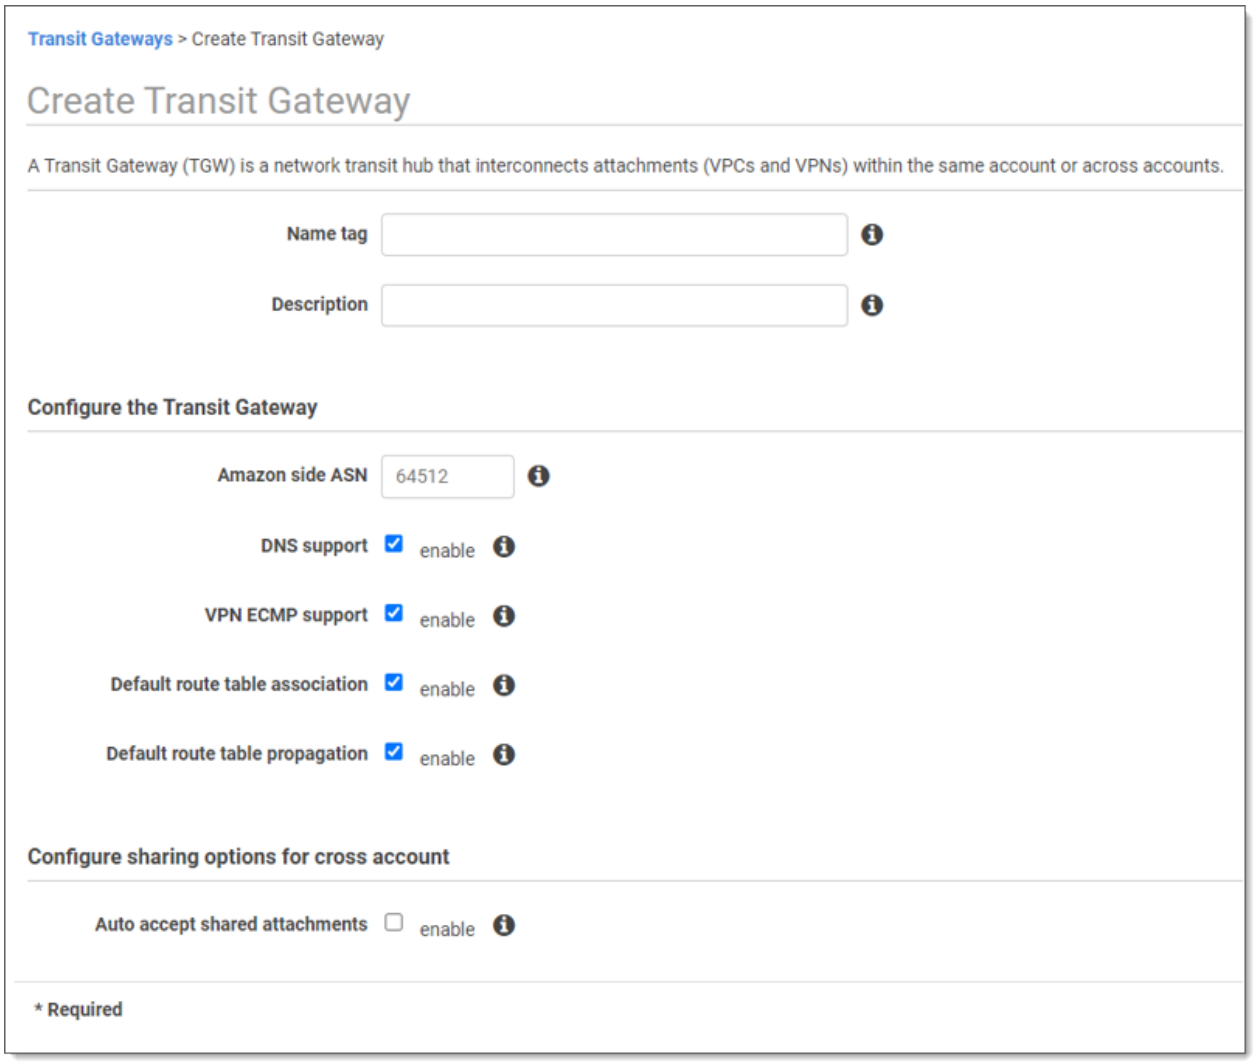
       NOTE: Tags enable you to provide additional information about your user or group for tracking and organizational purposes. Up to 50 tags are allowed.
9. Select **Next: Review**. This page displays the review of the profile you just created for your user. The **User Details**, **Permissions Summary**, and additional information such as tag, are shown.
10. Select **Create User**. The page should now show the following success message, along with **Access Key ID** and the **Secret Access Key** associated with your configured user. Copy and paste the Access Key ID and the Secret Access Key to a secure place for later use. You will need these when adding the AWS account on Orchestrator.

Create Transit Gateways

Next, you must create Transit Gateways (or select existing Transit Gateways you have already created) to associate with your AWS Network Manager, which you create in the steps below. Transit Gateways will terminate the Site-to-Site IPSec tunnels established from the EdgeConnect appliances in your network.

To create a new Transit Gateway, complete the following steps:

1. Navigate to the **Virtual Private Cloud (VPC) Dashboard** (**Services > Networking & Content Delivery**).
2. Click **Transit Gateways**, under **Transit Gateways** in the left menu.
3. Click **Create Transit Gateways**.



4. Complete the following fields to create the Transit Gateway.

Field	Description
Name Tag	Enter a name that represents your Transit Gateway.
Description	Enter a description to help identify your Transit Gateway. This is the description for the Name Tag.
Amazon side ASN	Autonomous System Number that represents your Transit Gateways in AWS. You can use an existing ASN assigned to your global network or a private ASN. See the range limitations in AWS.
DNS Support	Select this check box if you want to enable Domain Name System support for your VPC within your Transit Gateways.

Field	Description
VPN ECMP support	Select this check box if you want to enable Equal Cost Multi-Path routing support in your Transit Gateways. This allows traffic with the same source and destination to be sent across the same multiple paths.
Default Route Table Association	Select this check box if you want to automatically associate other Transit Gateways to the route table that this one is using.
Default Route Propagation	Select this check box if you want to automatically create other Transit Gateways with this same route table.
Auto-accept shared attachments	Select this check box if you want your transit gateways to automatically accept attachments associated with different accounts.

5. Click **Create Transit Gateway**. A success message should display along with your Transit Gateway ID.

Create a Network Manager

After you create your Transit Gateway, you must create a Global Network in AWS. A Global Network hosts your specified Transit Gateways. It is managed by the AWS Network Manager.

1. Navigate to the **VPC Dashboard**.
2. Click **Network Manager** under **Transit Gateways**.
3. Click **Create Global Network**.
4. Enter a **Name** and **Description** for your **Global Network**.
5. Click **Create**.

Orchestrator Configuration

After completing the AWS prerequisites, navigate to the **AWS Network Manager** tab in Orchestrator. There are seven buttons above the table on this tab that you use to complete the AWS and Orchestrator integration: **Subscription**, **Interface Labels**, **Tunnel Settings**, **VTI Subnet Pool**, **Segment & Zone Association**, **AWS Resources**, and **Appliance Association**.

Subscription

1. To begin, click the **Subscription** button.
2. Enter the **Access Key ID** and the **Secret Access Key** that reflect your user account in AWS. This is the Access Key ID and the Secret Access Key you copied earlier in the Create a User Profile in AWS section.
3. Click **Save** after you finish entering the information in the table below. The **AWS Reachability** field should reflect **Connected**.

Field	Description
AWS Reachability	Connection status of the AWS Network Manager to Orchestrator: **Connected** or **Not Connected**.
Access Key ID	Access Key given to you in AWS to log in to the AWS console.
Secret Access Key	Secret Access Key given to you in AWS to log in to the AWS console.
Polling Interval	Indicates how often Orchestrator should check for configuration changes in the AWS transit gateways or Network Manager. The default polling interval is ten minutes.

4. Click **Save**.

You now should have an established connection with Orchestrator to your AWS account.

Interface Labels

The Build Tunnels Using These Interfaces dialog box enables you to select the interfaces to build your tunnels to AWS.

1. Click the **Interface Labels** button. The Build Tunnels Using These Interfaces dialog box opens.
2. Drag the interface labels you want to apply from the column on the right into the **Primary** columns.
3. Click **Save**.

Tunnel Settings

The **Tunnel Settings** dialog box shows IKE and IPSec parameters used by Orchestrator when building Site-to-Site IPSec tunnels from the EdgeConnect appliances to the Transit Gateways. No changes are necessary for these parameters.

VTI Subnet Pool

In this dialog box, set the Subnet IP address and the mask for the AWS subnet pool. Enter the subnet IP address and the mask ID in the designated fields.

- Any updates to the subnet pool configuration results in service disruption.
- You can have duplicated ASNs if you have a site with the same name.

NOTE: This is an AWS-specific subnet pool. Therefore, every subnet IP address must start with **169.254** to be included in this pool.

Segment & Zone Association

You can apply configured segments to your VTI interfaces associated for AWS. Click the **Zone** icon and select the zone you want to apply from the drop-down list.

AWS Resources

Documentation for this section is under development.

Appliance Association

In this dialog box, you can choose which Transit Gateways you want to associate with your EdgeConnect appliances.

NOTE: You must first select the EdgeConnect appliances on the Orchestrator appliance tree, and then open the Network Manager Association tab to associate the appliances to your Transit Gateways.

1. Select or clear the check box next to the appliance you want to connect to or disconnect from the Network Manager.
2. See the following table for field descriptions.

Field	Description
Hostname	Host name of the appliance you want to connect to or disconnect from the Network Manager.
Transit Gateways Present	Lists the Transit Gateways that are already associated with the EdgeConnect appliances.
Transit Gateways Changes	Displays the EdgeConnect appliances that will be added or removed from the Transit Gateways.

3. Click **Save**.

   Orchestrator starts to establish the Site-to-Site IPSec tunnels from the EdgeConnect appliances to the selected Transit Gateways.

Verification

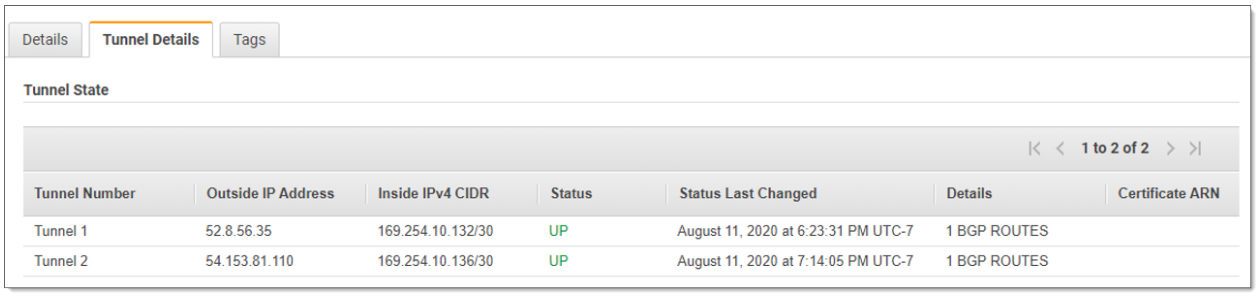
You can verify the stability and connectivity of your tunnels to the AWS Network Manager using the Connection Status column on the AWS Network Manager tab. This column shows the BGP Peer status. You can find additional details on the **Tunnels**, **VTI**, and **BGP** tabs.

Also, you can verify the AWS resources that Orchestrator created on the VPC Dashboard. To view the resources on the VPC dashboard, navigate back to the **Virtual Private Network** section in AWS and select **Customer Gateways** and **Site-to-Site VPN Connections**. On these tabs, you can confirm that the IPSec tunnels you created in Orchestrator are functioning correctly.

The tunnels should be in the ‘available’ state.

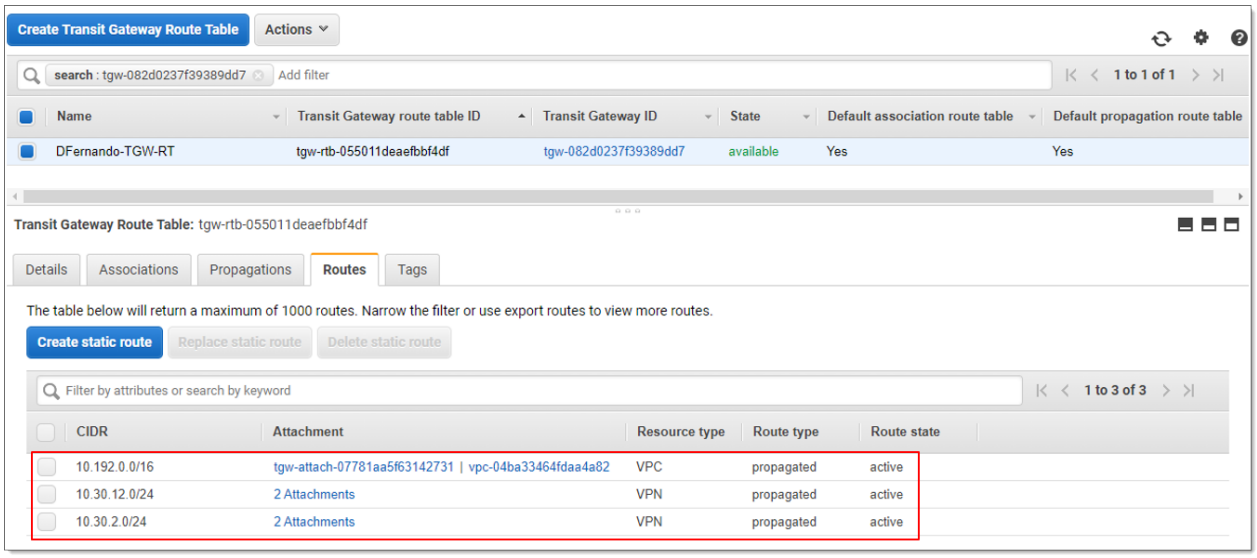


The IPSec tunnel statuses should be ‘UP’.



Route Tables and Static Routes

After the tunnels and the BGP sessions are established, the TGW route table shows the routes learned from the EdgeConnect devices. To create a route table for your transit gateways, navigate to the **VPC Dashboard** in AWS and click **Transit Gateway Route Tables** under **Transit Gateways**. To create a static route, select the transit gateway from the Route Table and navigate to the **Routes** tab.

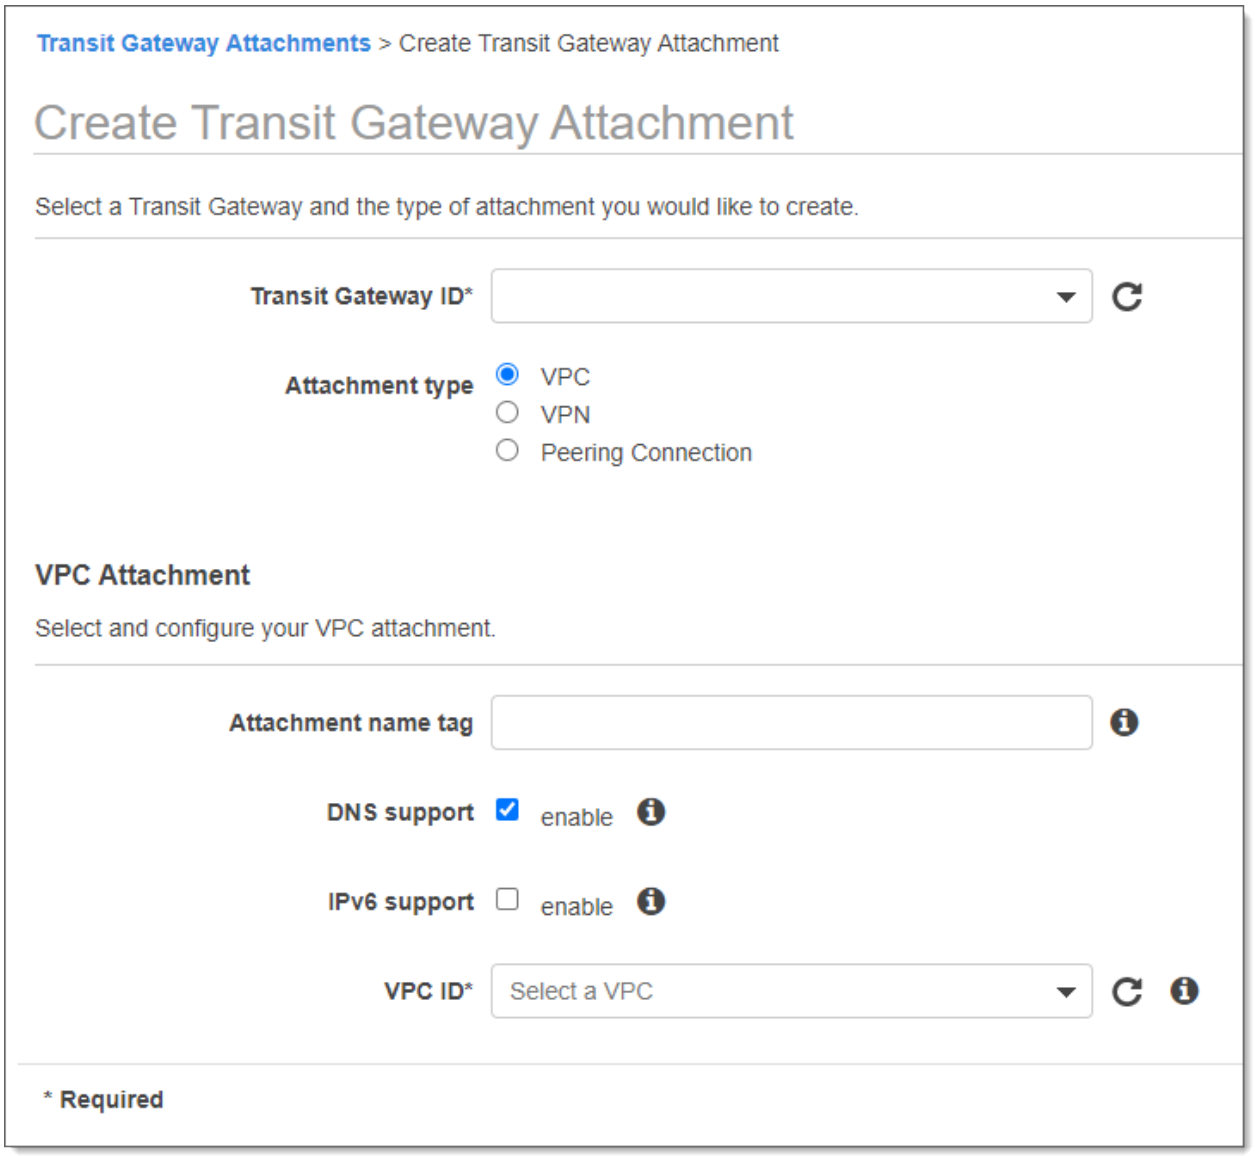


Complete the following fields, and then click **Create Static Route**.

Field	Description
CIDR	Specified range of IPv4 addresses for your VPC.
Blackhole	Enable if you want your matched traffic to be dropped.
Choose attachment	Choose the attachment for your static route.

Peering

To begin sending traffic from the spoke VPCs where your AWS workloads are running, you must peer the VPCs with the Transit Gateways. To peer your configured Transit Gateways, navigate back to your VPC dashboard in AWS and click **Transit Gateway Attachments** under **Transit Gateways**. Complete the following steps.



1. Select the check box next to the available transit gateways you want to peer.
2. Click **Create Transit Gateway Attachment**.
3. Choose the **Transit Gateway ID** from the drop-down menu.
4. For **Attachment Type**, select **Peering Connection**.
5. For **Attachment Name Tag**, enter text for identification purposes.
6. For **Account**, select the check box for **My Account**.
7. For **Region**, choose the destination region you want the BGP peering to connect with.

Azure Network Manager

Configuration > Cloud Services > Azure Network Manager

Microsoft Azure optimizes routing, automates large scale connectivity from various branches to Azure workloads, and provides unified network and policy management within Orchestrator.

The Azure Network Manager feature simplifies connectivity between EC-Vs in Azure and Azure Networking services. It enables seamless integration with services such as the Standard Internal Load Balancer, Azure Route Server, and Azure Virtual WAN Hub. This feature allows you to streamline the horizontal scaling of EC-Vs in Azure, thereby enhancing resilience and redundancy for Azure-based mission-critical workloads.

IMPORTANT: For more detailed information and configuration options, including integration with services such as the Standard Internal Load Balancer and Azure Route Server, see the HPE Aruba Networking EdgeConnect SD-WAN Virtual (EC-V) in Microsoft Azure Deployment Guide.

Before you begin Microsoft Azure Virtual WAN configuration in Orchestrator, you need to use the Azure Virtual WAN portal to authenticate and authorize Orchestrator in Azure. You need to create the service principal, which focuses on single-tenant application to run within only one organization. To get started, see Register a Microsoft Entra app and create a service principal.

Microsoft Azure Prerequisites

1. Create an application in Azure and note the following Subscription details from the Azure Active Directory:
   - Subscription ID
   - Directory (tenant) ID
   - Application (client) ID
   - Client secret value
2. Create a storage account in Azure and get the following:
   - Storage Account Name
   - Storage Access Key
3. Create a resource group.
4. Create Azure Virtual WANs with hubs from your resource groups.

Orchestrator Prerequisites

Complete the following tasks in Orchestrator:

1. Configure a VTI IP Pool.
   - Enter a valid IPv4 Subnet.

     NOTE: This is a unique address across the network. VTI interfaces created for Azure integration will be selected from this pool.

     NOTE: Azure VTI interface zone is set to WAN interface zone. Any change in deployment for the WAN interface zone is applied to Azure VTI as well.

WARNING: Any change in the VTI pool after it is configured is networking affecting. This operation should be performed during a maintenance window as it can take several hours for some Cloud services to complete.

2. Configure BGP ASN Global Pool.
   - Enter the start and end ranges for ASNs.
   - Add any reserved ASNs to exclude from being applied to appliances.

     NOTE: If not previously enabled, Orchestrator enables BGP.

Orchestrator Configuration

When you are finished with the Azure and Orchestrator prerequisites, navigate to **Configuration > Cloud Services > Azure Network Manager**.

Cloud Gateways

There are five buttons at the top of the table that are used to complete the Azure and Orchestrator integration for Cloud Gateways: **Subscription**, **LAN interface labels**, **Azure resources**, **Appliance association**, and **Pause Orchestration**.

Subscription

1. Click **Subscription**.
2. In the Subscription For Azure Cloud Gateways LAN-side Automation dialog box, select the Azure subscription you want to use.

   Under Connectivity status, Orchestrator displays the status for the load balancer, route server, and VWAN hub you connected through your Azure account.

   NOTE: If any modifications you make (such as creating or deleting subnets) in Azure are not reflected on the Orchestrator after you add the Azure subscription to Orchestrator, click the refresh icon to retrieve the latest Azure configuration on demand.
3. Click **Save**.

LAN interface labels

1. Click **LAN interface labels**.
2. In the Establish Connectivity Using These Network Interfaces dialog box, drag and drop the LAN label into the Primary column.
3. Click **Save**.

Azure resources

1. Click **Azure resources**.
2. In the Configure SD-WAN VNet Resources For Cloud-Deployed EdgeConnect Appliances dialog box, click **Add**.

3. In the Configure Azure Resources dialog box, enter the following information:

Field	Description
Rule name	Enter a descriptive name to help identify the EC-V gateways and their VNet.
Appliances	Select the EC-V gateways that should be associated with the ILB. You can either enter the appliance names or click **Use Tree Selection** to browse and select them.
Resource group	From the drop-down menu, select the Azure resource group where the EC-V gateways are deployed.
Region	Select the Azure region that corresponds to the deployment location of the EC-V gateways.
Virtual network	Choose the VNet where the EC-V gateways reside. This ensures that the appliances are correctly associated with the ILB within the specified VNet.

4. Click **Done**.

Appliance association

This step prompts you to select the EC-Vs and the services (vWAN hub, ARS, or ILB) that need to establish connectivity.

1. In the appliance tree in Orchestrator, select each EC-V that you want to establish connectivity with.
2. On the Azure Network Manager tab, click **Appliance Association**.
3. Depending on what you are trying to connect, click **Virtual WAN Hub**, **Azure Route Server**, or **Azure Standard Internal Load Balancer**. You will see all vWAN hubs, ARSs, and ILBs the Orchestrator has access to. If you do not see the service that you want to associate your EC-Vs with, check the permissions assigned to the Orchestrator's custom role on Azure Portal.
4. From the left-side panel, select **Add** for each EC-V you want to associate.
5. Click **Save**.

Pause Orchestration

Click this button to pause LAN-side orchestration. You will be prompted to confirm your selection.

Deploy Route Servers

You can create a new ARS or associate an existing one with EC-V gateways from Orchestrator. The following steps explain how to create a new ARS from Orchestrator within a transit VNet:

1. In Orchestrator, navigate to **Configuration > Cloud Services > Azure Network Manager**.

2. Under the Route Servers header, click **Deploy route server**.
3. On the Azure Route Server Deployment Configuration dialog box, enter the following settings:

Field	Description
Name	Enter a descriptive name for the ARS.
Azure account	Select the Azure subscription for deployment.
Resource group	Select the resource group where the ARS is created.
Region	Select the Azure region for the EC-V gateways.
Virtual network	Select the transit VNet where the EC-Vs are deployed. **NOTE:** For Orchestrator to establish BGP between the EC-V and the ARS, they both must reside within the same VNet.
Route server subnet	If your transit VNet already has a RouteServerSubnet of /26 or larger, it will be auto-selected. This is useful for users deploying the subnet themselves and using Orchestrator only for BGP setup. If not, you will be asked to provide a CIDR block so Orchestrator can create the subnet for ARS deployment.
Comment (Optional)	Add comments for tracking purposes.

4. Click **Deploy**.

   NOTE: Route Server creation can take up to 30 minutes. You can monitor progress in the audit logs in Orchestrator or the activity logs in the Azure Portal.
5. Confirm that the ARS resource appears in the Azure resource group, and check its status and assigned IP addresses.
6. Verify that the ARS is listed by clicking **Appliance association**, and then clicking **Azure Route Server**.

Deploy Load Balancers

You can deploy a new ILB or associate an existing one with EC-V gateways from Orchestrator. The following steps show how to create a new ILB from Orchestrator.

1. From the Azure Network Manager tab, click **Deploy load balancer**.
2. On the Azure Load Balancer Deployment Configuration dialog box, enter the following settings:

Field	Description
Name	Enter a descriptive name for the ILB.
Azure account	Select the Azure subscription for deployment.

Field	Description
Resource group	Select the resource group added to the Orchestrator.
Region	Select the Azure region that corresponds to the deployment location of EC-V gateways.
Virtual network	The VNet where EC-Vs are deployed. When deploying through Orchestrator, the ILB and EC-Vs must be in the same VNet. Deploying the ILB in a different VNet from the EC-Vs is not supported.
New subnet	If you want Orchestrator to create a new subnet for the ILB, select the check box. When you select the checkbox, two things happen: i. The Orchestrator displays all existing subnets within the VNet you selected. The Load balancer subnet CIDR field appears. ii. You must enter a subnet for the load balancer on this field. The subnet you enter must not overlap with any existing subnets in the virtual network and be /27 or bigger (such as /26 or /25). After the Orchestrator creates the ILB, it will use the fourth IP address of this subnet as the frontend IP address of the load balancer. For instance, if you enter 10.0.0.0/27 as the subnet mask, 10.0.0.4 is assigned as the frontend IP address of the load balancer. If the New subnet check box is not selected, you must choose an existing subnet in your VNet for the ILB. This option is available for users whose organizations do not allow a third-party application such as the Orchestrator to create new subnets in a VNet. Ensure that the subnet you select is /27 or bigger (such as /26 or /25) in size. The Orchestrator will not accept a /28 or smaller subnet. This option allows you to create a subnet manually but still use the Orchestrator to deploy the ILB within that subnet. When this option is selected, you must also enter a frontend IP address for the ILB, and the IP address must be available within the selected subnet.
TCP health probe port	Enter a port number that you want to use for the ILB health probe. Port 443 is recommended.
Comment (Optional)	Add a comment to identify your ILB deployment.

3. Click **Deploy**.

Orchestrator now deploys the ILB and creates a backend pool, a health probe (on TCP port 443 by default), and a load-balancing rule. The backend pool is empty because you have not associated any EC-V gateways to the ILB yet. After you associate EC-V gateways to the ILB, the Orchestrator adds the LAN interface IPs of the selected EC-Vs to the ILB's backend address pool.

On-prem Gateways

There are seven buttons at the top of the table that are used to complete the Azure and Orchestrator integration for On-prem Gateways: **Subscription**, **WAN Interface Labels**, **Tunnel Settings**, **VTI Subnet Pool**, **Zone/Role**, **Virtual WAN Association**, and **Pause Orchestration**.

To begin, click **Subscription**.

Subscription

1. In the Subscription for Azure dialog box, enter the following information to reflect your Azure portal account:

Field	Description
Azure reachability	Connection status of your account with Azure.
Subscription ID	ID of your subscription.
Directory (tenant) ID	Name of your Azure AD tenant.
Application (client) ID	Client ID of your Azure portal.
Client secret value	Secret key of your Azure application.
Storage account Name	Name of your storage account.
Storage account Key	Storage account key.
Storage URL	Storage account URL.*
Configuration polling interval	Indicates hows often Orchestrator should check for configuration changes in Azure. The default polling interval is 10 minutes.

2. Click **Save**. The Azure field should appear as *Connected*.

Storage URL

The Storage URL is present on the **Storage Accounts** tab in your Azure portal. Complete the following steps to obtain your storage account URL.

1. After your storage account is created in Azure, create a blob container.
2. Get the blob container URL.
3. Suffix the URL with a slash and add a file name in the **Storage URL** field.

   NOTE: Append the URL with a slash for the file name. Do not end the URL with a slash.

WAN Interface Labels

Select the order in which you want your interface labels to be used.

1. Click the **Interface Labels** button. The **Build Tunnels Using These Interfaces** displays.
2. Drag the Interface labels you want to use into the **Preferred Interface Label Order** column.
3. Click **Save**.

Tunnel Settings

The **Tunnel Settings** button opens the Tunnel Settings dialog box, which enables you to define the tunnels associated with Azure and Orchestrator. It is recommended that you use the default tunnel settings for General, IKE, and IPSec; however, you can modify any field. The tunnel settings are set using the default VPN configuration parameters received from virtual WAN APIs located in your Azure portal account.

In your Azure Portal Account, navigate to the Azure Configuration table. This table displays the VPN site created for Orchestrator appliances associated to Azure virtual WANs. Additionally, manually associate sites to your hubs in Azure.

1. Navigate to **Azure Virtual WAN**.
2. Select **Azure VPN site**.
3. Select **New Hub Association**.

VTI Subnet Pool

Enter a subnet IP and mask value in the provided fields, and then click **Save**.

Zone/Role

You can apply configured segments to your VTI interfaces associated for Azure. Click the **Zone** button and select the zone from the drop-down you want to apply.

Virtual WAN Association

Each appliance is associated with **one** virtual WAN. Use the Virtual Wan Association button to add or remove specific sites to your virtual WANs.

1. Click the **Virtual Wan Association** button.
2. Select an appliance from the tree in the left menu.
3. Select the check box to **Add** or **Remove** the appliance to your virtual WAN in Azure.

Pause Orchestration

Click this button to pause LAN-side orchestration. You will be prompted to confirm your selection.

Verification

The **Tunnel** page displays that Azure and Orchestrator have an established connection with Azure by displaying a tunnel status of **up - active**.

For more information about Azure configuration, visit the following link: https://docs.microsoft.com/en-us/azure/virtual-wan/virtual-wan-site-to-site-portal.

HPE SSE

Configuration > Cloud Services > HPE SSE

HPE SSE automates Secure Web Gateway (SWG) and SSE connector functionalities. SWG protects and governs internet-bound traffic, while the SSE connector enables secure, identity-based access to private applications.

NOTE: HPE SSE is not supported in air-gapped environments.

Click one of the following buttons to configure HPE SSE:

- Secure Web Gateway
- SSE Connector

Secure Web Gateway

HPE Aruba Networking SSE is a cloud security service. EdgeConnect traffic can be service chained to HPE SSE for additional security inspection. Orchestrator supports IPSec tunnel mode for HPE SSE.

IMPORTANT: By default, the maximum limit is 100 tunnels per HPE SSE tenant. If you want to increase the limit, you must contact HPE Aruba Networking support for assistance.

The following table describes the fields on the Secure Web Gateway view of the HPE SSE tab.

Field	Description
Appliance	Name of the appliance to connect to HPE SSE.
Interface Label	Interface label for the interfaces you want to connect to HPE SSE.
Location	Physical location of the appliance to connect to HPE SSE.
HPE SSE POP IPs	These are the HPE SSE endpoints to which the tunnels connect. This field is populated with discovered Public Service Edges based on the appliance's geographical location.
HPE SSE Deployment Status	Status of the HPE SSE deployment (**Creating**, **Pending**, or **Deployed**). *Deployed* indicates successful deployment.
Connection Status	Status of the HPE SSE connection based on tunnel and IP SLA statuses.

Configure Secure Web Gateway

Before you configure HPE SSE, you must create an HPE SSE account and have an HPE SSE tenant provisioned. Contact HPE Networking support for assistance with provisioning an HPE SSE tenant.

Subscription

Complete the Subscription dialog box to establish connectivity between Orchestrator and the HPE SSE (Axis) portal. The subscription is shared by Secure Web Gateway and SSE Connector.

1. In Orchestrator, navigate to **Configuration > Cloud Services > HPE SSE**.
2. Click **Subscription**.

   The Subscription dialog box opens. Leave the dialog box open, so you can paste your HPE SSE API token key in the API Token Key field.
3. In a new browser tab, go to https://auth.axissecurity.com/ and log in to your HPE SSE account.
4. From the Dashboard, click **Settings** and then click **Admin API**.

   The Admin API page opens.
5. Click **New API Token**.

   The New API Token dialog box opens.
6. Enter a Name for the new API token. The name should identify your Orchestrator.
7. Under Token Permissions, select **Read and Write**.
8. Under Token Scopes, select only **Tunnels** and **Locations**.
9. Under Token Expiration, enter **12**.
10. Click **Submit**.

    The New API Token dialog box opens and displays the token you created.
11. Copy the token.
12. In Orchestrator on the Subscription dialog box, paste the token into the API Token Key field.
13. In your HPE SSE account after you have copied the token, click **OK**.
14. In Orchestrator, enter the appropriate information in the remaining fields on the Subscription dialog box to reflect your HPE SSE account.

    The following table describes the fields.

Field	Description
HPE SSE	Indicates whether you are connected to your HPE SSE account.

Field	Description
API token name	Enter the name you assigned to the API token you created in your HPE SSE account. **NOTE:** The name should match exactly what you entered in the HPE SSE dashboard.
API token key	Enter (paste) the API token you created in your HPE SSE account. This token is used to access the HPE SSE APIs.
API domain	The domain name of the HPE SSE APIs that are used in tunnel creation. Leave the default setting.
Tunnel identifier	A unique identifier for the tunnel that is used when building the tunnel IKE identifiers. Enter the domain name for your company. For example, arubanetworks.com.
Orchestrator identifier	A unique identifier for the Orchestrator instance. This value is used to uniquely associate connectors created by this Orchestrator in the SSE (Axis) portal. By default, this is the configured hostname for the Orchestrator. You can change this to be any string, but it should be entered in the format of a hostname (containing no spaces or special characters).
Polling interval	Indicates how often Orchestrator should check for configuration changes in HPE SSE. The default polling interval is ten minutes.

15. Click **Save**. The HPE SSE field should indicate *Connected*.

Interface Labels

Select the WAN interfaces you want to use for HPE SSE internet traffic. You can specify primary and backup interfaces as described below. If a primary interface is unavailable, Orchestrator will use a backup interface if specified. Optionally, you can specify secondary interfaces as well. In this case, the fallback order is primary, secondary, and then backup.

1. On the HPE SSE tab, click **Interface Labels**.

   The Build HPE SSE Tunnels Using These Interfaces dialog box opens.
2. Drag the interfaces you want to use from the right side of the dialog box to the Primary and Backup areas. The interfaces are grayed out until you move them into the areas.
3. If you want to specify secondary interfaces, click **Show Secondary** to display the Secondary area, and then drag the appropriate interfaces to this area.
4. Click **Save**.

WARNING: This is service affecting. Any changes to the interface selection can cause previously built tunnels to be deleted and rebuilt.

Tunnel Settings

The **Tunnel Settings** button opens the HPE SSE Tunnel Setting dialog box, enabling you to define the tunnels associated with HPE SSE and EdgeConnect. The Mode field on the General tab allows you to select **IPSec** as the tunnel protocol for the specified WAN interface label. Use HPE SSE defaults for tunnel settings defined by the system.

NOTE: You can configure General, IKE, and IPSec tunnel settings. Settings are automatically generated, but you can change them if you want to.

IP SLA

Configure IP SLA for HPE SSE tunnels. This configuration ensures tunnel connectivity and internet availability between HPE SSE and Orchestrator. If the tunnel cannot reach HPE SSE, the tunnel is considered DOWN.

IMPORTANT: You must configure a loopback interface and a unique LAN-side label (such as "LOOPBACK") for the orchestrated loopback interface before you can set up IP SLA for HPE SSE tunnels. See Loopback Orchestration and Interface Labels for more information.

1. On the HPE SSE tab, click **IP SLA**.

   The HPE SSE Configuration dialog box opens.
2. If all fields are dimmed, click **Enable IP SLA rule orchestration**.
3. Select an orchestrated loopback label from the **Source Interface** field.
4. (Optional) In the **Route Label** field, select a route label to apply to this orchestration. Orchestrator automatically revokes the associated static default route when the IP SLA rule goes down, without requiring a separate standalone IP SLA rule. Depending on the IP SLA Condition set for the route label, the route is marked DOWN if *any* of the associated IP SLA rules are DOWN or only when *all* associated IP SLA rules are DOWN. For more information, see Route Labels.
5. Accept the default values for the remaining fields and click **Save**.

   Orchestrator builds the tunnels.

Sub-Locations

Sub-locations are a mechanism to configure and deploy different security policies to different types of traffic, at scale. When configuring a sub-location, you specify a subnet range to which the sub-location applies. Orchestrator then creates a corresponding sub-location in HPE SSE using that subnet range, and EdgeConnect appliances automatically provision the sub-location names.

From HPE SSE, you can apply policy rules to sub-locations. The policy rules are applied to all appliances that are configured as part of a sub-location regardless of physical location.

1. On the HPE SSE tab, click **Sub-Locations**.

   The HPE SSE Sub Locations dialog box appears.

2. Click **Add**.

   The Sub-Location Match Criteria dialog box opens. Enter the appropriate information for the following fields.

3. Enter a name for the sub-location in the Name field. This name will also be used for the corresponding sub-location in HPE SSE.
4. In the Appliances field, do one of the following to specify appliances to which the sub-location applies:

   - Start typing in the field and select “any”.
   - Enter “group” and select the name of an appliance group from the list.
   - To specify the appliances currently selected in the appliance tree, click **Use Tree Selection**. The appliance names appear beneath the Appliances field.

5. In the Internal IPs field, do one of the following to specify the subnet range for the sub-location:

   - Enter the name of a configured LAN label, firewall zone, or address group.
   - Enter an IP address or IP address range and click **+Add**.

6. Click **Save**.

   The Sub-Location Match Criteria dialog box closes.

7. Click **Save**.

HPE SSE POP Override

You can override the automatically selected endpoints for specific sites. You have the option to add this exception to one or more sites within your network.

1. On the HPE SSE tab, click **HPE SSE POP Override**.

   The HPE SSE POP Override dialog box opens.

2. Enter the appliance name, the interface label, and the primary and secondary FQDNs or IP addresses. Orchestrator will build tunnels to those endpoints.

Field	Description
Appliance	Appliance for which to override HPE SSE endpoints.
Interface Label	Interface label from which tunnels are built.
Primary FQDN or IP	FQDN or IP address of the primary HPE SSE endpoint.
Secondary FQDN or IP	FQDN or IP address of the secondary HPE SSE endpoint.

3. Click **Save**.

HPE SSE Association

The final step to configure the integration in Orchestrator is to associate EdgeConnect appliances to HPE SSE.

1. In the Orchestrator appliance tree, select one or more appliances to associate with HPE SSE.
2. On the HPE SSE tab, click **HPE SSE Association**.

   The HPE SSE Appliance Association dialog box opens.
3. In the table, select one or more appliances you want to associate with HPE SSE, and then select the **Add** check box.

   Select the **Remove** check box to remove HPE SSE association from selected appliances in the table.
4. Verify the changes, and then click **Save**.

Pause Orchestration

When troubleshooting, you can click **Pause Orchestration** and then click **Save** to pause orchestration. To restart, click **Resume Orchestration**.

Redirect HPE SSE Application Traffic via Break Out Locally

Traffic categorized as HPE SSE, which includes HPE SSE agent traffic (formerly known as Axis Security client), should break out locally, so it is not redirected to the HPE SSE tunnel and does not consume SASE secure web gateway bandwidth. To do this, you need to select the HPE SSE service in at least one Business Intent Overlay (BIO) Breakout Traffic Policy to steer traffic to it.

1. Navigate to **Configuration > Overlays & Security > Business Intent Overlays**.
2. Click an overlay at the top of the BIO priority to use for breakout traffic to HPE SSE.

   The Overlay Configuration dialog box opens.
3. Click the **Breakout Traffic to Internet & Cloud Services** tab.
4. Drag **HPE SSE Cloud** from the Available Policies column to the Preferred Policy Order column.
5. Drag **Break Out Locally** from the Available Policies column to the Preferred Policy Order column and make sure that it appears below HPE SSE Cloud.
6. Continue to Set Match Criteria.

Set Match Criteria

In the Overlay ACL, you need to set match criteria to steer the HPE SSE application traffic appropriately.

NOTE: Modify the top BIO to ensure that “HPE-SSE” application traffic is matched at the top of the BIO list. You can easily match HPE SSE flows using the built-in “HPE-SSE” application. This built-in application matches various “axisapps.io” domains and known public IPs used by HPE SSE. See Application Definitions Tab.

1. In the Match field, ensure that Overlay ACL is selected from the menu, and then click the edit icon.

   The Associate ACL dialog box opens.
2. Click **Add Rule**.
3. In the Match Criteria field click the edit icon.

   The Match Criteria dialog box opens.
4. Select the **Application** check box, enter “HPE-SSE” in the field, and select **HPE-SSE** from the drop-down menu.
5. Ensure that the Permit field shows “permit”.

   NOTE: If the overlay at the top of the BIO priority list is breaking out to the HPE SSE Cloud path, in the Permit field select **deny** from the menu, so the traffic for the HPE SSE application goes through another BIO policy that does break out locally.
6. Click **Save** to close the Match Criteria dialog box.
7. Click **Save**.
8. Click **OK**.

Verify HPE SSE Deployment

After HPE SSE is configured, deployment will begin automatically. Navigate to the HPE SSE tab to verify successful deployment. The HPE SSE Deployment Status column should have a green status of *Deployed*, and the Connection status column should have a green status of *Up*. The Connection Status column indicates the status of the HPE SSE connection based on tunnel and IP SLA statuses.

NOTE: HPE SSE is deployed and orchestrated for an appliance based on the HPE SSE Appliance Association dialog box. Business Intent Overlays (BIOs) are used to configure breakout internet policies to HPE SSE. This is used for automatic load distribution and failover.

You can also verify that your HPE SSE tunnels have been successfully deployed on the Tunnels tab. The Passthrough Tunnel column should list your HPE SSE tunnels, and the Status column should have a green status of *up – active*.

You can view the Audit Log to check for orchestration errors. Navigate to **Orchestrator > Audit Logs** and enter `hpesse` in the search field above the table.

SSE Connector

The SSE Connector enables EdgeConnect appliances to integrate with the HPE SSE cloud security service. The SSE Connector provides secure, identity-based access to private applications. It is supported in both Edge HA and Traditional HA deployments. Each appliance can have only *one connector installed*.

NOTE: The SSE connector is not supported on EC-US models or older EC-XS models.

NOTE: On EdgeConnect appliances with fewer than 16 cores, the WAN Optimization feature (formerly called Boost) is disabled when the SSE Connector is installed. You must reboot to apply this change.

The following table describes the fields on the SSE connector view of the HPE SSE tab.

Field	Description
Appliance	Name of the appliance to connect to HPE SSE.
Eligible	Indicates whether the appliance meets the minimum hardware and platform requirements to support the SSE connector. Any supported EC or EC-V with 4 or more cores can support the SSE connector.
Admin	Indicates whether the SSE connector is administratively enabled or disabled in the Axis portal.
Connector name	Name of the SSE connector created for the appliances. This name is automatically generated by Orchestrator using the Orchestrator identifier and the appliance name to create a unique identifier in the Axis portal.
Connector zone	SSE connector zone associated with the connector. This zone is required for connector operation. It is derived from the subscription configuration and site or cluster information, if available. If no connector zone is specified and no site or cluster information is available, Orchestrator automatically creates a default connector zone using the connector name. This ensures that the connector zone is unique and valid.
Version	Software version of the SSE connector currently installed on the appliance. If no connector is installed, this field is set to “None”.
Public IP	Public IP that is automatically populated after the connector is created and successfully registered with the SSE cloud.
Status	Current operational state of the connector on the appliance (for example, “Up” or “No connector installed”). This column reflects real-time connector availability and connectivity to the Axis portal.
Details	Click to display a detailed status and diagnostics view for the connector. This view provides installation and operational information, including download progress, image verification, installation steps, connection status, and error messages. The Details view is the primary troubleshooting interface for connector install, upgrade, and uninstall operations. **NOTE:** Closing the Details dialog box does not stop an active operation. Installation or uninstall operations continue in the background.
Actions	Hover to view available actions for this row (edit, delete, and so on).

Configure SSE Connector

Before you configure HPE SSE, you must create an HPE SSE account and have an HPE SSE tenant provisioned. Contact HPE Networking support for assistance with provisioning an HPE SSE tenant.

Before you begin:

- Ensure that you have access to the HPE SSE (Axis) Admin Portal. You will use this portal to create API tokens required for HPE SSE configuration.

- Plan a maintenance window for installing the HPE SSE connector, as installation may temporarily interrupt connector traffic and may require an appliance restart on low-end platforms.

Subscription

Complete the Subscription dialog box to establish connectivity between Orchestrator and the HPE SSE (Axis) portal.

TIP: The subscription is shared by Secure Web Gateway and SSE Connector. If the HPE SSE status on the Subscription dialog box indicates Connected and your subscription includes API tokens with permissions for Connector Zones and Connectors, you can skip this step. Otherwise, update the subscription to include the required tokens.

1. In Orchestrator, navigate to **Configuration > Cloud Services > HPE SSE**.
2. Click **SSE Connector**.
3. Click **Subscription**.

   The Subscription dialog box opens. Leave the dialog box open, so you can paste your HPE SSE API token key in the API Token Key field.
4. In a new browser tab, go to https://auth.axissecurity.com/ and log in to your HPE SSE account.
5. From the Dashboard, click **Settings** and then click **Admin API**.

   The Admin API page opens.
6. Click **New API Token**.

   The New API Token dialog box opens.
7. Enter a name for the new API token. The name should identify your Orchestrator.
8. Under Token Permissions, select **Read and Write**.
9. Under Token Scopes, select only **Connector Zones** and **Connectors**.
10. Under Token Expiration, enter **12**.
11. Click **Submit**.

    The New API Token dialog box opens and displays the token you created.
12. Copy the token.
13. In Orchestrator, paste the token into the **API Token Key** field on the Subscription dialog box.
14. In your HPE SSE account after you have copied the token, click **OK**.
15. In Orchestrator, enter the appropriate information in the remaining fields on the Subscription dialog box to reflect your HPE SSE account.

    The following table describes the fields.

Field	Description
HPE SSE	Indicates whether you are connected to your HPE SSE account.
API token name	Enter the name you assigned to the API token when you created it in your HPE SSE account. **NOTE:** The name should exactly match what you entered in the HPE SSE dashboard.
API token key	Enter (paste) the API token key you created in your HPE SSE account. This token is used to access the HPE SSE APIs.
API domain	Domain name of the HPE SSE APIs that are used in tunnel creation. Do not change the default setting.
Tunnel identifier	Unique identifier for the tunnel that is used when building the tunnel IKE identifiers. Enter the domain name for your company. For example, hpearubanetworks.com.
Orchestrator identifier	Unique identifier for the Orchestrator instance. This value is used to uniquely associate connectors created by this Orchestrator in the SSE (Axis) portal. By default, this is the configured hostname for the Orchestrator. You can change this to be any string, but it should be entered in the format of a hostname (containing no spaces or special characters).
Polling interval	Indicates how often Orchestrator should check for configuration changes in HPE SSE. The default polling interval is ten minutes.

16. Click **Save**. The HPE SSE field should change to *Connected*.

SSE Connector Interface Orchestration

A connector interface is required to pass traffic between EdgeConnect appliances and the SSE connector. A single subnet is specified for all appliances using the “Global” region. These SSE connector interfaces (“veth” interfaces) are required for connecting EdgeConnect appliances to the SSE connector. The SSE connector software is installed and managed separately.

The following table describes the fields on the SSE connector interface orchestration dialog box.

Field	Description
Segment	Network segment associated with the connector interface. After this value is created, it cannot be modified. If the connector is deployed to a non-default segment, be sure to configure the appropriate intersegment rules and firewall policies to allow connector traffic. You can match traffic based on Zone, Role, or the configured veth pool range. Additionally, all traffic to and from the HPE SSE cloud services can be identified and matched using the “HPE-SSE” application name. **NOTE:** All traffic to and from the connector is subject to existing policies, similar to any other LAN- or WAN-side flows.
Region	In Orchestrator release 9.7.0, the region setting is fixed to “Global” and cannot be changed.
Label	Label to identify the interface for organizational or policy purposes.
Zone	Security zone assigned to the connector interface.
Role	Role applied to traffic originating from the connector interface.
IP pool	IP address pool used for connector interfaces. The pool size determines how many appliances can be supported. Each appliance consumes two IP addresses from the SSE connector interface pool.
Allocated / Total	Shows how many IPs are currently assigned versus the total number of IPs available in the pool.
Actions	Hover to view available actions for this row (edit, delete, and so on).

Add SSE Connector Interface Pools

IMPORTANT: Plan the SSE connector interface subnet and segment before you begin. You cannot change the IP subnet or the segment after the SSE connector is installed. If you need to change these values, you must uninstall the SSE connector from all associated appliances, remove the SSE connector interface, and then recreate the SSE connector interface orchestration. SSE connector interfaces will appear as “veth20000”.

IMPORTANT: Ensure that among the top three DNS Server IPs, at least one is an IPv4 address and is reachable from the connector segment. This can be achieved by either placing the DNS server within the same segment as the connector or adding an inter-segment rule from the connector segment to the DNS server’s segment.

NOTE: If the DNS Server IP is reachable only via the SD-WAN fabric, then the veth IP subnet should be configured as “subnet shared” with the peer to ensure proper return traffic flow.

NOTE: Management Services configuration is not applicable for packets generated from the “ssecon” namespace.

The following table describes the fields on the SSE connector interface pool dialog box.

Field	Description
Segment	Network segment associated with the connector interface. You cannot modify this value. All interfaces in the region inherit the same segment.
Label	Label to identify the interface for organizational or policy purposes. All interfaces in the region inherit the same label. This field is optional.
Zone	Security zone assigned to the connector interface. All interfaces in the region inherit the same zone.
Role	Role applied to traffic using the connector interface. All interfaces in the region inherit the same role.
Region	Region is a logical grouping used for SSE connector interface orchestration. In Orchestrator release 9.7.0, the region setting is fixed to “Global” and cannot be changed.
IP pool	Virtual interface automatically created on each appliance to enable communication between the appliance and the SSE connector. The interface is assigned IP addresses from an SSE connector interface pool. The SSE connector interface pool can be defined in the Global region. A pool defined in the Global region will be used across multiple regions. Support for creating SSE connector interface pools for specific regions will be in future releases.

To add an SSE connector interface pool:

1. Click **Add**.

   The SSE connector interface pool dialog box opens.

2. From the **Segment** field, select the segment for this connector interface. The list contains existing segments defined in Orchestrator. If you select Default, the SSE connector interface is placed in the default network segment (VRF) on the appliance.

3. From the **Label** field, select the label for this connector interface. The list contains existing labels defined in Orchestrator. If you select NONE, no label is applied to the SSE connector interface. The interface is created without label-based classification and will not be included in policies that depend on interface labels.

   NOTE: Do not assign the same label to the connector interface as any other LAN-side interface. The connector should have its own dedicated label, role, and zone. Sharing labels with loopback interfaces or other LAN interfaces can lead to policy conflicts.

4. From the **Zone** field, select the zone for this connector interface. The list contains existing zones defined in Orchestrator. If you select Default, the SSE connector interface is assigned to the default security zone defined in Orchestrator. The interface inherits the policies and access controls associated with that zone.

5. In the **Role** field, select the role to assign to this interface. If you select Default, the SSE connector interface uses the default role defined in Orchestrator. The interface inherits the traffic policies and behavior associated with that role.
6. In the **Region** field, select the region you want Orchestrator to create an SSE connector interface for. In release 9.7.0, the region setting is fixed to "Global" and cannot be changed.
7. In the **IP pool** field, enter the IPv4 subnet in CIDR notation (for example, /16). The subnet size determines how many appliances the fabric can support.

   NOTE: Ensure that the pool is large enough. Each appliance requires two IP addresses from the pool, and the subnet size determines how many appliances a region can have.
8. Click **Save**.

   Orchestrator automatically creates the SSE connector interface on each appliance in the region. You may see new virtual interfaces or assigned IP addresses before the SSE connector is installed or active.
9. Click **Close**.
10. Continue to SSE connector interfaces to confirm that the interface orchestration completed successfully. The configuration is successful when the SSE connector interface appears in the SSE Connector Interfaces view with Admin set to *Enabled* and Status set to *Up*.

SSE Connector Interfaces

Use the SSE connector interfaces dialog box to view and validate that interface orchestration completed successfully.

The following fields describe the connector interface created on each appliance.

Field	Description
Appliance	Appliances that are selected in the appliance tree.
Segment	Inherited from interface orchestration configuration.
Zone	Inherited from interface orchestration configuration.
Role	Inherited from interface orchestration configuration.
Label	Inherited from interface orchestration configuration.
Interface	Automatically generated virtual interface created for the connector (for example, veth20000).
IP/Mask	A /31 address block assigned from the SSE connector interface. Each appliance is allocated two IP addresses within this range.
Admin	Indicates whether Orchestrator has configured and enabled this interface.
Status	Indicates whether the interface is up and operational on the appliance.

If the interface status is UP, continue to SSE connector management to download/install the HPE SSE connector.

SSE Connector Management

The final step to configure the SSE connector is to download and install the SSE connector on the appliances. You can either download only or download and install the SSE connector image to appliances. Traffic on the selected appliances is interrupted when the SSE connector is installed. HPE Networking recommends that you install connectors in a maintenance window.

NOTE: A reboot is required when installing or uninstalling the Connector on any appliance with 16 cores or fewer. Additionally, WAN Optimization (formerly called Boost) is disabled when the connector is installed on EdgeConnect appliances with less than 16 cores.

The following tables explain the fields on this dialog box.

Connector Images Table

Field	Description
Connector name	Auto-generated name created by Orchestrator using the Orchestrator identifier and appliance name. This name uniquely identifies the connector in the Axis portal.
Version	Connector image version available for installation.
Build date	Date on which the connector image was built.
Release notes	Click the information icon (**i**) to download the release notes for the selected SSE connector image.
Certified appliance version	Indicates whether the connector image is certified for the current ECOS version. HPE Networking tests and certifies the current HPE SSE connector version at the time of the EdgeConnect software release.

Target Appliances Table

Field	Description
Appliance	The appliance on which the connector image will be installed, upgraded, or removed.
Site/Cluster name	The site or cluster with which the appliance is associated. This field is blank if no site or cluster is associated with the appliance.
Status	The current connector state of the appliance (for example, no connector installed, connector installed and running, or an error message). This field updates dynamically.

Field	Description
Progress	The real-time progress of the install, upgrade, or uninstall operation.
Image certified	Indicates whether the selected connector image is certified for the ECOS version running on the appliance.
Restart required	Indicates whether installing a connector image requires a reboot (appliances with 16 cores or less).
Actions	Hover to view available actions for this row (edit, remove, and so on).

Download and Install an SSE Connector Image

To download or download and install an SSE connector image to appliances:

1. In the appliance tree, select the appliances you want to download and install the SSE connector image to.
2. Click **SSE connector management**.

   The SSE connector management dialog box opens.
3. In the Target appliances table, verify that the appliances you want to install the SSE connector on are listed in the table. The Restart required column indicates whether a restart is required after the connector is installed.

   IMPORTANT: Check the Image certified column in the Target appliances table to verify that the image is certified for use with the appliance's ECOS version.
4. If you do not want to install the connector image on a target appliance, hover your cursor in the Actions column of that appliance and click the trash can icon.

   Orchestrator removes the appliance from the Target appliances table.
5. From the Connector images table, select the connector image you want to download to the appliances. Use the Image certified column to determine which images are certified for your ECOS version.

   TIP: Click **Refresh connector images** to update the list of available connector images.
6. Click the information icon (**i**) in the Release notes column to download the release notes associated with the connector image. The release notes provide critical information concerning the connector installation or upgrade. Review the release notes before you proceed.
7. Select one of the following options:

   - **Download**: Downloads the SSE connector image to the selected appliances but does not install it. This option does not require a maintenance window.

- **Install already downloaded image**: Installs the SSE connector image that has already been downloaded to the selected appliances. This option requires a maintenance window only for low-end appliances (fewer than 16 CPU cores), that reboot during installation. High-end appliances do not require a reboot, but connector traffic is briefly interrupted during installation.
- **Download and install**: Downloads and installs the SSE connector image on the selected appliances. This option requires a maintenance window only for low-end appliances (fewer than 16 CPU cores), that reboot during installation. High-end appliances do not require a reboot, but connector traffic is briefly interrupted during installation.

NOTE: To reduce the risk of connectivity disruption, HPE Networking recommends downloading the connector image first and performing the installation in a separate step.

IMPORTANT: Appliances with fewer than 16 CPU cores require a reboot during connector install or uninstall. This automatic reboot temporarily interrupts traffic on the appliance. Perform these operations during a maintenance window. Appliances with 16 CPU cores or more do not require a reboot during connector install or uninstall. However, connector traffic is still briefly interrupted. Additionally, WAN Optimization (formerly called Boost) is disabled when the Connector is installed on EdgeConnect appliances with less than 16 cores.

8. Click **Install** to initiate the selected action (download, install, or both).

   The Progress column in the Target appliances table indicates the progress of the download or installation.
9. If a restart is required, a Connector install confirmation dialog box opens. Confirm the restart.
10. Click **Install** to continue.

The Progress column in the Target appliances table updates in real time to show the current progress. You can click **Cancel** to close this dialog box without interrupting the download or installation.

Uninstall a Connector Image

To uninstall a connector:

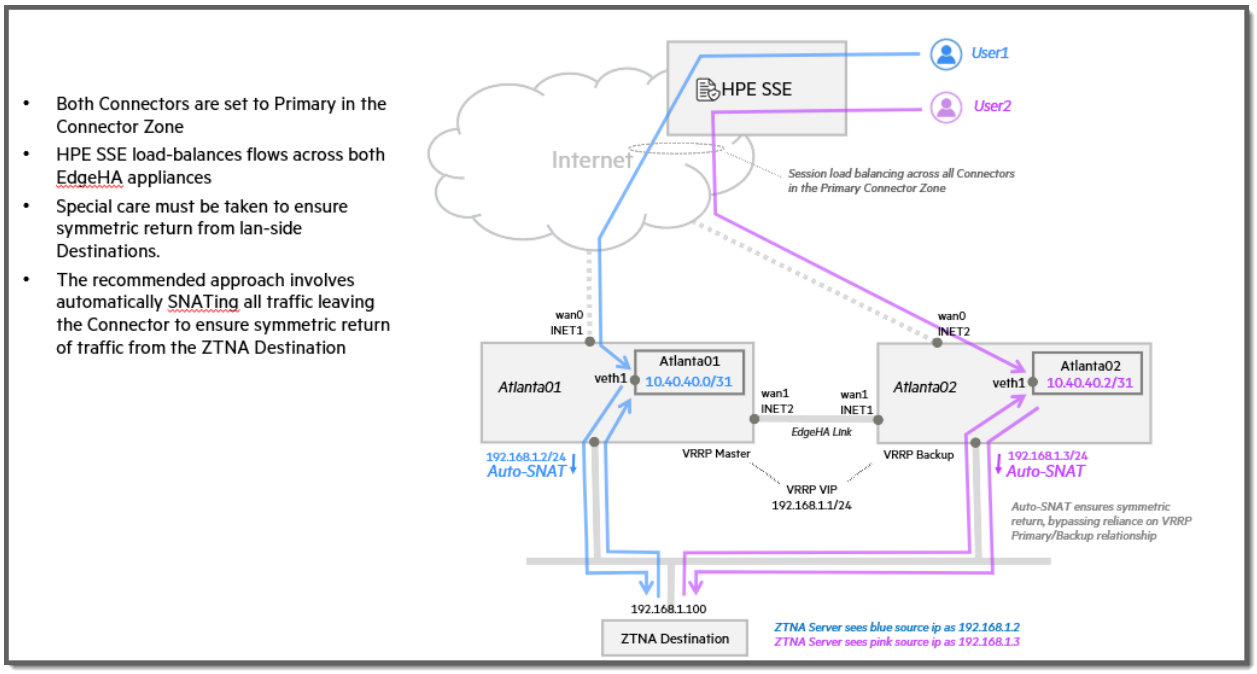
1. Hover over the Actions column for the appliance whose connector you want to uninstall and click the trash can icon.

   The Uninstall connector dialog box opens.

   IMPORTANT: Uninstalling the connector interrupts connector traffic on all appliances. Low-end appliances (fewer than 16 CPU cores) are automatically restarted during uninstall, which temporarily disrupts all traffic on the appliance. HPE Networking recommends uninstalling connectors during a maintenance window.
2. Click **Uninstall**. The uninstall operation starts immediately and cannot be cancelled.

EdgeHA Functionality and Connector Traffic Flow

The following diagram illustrates how connector traffic is handled within the EdgeConnect appliance and forwarded to the SSE service. Use this diagram to understand how interfaces, zones, and roles influence traffic flow and policy enforcement, and to assist with troubleshooting connectivity issues.



Troubleshoot SSE Connector

The ZTNA Connector button on the Flows tab allows you to filter and identify traffic flows associated with the HPE SSE connector. You can use this feature to:

- Verify that ZTNA traffic is reaching the connector.
- Confirm that traffic is being forwarded to the SSE cloud.
- Troubleshoot connectivity issues with ZTNA-protected applications.

To view ZTNA traffic:

1. Navigate to **Monitoring > Tunnel Bandwidth > Flows**.
2. Click **More options**.
3. From the appliance tree, select the appliance or appliances you want to monitor.
4. Click the **ZTNA connector** button to filter the traffic so that only SSE connector traffic is displayed.
5. In the Details column, click the information icon (**i**) to display the details for a specific ZTNA connector flow.

6. Confirm that the Source IP corresponds to the SSE connector interface IP and that the Destination IP corresponds to either the protected application or the SSE cloud PoP, depending on the flow direction.

   If the endpoints are not communicating, validate that the SSE connector is installed and that the connector interface Admin state is enabled and the Status state is UP.

   Ensure that the EdgeConnect firewall policies are properly configured to allow the ZTNA connector traffic. As a best practice, assign a specific role to the connector's veth interface and then permit that role across the entire fabric.

7. On the top of the dialog box, click **NAT**.
8. Confirm that the implicit SNAT is being applied to the connector to ZTNA-protected application traffic as expected.

   If the source IP addresses match the expected values, the SNAT is functioning correctly. If not, verify that the Original Source IP is the SSE connector interface IP and that the Translated Source IP is the appliance's LAN interface IP.

Microsoft Office 365

Configuration > Cloud Services > Microsoft Office 365

Ensure that your overlays have the following options configured to preserve the Works with Office 365 default applications. The table below indicates the default overlays, applications, and preferred policy order configured on the **Business Intent Overlays** tab within Orchestrator. The overlay name indicated in the table below is the default that ships with Orchestrator. This can be modified with user configuration.

NOTE: Skype for Business, SharePoint Online, and Office 365 Exchange **must** break out locally.

Overlay	Application	Preferred Policy Order (Breakout Traffic to Internet & Cloud Services)	What It Matches
Real-Time	Skype for Business		Microsoft Office 365 **Optimize** and **Allow** categories for the respective applications
CriticalApp	SharePoint Online, Office 365 Exchange		Microsoft Office 365 **Optimize** and **Allow** categories for the respective applications

Overlay	Application	Preferred Policy Order (Breakout Traffic to Internet & Cloud Services)	What It Matches
Default	For everything	Any policy order except “Drop”	Matches Microsoft Office 365 **Default** categories Office365 Common applications **NOTE:** Do not specify other individual Office applications in this group or overlay.

For more information about applications that work with Office 365, go to Microsoft 365 & Security for Partners.

Zscaler Internet Access

Configuration > Cloud Services > Zscaler Internet Access

Zscaler Internet Access (ZIA) is a cloud security service. EdgeConnect traffic can be service chained to Zscaler for additional security inspection. Orchestrator supports IPSec and GRE tunnel modes for Zscaler.

NOTE: GRE tunnels are not formed across an EdgeHA link.

NOTE: Zscaler’s term for *ZEN* is now *Service Edge*.

WARNING: If two or more Orchestrators are connected to a single Zscaler account, you must set the *MultipleOrchestratorsForOneZscalerAccount* property to “true” on the Advanced Orchestrator Properties dialog box for both Orchestrators. This is service affecting. When you change this property, all Zscaler artifacts will be rebuilt to include the UUID for each Orchestrator connected to the Zscaler account.

The following table describes the fields on the Zscaler Internet Access tab.

Field	Description
Appliance	Name of the appliance to connect to Zscaler.
Interface Label	Interface label for the interfaces you want to connect to Zscaler.
Mode	Tunnel mode (**IPSec** or **GRE**) for Zscaler. The default mode is IPSec.

Field	Description
Gateway Options	A feature that enables you to configure sub-locations and various rules for your sub-locations. Gateway Options is an optional add-on.
Bandwidth	Upload and download bandwidth speeds (in Mbps) to and from Zscaler.
Zscaler Deployment Status	Status of the Zscaler deployment (**Creating**, **Pending**, or **Deployed**). *Deployed* indicates successful deployment.
Zscaler Service Edges	These are the Zscaler endpoints to which the tunnels connect. This field is populated with discovered Public Service Edges based on the appliance's geographical location.
Connection Status	Status of the Zscaler connection based on tunnel and IP SLA statuses.
Zscaler ZDX	When configured, click the link to open a new tab to your ZDX web portal.

Configure Zscaler

Before you configure Zscaler, you must create a Zscaler account and ensure that you have an established connection with Zscaler.

NOTE: Ensure that both IPSec and GRE services are enabled in your Zscaler subscription so that Orchestrator can download data appropriately from Zscaler.

Subscription

1. Go to https://help.zscaler.com/zia/sd-wan-api-integration and follow the steps to configure your Zscaler account.
2. After configuring your Zscaler account, navigate to the Zscaler Internet Access tab in Orchestrator (**Configuration > Cloud Services > Zscaler Internet Access**).
3. Click **Subscription**.

   The Subscription dialog box opens.
4. Enter the appropriate information to reflect your Zscaler account.

   The following table describes the fields.

Field	Description
Zscaler	Indicates whether you are connected to your Zscaler account.
Zscaler Cloud	Zscaler cloud URL. For example, zsapi.zscalerthree.net.

Field	Description
Partner Username	Partner administrator username you created when configuring Zscaler.
Partner Password	Partner administrator password you created when configuring Zscaler.
Partner Key	Partner key you created when configuring your Zscaler account. Select Silver Peak from the list of partners.
Domain	Domain provisioned in Zscaler for your enterprise.
SubCloud ID	*(Optional)* A subcloud can be a subset of ZIA Public Service Edges, a subset of Private Service Edges, a subset of PZENs, or a subset of both ZIA Public Service Edges and Private Service Edges or PZENs. If you subscribe to any of these services, you must specify in this field the name of your subcloud (for example, Americas) to obtain a full list of Service Edges for your organization. **WARNING:** Because this is service affecting, configure this ID during a maintenance window only. This will cause previously built tunnels to be deleted and rebuilt.
Link to Zscaler ZDX	*(Optional)* Provides direct browser access to the Zscaler Digital Experience (ZDX) monitoring service through a popout URL on the Zscaler Internet Access tab or in the appliance tree. To enable this hyperlink, switch on the toggle and enter your ZDX URL. **NOTE:** This URL could be customized for your ZDX web portal. Confirm the correct URL in your ZDX account.
Configuration Polling Interval	Indicates how often Orchestrator should get “other” sublocations of VPN locations from Zscaler. The default polling interval is ten minutes.

5. Click **Save**. The Zscaler field should indicate *Connected*.

Interface Labels

Select the WAN interfaces you want to use for Zscaler internet traffic. You can specify primary and backup interfaces as described below. If a primary interface is unavailable, Orchestrator will use a backup interface if specified. Optionally, you can specify secondary interfaces as well. In this case, the fallback order is primary, secondary, and then backup.

1. On the Zscaler Internet Access tab, click **Interface Labels**.

   The Build Zscaler Tunnels Using These Interfaces dialog box opens.

2. Drag the interfaces you want to use from the right side of the dialog box to the Primary and Backup areas. The interfaces are grayed out until you move them into the areas.

3. If you want to specify secondary interfaces, click **Show Secondary** to display the Secondary area, and then drag the appropriate interfaces to this area.
4. Click **Save**.

WARNING: This is service affecting. Any changes to the interface selection can cause previously built tunnels to be deleted and rebuilt.

Tunnel Settings

The **Tunnel Settings** button opens the Zscaler Tunnel Setting dialog box, enabling you to define the tunnels associated with Zscaler and EdgeConnect. The Mode field on the General tab allows you to select **IPSec** or **GRE** as the tunnel protocol for the specified WAN interface label. Use Zscaler defaults for tunnel settings defined by the system.

NOTE: For IPSec mode, you can configure General, IKE, and IPSec tunnel settings. For GRE mode, you can configure General tunnel settings. Settings are automatically generated, but you can change them if you want to.

Service Edge Override

You can override the automatically selected Service Edge pair for specific sites. You have the option to add this exception to one or more sites within your network.

NOTE: Orchestrator does not support Service Edge Override for GRE tunnels.

1. On the Zscaler Internet Access tab, click **Service Edge Override**.

   The Service Edge Override dialog box opens.
2. Enter the appliance name, the interface label, and the primary and secondary IP addresses. Orchestrator will build tunnels to those Service Edges.

Field	Description
Appliance	Appliance for which to override Zscaler Service Edges.
Interface Label	Interface label from which tunnels are built.
Primary IP	IP address of the primary Zscaler Service Edge.
Secondary IP	IP address of the secondary Zscaler Service Edge.

3. Click **Save**.

IP SLA

Configure IP SLA for Zscaler tunnels. This configuration ensures tunnel connectivity and internet availability between Zscaler and Orchestrator. If the tunnel cannot reach Zscaler, the tunnel is considered DOWN.

1. On the Zscaler Internet Access tab, click **IP SLA**.

   The Zscaler IP SLA Configuration dialog box opens.

2. If all fields are dimmed, click **Enable IP SLA rule orchestration**.
3. Complete the following fields.

Field	Description
Monitor	Ping or HTTP/HTTPS.
Address	URL to the Zscaler endpoint that the IP SLA subsystem will ping. You can configure up to three addresses.
Source Interface	Select an orchestrated loopback label.
Route Label	(Optional) Select a route label to apply to this orchestration. Orchestrator automatically revokes the associated static default route when the IP SLA rule goes down, without requiring a separate standalone IP SLA rule. Depending on the IP SLA Condition set for the route label, the route is marked DOWN if *any* of the associated IP SLA rules are DOWN or only when *all* associated IP SLA rules are DOWN. See Route Labels.

4. Accept the default values for the remaining fields and click **Save**.

   Orchestrator builds the tunnels.

Country / Timezone

You can use the Zscaler Country / Timezone dialog box to configure standard ISO Country Codes to Zscaler Country Enums and standard Time Zones to Zscaler Time Zone Enums. On the Zscaler Internet Access tab, click **Country / Timezone** to open the dialog box. Make changes, and then click **Save**.

NOTE: If the Zscaler VPN Location request fails with an invalid request body, you can use this dialog box to change the ISO Country Code to the correct Zscaler Country Enums. The Zscaler enum list is available in the Zscaler documentation and this Zscaler Trust post.

Gateway Options

You can configure gateway options and rules for Zscaler sub-locations. Orchestrator uses location and sub-locations to better define a branch site in the Zscaler cloud. Sub-locations are LAN-side segments within each branch. They can be identified by LAN interfaces, zones, or a collection of LAN subnets.

IMPORTANT: After you configure gateway options and sub-locations as described in this section, you can enable orchestration of these configurations to the Zscaler portal by moving the **Orchestrator Gateway Options** toggle to the right. Be aware that this overwrites existing orchestrated location and sub-location configurations on the Zscaler portal. Any locations created manually on the portal are not overwritten. This action is service affecting, so do this during a maintenance window.

Configure Gateway Options

To configure gateway options:

1. On the Zscaler Internet Access tab, click **Gateway Options**.

   The Zscaler Gateway Options dialog box opens.

2. Click **Add**.

   The Location / Sub-Location Match Criteria dialog box opens.

3. Enter a name for the new rule in the **Rule Name** field.

   WARNING: If two rules have the same sub-location name or IP address, Orchestrator picks the first match and considers the order of the rules.

4. Specify a location by entering an appliance name, region, or group in the **Appliances** field.

5. Enter the WAN label in the **Location Label** field.

6. If you select the **Sub-Location** check box:

   1. Enter the sub-location name in the **Name** field.
   2. Enter the subnet address (LAN label, Firewall Zone, or subnet) in the **Internal IPs** field.

7. Click **Save**.

   NOTE: Sub-locations can be applied to all WAN links selected in the Build Tunnels Using These Interfaces dialog box (accessed by clicking the Interface Label button on the Zscaaler Internet Access tab).

If you select the **Show sub-locations** check box on the Zscaler Internet Access tab, the sub-locations configured in Gateway Options appear in the Zscaler table.

Configure Bandwidth Control

You can set up bandwidth controls for your Zscaler sub-locations configured in Gateway Options. Select from bandwidth control options that use fixed amounts of bandwidth, inherit bandwidth values from parent locations, or use percentages of deployment bandwidth.

1. On the Zscaler Internet Access tab, click **Gateway Options**.

   The Zscaler Gateway Options dialog box opens.

2. In the table, locate the rule name row for which you want to configure bandwidth control, and then click the linked text in the Gateway Options column.

   The Zscaler Gateway Options & Bandwidth Control dialog box opens.

3. Select one of the following options from the **Bandwidth Control** drop-down list:

Bandwidth Control Option	Description
OFF	Do not use bandwidth control. This is the default setting.
Fixed bandwidth	Use fixed amounts of bandwidth for the sub-location. Specify amounts for download and upload in Mbps.
Inherit (parent) location bandwidth	Inherit the parent location's bandwidth values.
Use deployment WAN label bandwidth	Use percentages of the deployment WAN label's bandwidth. Specify amounts for download and upload as percentages. Each specified percentage cannot exceed 100%. Orchestrator will automatically translate percentages into Mbps and send them to Zscaler. Sub-locations will use these values as percentages of deployment bandwidth.

4. Click **Save**.

   The Change Gateway Options dialog box opens.

   WARNING: Changing Gateway Options is service affecting. Make changes during a maintenance window.

5. Click **Change Gateway Options**.

   Your changes are applied to Orchestrator and Zscaler. This process takes time to complete.

Zscaler Association

The final step to configure the integration in Orchestrator is to associate EdgeConnect appliances to Zscaler.

1. In the Orchestrator appliance tree, select one or more appliances to associate with Zscaler.
2. On the Zscaler Internet Access tab, click **Zscaler Association**.

   The Zscaler Appliance Association dialog box opens.

3. In the table, select one or more appliances you want to associate with Zscaler, and then select the **Add** check box.

   Select the **Remove** check box to remove Zscaler association from selected appliances in the table.

4. Verify the changes, and then click **Save**.

Pause Orchestration

When troubleshooting, you can click **Pause Orchestration** and then click **Save** to pause orchestration. To restart, click **Resume Orchestration**.

Using Zscaler for Breakout Traffic

Finally, you need to select the Zscaler service in at least one Business Intent Overlay Breakout Traffic Policy to steer traffic to it.

1. Navigate to the Business Intent Overlays tab in Orchestrator (**Configuration > Overlays & Security > Business Intent Overlays**).
2. Click the overlay that breaks out traffic to Zscaler.

   The Overlay Configuration dialog box opens.
3. Click the **Breakout Traffic to Internet & Cloud Services** tab.
4. Drag **Zscaler Cloud** from the **Available Policies** column to the **Preferred Policy Order** column.

Verify Zscaler Deployment

After Zscaler Internet Access is configured, deployment will begin automatically. Navigate to the Zscaler Internet Access tab to verify successful deployment. The Zscaler Deployment Status column should have a green status of *Deployed*, and the Connection status column should have a green status of *Up*. The Connection Status column indicates the status of the Zscaler connection based on tunnel and IP SLA statuses.

NOTE: Zscaler is deployed and orchestrated for an appliance based on the Zscaler Appliance Association dialog box. Business Intent Overlays (BIOs) are used to configure breakout internet policies to Zscaler. This is used for automatic load distribution and failover.

You can also verify that your Zscaler tunnels have been successfully deployed on the Tunnels tab. The Passthrough Tunnel column should list your Zscaler tunnels, and the Status column should have a green status of *up – active*.

You can view the Audit Log to check for orchestration errors. Navigate to **Orchestrator > Audit Logs** and enter `zscaler` in the search field above the table.

Netskope

Configuration > Cloud Services > Netskope

Netskope is a cloud security service. EdgeConnect traffic can be chained to Netskope for additional security inspection. Orchestrator supports IPSec tunnel mode for Netskope.

NOTE: Be aware that design changes that occur in the Netskope application (especially the user interface) could affect instructions provided in this topic. Therefore, these instructions are provided as guidelines rather than precise steps.

IMPORTANT: If you have Netskope running through Service Orchestration, you must take down the manual tunnels *before* enabling the API through the **Configuration > Cloud Services > Netskope** feature.

The following table describes the fields on the Netskope tab.

Field	Description
Appliance	Name of the appliance to connect to Netskope.
Interface Label	Interface label for the interfaces you want to connect to Netskope.
Mode	Tunnel mode (**IPSec**) for Netskope.
Netskope Deployment Status	Status of the Netskope deployment (**Creating**, **Pending**, or **Deployed**). *Deployed* indicates successful deployment.
Netskope Service Edges	These are the Netskope endpoints to which the tunnels connect. This field is populated with discovered Public Service Edges based on the appliance's geographical location.
Connection Status	Status of the Netskope connection based on tunnel and IP SLA statuses.

Configure Netskope

Before you configure Netskope, you must create a Netskope account and ensure that you have an established connection with Netskope.

Subscription

1. Go to https://docs.netskope.com/en/rest-api-v2-overview-312207.html and follow the steps to configure your Netskope account.

   When you create your REST API token in the steps, add the following endpoints and assign the indicated privileges.

Endpoint	Privileges
/api/v2/steering/ipsec/pops	Read
/api/v2/steering/ipsec/tunnels	Read + Write

2. After configuring your Netskope account, navigate to **Configuration > Cloud Services > Netskope**.
3. Click **Subscription**.

   The Subscription dialog box opens.
4. Enter the appropriate information to reflect your Netskope account.

   The following table describes the fields.

Field	Description
Netskope	Indicates whether you are connected to your Netskope account.
API Token Name	The API Token name you created when configuring Netskope.
API Token Key	The API Token key you created when configuring your Netskope account.
Domain	Domain provisioned in Netskope for your enterprise.
Configuration Polling Interval	Indicates how often Orchestrator should check for configuration changes in Netskope. The default polling interval is ten minutes.

5. Click **Save**. The Netskope field should indicate *Connected*.

Interface Labels

Select the WAN interfaces you want to use for Netskope internet traffic. You can specify primary and backup interfaces as described below. If a primary interface is unavailable, Orchestrator will use a backup interface if specified. Optionally, you can specify secondary interfaces as well. In this case, the fallback order is primary, secondary, and then backup.

NOTE: When two or more labels are configured and active at the same level (primary, secondary, or backup) new flows will be load balanced across the Netskope tunnels based on current available bandwidth for the label. Labels with more available bandwidth will receive more flows than labels with less available bandwidth.

1. On the Netskope tab, click **Interface Labels**.

   The Build Tunnels Using These Interfaces dialog box opens.

2. Drag the interfaces you want to use from the right side of the dialog box to the Primary and Backup areas. The interfaces are grayed out until you move them into the areas.
3. If you want to specify secondary interfaces, click **Show Secondary** to display the Secondary area, and then drag the appropriate interfaces to this area.
4. Click **Save**.

WARNING: This is service affecting. Any changes to the interface selection can cause previously built tunnels to be deleted and rebuilt.

Tunnel Settings

The **Tunnel Settings** button opens the Netskope Tunnel Setting dialog box, enabling you to define the tunnels associated with Netskope and EdgeConnect. Use Netskope defaults for tunnel settings defined by the system.

NOTE: You can configure General, IKE, and IPSec tunnel settings. Settings are automatically generated, but you can change them if you want to.

Netskope Service Edge Override

This section is for future content.

IP SLA

Configure IP SLA for Netskope tunnels. This configuration ensures tunnel connectivity and internet availability between Netskope and Orchestrator. If the tunnel cannot reach Netskope, the tunnel is considered DOWN.

1. On the Netskope tab, click **IP SLA**.

   The Netskope IP SLA Configuration dialog box opens.
2. If all fields are dimmed, click **Enable IP SLA rule orchestration**.
3. Select an orchestrated loopback label from the **Source Interface** field.

   Note: When IP SLA is enabled for Netskope, Orchestrator automatically sets the Monitor field to Ping and uses the IP SLA targets specified by Netkope in the RESTv2 API response. Each Netskope POP uses a unique IP SLA target. The auto-IP SLA target typically ends in `.216`. For example, `10.162.6.216` (LON1) and `10.177.6.216` (LON2).
4. (Optional) In the **Route Label** field, select a route label to apply to this orchestration. Orchestrator automatically revokes the associated static default route when the IP SLA rule goes down, without requiring a separate standalone IP SLA rule. Depending on the IP SLA Condition set for the route label, the route is marked DOWN if *any* of the associated IP SLA rules are DOWN or only when *all* associated IP SLA rules are DOWN. For more information, see Route Labels.
5. Accept the default values for the remaining fields and click **Save**.

   Orchestrator builds the tunnels.

Netskope Association

The final step to configure the integration in Orchestrator is to associate EdgeConnect appliances to Netskope.

1. In the Orchestrator appliance tree, select one or more appliances to associate with Netskope.
2. On the Netskope tab, click **Netskope Association**.

   The Netskope Appliance Association dialog box opens.
3. In the table, select one or more appliances you want to associate with Netskope, and then select the **Add** check box.

   Select the **Remove** check box to remove Netskope association from selected appliances in the table.
4. Verify the changes, and then click **Save**.

Pause Orchestration

When troubleshooting, you can click **Pause Orchestration** and then click **Save** to pause orchestration. To restart, click **Resume Orchestration**.

Using Netskope for Breakout Traffic

Finally, you need to select the Netskope service in at least one Business Intent Overlay Breakout Traffic Policy to steer traffic to it.

1. Navigate to the Business Intent Overlays tab in Orchestrator (**Configuration > Overlays & Security > Business Intent Overlays**).
2. Click the overlay that breaks out traffic to Netskope.

   The Overlay Configuration dialog box opens.
3. Click the **Breakout Traffic to Internet & Cloud Services** tab.
4. Drag **Netskope** from the **Available Policies** column to the **Preferred Policy Order** column.

Verify Netskope Deployment

After Netskope is configured, deployment will begin automatically. Navigate to the Netskope tab to verify successful deployment. The Netskope Deployment Status column should have a green status of *Deployed*, and the Connection status column should have a green status of *Up*. The Connection Status column indicates the status of the Netskope connection based on tunnel and IP SLA statuses.

NOTE: Netskope is deployed and orchestrated for an appliance based on the Netskope Appliance Association dialog box. Business Intent Overlays (BIOs) are used to configure breakout internet policies to Netskope. This is used for automatic load distribution and failover.

You can also verify that your Netskope tunnels have been successfully deployed on the Tunnels tab. The Passthrough Tunnel column should list your Netskope tunnels, and the Status column should have a green status of *up – active*.

You can view the Audit Log to check for orchestration errors. Navigate to **Orchestrator > Audit Logs** and enter `Netskope` in the search field above the table.

Service Orchestration

Configuration > Cloud Services > Service Orchestration

To watch a video of this feature, see How to Integrate with Third-Party Service Providers.

Use the Service Orchestration tab to automate the integration of third-party service providers without an API. Service Orchestration automates the creation and deployment of IPSec tunnels and IP SLA probes and manages the lifecycle of the tunnels and probes.

Service Orchestration creates a local tunnel identifier (IKE ID) for each tunnel to the third-party service provider. After the tunnels are created, complete the integration on the third-party service provider's site by replacing the source identity values with the local tunnel identifiers (IKE IDs) that Orchestrator created for each endpoint.

Prerequisites

- You must have loopback interfaces configured to use the Service Orchestration feature.
- Service Orchestration supports third-party service providers that use IPSec IKEv2 endpoints.
- You will need the following information from the third-party service provider for each endpoint you want to add:
  - Endpoint name
  - IP address
  - Probe address
  - Probe type (Ping or HTTP/HTTPS)

Set Up a New Service

To set up a new third-party service:

1. Click **+Add Service** and complete the following fields.

Field	Description
Name	Name of the new service.
Prefix	A prefix to assign to all tunnels for this service. Orchestrator will use this prefix to filter tunnels and IP SLAs.

2. Click **Save**.

   A new tab is created on the Service Orchestration page.

   TIP: To edit or delete a service, click the edit icon next to the service name.
3. Select the tab for the new service and follow the steps below to integrate this new service.

Remote Endpoint Configuration

Add the remote endpoints for the third-party service provider. You can add one endpoint at a time or add endpoints in bulk by importing the information from a CSV file.

Add Endpoints One at a Time

1. Click **Remote Endpoint Configuration**.

   The Add Remote Endpoints dialog box opens.
2. Click **+Remote Endpoint**.
3. Complete the following fields. Press the **Tab** key to navigate to the next field.

Field	Description
Name	Name of the third-party service provider endpoint. **IMPORTANT:** If an endpoint name is decommissioned or modified, you must update the value in this table.
IP Address	IP address of the third-party service provider endpoint. If you do not have the IP address, you can use the FQDN. **IMPORTANT:** If an IP address is decommissioned or modified, you must update the value in this table.
Interface Label	The interface labels that can be provisioned for this endpoint. Only labels in this list will be provisioned. **HINT:** Click **Interface Label Default** to reset the Interface Label for every endpoint in the table to the default value of **Any**.
Pre-shared Key	The pre-shared key for the endpoint. To display the pre-shared key, click anywhere in the field. Do one of the following: Edit this field for each endpoint. This value can be an ASCII string, a hex-encoded string (if it has a 0x prefix), or a base64-encoded string (if it has a 0s prefix). Click **PSK Default** to create and save a pre-shared key. Every endpoint will use the pre-shared key you create. Because traffic going to these endpoints is encrypted, it will not compromise security to use the same pre-shared key for each endpoint.
Probe Address	The third-party service provider endpoint that the IP SLA subsystem will ping. You can obtain the probe address from the third-party service provider. **IMPORTANT:** Orchestrator will prefill the Address field in the IP SLA Settings dialog box with this value. If you delete the value in the Probe Address field in this table, Service Orchestration will ping the value specified in the Address field in the IP SLA Settings dialog box.

4. Click **Save**.
5. Repeat steps 2 - 4 for each endpoint you want to add.
6. After your endpoints are created, enter the probe address and a backup remote endpoint for each endpoint you defined.

Field	Description
Probe Address	The third-party service provider endpoint that the IP SLA subsystem will ping. You can obtain the probe address from the third-party service provider. **IMPORTANT:** Orchestrator will prefill the Address field in the IP SLA Settings dialog box with this value. If you delete the value in the Probe Address field in this table, Service Orchestration will ping the value specified in the Address field in the IP SLA Settings dialog box.
Backup Remote Endpoint	Enter the third-party service provider endpoint that you want to use as a backup tunnel. For example, ATL1-Atlanta could use DFW1-Dallas as a backup remote endpoint. If you leave this field empty, the endpoint will not have a backup tunnel. The BIO determines how traffic will be handled if a single or single and backup tunnel go down.

TIP: To delete an endpoint, click the **X** in the last column in the table.

7. Click **Save**.

   Updates are orchestrated immediately.

Add Endpoints in Bulk

1. Click **Remote Endpoint Configuration**.

   The Add Remote Endpoints dialog box opens.

2. Click **Import** to import a list of remote endpoints from a CSV file. The CSV file must contain columns for name, IP address, interface label, pre-shared key, probe address, and backup remote endpoint, in that order.

   NOTE: Remove any header rows before you import the file.

3. Click **Choose File**.
4. Navigate to the file, select the file, and then click **Open**.
5. Click **Save**.

   Updates are orchestrated immediately.

Bulk Edits

To make bulk edits to the table:

1. Click **Export**.
2. Open the CSV file and delete the three header rows.
3. Modify, save, and close the file.

4. Click **Import**, and then click **Choose File**.
5. Locate and select the file, and then click **Open**.

   Orchestrator updates the table.
6. Click **Save**.

Interface Labels

Select the Primary and Backup interface labels for your traffic. Backup interface labels will be used if the primary interface labels are unreachable.

1. Click **Interface Labels**.

   The Build Tunnels using these Interfaces dialog box opens.
2. Drag the interface labels you want to use into the Primary area. (The Peer/Service names in the Tunnels table will be XXX_Primary_1 and XXX_Primary_2.)
3. Drag the interface labels you want to use into the Backup area. (The Peer/Service names in the Tunnels table will be XXX_Backup_1 and XXX_Backup_2.)
4. Drag the interface labels up or down to reorder the list as necessary.
5. Click **Save**.

Tunnel Settings

1. Click **Tunnel Settings** to configure the tunnel settings.

   The Tunnel Settings dialog box opens. The General tab is displayed with the Mode field set to IPSec.
2. Complete the following fields as required for security service.

Field	Description
Mode	Indicates that the tunnel protocol is IPSec. You cannot edit this field.
IPSec Suite B Preset	Select an IPSec Suite B Preset if required by the security service (**GCM-128**, **GCM-256**, **GMAC-128**, or **GMAC-256**). The default setting is None. If IPSec Suite B Preset is set to None, no preset is selected, but GCM and GMAC algorithms are available to set independently. If an IPSec Suite B preset is selected, various settings on the IKE and IPSec tabs are configured automatically based on the selected preset.
Auto max BW enabled	When enabled, allows the appliances to auto-negotiate the maximum tunnel bandwidth. Enabled by default.

3. Click the **IKE** tab, and then complete the following fields.

Field	Description
IKE Version	IKE v2. You cannot edit this field.
Preshared Key	Pre-shared key used for IKE authentication. This key is generated dynamically.
Authentication algorithm	Authentication algorithm used for IKE security association (SA). Authentication algorithm can be set to **SHA1**, **SHA2-256**, **SHA2-384**, **SHA2-512**, or **NULL**.
Encryption algorithm	Encryption algorithm used for IKE security association (SA). Encryption algorithm can be set to **AES-128**, **AES-256**, **AES-GCM-128**, **AES-GCM-256**, or **NULL**.
Diffie-Hellman group	Diffie-Hellman Group used for IKE security association (SA) negotiation. If the IPSec Suite B Preset field on the General tab is set to None, you can select the appropriate group. Available groups are 14 through 21, 26, and 31. If the IPSec Suite B Preset field is set to any other setting, this field is automatically set to the appropriate group.
Rekey interval/lifetime	Rekey interval/lifetime of IKE security association (SA) in minutes. The default is 480 minutes.
Dead peer detection	**Delay time**: The interval (in seconds) to check the lifetime of the IKE peer. **Retry count**: The number of times to retry the connection before determining that the connection is dead. This field is not editable.
Phase 1 mode	Exchange mode for the IKE security association (SA) negotiation. This field is automatically set to Aggressive. This field is not editable.

Field	Description
Local IKE identifier	By default, the Service Orchestration feature creates local IKE IDs using the following fixed format: *%hostname%_%label%@%tunnel_dst%* You can create custom local IKE IDs by specifying one or more of the following macros: *%hostname%* Appliance host name *%label%* Interface label name *%tunnel_source_ip%* Tunnel source IP address *%tunnel_dst%* Tunnel destination IP address/FQDN *%appliance_key%* Appliance key For example, to create an IKE ID that contains an email domain, enter *%hostname%_%label%@customerdomain.com* which results in the following configuration: ECV-A_MPLS@customerdomain.com **IMPORTANT:** The custom local IKE ID cannot exceed 64 characters.

Field	Description
Remote IKE identifier	By default, the Service Orchestration feature creates remote IKE IDs using the following fixed format: *%tunnel_dst_ip%* You can create custom remote IKE IDs by specifying one or more of the following macros: *%hostname%* Appliance host name *%label%* Interface label name *%tunnel_source_ip%* Tunnel source IP address *%tunnel_dst_fqdn%* Tunnel destination FQDN *%tunnel_dst_ip%* Tunnel destination IP address *%appliance_key%* Appliance key For example, to create an IKE ID that contains an email domain, enter *%hostname%_%label%@customerdomain.com* which results in the example configuration: ECV-A_MPLS@customerdomain.com **IMPORTANT:** The custom remote IKE ID cannot exceed 64 characters.

4. Click the **IPSec** tab, and then complete the following fields:

Field	Description
Authentication algorithm	Authentication algorithm used for the IPSec security association (SA). Authentication algorithm can be set to **SHA1**, **SHA2-256**, **SHA2-384**, **SHA2-512**, **AES-GCM-128**, **AES-GCM-256**, or **NULL**.
Encryption algorithm	Encryption algorithm used for the IPSec security association (SA). Encryption algorithm can be set to **AES-CBC-128**, **AES-CBC-256**, **AES-GCM-128**, **AES-GCM-256**, or **NULL**.
IPSec anti-replay window	Select a size from the drop-down list or **Disable** to disable the IPSec anti-replay window. If a size is selected, protection is provided against an attacker duplicating encrypted packets by assigning a unique sequence number to each encrypted packet.

Field	Description
Rekey interval/lifetime	Rekey interval/lifetime of the IPSec security association (SA) in minutes. The default is 120 minutes.
Perfect forward secrecy group	Diffie-Hellman group used for IPSec security association (SA) negotiation. Based on the setting of the IPSec Suite B Preset field on the the General tab, this field is set to the following Diffie-Hellman group: For None: 14 (by default) For GCM-128 or GMAC-128: 19 For GCM-256 or GMAC-256: 20

5. Click **Save**.

TIP: Click **Use Default** to reset all tunnel settings to the global defaults for Service Orchestration.

IP SLA Settings

1. Click **IP SLA Settings**.

   The IP SLA Settings dialog box opens.
2. If all fields are dimmed, click **Enable IP SLA rule orchestration**.
3. Complete the following fields.

Field	Description
Monitor	Ping or HTTP/HTTPS.
Address	The third-party service provider endpoint that the IP SLA subsystem will ping. Orchestrator prefills the Address field with the value from the Remote Endpoint Configuration table. You can configure up to three addresses.
Source interface	Select an orchestrated loopback label.
Route Label	(Optional) Select a route label to apply to this orchestration. Orchestrator automatically revokes the associated static default route when the IP SLA rule goes down, without requiring a separate standalone IP SLA rule. Depending on the IP SLA Condition set for the route label, the route is marked DOWN if *any* of the associated IP SLA rules are DOWN or only when *all* associated IP SLA rules are DOWN. See Route Labels.

4. Accept the default values for the remaining fields, and then click **Save**.

Orchestrator builds the tunnels.

Pause Orchestration (Optional)

When troubleshooting, you can click **Pause Orchestration** and then click **Save** to pause the service orchestration. To restart the service orchestration, click **Resume Orchestration**.

BIO Breakout

By default, the tunnels associated with a third-party service provider will be available for BIOs. You can upload an icon to display on the Business Intent Overlays tab.

NOTE: Supported file types include PNG, JPEG, SVG, and WEBP. The recommended dimensions are 60 x 20 pixels.

1. Click **BIO Breakout**.

   The Configure BIO Breakout dialog box opens.
2. Click **Upload Service Icon**.
3. Locate and select the file, then click **Open**.
4. Click **Save**.

   This icon will display next to the service name on the Business Intent Overlays tab.

If you do not want this third-party provider to be available for BIOs, do the following:

1. Click **BIO Breakout**.

   The Configure BIO Breakout dialog box opens.
2. Clear the **BIO Breakout** check box.
3. Click **Save**.

Remote Endpoint Association

The final step to configure the integration in Orchestrator is to associate EdgeConnect appliances with remote endpoints. Use this page to add or remove endpoints from an appliance. It is recommended that you associate one remote endpoint per EdgeConnect appliance.

1. In the Orchestrator appliance tree, select one or more appliances to associate with the third-party service provider remote endpoints.
2. Click **Remote Endpoint Association**.

   The Associate an Appliance to Remote Endpoints dialog box opens.
3. Select the **Add** or **Remove** check box next to the endpoints you want to associate with the selected appliances. Be sure to add the endpoints that are geographically closest to the appliances.
4. Verify the proposed changes to remote endpoints in the table to the right, and then click **Save**.

Add Tunnel Local Identifiers to the Third-Party Service Provider

After the Service Orchestration integration is complete in Orchestrator, you must add the local tunnel identifiers (IKE IDs) to the third-party service provider. You can simplify this process by exporting the third-party service provider configuration to a CSV file. The exported file contains all of the configuration details in the table on the third-party service provider page for all selected appliances, including IKE IDs.

NOTE: The default tunnel local identifier value is a fixed format: *hostname_labelname@IPaddress*. For example, EAST3-AWS_INETA@192.x.x.xxx.

If you created a custom IKE ID, the local tunnel identifier value will follow the format you defined in the IKE identifier field on the Tunnel Settings dialog box.

1. In the Orchestrator appliance tree, select all appliances associated with third-party service provider remote endpoints.
2. On the third-party service provider page on the Service Orchestration tab, click **Export** to save the contents of the table to a CSV file.
3. Log in to the third-party service provider.
4. In the IPSec/Location configuration panel, replace the Source Identity values with the corresponding Tunnel Local Identifiers (IKE IDs) created by Orchestrator.

Verification

After the third-party service provider is configured and the third-party service provider policy is applied successfully in the BIO, deployment will begin automatically. Go to the third-party service provider tab and view the Connection Status column to verify that the deployment was successful.

Deploy Cloud Hubs

You can deploy one or more EdgeConnect Virtual (EC-V) appliances in supported platforms. At this time, AWS, Azure, and GCP are supported.

Before you begin, complete the following tasks:

1. On the AWS dashboard, Azure portal, or GCP dashboard, create an Identity and Access Management (IAM) user account with required permissions for Orchestrator to create resources. A dedicated IAM user account for Orchestrator is recommended.
    1. Create a policy that contains all permissions the Orchestrator requires to create an EC-V.
    2. Attach the policy to the Orchestrator's IAM user account.
    3. Download the Security credentials of the Orchestrator's IAM user account.

2. If you are deploying EV-Cs in AWS, on the EC2 dashboard, create a key pair to assign to the EC-V. You will need this key pair if you want to SSH into the EC-V after the deployment.

After creating the IAM account, click **New Deployment** on the Cloud Hubs in AWS, Cloud Hubs in Azure, or Cloud Hubs in GCP tab to configure and deploy one or more EC-V cloud instances.

After deploying an EC-V in the cloud, navigate to the Discovered Appliances page in Orchestrator to view the deployment status. If the EC-V is still being deployed, the status in the Approve column will indicate Configuring. It takes approximately ten minutes to deploy and configure a cloud EC-V. Click **Refresh Discovery Information** to determine whether the appliance is ready to be approved into the SD-WAN fabric.

When configuration is complete and the green Approve button appears, the EC-V is fully configured in Inline Router mode with mgmt0, wan0, and lan0 MAC addresses assigned. While adding the EC-V, the Deployment Profile page will show LAN IP address, WAN IP address, WAN interface firewall mode, and WAN bandwidth value assigned by Orchestrator.

You can upgrade the appliance software version on a cloud EC-V after approving and adding it to the SD-WAN fabric.

After a cloud EC-V has been deployed, you can add another EC-V into the same deployment. The new EC-V will use the same settings from the existing deployment configuration such as account, region, VPC, key pair, and instance type. You can deploy the new instance into an Availability Zone that is already used by an existing appliance or a new Availability Zone.

Cloud Hubs in AWS

Configuration > Cloud Services > IaaS > Deploy Cloud Hubs in AWS

The Cloud Hubs in AWS tab provides the AWS account details and EC-V deployment configuration details for all cloud EC-Vs that have been deployed.

Use this tab to:

- Create and modify AWS accounts
- Deploy EC-Vs in the AWS cloud
- Remove an AWS cloud deployment

NOTE: Before you can deploy EC-Vs to the AWS cloud, you must perform several tasks in AWS. For more information, see AWS Account Configuration.

The following table describes each field on this tab.

Field	Description
Name	Name given on the deployment configuration page.
VPC	CIDR block used for deployment.
Account	Name of the AWS account that was used to deploy the EC-Vs.

Field	Description
Instances	Number of EC-V instances in the deployment. To add one or more EC-Vs to the deployment, click **+Add**. In the New Instance on AWS dialog box, select the availability zone to use and any optional tags to apply to the new instance. **Max** indicates that the maximum number of instances have been created for the VPC CIDR block.
Status	Status of the deployment. If more information is available, an info icon is displayed. **NOTE:** If the deployment was incomplete, the info dialog contains a link to download the log file and steps to resolve the issue.
Terminate	To permanently delete a deployment, click **Terminate**. This action deletes all resources associated with the EC-Vs, including all EC2 resources.
Deployment Info	Click the info icon in this column to view deployment and instance details, including the IP addresses associated with the mgmt0, wan0, and lan0 interfaces.
Resources	Click the info icon in this column to view details about each AWS resource that Orchestrator created during the deployment. This information is helpful when, for example, you need to identify the IP address of a security group to add a user to.
Comment	Comments that were added to the deployment when the EC-V was created. To edit the comment, click the edit icon.

Create or Modify an AWS Account

To create or modify an AWS account to Orchestrator:

1. Click **AWS Accounts**.

   The AWS Accounts dialog box opens.
2. Click **New AWS Account** or click the edit icon next to the account you want to edit.

   The AWS Account Configuration dialog box opens.
3. Complete or modify the elements as necessary.

Deploy a New EC-V

Click **New Deployment** to deploy one or more EC-V instances in AWS.

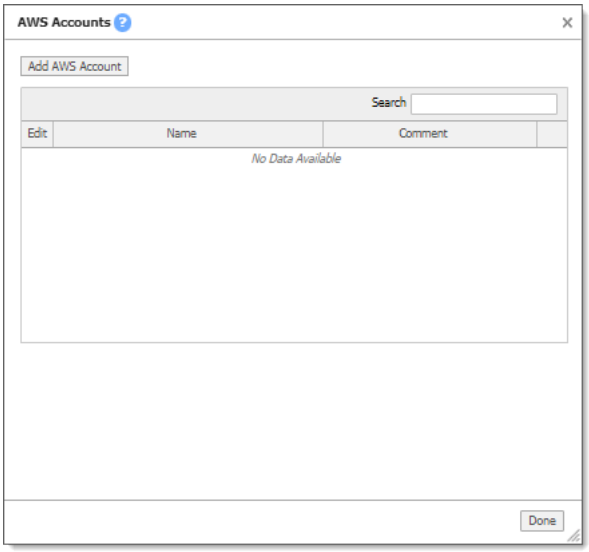
Remove an EC-V

If a deployment does not complete or you no longer want the EC-V in the AWS cloud, you can remove the deployment and all associated artifacts.

To remove a deployment, locate the deployment you want to remove, and then click **Terminate** in the desired row.

AWS Accounts

The AWS Accounts dialog box lists all of the AWS accounts that have been added.



- Click **Add AWS Account** to create a new account for EC-V deployments.
- Click the edit icon next to an existing account to modify that account's details.

NOTE: You cannot modify accounts that have active deployments.

AWS Account Configuration

Complete the following steps to create an AWS IAM user account with the required permissions for creating EC-V instances in AWS.

Create a Policy with Required Permissions

1. Log in to the AWS Dashboard.
2. On the Find Services search menu, enter *IAM* to open the Identity and Access Management (IAM) page.
3. Under Access Management, click **Policies**. The Policies page opens.
4. Click **Create policy** and click the **JSON** tab.
5. Delete the existing text.

6. Go to Permissions required for deploying Cloud EC-Vs and Cloud network services from Orchestrator, click **Permissions required to deploy Cloud Hubs in AWS in a Transit Virtual Private Cloud (VPC)**, and then copy and paste the JSON policy text into the editor.
7. Click **Next: Tags**.
8. (Optional) Add metadata to the policy by attaching tags as key-value pairs.
9. Click **Next: Review**.
10. On the Review policy page, enter a name and optional description for the new policy.
11. Review the policy summary to see the permissions granted by your policy, and then click **Create policy** to save your work.

Attach Policy to the Orchestrator IAM User Account

1. Click **Users > Add user**. The Add user page opens.
2. Enter a user name in the User name field (for example, ArubaOrchestrator).
3. Under Access type, select **Programmatic access**, and clear the AWS Management Console access check box.
4. Click **Next: Permissions**.
5. Under Set Permissions, click **Attach existing policies**.
6. Select the Policy document you created from the list, and then click **Next: Review**.
7. Under Permissions summary, click **Add permissions**.

Download Orchestrator IAM User Account Credentials

1. On the Users page, click the **Security credentials** tab.
2. Download or copy and paste the Access key ID and Secret key ID to a secure place for later use.

Create a Key Pair to Assign to EC-Vs

Review the instructions on this page to create a key pair on the AWS region where you plan to deploy the EC-V.

Subscribe to EdgeConnect SD-WAN Product on the AWS Marketplace Portal

Before you deploy an EdgeConnect SD-WAN instance from the Orchestrator in AWS, you must subscribe to the EdgeConnect SD-WAN product on the AWS Marketplace portal:

1. Navigate to the EdgeConnect SD-WAN product page: AWS Marketplace: Silver Peak Unity EdgeConnect for AWS
2. Click **Continue to Subscribe**. If prompted, log in to your AWS account.
3. Under Terms and Conditions, click **Accept Terms**. The subscription might take a few minutes to process.
4. After the Thank you message appears, proceed with launching your instance.

Add the AWS Account to Orchestrator

Complete the following fields for Orchestrator, and then click **Save** when finished.

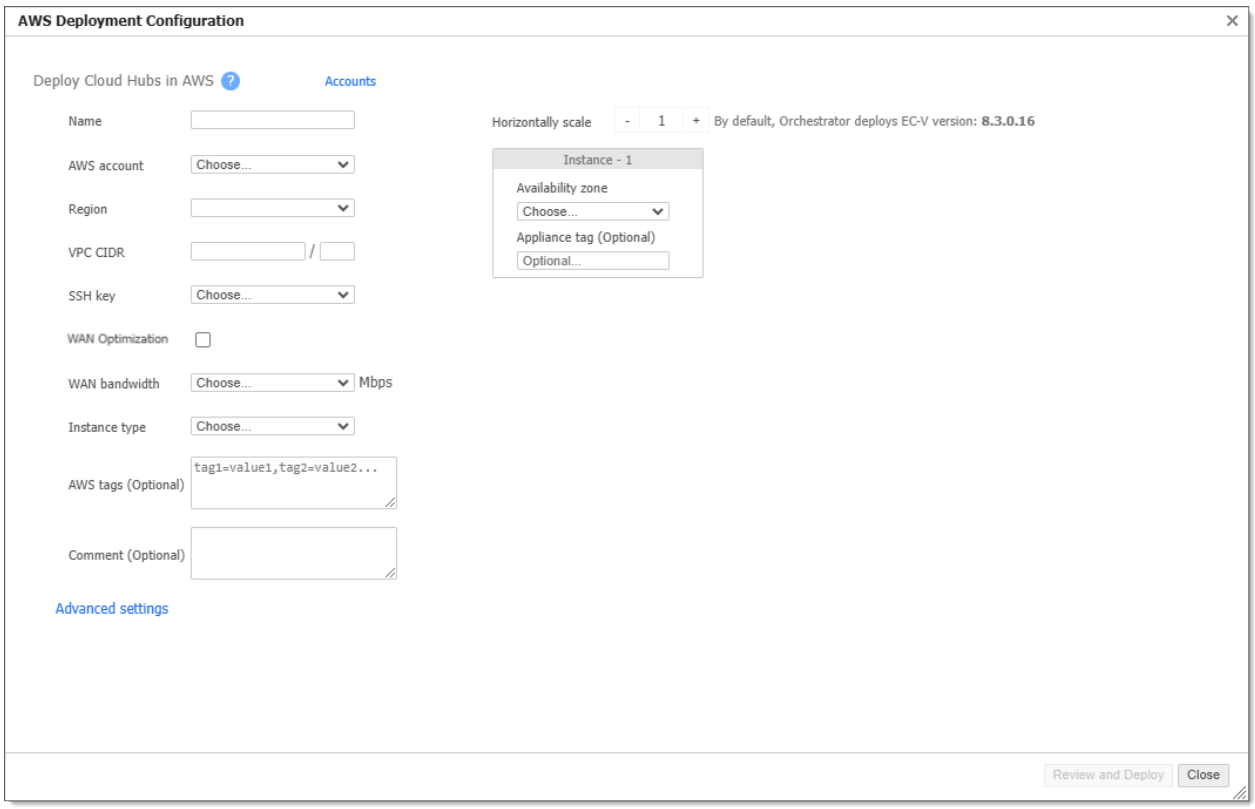
Field	Description
Name	Enter a unique name. If you have multiple AWS accounts, you must enter a unique name for each account.
Access Key	Enter the Orchestrator IAM user's Access Key ID that you saved earlier.
Secret Key	Enter the Orchestrator IAM user's Secret Key ID that you saved earlier.
Comment	Enter a comment that provides any additional information about the AWS account.

Orchestrator validates the account information. This takes approximately 45 seconds.

AWS Deployment Configuration

Use the AWS Deployment Configuration page to create one or more EC-V instances in an AWS region.

NOTE: If you do not have an AWS account configured in Orchestrator, the AWS Deployment Configuration dialog box is blank. Click the **Accounts** link to create an AWS account.



Field	Description
Name	Enter a name for the deployment. This name is used only for identifying the deployment. A deployment consists of one or more EC-Vs that an Orchestrator creates in an AWS Virtual Private Cloud (VPC). Only alphanumerical letters and hyphens are allowed in the deployment name. The maximum allowed length is 20 characters.
AWS Account	Select an AWS account to use for deploying the EC-V.
Region	Select an AWS region where you want to deploy the EC-V.
VPC CIDR	Enter a VPC Classless Inter-Domain Routing (CIDR) block. The smallest supported CIDR block is /24 and the largest supported CIDR block is /16. Orchestrator creates all AWS resources required for the EC-V deployment within this VPC. For each EC-V you deploy, Orchestrator creates three subnets that are /28 in size. In other words, if you deploy two EC-Vs, Orchestrator creates six subnets in total. This is true even if both EC-Vs are created in a single Availability Zone.
SSH Key	Select an existing AWS key pair to assign to the EC-V. A key pair must be created prior to the deployment.
WAN Optimization (Optional)	WAN Optimization requires additional resources on an AWS EC2 instance. After WAN Optimization and an appropriate WAN Bandwidth value are selected, Orchestrator displays the appropriate AWS instance types for the deployment on the Instance Type drop-down menu. **NOTE:** Selecting the WAN Optimization check box does not enable WAN Optimization on the EC-V. It only allows Orchestrator to display appropriate AWS instance types that can support WAN Optimization for the selected WAN bandwidth. To enable WAN Optimization on the EC-V, go to the Deployment page and the Business Intent Overlay (BIO) page after the deployment is complete.
WAN Bandwidth	The Bandwidth drop-down list displays the current EdgeConnect license tiers. After you select a WAN Bandwidth value, Orchestrator displays the appropriate AWS instance types for the deployment on the Instance Type drop-down menu.
Instance Type	Based on your selection of WAN Optimization and WAN Bandwidth values, Orchestrator displays the appropriate AWS instance types on this drop-down menu. **NOTE:** The Elastic Network Adapter (ENA) driver on EC-V supports only IPv4 traffic. To send IPv6 traffic or establish IPv6-based tunnels, deploy EC-V on an AWS C4 instance.

Field	Description
AWS Tags (Optional)	Any comma-separated tags entered here are applied to all AWS resources that Orchestrator creates while deploying the EC-V. If you do not enter any tags, Orchestrator automatically creates a unique tag for each AWS resource that it creates while deploying the EC-V. This AWS tag is created to identify each resource created by Orchestrator. The tag is formatted as follows: sp-automated-*deployment name*-instance-*index-resource name*.
Comment (Optional)	Enter an optional comment if you want to attach any additional details for the deployment.
Advanced Settings	**Custom AMI ID:** If you want to deploy the EC-V with a specific public or private image, provide the AMI ID. You can obtain the AMI ID from the AWS console. Leave this field blank to allow Orchestrator to deploy the EC-V with the base AMI obtained from the AWS Marketplace.
Horizontally Scale	You can deploy multiple EC-Vs by clicking **+** and selecting the Availability Zone for each EC-V. If the selected region supports multiple Availability Zones, each Availability Zone is shown on the drop-down menu. When deploying multiple EC-Vs, it is best practice to deploy each EC-V in a unique Availability Zone.
Appliance Tag (Optional)	Enter an Appliance Tag on this field if you want to assign a pre-configuration file to the deployment. If this field is left blank, Orchestrator will automatically assign an Appliance Tag for its own configuration purposes.

When you have completed all of the required fields, click **Review and Deploy**. Review the configuration summary, and click **Deploy** to create the EC-V instances.

Deploy Cloud Hubs in Azure

Configuration > Cloud Services > IaaS > Deploy Cloud Hubs in Azure

The Deploy Cloud Hubs in Azure tab provides the Azure account details and EC-V deployment configuration details for all Azure cloud EC-Vs that have been deployed.

IMPORTANT: For more detailed information and configuration options, including integration with services such as the Standard Internal Load Balancer and Azure Route Server, see the HPE Aruba Networking EdgeConnect SD-WAN Virtual (EC-V) in Microsoft Azure Deployment Guide.

NOTE: Before you can deploy EC-Vs to the Azure cloud, you must perform several tasks on the Azure portal. For more information, see Azure Subscriptions.

NOTE: EC-Vs that are deployed manually in Azure will not be displayed in Orchestrator.

Use this tab to:

- Create and modify Azure subscriptions
- Deploy EC-Vs in the Azure cloud, within either a Transit VNet or a vWAN hub
- Remove an Azure cloud deployment

  NOTE: When you remove a deployment, all EC-Vs in the deployment will be deleted.

The following table describes each field on this tab.

Field	Description
Deployment name	Name given on the deployment configuration page.
Subscription	Name of the Azure account used to deploy the EC-Vs.
Region	Region of the EC-V deployment.
Resource group	Name of the Azure Resource Group used for the EC-V deployment.
Platform	Type of platform (azure or azure_nva) used for the EC-V deployment.
Virtual network	CIDR block used for deployment.
Instances	Number of EC-V instances in the deployment. To add one or more EC-Vs to the deployment, click **+Add**. In the New Instance in Azure dialog box, select the Availability Zone to use and an optional tag to apply to the new instance. **Max** indicates that the maximum number of instances have been created for this deployment. If the region you selected does not support Availability Zones, the New Instance in Azure dialog box does not display an Availability Zone menu.
Status	Status of the deployment. If more information is available, an information icon is displayed. **NOTE:** If the deployment was incomplete, the info dialog contains a link to download the log file and steps to resolve the issue.
Terminate	To permanently delete a deployment, click **Terminate**. This action deletes all resources associated with the EC-Vs, including all Azure resources. If you created more than one EC-V in the deployment, all EC-Vs will be deleted when you click **Terminate**. The Resource Group that was used for the deployment will *not* be deleted.
Deployment info	Click the info icon in this column to view deployment and virtual machine details.

Field	Description
Resources	Click the info icon in this column to view details about each Azure resource that Orchestrator created during the deployment.
Comment	Comments that were added to the deployment when the EC-V was created. To edit the comment, click the edit icon.

Create or modify an Azure subscription

Click **Azure Subscriptions** to create or modify an Azure subscription to Orchestrator.

Deploy a new EC-V

You can deploy a new EC-V within either a Transit Virtual Network (VNet) or a Virtual WAN (vWAN) Hub as a Managed NVA.

When deciding between the two deployment methods, there are several factors to consider. Deploying an EC-V inside a transit VNet offers more flexibility and control of the EC-V deployment. Managed NVA includes limitations such as not being able to add or delete network interfaces on the EC-V, not having access to SSH, and others. For more information on this decision and design specifications of each deployment, see the HPE Aruba Networking EdgeConnect SD-WAN Virtual (EC-V) in Microsoft Azure Deployment Guide.

Click **Deploy Cloud Hubs in Transit VNet** or **Deploy Cloud Hubs in Virtual WAN Hub** to deploy one or more EC-V instances in Azure.

Remove an EC-V

If a deployment does not complete or you no longer want the EC-V in the Azure cloud, you can remove the deployment and all associated artifacts.

To remove a deployment, locate the deployment you want to remove, and then click **Terminate** in the desired row.

Azure subscriptions

Before deploying an EC-V in Azure, you must configure an Azure subscription and add it to Orchestrator.

The Azure Subscriptions dialog box lists all the Azure subscriptions that have been added to Orchestrator.

Add new Azure subscription

To add a new Azure subscription, click **Add Azure Subscription**.

Edit an existing Azure subscription

To edit an existing Azure subscription:

1. Hover over the Actions column of the subscription you want to edit, and then click the edit icon.

   The Azure Subscription Configuration dialog box appears.

   NOTE: You cannot modify subscriptions that have active deployments.
2. Modify the elements as necessary.
3. Click **Save**.

   Orchestrator validates the subscription information.
4. Click **Close**.

Deploy cloud hubs in Transit VNet

Use the Deploy Cloud Hubs in Transit VNet dialog box to create one or more EC-V instances in Azure.

Configure the Azure subscription

Before deploying an EC-V via Orchestrator, you must configure your Azure subscription by accepting Marketplace terms for EdgeConnect to enable programmatic deployment, creating an app registration (service principal), resource group, and custom role, and assigning that role to the resource group.

Accept Azure Marketplace image terms

Accepting Azure Marketplace's image terms for EdgeConnect is required for the Orchestrator to automatically deploy an EdgeConnect image from the Azure Marketplace. You only need to do this once per Azure subscription.

1. Log in to the Azure Portal.
2. Under Azure services, click **+ Create a resource**.
3. On the Create a resource page, enter *edgeconnect*, and then click **Silver Peak Unity EdgeConnect**.
4. On the Plan drop-down menu, select the latest EdgeConnect version, and then click **Get started**.
5. On the Configure Programmatic Deployment page, select **Enable** next to the subscription ID that you want to use to deploy the EdgeConnect VMs.
6. Click **Save**.

   A message at the top of the page notifies you when configuration updates are complete.

Create a new app registration

Azure app registrations are used to create service principals, which are identities used by applications or services to access Azure resources securely and with restricted permissions. The Orchestrator uses the app registration to authenticate itself with Azure. An app registration allows you to assign least-privilege access using custom roles. It avoids sharing credentials and exposing elevated privileges unnecessarily.

To create a new app registration for the Orchestrator:

1. Log in to the Azure Portal.
2. In the main search menu, enter *app registrations*, and then click **App registrations**.
3. Click **+ New registration**.
4. On the Register an application page, in the Name field, enter a user-facing display name for the application.
5. Under Supported account types, select **Accounts in this organizational directory only (Default Directory only - single tenant)**.
6. *(Optional)* Enter a redirect URI.
7. Click **Register**.

   NOTE: Note the application (client) ID and directory (tenant) ID. You need these IDs when you add the subscription details in Orchestrator.
8. Under Manage, click **Certificates & secrets**.
9. Click **New client secret**.
10. Enter a Description and Expiration Date.
11. Click **Add**.

    A new client secret is created.
12. Copy the text in the Value column.

    IMPORTANT: This text can only be viewed immediately after creation. Be sure to save the secret before leaving the page.
13. On the main search menu bar, enter *subscription*, and then press **Enter**.
14. Copy the subscription ID.

Create a new resource group

Creating a new resource group on the Azure Portal is considered a best practice. This ensures that the SD-WAN Orchestrator only has access to that resource group to deploy EC-Vs. However, it is possible to deploy one or more EC-Vs into an existing resource group that contains other Azure resources.

To create a new resource group:

1. On the main search menu, enter *resource group*, and then click **Resource groups**.
2. Click **+ Create**.

3. On the Create a resource group page, select the subscription that you want to use to create the resource group.
4. Enter a name for the resource group, and then select a region.
5. Click **Review + create**.
6. Click **Create**.

Create a custom role

To create custom roles, you must have owner or user access administrator permissions. There are multiple ways to create a custom role. The following steps create a custom role from within the resource group that you created.

1. Select the resource group you created in Create a new resource group, and then click **Access control (IAM)**.
2. Click **Add**, and then click **Add custom role**.

   The Custom Roles editor appears (the Basics tab is displayed).
3. In the Custom role name field, enter a name for the custom role. The name must be unique for the Azure AD directory. The name can include letters, numbers, spaces, and special characters.
4. *(Optional)* In the Description field, enter a description for the custom role. The description displays in the tool tip for the custom role.
5. Accept the default value for the Baseline permissions, and then click the **JSON** tab.
6. Click **Edit**.
7. Go to https://arubanetworking.hpe.com/techdocs/SilverPeak/files/cloud-ecv/cloud_ecv_json.htm and click **Permissions required to deploy Cloud Hubs in Azure in a Transit Virtual Network (VNet)**.
8. Copy and paste the list of Azure permissions within the square brackets next to actions (line 10).
9. Click **Save**.
10. Click the **Assignable scopes** tab, and then verify that the resource group you created is added as an assignable scope and Type is set to the resource group.
11. Click the **Permissions** tab, and then verify that the permissions, descriptions, and permission types you added are listed.
12. Click **Review + create**.
13. Click **Create**.

    A message displays to confirm that you have successfully created your custom role.

Assign the custom role to the resource group

Assigning the custom role to the resource group ensures that the Orchestrator has the exact set of permissions needed to deploy only the resources within that resource group, in line with the principle of least privilege. However, if you want to deploy EdgeConnect gateways into multiple resource groups, you can assign the custom role at the Azure subscription level. This allows the Orchestrator to list multiple resource groups and lets you select the one you want for deployment. In this example, the custom role is assigned only to the specific resource group.

1. Navigate to the resource group you created, and then click **Access control (IAM)**.

   NOTE: If you just completed the previous task of creating a custom role, the Access control (IAM) page is already open.

2. Click **Add**, and then click **Add role assignment**.

   The Role assignment page appears.

3. On the Role tab, enter the name of your custom role.

   NOTE: If the role you created is not displayed, refresh the page.

4. Select the custom role, and then click **Next**.

   The Members tab appears.

5. Ensure that **User, group, or service principle** is selected, and then click **+ Select members**.

   The Select members page appears.

6. Enter the name of your app registration (service principle), select your app, and then click **Select**.

   Your app is added under Members.

7. Click **Review + assign**.

8. Click **Review + assign** again.

Add Azure subscription details to Orchestrator

Add your Azure subscription details such as subscription ID, directory (tenant) ID, application (client) ID, and client secret value to Orchestrator. You only need to do this once per Azure subscription unless you are a large enterprise with multiple teams that want to use different resource groups to segregate deployments.

To add the Azure subscription to Orchestrator:

1. Log in to Orchestrator.
2. Navigate to **Configuration > Cloud Services > IaaS > Deploy Cloud Hubs in Azure**.
3. Click **Azure Subscriptions**.
4. Click **Add Azure subscription**.

5. Enter a Name, Subscription ID, Directory (tenant) ID, Application (client) ID, and Client secret value for the Azure subscription.

   NOTE: You copied the directory (tenant) ID, application (client) ID, and client secret in Create a new app registration.

   NOTE: If you copy and paste the subscription ID, Azure might add a blank space to the beginning of the subscription ID. Be sure to remove all spaces from your subscription ID.

6. Click **Save**.

   Orchestrator validates the subscription information.

NOTE: If you have multiple Azure subscriptions, you can load the credentials for each subscription.

Deploy EC-V into a transit VNet from Orchestrator

To deploy EC-V into a transit VNet from Orchestrator:

1. On the Deploy Cloud Hubs in Azure tab, click **Deploy Cloud Hubs in Transit VNet**.

   The Deploy Cloud Hubs in Transit VNet dialog box appears.

2. To instantiate EC-V instances, enter the deployment details in the Orchestrator's Deploy Cloud Hubs in Transit VNet dialog box, as noted in the table below.

   NOTE: If you do not have an Azure subscription configured in Orchestrator, the Deploy Cloud Hubs in Transit VNet dialog box is blank. Click the **Subscriptions** link to go to the Azure subscription page and enter your Azure credential.

Field	Description
Name	Enter a name for the deployment. This name is used only for identifying the deployment. A deployment consists of one or more EC-Vs that an Orchestrator creates in an Azure Virtual Network. Only alphabetical letters and hyphens are allowed in the deployment name. The maximum allowed length is 20 characters.
Azure account	Select an Azure account to use for deploying the EC-V.
Resource group	Select an Azure resource group to use for deploying the EC-V.
Region	Select an Azure region where you want to deploy the EC-V.

Field	Description
Virtual network	Select **Create new virtual network** or **Existing virtual network**. **Create new virtual network:** If selected, Orchestrator creates a new VNet for the EC-V. Orchestrator creates three subnets (MGMT0, WAN0, and LAN0) for each EC-V you deploy. For example, if you deploy two EC-Vs, Orchestrator creates six subnets. **Existing virtual network:** If selected, Orchestrator allows you to select an existing VNet and subnets for MGMT0, WAN0, and LAN0 interfaces. **NOTE:** If you select **Existing virtual network** and you deploy multiple ECVs using the Horizontally scale setting, the MGMT0, WAN0, and LAN0 interfaces for each EC-V are created using the subnets you select in the Available subnets fields. For example, if you deploy two EC-Vs, the MGMT0 interface for each EC-V is created on the MGMT0 subnet you select in the mgmt0 field, the WAN0 interface for each EC-V is created on the subnet you select in the wan0 field, and the LAN0 interface for each EC-V is created on the subnet you select in the lan0 field. For a two-EC-V (or a multiple EC-V) deployment, you need three subnets. If you select **Create new virtual network**, Orchestrator creates six new subnets for a two-EC-V deployment. If you select **Existing virtual network**, the VNet that you select must be created within the same resource group as the EC-V deployment. If your VNet is created in a different resource group, the Orchestrator does not display that VNet under the Existing virtual network option.
Virtual network CIDR	If you select the Create new virtual network option, you need to enter a Virtual Network Classless Inter-Domain Routing (CIDR) block. Orchestrator uses this CIDR block to create a new VNet. The smallest supported CIDR block is /24 and the largest supported CIDR block is /16. Orchestrator creates all Azure resources required for the EC-V deployment within this virtual network. For each EC-V you deploy, Orchestrator creates three subnets that are /28 in size. In other words, if you deploy two EC-Vs, Orchestrator creates six subnets in total. This is true even if both EC-Vs are created in a single Availability Set or Availability Zone.
Choose virtual network	If you selected Existing virtual network, enter the name of the network in this field.
Available subnets	If you selected Existing virtual network, enter a subnet for each network interface.

Field	Description
WAN optimization	After WAN optimization and an appropriate WAN bandwidth value are selected, Orchestrator displays the appropriate Azure instance types for the deployment on the Instance Type menu. **NOTE:** Selecting WAN optimization does not enable WAN optimization on the EC-V. It only allows Orchestrator to display appropriate Azure instance types that can support WAN optimization for the selected WAN bandwidth. To enable WAN optimization on the EC-V, go to the Deployment page and the Business Intent Overlay (BIO) page after the deployment is complete.
WAN bandwidth	The WAN bandwidth list displays the current EdgeConnect license tiers. After you select a WAN bandwidth value, Orchestrator displays the appropriate Azure instance types for the deployment in the Instance type list.
Instance type	Based on your selected WAN optimization and WAN bandwidth values, Orchestrator displays the appropriate instance types.
Availability option	Select **Availability Set** or **Availability Zone**. Some regions only support Availability Set. HPE Aruba Networking recommends selecting Availability Zone, if available.
SSH public key	Generate a public key with an application, such as PuTTYgen, and then input the value here. **IMPORTANT:** EdgeConnect only supports single-line SSH public keys. *Do not use multi-line SSH public keys.* Additionally, use an EdDSA (ED25519) key pair. **NOTE:** Save the private key file. If you need to log in via SSH to the appliance after it is deployed, you will need this key.
Azure tags (Optional)	Any comma-separated tags entered here are applied to all Azure resources that Orchestrator creates while deploying the EC-V. If you do not enter any tags, Orchestrator automatically creates a unique tag for each Azure resource that it creates while deploying the EC-V. This Azure tag is created to identify each resource created by Orchestrator. The tag is formatted as follows: `sp-automated-deployment name-instance-index-resource name`.
Comment (Optional)	Enter a comment if you want to attach any additional details for the deployment.
Advanced settings	**Custom VHD**: Leave this field blank unless you have an EdgeConnect VHD that you want to use for the deployment. When this field is blank, the Azure Marketplace image is deployed.

Field	Description
Horizontally scale	You can deploy multiple EC-Vs by clicking **+** and selecting the Availability Set or Availability Zone for each EC-V. If the selected region supports multiple availability zones, each availability zone appears on the menu. You can deploy up to five EC-Vs with a CIDR block of /24. If you need to deploy more than five EC-Vs within a single virtual network, select a virtual network CIDR block that is bigger than /24, such as /23 or /22. The maximum number of EC-Vs you can deploy within a single network is 20.
Appliance tags (Optional)	Enter an appliance tag. If this field is left blank, Orchestrator automatically assigns an appliance tag for its own configuration purposes.
Availability zone	Enter the Azure availability zone for the EC-V. **NOTE:** This field only displays if the region supports availability zones.

3. When you have completed all the required fields, click **Review and deploy**.
4. Review the configuration summary, and then click **Deploy** to create the EC-V instances.

   NOTE: After a few minutes, click the refresh icon next to the Deploy Cloud Hubs in Azure label to check the status. If your EC-V deployment fails, the status will be shown as *Incomplete* in the Cloud Hubs in Azure table. Click the info icon to download the log file. The reason for the failure is typically at the bottom of the log file.

   - To remove (delete) a deployment after a failure, click **Terminate**. All Azure resources created by the Orchestrator are removed. If termination fails, it indicates that an Azure resource is blocking the deletion. To view details of this Azure resource, download the log file and check the last error message.
   - To proceed with the termination, you can manually delete this resource on the Azure Portal.

   If your deployment succeeds, the status will be shown as *Deployed*.
5. After the deployment succeeds, the newly deployed EC-Vs appear on the Discovered Appliances tab in Orchestrator. Please allow at least 10 minutes for EC-Vs to be discovered in Orchestrator. On the Discovered Appliances tab, the Approve button appears after the EC-V is fully configured.

Deploy Cloud Hubs in Virtual WAN Hub

This section describes how to deploy EC-V gateways directly within an Azure vWAN hub and establish BGP. This design is also referred to as *Managed NVA* or *NVA in vWAN hub*. To learn more about this design, see the HPE Aruba Networking EdgeConnect SD-WAN Virtual (EC-V) in Microsoft Azure Deployment Guide.

Configure the Azure subscription

Before deploying Managed NVA via Orchestrator, you must configure your Azure subscription by creating an app registration (service principal), a resource group, and a custom role, and assigning that role to the resource group.

Create a new app registration

Azure app registrations are used to create service principals, which are identities used by applications or services to access Azure resources securely and with restricted permissions. The Orchestrator uses the app registration to authenticate itself with Azure. An app registration allows you to assign least-privilege access using custom roles. It avoids sharing credentials or exposing elevated privileges unnecessarily.

To create a new App registration for the Orchestrator:

1. Log in to the Azure Portal.
2. In the main search menu, enter *app registrations*, and then click **App registrations**.
3. Click **+ New registration**.
4. On the Register an application page, in the Name field, enter a user-facing display name for the application.
5. Under Supported account types, select **Accounts in this organizational directory only (Default Directory only - single tenant)**.
6. *(Optional)* Enter a redirect URI.
7. Click **Register**.

   NOTE: Note the application (client) ID and directory (tenant) ID. You need these IDs when you add the subscription details in Orchestrator.
8. Under Manage, click **Certificates & secrets**.
9. Click **New client secret**.
10. Enter a Description and Expiration Date.
11. Click **Add**.

    A new client secret is created.
12. Copy the text in the Value column.

    NOTE: This text can only be viewed immediately after creation. Be sure to save the secret before leaving the page.
13. On the main search menu bar, enter *subscription*, and then press **Enter**.
14. Copy the subscription ID.

Create a new resource group

Creating a new resource group on the Azure Portal is considered the best practice. This ensures that the SD-WAN Orchestrator only has access to that resource group to deploy EC-Vs. However, it is possible to deploy one or more EC-Vs into an existing resource group that contains other Azure resources.

To create a new resource group:

1. From the main search menu in Azure Portal, enter *resource group*, and then click **Resource groups**.
2. Click **+ Create**.
3. On the Create a resource group page, select the subscription that you want to use to create the resource group.
4. Enter a name for the resource group, and then select a region.
5. Click **Review + create**.
6. Click **Create**.

Create a custom role

To create custom roles, you must have Owner or User Access Administrator permissions. There are multiple ways to create a custom role. The following steps create a custom role from within the resource group you created.

1. Select the resource group you created in the previous section, and then click **Access control (IAM)**.
2. Click **Add**, and then click **Add custom role**.

   The Custom Roles editor appears (the Basic tab is displayed).
3. In the Custom role name field, enter a name for the custom role. The name must be unique for the Azure AD directory. The name can include letters, numbers, spaces, and special characters.
4. *(Optional)* In the Description field, enter a description for the custom role. The description displays in the tool tip for the custom role.
5. Accept the default value for the Baseline permissions, and then click the **JSON** tab.
6. Click **Edit**.
7. Go to https://arubanetworking.hpe.com/techdocs/SilverPeak/files/cloud-ecv/cloud_ecv_json.htm, and then click **Permissions required to deploy EC-Vs inside an Azure Virtual WAN Hub-Managed NVA and establish BGP connectivity**.
8. Copy the list of Azure permissions, and then paste the list within the square brackets under Actions (line 10).
9. Click **Save**.
10. Click the **Assignable scopes** tab, and then verify that the resource group you created is added as an assignable scope and Type is set to the resource group.

11. Click the **Permissions** tab, and then verify that the permissions, descriptions, and permission types you added are listed.
12. Click **Review + create**.
13. Click **Create**.

A message displays to confirm that you have successfully created your custom role.

Assign the custom role to the resource group

Assigning the custom role to the resource group ensures that the Orchestrator has the exact set of permissions needed to deploy only the resources within that resource group, in line with the principle of least privilege. However, if you want to deploy EdgeConnect gateways into multiple resource groups, you can assign the custom role at the Azure subscription level. This allows the Orchestrator to list multiple resource groups and lets you select the one you want for deployment. This example assigns the custom role only to the specific resource group.

1. Navigate to the resource group you created, and then click **Access control (IAM)**.

   NOTE: If you just completed the previous task of creating a custom role, the Access control (IAM) page is already open.
2. Click **Add**, and then click **Add role assignment**.
3. On the Role tab, enter the name of your custom role.

   NOTE: If the role you created is not displayed, refresh the page.
4. Select the custom role, and then click **Next**.
5. On the Members tab, ensure that User, group, or service principle is selected, and then click **+ Select members**.
6. On the Select members tab, enter the name of your App registration (Service Principle), select your app, and then click **Select**.

   Your app is added under Members.
7. Click **Review + assign**.
8. Click **Review + assign** again.

Add Azure subscription details on Orchestrator

This section explains how to add your Azure subscription details—such as Subscription ID, Directory (tenant) ID, Application (client) ID, and Client secret value—to Orchestrator. You only need to do this once per Azure subscription unless you are a large enterprise with multiple teams that want to use different resource groups to segregate deployments.

To add the Azure subscription to Orchestrator:

1. Log in to Orchestrator, and then navigate to **Configuration > Cloud Services > IaaS > Deploy Cloud Hubs in Azure**.
2. Click **Azure Subscriptions**, and then click **Add Azure subscription**.

3. Enter a Name, Subscription ID, Directory (tenant) ID, Application (client) ID, and Client secret value for the Azure subscription.

   NOTE: You copied the directory (tenant) ID, application (client) ID, and client secret in Create a new app registration.

   NOTE: If you copy and paste the subscription ID, Azure might add a blank space to the beginning of the subscription ID. Be sure to remove all spaces from your subscription ID.

4. Click **Save**.

Orchestrator validates the subscription information.

NOTE: If you have multiple Azure subscriptions, you can load the credentials for each subscription.

Deploy Managed NVA

To deploy Managed NVA from Orchestrator:

1. Log in to the Orchestrator, and then navigate to **Configuration > Cloud Services > IaaS > Deploy Cloud Hubs in Azure**.
2. Click **Deploy Cloud Hubs in Virtual WAN Hub**.
3. Enter the deployment details shown below.

   NOTE: If you do not have an Azure subscription configured in Orchestrator, the Deploy Cloud Hubs in Virtual WAN Hub dialog box is blank. Click the **Subscriptions** link to go to the Azure subscription page and enter your Azure credential.

Field	Description
Name	Enter a name for the deployment. This name is used only for identifying the deployment. A deployment consists of two EC-Vs created within a vWAN Hub. Only alphabetical letters and hyphens are allowed in the deployment name. The maximum allowed length is 20 characters.
Azure account	Select an Azure subscription to use for deploying the Managed NVA.
Resource group	Select an Azure resource group to use for deploying the Managed NVA.
Region	Select an Azure region where you want to deploy the Managed NVA.
Virtual WAN hub	Select **Create new virtual network** or **Existing virtual network**. Select your vWAN Hub. If no vWAN Hubs are shown, the region you selected does not have any vWAN Hubs created, or you have not entered the required permissions to your custom role.

Field	Description
Scale units	The scale unit determines the resource allocation on the Managed NVA. The higher the scale unit, the greater the amount of traffic that can be handled.
EdgeConnect ASN	Enter an ASN between 64512 and 65534, excluding the following ASNs: i. ASNs reserved by Azure: 1. Public ASNs: 8074, 8075, 12076 2. Private ASNs: 65515, 65517, 65518, 65519, 65520 ii. ASNs reserved by IANA: 1. 23456, 64496-64511, 65535-65551 **NOTE:** vWAN Hub only supports 16-bit ASNs. As a result, do not enter a 32-bit ASN. The ASN you enter here will be assigned to both EC-Vs.
SSH public key	Generate a public key with an application, such as PuTTYgen, and then input the value here. **IMPORTANT:** EdgeConnect only supports single-line SSH public keys. *Do not use multi-line SSH public keys.* Additionally, use an EdDSA (ED25519) key pair. **NOTE:** Save the private key file. If you need to log in via SSH to the appliance after it is deployed, you will need this key.
Comment (Optional)	Enter a comment if you want to attach any additional details for the deployment.

4. Click **Review and deploy**.
5. Review the configuration summary, and then click **Deploy** to create the EC-V instances.

   NOTE: To check the status after a few minutes, click the refresh icon next to the Deploy Cloud Hubs in Azure button.
6. If your EC-V deployment fails, the status appears as *Incomplete* in the Cloud Hubs in Azure table. To download the log file, click the info icon. The reason for the failure is typically at the bottom of the log file.

   - To remove (delete) a deployment after a failure, click **Terminate**. All Azure resources created by the Orchestrator will be removed. If termination fails, an Azure resource is blocking the deletion.

- To view details of this Azure resource, download the log file and then check the last error message. To proceed with the termination, you can manually delete this resource on the Azure Portal.

If your deployment succeeds, the status appears as *Deployed*.

7. After the deployment succeeds, the newly deployed EC-Vs appear on the Discovered Appliances tab in Orchestrator. Please allow at least 10 minutes for EC-Vs to be discovered in Orchestrator. On the Discovered Appliances tab, the Approve button appears after the EC-V is fully configured.

Cloud Hubs in GCP

Configuration > Cloud Services > IaaS > Deploy Cloud Hubs in GCP

The Cloud Hubs in GCP tab provides the Google Cloud Platform account details and EC-V deployment configuration details for all cloud EC-Vs that have been deployed.

Use this tab to:

- Add and modify GCP accounts
- Deploy EC-Vs in GCP
- Manage EC-V instances deployed in GCP

NOTE: Before you can deploy EC-Vs to the GCP cloud, you must perform several tasks in GCP. For more information, see GCP Account Configuration.

The following table describes each field on this tab.

Field	Description
Project	Project ID for the EC-V deployment.
Region	GCP Region in which the EC-V instance was deployed.
Zone	GCP zone in which the EC-V was deployed.
Tag	Appliance tag for an EC-V instance. If not configured, this value is automatically assigned.
Status	Status of an EC-V instance deployed in GCP. If more information is available, an info icon is displayed. **NOTE:** If the deployment failed for the selected instance, the info dialog contains a link to download the log file and steps to resolve the issue.
Terminate	To terminate an EC-V instance in GCP, click **Terminate**. This action deletes all resources associated with the selected EC-V instance.
Deployment Info	Click the info icon in this column to view details of the selected EC-V instance, including the IP addresses associated with the mgmt0, wan0, and lan0 interfaces.

Field	Description
Resources	Click the info icon in this column to view details about all GCP resources that Orchestrator created during the deployment for the selected EC-V instance.
Info	The appliance hostname after it has been approved and added to Orchestrator. If an appliance has not been approved, this column will be blank.
Comment	Comments that were added to the deployment when the EC-V was created. To edit the comment, click the edit icon.

Add or Modify a GCP Account

To add or modify a GCP account to Orchestrator:

1. Click **GCP Service Accounts**.

   The GCP Service Accounts dialog box opens.
2. Click **New GCP Account** or click the edit icon next to the account you want to edit.

   The GCP Account Configuration dialog box opens.
3. Complete or modify the elements as necessary.

Deploy a New EC-V

To deploy one or more EC-V instances in GCP, click **New Deployment**.

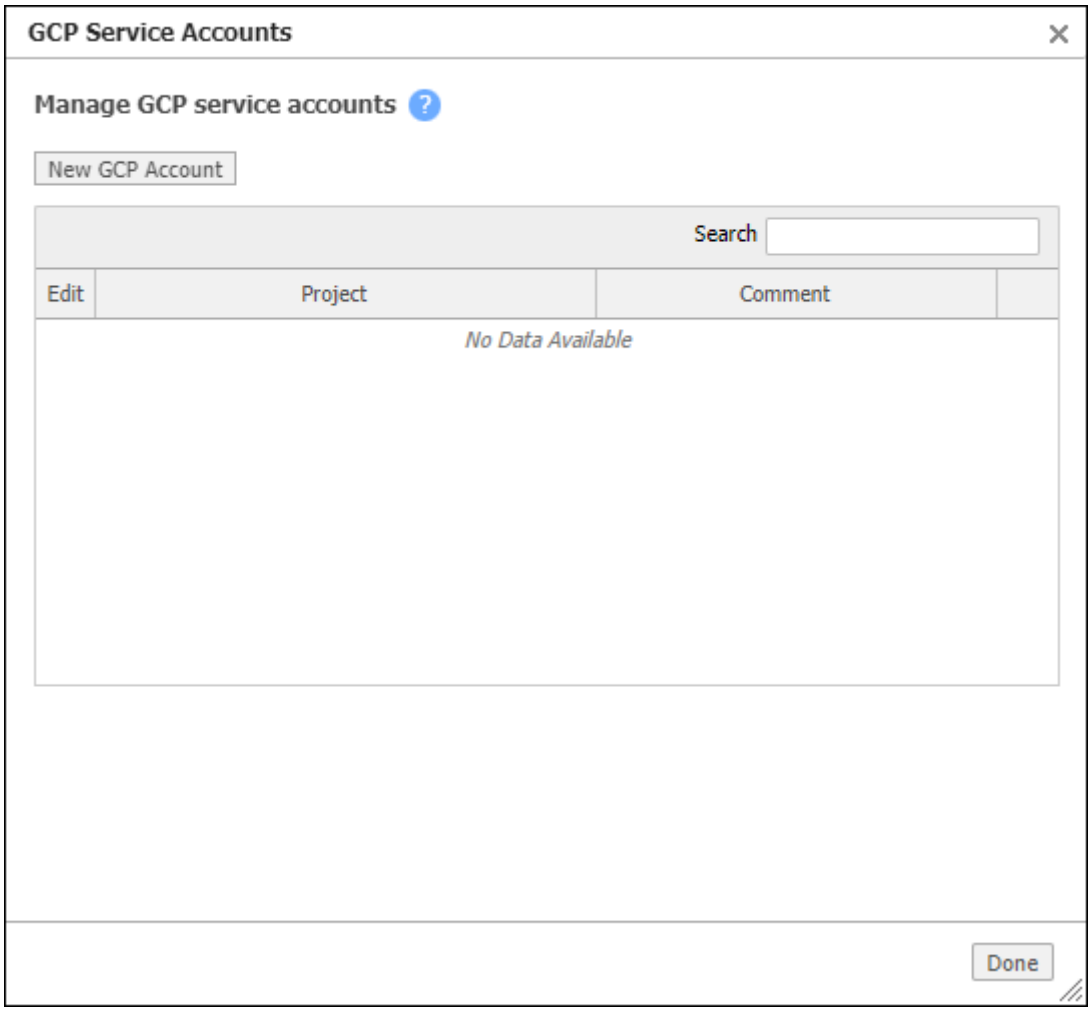
Manage an EC-V

If a deployment does not complete or you no longer want the EC-V in the GCP cloud, you can remove the deployment and all associated artifacts.

To remove a deployment, locate the deployment you want to remove, and then click **Terminate** in that row.

GCP Accounts

The GCP Service Accounts dialog box lists all of the GCP accounts that have been added.



- To create a new account for EC-V deployments, click **New GCP Account**.
- To modify an existing account's details, click the edit icon next to the account.

GCP Account Configuration

Complete the following steps to create an GCP user account with the required permissions for creating EC-V instances in GCP.

Create a GCP Project

1. Log in to GCP (https://cloud.google.com).
2. On the right-side menu, hover over **IAM & Admin**, and then click **Create a Project**.
3. Enter the required information, and then click **Create**.

Enable GCP APIs

This section explains how to enable two different APIs required for creating EC-V instances in GCP.

Enable Compute Engine API

1. On the right-side menu, hover over **APIs & Services**, and then click **Library**.
2. In the search bar, search for **Compute Engine API**.
3. Click **Compute Engine API**.
4. Click **Enable**.

Enable Google Cloud Resource Manager API

1. On the right-side menu, hover over **APIs & Services**, and then click **Library**.
2. In the search bar, search for **Cloud Resource Manager API**.
3. Click **Cloud Resource Manager API**.
4. Click **Enable**.

Create a Custom Role

This section explains how to create a custom role and attach permissions in GCP. There are two ways to do this. Select one of the following options (you do not need to do them both):

Create a Custom Role Using Google Cloud Shell (Recommended)

Create a Custom Role Using Google Cloud Console

Create a Custom Role Using Google Cloud Shell (Recommended)

1. In Google Cloud, click the **Activate Cloud Shell** icon on the top menu bar.
2. Paste the following into the command line, making sure to replace the ** variable with the project name you created above:

```
gcloud iam roles create CustomRoleForArubaOrchestrator --project=<PROJECT NAME>
    --title="Custom Role for SD-WAN Orchestrator" --description="Custom Role for
    SD-WAN Orchestrator." --permissions="compute.disks.create,compute.firewalls.
    create,compute.firewalls.delete,compute.firewalls.get,compute.images.get,
    compute.instances.create,compute.instances.delete,compute.instances.get,
    compute.instances.setMetadata,compute.networks.create,compute.networks.delete
    ,compute.networks.get,compute.networks.updatePolicy,compute.regions.list,
    compute.subnetworks.create,compute.subnetworks.delete,compute.subnetworks.get
    ,compute.subnetworks.use,compute.subnetworks.useExternalIp,compute.zones.get"
     --stage=GA
```

3. Exit Cloud Shell and verify that **Custom Role for SD-WAN Orchestrator** is enabled in the list of roles for your project.

Create a Custom Role Using Google Cloud Console

1. On the right-side menu, hover over **IAM & Admin**, and then click **Roles**.
2. Click **Create Role**.
3. Fill in the Title and ID fields.
4. In the Role launch stage field, select **General Availability**.
5. Click **Add Permissions**.
6. Add the following permissions to the role:

   - compute.disks.create
   - compute.firewalls.create
   - compute.firewalls.delete
   - compute.firewalls.get
   - compute.images.get
   - compute.instances.create
   - compute.instances.delete
   - compute.instances.get
   - compute.instances.setMetadata
   - compute.networks.create
   - compute.networks.delete
   - compute.networks.get
   - compute.networks.updatePolicy
   - compute.regions.list
   - compute.subnetworks.create
   - compute.subnetworks.delete
   - compute.subnetworks.get
   - compute.subnetworks.use
   - compute.subnetworks.useExternalIp
   - compute.zones.get
   - compute.globalOperations.get

   NOTE: These permissions cannot be batch-added in GCP. For each entry, you must search for the permission, select the check box next to it, and then click **Add**.
7. On the Create Role page, verify the assigned permissions, and then click **Create**.

Create a GCP Service Account

1. On the right-side menu, hover over **IAM & Admin**, and then click **Service Accounts**.
2. Click the project created in the previous steps.
3. Click **Create Service Account**.
4. Enter the required information, and then click **Create and Continue**.
5. In the Select a role field, select **Custom Role for SD-WAN Orchestrator**, and then click **Continue**.
6. Click **Done**.

Create a Service Account Key Pair

1. On the right-side menu, click **Service Accounts**.
2. From your project home page, click your service account from the list.
3. Click the **Keys** tab.
4. From the drop-down menu, click **Create new key**.
5. Leave the key type as **JSON**, and then click **Create**.

   The .json file saves to your system.

Add the GCP Account to Orchestrator

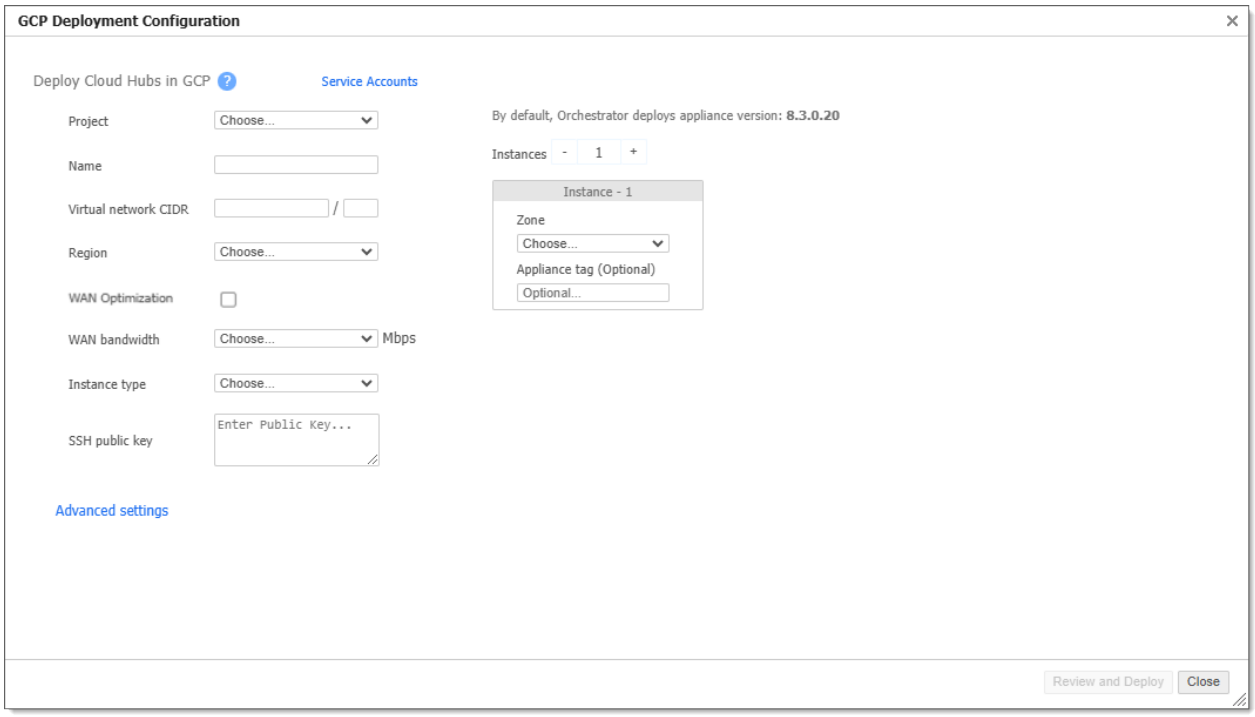
1. In Orchestrator, navigate to **Configuration > IaaS > Cloud Hubs in GCP**.
2. Click **GCP Service Accounts**.
3. Click **New GCP Account**.
4. Paste all content from the .json file saved in the previous section into the Key field, and then click **Save**.

   Orchestrator validates the account information. This takes approximately 45 seconds.

GCP Deployment Configuration

Use the GCP Deployment Configuration page to create one or more EC-V instances in a GCP region.

NOTE: If you do not have a GCP account configured in Orchestrator, the GCP Deployment Configuration dialog box is blank. To create a GCP account, click the **Service Accounts** link.



Field	Description
Project	Select the GCP project created earlier in this procedure.
Name	Enter a name for the deployment. This name is used only for identifying the deployment. A deployment consists of one or more EC-Vs that an Orchestrator creates in GCP. Only alphanumerical letters and hyphens are allowed in the deployment name. The maximum allowed length is 20 characters.
Virtual network CIDR	Enter a Virtual Classless Inter-Domain Routing (CIDR) block. The CIDR block must be at least /16. Orchestrator carves out three x /26 global subnets (mgmt, wan, lan) from Virtual Network CIDR for each region. A /16 CIDR supports deploying in 300 regions. **NOTE:** You only need to enter this value once per GCP project.
Region	Select the GCP region where you want to deploy the EC-V.

Field	Description
WAN Optimization *(Optional)*	WAN Optimization requires additional resources on a GCP instance. After WAN Optimization and an appropriate WAN Bandwidth value are selected, Orchestrator displays the appropriate GCP instance types for the deployment on the Instance Type drop-down menu. **NOTE:** Selecting the WAN Optimization check box does not enable WAN Optimization on the EC-V. It only allows Orchestrator to display appropriate GCP instance types that can support BWAN Optimization for the selected WAN bandwidth. To enable WAN Optimization on the EC-V, go to the Deployment page and the Business Intent Overlay (BIO) page after the deployment is complete.
WAN bandwidth	The Bandwidth drop-down list displays the current EdgeConnect license tiers. After you select a WAN Bandwidth value, Orchestrator displays the appropriate GCP instance types for the deployment on the Instance Type drop-down menu.
Instance type	Based on your selection of WAN Optimization and WAN Bandwidth values, Orchestrator displays the appropriate GCP instance types on this drop-down menu.
SSH public key	Enter the SSH public key for the deployment.
Instances	**Zone:** You can deploy multiple EC-Vs by clicking **+** and selecting the Zone for each EC-V. If the selected region supports multiple zones, each zone is shown on the drop-down menu. When deploying multiple EC-Vs, it is best practice to deploy each EC-V in a unique zone. **Appliance tag *(Optional)*:** Enter an Appliance Tag in this field if you want to assign a pre-configuration file to the deployment. If this field is left blank, Orchestrator will automatically assign an Appliance Tag for its own configuration purposes.
Advanced Settings	**Custom image:** If you want to deploy the EC-V with a specific public or private image, specify the image ID here. You can obtain the image ID from the GCP console. Leave this field blank to allow Orchestrator to deploy the EC-V with the base image obtained from GCP.

When you have completed all of the required fields, click **Review and Deploy**. Review the configuration summary, and then click **Deploy** to create the EC-V instances.

Administration

The menus under **Administration** are related to appliance administration. They include general settings, software management, and tools for troubleshooting and maintenance, and are organized as follows:

- General Settings
- Software
- Tools

Administration > General Settings

The options under **Administration > General Settings** focus on how to apply and manage the core settings for Orchestrator appliances, including user accounts, date/time, flows, certificates, SNMP, and more.

Appliance Users

Administration > General Settings > Users & Authentication > Appliance Users

Use this tab to view and manage user accounts on your network's appliances. This feature requires admin permissions. All user accounts are listed regardless of whether they are enabled. The table lists two or more user accounts for each appliance in your network or those specifically selected in the appliance tree.

- Each appliance automatically provides two default user accounts: *admin* and *monitor*. You cannot remove them. However, you *can* modify and enable/disable these built-in accounts.
- You can create user accounts for a specific appliance. Click the edit icon associated with the appliance to access the Appliance Users dialog box, and then click **Add**. For details, see Appliance Users Dialog Box below.
- You can also create additional *default* user accounts that apply to all appliances in your network by using the User Management template. Click **Manage Users with Templates** to access this template. For more information, see User Management Template.

Each user account is assigned a capability:

- *monitor* – Supports reading and monitoring all data, in addition to performing all actions. This is equivalent to *enable* mode privileges in the Command Line Interface (CLI).
- *admin* – Supports full privileges, along with permissions to add, modify, and delete. This is equivalent to *configuration* mode privileges in the CLI.

Click **Export** to download user account data to a CSV file.

Appliance Users Dialog Box

Use this dialog box to view and manage user accounts on an appliance in your network and to view user sessions currently active on that appliance.

The Appliance Users table lists all accounts configured on the appliance, regardless of whether they are enabled. You can modify and enable/disable any user account, but you cannot delete the built-in *admin* and *monitor* user accounts.

- To add an account, click **Add**. A row is added to the table. Click in the table cells to enter a username, select admin or monitor capability, and enter and confirm a password. User credentials are case-sensitive. To enable the user account, select the **Enabled** check box. Finally, click **Save**.
- To delete an account, click **X** in the last column, and then click **Save**.

The Active Sessions table lists users currently accessing the appliance and includes details, such as IP address location and login time. You can end a user session by clicking **X** in the last column.

Auth/RADIUS/TACACS+ Tab

Administration > *General Settings* > *Users & Authentication* > *Auth/RADIUS/TACACS+*

This tab displays the configured settings for authentication and authorization.

If the appliance relies on either a RADIUS or TACACS+ server for those services, those settings are also reported.

All settings are initially applied by means of the Auth/RADIUS/TACACS+ configuration template.

Authentication and Authorization

Authentication and Authorization Fields

Field	Description
Appliance	Name of the appliance selected.
Authentication Order	When it is possible to validate against more than one database (local, RADIUS server, TACACS+ server), Authentication Order specifies which method to try in what sequence: **Authentication Order First**, **Order Second**, and **Order Third**. . **NOTE:** If the Authentication Order field is set to Remote First, and if a password is configured for the CLI **enable** command, add a user named “enableuser” on the remote server and set the password to be identical to the one configured locally.

Field	Description
Authorization Map Order	Map ordering determines which server is used first. Select the map ordering from the drop-down list: **Local-Only**, **Remote-First**, and **Remote-Only**. The default (and recommended) value is **Remote-First**.
Authorization Default Role	Default role assigned for authorization. The default (and recommended) value is **admin**.
Authentication	Process of validating that the end user, or a device, is who they claim to be.
Authorization	Action of determining what a user is allowed to do. Generally, authentication precedes authorization.
Map Order	Default (and recommended) value is **Remote First**.

RADIUS and TACACS+

RADIUS and TACACS+ Server Fields

Field	Description
Server Type	RADIUS or TACACS+.
Auth Port	For RADIUS, the default value is **1812**. For TACACS+, the default value is **49**.
Auth Type	TACACS+ The options are **pap** or **ascii**.
Timeout	If a logged-in user is inactive for an interval that exceeds the inactivity time-out, the appliance logs them out and returns them to the login page. You can change that value, as well as the maximum number of sessions, in the Session Management template.
Retries	Number of attempts allowed before lockout.
Enabled	Whether or not the server is enabled.

Auth/RADIUS/TACACS+ Edit Row

Select the Authentication Order and Authorization information in this dialog box. You can also add a RADIUS and TACACS+ Server by clicking **Add** under each section.

Authentication Order

Choose which authentication database you want to be **First**, **Second**, and **Third** from the designated drop-down lists.

Authorization Information

Select the Map Order and the Default Role from the designated drop-down lists.

Authentication and Authorization

Authentication and Authorization Fields

Field	Description
Authentication	Process of validating that the end user, or a device, is who they claim to be.
Authorization	Action of determining what a user is allowed to do. Generally, authentication precedes authorization.
Authentication Order	When it is possible to validate against more than one database (local, RADIUS server, TACACS+ server), **Authentication Order** specifies which method to try in what sequence. Default is Local-first.
Map Order	Default (and recommended) value is **Remote First**.
Default Role	Default (and recommended) value is **admin**.

RADIUS and TACACS+

RADIUS and TACACS+ Server Fields

Field	Description
Order	Method RADIUS and TACAC+ specifies first– **local first**.
Auth Port	For RADIUS, the default value is **1812**. For TACACS+, the default value is **49**.
Auth Type	RADIUS The options are **pap** or **chap**. [TACACS+] The options are **pap** or **ascii**.
Enabled	Whether or not the server is enabled.
Retries	Number of attempts allowed before lockout.
Server Type	RADIUS or TACACS+.
Timeout	If a logged-in user is inactive for an interval that exceeds the inactivity time-out, the appliance logs them out and returns them to the login page. You can change that value, as well as the maximum number of sessions, in the Session Management template.

Use Appropriate RADIUS Configuration Options

Follow the steps below to use the CHAP Protocol option for RADIUS authentication. Doing so avoids potential security vulnerabilities.

1. In Orchestrator, assure that CHAP is selected for RADIUS authentication. Go to **Orchestrator > Authentication**.

   The **Authentication** dialog box opens.
2. Select the RADIUS server and click edit to open the **Remote Authentication Server** dialog box.
3. For **Authentication Type**, choose CHAP from the drop-down list.
4. Click **Save** to save your settings.
5. On the RADIUS server, assure that the *clients.conf* file contains the following command: *require_message_authenticator = yes*

   As an example of where to find this file, on a FreeRADIUS server, the path is here: */etc/raddb/clients.conf*
6. Verify that this configuration is in effect by using TCP dump in Orchestrator or via the CLI on your EdgeConnect. You should see that the CHAP message authenticator is included in packets that are exchanged with the RADIUS server.

   NOTE: If this configuration is not in effect, the RADIUS server will still work and a security vulnerability will exist.

Appliance IP Allow List

Administration > General Settings > Users & Authentication > Appliance IP Allow List

Use this tab to configure the Appliance IP Allow List. The Appliance IP Allow List restricts which IP addresses can establish SSH, SNMP and HTTPS connections to an appliance. This blocks access from unknown or potentially malicious IP addresses by ensuring that only specified IP addresses are permitted to initiate sessions. Some examples of specified IP addresses one might choose to allow are those from corporate offices or VPNs.

Appliance IP Allow List Settings

For a target appliance selected in the appliance tree, the configuration options below are available.

Field	Description
Enable	Enable/Disable the IP Allow List feature. The status indicator displays whether the IP Allow List feature is enabled or disabled for the selected appliance.
IP/Mask	Specify a valid IPv4 Remote Address by setting the Mask which determines the range of IP addresses allowed. Possible values are 24 or 32. If the last octet of the IP address is zero, use a mask of 24. Use a mask of 32 for absolute IP addresses.
Authentication Type	Choose one or more protocols (SSH, HTTPS, SNMP) for which the IP entry applies.

Field	Description
Comments	Optionally annotate entries to clarify purpose and ownership for audit or documentation purposes.

Upon saving changes, the configuration tab auto-closes to confirm successful application.

Note the following conditions and guidelines.

- Ensure the IP Allow List feature is enabled and protocols are correctly assigned before relying on its enforcement. Entries may appear even when the feature is disabled.
- Adding a new entry via the CLI with different protocols will override existing protocol settings for that IP.
- Review entries periodically to maintain alignment with security policies.

Appliance Password Settings

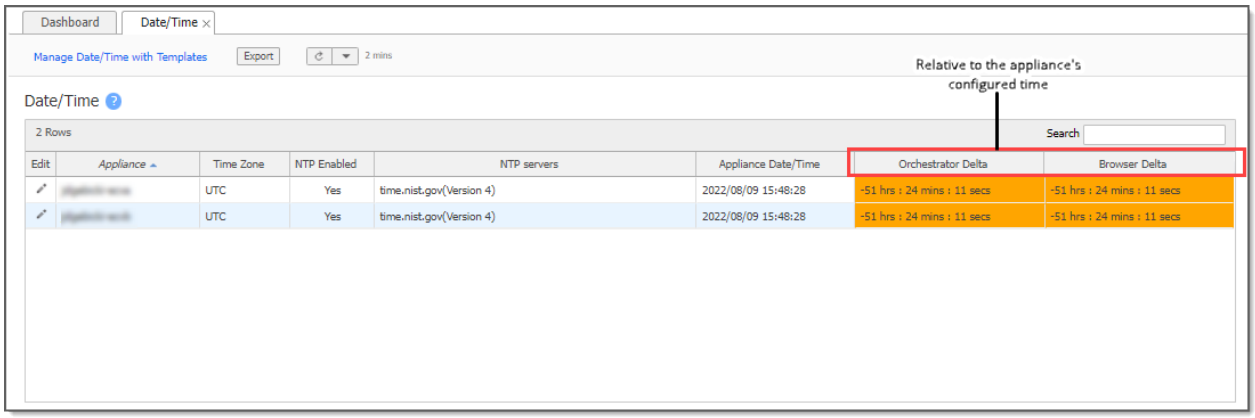
Administration > General Settings > Users & Authentication > Appliance Password Settings

This topic is for future content.

Date/Time Tab

Administration > General Settings > Setup > Date/Time

The Date/Time tab indicates the time zone and Network Time Protocol (NTP) settings for one or more appliances selected in the appliance tree. It can also indicate time discrepancies between your devices (appliances, the Orchestrator Server, and your browser). If the dates and times of devices are not synchronized, charts and report statistics will have varying timestamps for the same data.



Time differences shown for Orchestrator and your browser (client machine) are relative to the respective appliances. Appliance times should be within a minute of Orchestrator and browser times. If exceeded, an appliance alarm is issued.

TIP: To synchronize times, it is recommended that you configure the appliance, the Orchestrator server, and your client machine to use an NTP (Network Time Protocol) server.

To configure a synchronized date and time for an appliance, click the edit icon in the row associated with the appliance. The Date/Time Settings dialog box opens.

Date/Time Settings Dialog Box

Use the Date/Time Settings dialog box to configure a synchronized date and time across your appliances, Orchestrator, and browser (client machine).

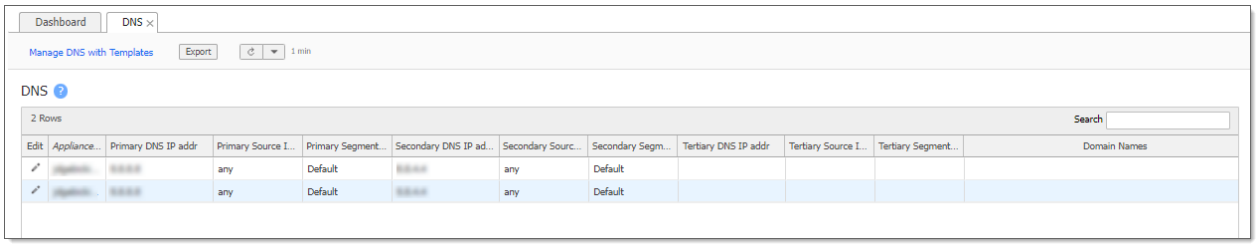
1. Select the appropriate time zone for your network from the **Time zone** drop-down list.
2. Select one of the following time/date configuration options:
   - **Manual** – Set the Date and Time fields manually.
   - **NTP Time Synchronization** – To specify an NTP server for automatic time/date synchronization, click **Add**. A grid row is created. In the **Server** field, type the URL for the NTP server to use. In the **Version** field, select the NTP version to use (version 3 or 4).

     NOTE: It is generally recommended that you specify two or more NTP servers for timekeeping accuracy and reliability.
3. Click **Save**.

DNS (Domain Name Servers) Tab

Administration > General Settings > Setup > DNS

This tab lists the Domain Name Servers that the appliances reference.



A Domain Name Server (DNS) uses a table to map domain names to IP addresses so you can reference locations by a domain name, such as *mycompany.com*, instead of using the IP address.

Each appliance can support up to three name servers.

Field	Description
Appliance	Name of the appliance.
Primary DNS IP addr	IP address of the DNS the system uses first.
Secondary DNS IP addr	IP address of the DNS the system uses second.
Tertiary DNS IP addr	IP address of the DNS the system uses last.

To add the three domain name servers, click the **Edit** icon.

DNS (Domain Name Servers) Edit Row

On this dialog box, you can configure up to three name servers. Enter the three server DNS IP addresses, and then click **Add** to apply the name to the domain.

SNMP

Administration > General Settings > Setup > SNMP

EdgeConnect appliances support Management Information Base (MIB-II) as described in RFC 1213 for cold start traps, warm start traps, and EdgeConnect private MIBs. Appliances issue an SNMP trap during reset when loading a new image, recovering from a crash, or rebooting.

An appliance sends a trap whenever an alarm is raised or cleared. Traps contain additional information about alarms, including severity, sequence number, a text-based description of the alarm, and the time the alarm was created. To download a .zip file of standard and proprietary MIB files you can load into the MIB browser you use, see MIB Files.

The SNMP tab lists Simple Network Management Protocol (SNMP) configuration summaries for appliances selected in the appliance tree.

Field	Description
Appliance	Name of the appliance.
Enable SNMP Agent	Indicates whether the SNMP agent is enabled on the appliance.
Enable SNMP Traps	Indicates whether the SNMP agent on the appliance is enabled to send traps to configured receivers.
Enable SNMP V1/V2	Indicates whether the SNMP agent on the appliance is enabled to send traps to configured receivers.
Enabled V3 Users	Indicates SNMP v3 users enabled on this appliance.
Trap Receiver 1	Configuration details for the first trap receiver if configured. Up to three trap receivers can be configured. Details include the IP address of the host receiving traps, the SNMP version of the receiver, and whether the receiver is enabled.

Field	Description
Trap Receiver 2	Configuration details for the second trap receiver if configured.
Trap Receiver 3	Configuration details for the third trap receiver if configured.

To modify the SNMP configuration for a listed appliance, click the edit icon associated with the appliance. The SNMP dialog box opens. For details, see Modify SNMP Configuration below.

NOTE: You can also use a template to manage SNMP configurations. Click the **Manage SNMP with Templates** link.

Click **Export** to save the SNMP data to a CSV file.

Modify SNMP Configuration

Use the SNMP dialog box to configure the appliance's SNMP agent and trap receivers.

1. To enable SNMP on this appliance, select the **Enable SNMP** check box. Additional areas on the dialog box appear (SNMP V1/V2, SNMP V3, and Trap Receivers).
2. To enable the SNMP agent on the appliance to send traps to configured receivers, select the **Enable SNMP Traps** check box.
3. In the **Default Trap Community** field, specify the string the trap receiver should use to accept traps sent to it. The default value is "public". You can modify this value.
4. Configure SNMP v1/v2, SNMP v3, and trap receivers, as described below.

SNMP V1/V2

Complete the following fields for SNMP v1 and v2c.

Field	Description
Enable SNMP	Enables the SNMP agent on the appliance to send traps to configured receivers.
Read-Only Community	The SNMP application needs to present this text string (secret) to poll the appliance's SNMP agent. The default value is "public". You can modify this value.

SNMP V3

For additional security, configure SNMP v3 if you want to authenticate without using clear text. To add an SNMP v3 user, click **Add** above the SNMP V3 table, and then complete the following fields.

Field	Description
Enabled	Select this check box to enable the selected user. Clear it to disable the user and maintain the configuration.
Username	Enter the username to identify the SNMP v3 user.
Authentication Type	Select the authentication type to use for SNMP requests from the user. **NOTE:** Authentication type is required. SHA-1 is the only supported algorithm.
Authentication Password	Enter a password that the SNMP agent can use to authenticate requests sent by the user. **NOTE:** The password must have at least 20 characters.
Privacy Type	Select the encryption type to use for encrypting requests from the user. **NOTE:** Encryption is required. AES-128 is the only supported algorithm.
Privacy Password	Enter a password (key) to use for encrypting requests sent by the user. **NOTE:** The password must have at least 20 characters.

To remove an SNMP v3 user, click the corresponding delete icon (**X**) in the last column.

Trap Receivers

To configure a trap receiver, click **Add** above the Trap Receivers table, and then complete the following fields.

NOTE: You can configure up to three trap receivers per appliance.

Field	Description
Host	IP address of the host to which to send traps.
Version	Select the SNMP version of the trap receiver.
Community/Username	For v1 and v2c, enter the community string the receiver should use to accept traps. If left blank, the default community string (public) is used. If a different community string is configured on the trap receiver, enter it in this field. For v3, specify the SNMP v3 user sending traps to the receiver.
Enabled	Select this check box to enable the receiver. Clear it to disable the receiver and maintain the configuration.

To remove a receiver, click the corresponding delete icon (**X**) in the last column.

Flow Export Tab

Administration > General Settings > Setup > Flow Export

This tab summarizes how the appliances are configured to export statistical data to NetFlow and IPFIX collectors. The Flow Exporting Enabled setting allows the appliance to export the data to collectors. The appliance exports flows against two virtual interfaces—sp_lan and sp_wan—that accumulate the total of LAN-side and WAN-side traffic, regardless of physical interface.

To open the Flow Export Configuration dialog box, click the **Edit** icon.

Custom Information Elements

The following tables describe the Custom Information Elements.

Data Type: ipv4Address

Custom IE Name and Implementation Description	Semantics	Units	Field Length (bytes)	Enterprise ID
clientIPv4Address TCP: source ipv4 address of SYN initiator is the client. UDP: source ipv4 address of the first packet is the client.	default		4	1
serverIPv4Address TCP: destination ipv4 address of SYN initiator is the client. UDP: destination ipv4 address of the first packet is the client.	default		4	2
connectionInitiator TCP: source ipv4 address of SYN initiator is the connection initiator. UDP: source ipv4 address of the first packet is the connection initiator.	default		4	7

Data Type: unsigned8

Custom IE Name and Implementation Description	Semantics	Units	Field Length (bytes)	Enterprise ID
connectionNumberOfConnections Number of TCP connections (3-way handshake) or UDP sessions established.	totalCounter		1	9
connectionServerResponsesCount Currently 1.	totalCounter		1	10
connectionTransactionCompleteCount Currently 1.	totalCounter		1	21

Data Type: unsigned32

Custom IE Name and Implementation Description	Semantics	Units	Field Length (bytes)	Enterprise ID
connectionServerResponseDelay TCP: Round-trip time between SYN and SYN-ACK. UDP: Round-trip time between first onward and return packet.		microseconds	4	11
connectionNetworkToServerDelay TCP: Round-trip time between SYN and SYN-ACK. UDP: Round-trip time between first onward and return packet. It is also called Server Network Delay (SND).		microseconds	4	12
connectionNetworkToClientDelay TCP: Round trip between SYN-ACK and ACK. UDP: Round-trip time between first response and second request packet. It is also called Client Network Delay (CND).		microseconds	4	13

Custom IE Name and Implementation Description	Semantics	Units	Field Length (bytes)	Enterprise ID
connectionClientPacketRetransmissionCou	totalCounter		4	14
Currently 1.				
connectionClientToServerNetworkDelay		microseconds	4	15
Network Time/Network Delay is known as the round-trip time that is the summation of CND and SND. It is also called Network Delay (ND).				
connectionApplicationDelay		microseconds	4	16
TCP: Round-trip time between SYN and SYN-ACK.				
UDP: Round-trip time between first onward and return packet.				
connectionClientToServerResponseDelay		microseconds	4	17
The round-trip time that is the summation of CND and SND.				
connectionTransactionDuration		microseconds	4	18
The flow displays the time difference between the first and last packet.				
connectionTransactionDurationMin		microseconds	4	19
The flow displays the time difference between the first and last packet.				
connectionTransactionDurationMax		microseconds	4	20
The flow displays the time difference between the first and last packet.				

Data Type: unsigned64

Custom IE Name and Implementation Description	Semantics	Units	Field Length (bytes)	Enterprise ID
connectionServerOctetDeltaCount Server initiated byte count. If flow is lan to wan, Lan-Tx byte counter. If flow is wan to lan Lan-Rx byte counter.	deltaCounter	octets	8	3
connectionServerPacketDeltaCount Server initiated byte count. If flow is lan to wan, Lan-Tx byte counter. If flow is wan to lan Lan-Rx byte counter.	deltaCounter	packets	8	4
connectionClientOctetDeltaCount Server initiated byte count. If flow is lan to wan, Lan-Tx byte counter. If flow is wan to lan Lan-Rx byte counter.	deltaCounter	octets	8	5
connectionClientPacketDeltaCount Server initiated byte count. If flow is lan to wan, Lan-Tx byte counter. If flow is wan to lan Lan-Rx byte counter.	deltaCounter	packets	8	6

Data Type: String

Custom IE Name and Implementation Description	Semantics	Units	Field Length (bytes)	Enterprise ID
applicationHttpHost HTTP destination domain name.	default		variable length	8
applicationCategory Application group.	default		variable length	27
from-zone (Source Zone) name for the flow when ZBF is configured.	default		variable length	22

Custom IE Name and Implementation Description	Semantics	Units	Field Length (bytes)	Enterprise ID
to-zone (Destination zone) name for the flow when ZBF is configured.			variable length	23
tag User-specified readable string/tag that can be specified when the ZBF rule is configured. If “tag” is not specified, an automatic tag will be created and exported. The automatic/default tag is constructed by concatenating <from-zone>_<to-zone>_<rule priority>. For example, “lan-zone_corp-zone_10000”.	default		variable length	24
overlay Overlay name the zone belongs to.	default		variable length	25
direction Direction of the flow: outbound or inbound.	default		variable length	26

Flow Export Edit Row

The following table describes the Flow Export configuration options.

Field	Description
Enable Flow Exporting	Move the toggle to enable or disable flow exporting.
Active Flow Timeout	Amount of time an active flow has been timed out (in minutes).
IPFIX Template Timeout	Resending of templates based on a timeout.
Traffic Type	Check as many of the traffic types as you want. The default is **WAN TX**.
Information Elements	Check **Firewall Zones**, **Application Performance**, or both.

- If you check **Firewall Zones**:
  - Orchestrator generates data based specifically on the zone-based firewalls associated with the specified flow.
  - For example: Host Name, From Zone, To Zone, Tag, Action, Direction, and so forth.

- If you check **Application Performance**:
  - Orchestrator generates data based specifically on the application performance associated with each flow.
  - For example: clientIPv4Address, serverIPv4Address, connectionInitiator, applicationHttpHost, and so forth.
  - These interfaces appear in SNMP and are, therefore, “discoverable” by NetFlow and IPFIX collectors.
  - The **Collector’s IP Address** is the IP address of the device to which you are exporting the NetFlow/IPFIX statistics. The default Collector Port is **2055**.
- For more information about IPFIX and the associated Custom Information Elements (IEs), see Custom Information Elements.

Logging Tab

Administration > General Settings > Setup > Logging

The Logging tab summarizes the following configured logging parameters:

- *Log Settings* refers to local logging.
- *Log Facilities Configuration* refers to remote logging.

The logs keep track of alarms, events, and any other issues involving your appliances.

The following table provides more details.

Field	Description
Appliance	Name of the appliance associated with the recorded logs.
Minimum Severity	Minimum severity level the issue is recorded as. For descriptions of levels, see Severity Levels.
Log File Size Threshold	Set threshold configured for the log size limit.
Number of Logs to Keep	Maximum number of logs to keep for the appliance.
System	Assigned log facility for System.
Audit	Assigned log facility for Audit.
Firewall	Assigned log facility for Firewall.
Ids	Assigned log facility for IDS.
Log Stateful WAN Drops	Enable log information for discarded inbound packets, even at high traffic rates, for WAN-side interfaces running in stateful, stateful+SNAT, or hardened modes. Drops are logged to the firewall log, with the description of `Inbound drop on stateful wan interface`.

Field	Description
Anonymize IPs	True or false. Indicates if IP addresses are anonymized in log messages or not.
Bit Masking	If Anonymize IPs is enabled, this indicates how bit masking is applied to IP addresses in log messages (options: **Mask All, /8, /16, or /24**).
Jsonify	True or false. Indicates that firewall policy and IDS/IPS log messages are converted to JSON format when exported.
Remote Receiver	IP address of the remote receiver applicable to the log file.
Remote Receiver Minimum Severity	Lowest level of severity logged for the remote log receiver. For details about severity levels, see the "Severity Levels" section below this table.
Facility	Log facility used for the remote log receiver.

To edit the logging configuration for one of the listed appliances, click the edit icon in the left column of the table. The Logging dialog box opens. For details, see Logging Dialog Box

Severity Levels

In order of decreasing severity, the levels are as follows:

Severity Level	Description
Emergency	System is unusable.
Alert	Includes all alarms the appliance generates: **CRITICAL**, **MAJOR**, **MINOR**, and **WARNING**.
Critical	Critical event.
Error	An error. This is a non-urgent failure.
Warning	A warning condition. Indicates an error will occur if action is not taken.
Notice	A normal, but significant, condition. No immediate action required.
Info	Informational. Used by Support for debugging.
Debug	Used by Support for debugging.
None	This indicates that no events are logged.

These are related to event logging levels, not alarm severities, even though some naming conventions overlap. Events and alarms have different sources. Alarms, when they clear, list as the ALERT level in the Event Log.

Remote Logging

- You can configure the appliance to forward all events, at and above a specified severity, to a remote syslog server.

- A syslog server is independently configured for the minimum severity level that it will accept. Without reconfiguring, it might not accept as low a severity level as you are forwarding to it.
- Each message/event type (System / Audit / Firewall / Ids) is assigned to a syslog facility level (local0 to local7).

Logging Dialog Box

Use this dialog box to configure log settings and log facilities. You can also add remote log receivers.

WARNING: Appliance logging levels should only be set to "Notice" unless TAC asks you to set it differently. This applies to both the Minimum severity level field in the Log Settings area of this dialog box and the Minimum Severity field in the Remote Log Receivers area. Be aware that setting this level to "Debug" will generate logs for all modules that are turned on, which causes the packet processing engine to spend excessive time logging instead of forwarding packets.

Log Settings

Setting	Description
Minimum severity level	Minimum severity level that the system will log. (See the WARNING note above.) For details about severity levels, see Severity Levels.
Start new file when log reaches	Enter the maximum size (in MB) for a log file. Orchestrator generates a new file when this maximum size is reached. Specify a size from 1 to 50.
Keep at most log files	Maximum number of log files to allow to be stored. Specify a value from 1 to 100.
Log stateful wan-interface drops	Select to log information for discarded inbound packets, even at high-traffic rates. **NOTE:** Enabling this option may impact system performance.
Anonymize IPs	Click the check box to anonymize IP addresses in log messages.
Bit Masking	If Anonymize IPs is enabled, select how bit masking is applied to IP addresses in log messages (options: **Mask All, /8, /16, or /24**).

Setting	Description
Jsonify	Click the check box to convert firewall policy and IDS/IPS log messages to JSON format when exported. **NOTE:** When you click the Anonymize IPs check box, the Jsonify check box is automatically selected.

Log Facilities Configuration

Select the log facilities you want the System, Audit, Firewall, and IDS/IPS Events logs to use. You can choose between Local0 and Local7 for each.

NOTE: The log facilities you select for System, Audit, Firewall, and IDS/IPS Events must be uniquely assigned; they cannot overlap. For example, System can be assigned to local2 and Audit to local3, but both cannot be assigned to local2.

Remote Log Receivers

Follow these instructions to add a remote receiver for an appliance syslog server that uses an end entity certificate.

NOTE: To use an end entity certificate, you must first create an end entity certificate for use. To do this, see End Entity Certificates Tab.

1. Navigate to **Administration > General Settings > Setup > Logging**.
2. Click the edit icon next to the appliance for which you want to configure a receiver.

   The Logging dialog box opens.
3. Under Remote Log Receivers, click **Add** and then configure the following information.

Field	Description
IP Address	Enter the IP address for the remote receiver.
Port	Enter the port number of the remote syslog server. The default for TCP SSL is 6514.
Protocol	Select **TCP SSL**.
Minimum Severity	Select the minimum severity level of messages you want to log. (See the WARNING note above.) For details about severity levels, see Severity Levels. For Common Criteria mode, **Debug** should be used to ensure all logs are sent to the syslog receiver.
Facility	Select **all**, **local1**, **local2**, **local3**, **local4**, **local5**, **local6**, or **local7**.

4. In the Client Certificate column, click **Add**.

   The Add Remote Receiver SSL Certificate dialog box opens.

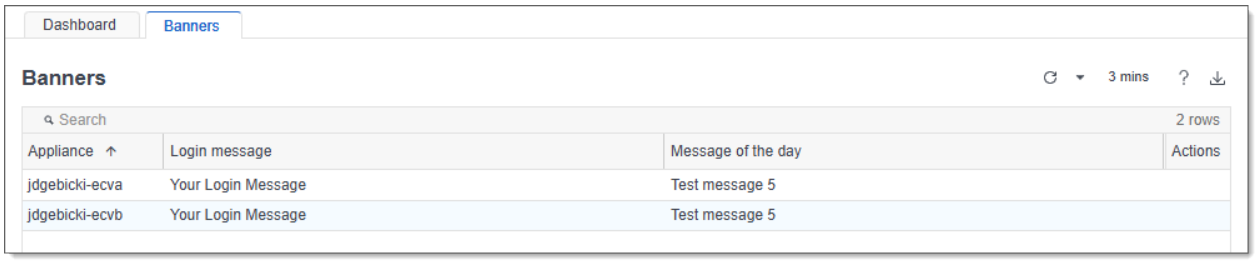
5. Click **Use End Entity Certificate** and then select the end entity certificate from the End Entity Certificate drop-down menu.
6. Click the cell in the Verify column to display a check box, and then click the check box to verify the server certificate.
7. Click **Add**.
8. Click **Save**.

For information about remote log receivers, including how to add and configure a receiver, see Remote Log Receivers.

Banners Tab

Administration > General Settings > Setup > Banners

The Banners tab lists the banner messages on each appliance.



Each appliance can have two banner messages:

- The *Login Message* appears before the login prompt.
- The *Message of the Day* appears after a successful login.

To enter your banner messages, hover over the Actions column and click the edit icon.

Banners Edit Row

Enter your messages in the boxes, and then click **Save**.

HTTPS Certificate Tab

Administration > General Settings > Setup > HTTPS Certificate

On this tab, you can view the HTTPS server certificate for each appliance. To edit, add, or assign an end entity certificate for a specific appliance, click the edit icon next to the appliance for which you want to add a certificate.

HTTPS Certificate Dialog Box

On this dialog box you select the type of certificate to use with the appliance. By default, EdgeConnect appliances present a self-signed server certificate to any client opening a TLS connection to the appliance web UI. To ensure secure communications, TLS clients will cryptographically verify that a trusted Certificate Authority (CA) issued the EdgeConnect certificate. If you use the default option, **Self-Signed Certificate**, browsers will show this as not secure; most enterprise IT departments will not allow this. Enterprises must set up an HTTPS server certificate for their EdgeConnect appliances if they intend to use the EdgeConnect web UI directly. However, it is highly recommended to perform all configuration through Orchestrator.

There are three ways to set up an HTTPS server certificate for EdgeConnect appliances.

- **Use an EST server and globally orchestrated end entity profiles to automate certificate enrollment.** This is the recommended option. For more information about this method, see End Entity Certificates.

  NOTE: Configuration for this method is not done on this tab.

- **Manually create a Certificate Signing Request (CSR) in Orchestrator.** As part of this process, Orchestrator creates the public key private key pair. The user downloads and submits the CSR for signing by a Certificate Authority (CA). The signed certificate is then uploaded in Orchestrator for use in one of several applications. The end entity certificate contains a label, which is significant to Orchestrator and allows this certificate to be used by referring to its label. You must repeat this process for each EdgeConnect appliance.

  To use an end entity certificate obtained by manually creating a CSR in Orchestrator:

  NOTE: To use an end entity certificate, you must first create an end entity certificate for use. To do this, see End Entity Certificates Tab.

  NOTE: This must be performed one appliance at a time.

  1. Navigate to **Administration > General Settings > Setup > HTTPS Certificate**.
  2. Click the edit icon next to the appliance for which you want to add a certificate.
  3. Click **End Entity Certificate** and then select the end entity certificate from the drop-down menu.
  4. Click **Save**.

- **Use a Custom Certificate.** This requires everything to be done externally including creating the public key private key pair and creating the CSR. This legacy method is not recommended.

  To use a Custom Certificate (legacy method):

  1. Consult with your IT security team to generate a certificate signing request (CSR), and then submit it to your organization's chosen SSL Certificate Authority (CA).
     - Examples of Certificate Authorities include GoDaddy, Verisign, Comodo, Symantec, Microsoft Entrust, GeoTrust, and so forth.
     - All certificate and key files must be in PEM format.

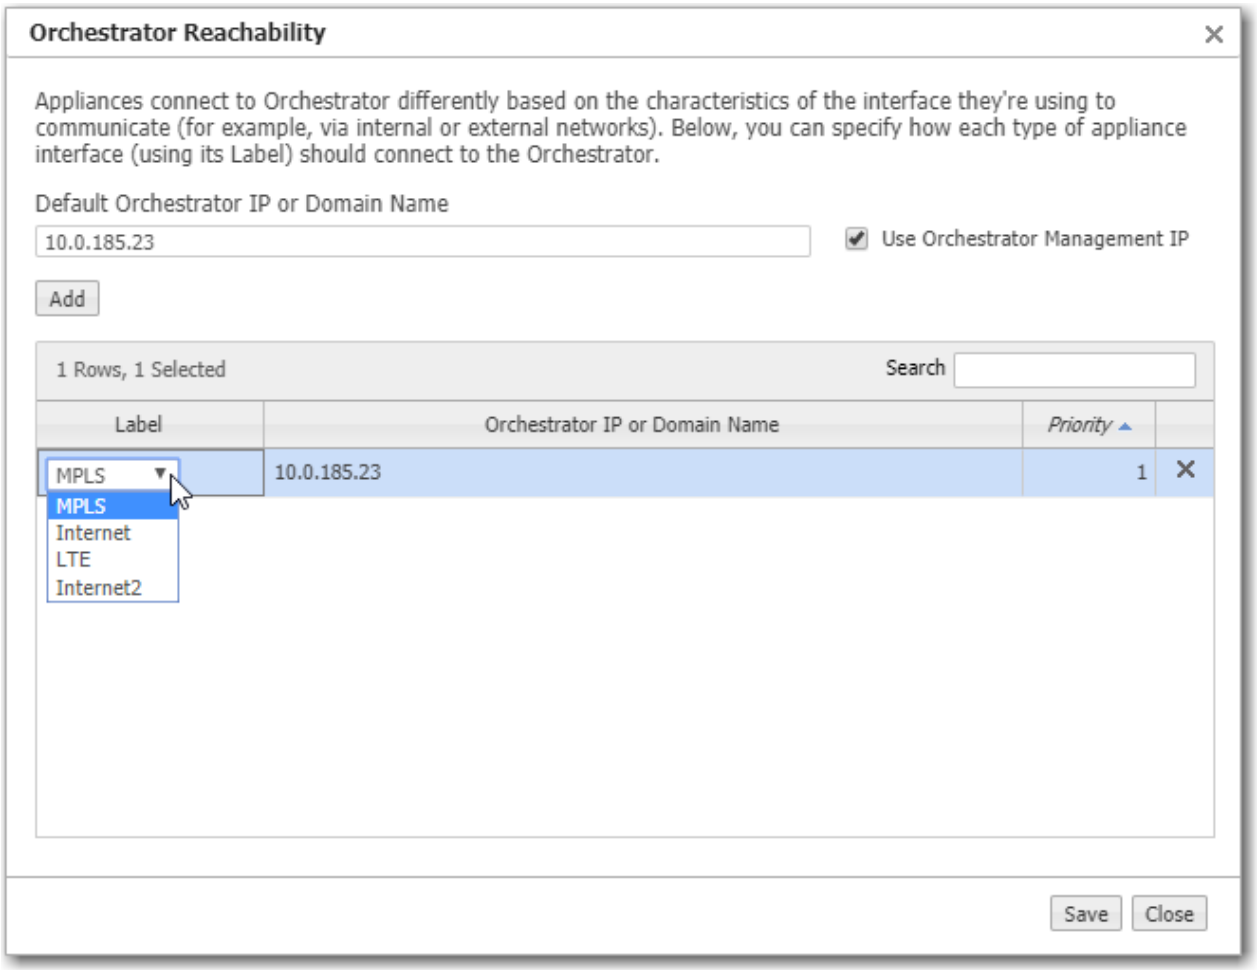
2. After the Certificate Authority provides a CA-verified certificate, navigate to **Administration > General Settings > Setup > HTTPS Certificate**.
3. Click the edit icon next to the appliance for which you want to add a certificate.
4. Click **Custom Certificate**, and then click **Upload and Replace**.

   The Add HTTPS Certificate dialog box appears.
5. If your IT security team advises the use of an Intermediate CA, upload an Intermediate Certificate File. Otherwise, skip this file.
6. Upload the Certificate File from the CA.
7. Upload the Private Key File that was generated as part of the CSR.
8. Click **Add** to close the Add HTTPS Certificate dialog box.
9. Click **Save**.

Orchestrator Reachabililty Tab

Administration > General Settings > Setup > Orchestrator Reachability

You can specify how each appliance connects to Orchestrator by designating one of its interface Labels.



Custom Appliance Tags

Administration > General Settings > Setup > Custom Appliance Tags

Use the Custom Appliance Tags tab to create and assign administrative information to your appliances. You can store up to eight key/value pairs for each appliance.

Service providers, backend IT systems, and third-party partners can use these tags to store account information, contact phone numbers, license information, and so forth. For example, the first tag for an appliance could have the key “customer” and “ABC Company” as its value. The second tag could have the key “accountid” and the value “19283270”.

Custom appliance tags can be retrieved from the Orchestrator REST API using /customApplianceTags/getCustomizableBindings.

NOTE: Orchestrator does not use these tags for any other purpose, such as for filtering or reporting data.

The tab lists eight tag rows for each appliance selected in the appliance tree. To create custom tags for an appliance listed on the tab:

1. Click the edit icon for one of the tag rows associated with the appliance.

   The Custom Appliance Tags dialog box opens. The table lists eight tag rows you can use for the appliance.

2. For each tag you want to create for the appliance, type the name for the key in the **Key** field, and then type the value you want to assign to the key in the **Value** field.
3. When you finish specifying tags for the appliance, click **Save**.

Administration > Software

The options under **Administration > Software** focus on software-related tasks such as managing system information with templates, upgrading appliance software, configuring and restoring backups, and removing appliances from Orchestrator.

System Information

Administration > Software > Upgrade > System Information

You can manage system information with templates except for Deployment Mode, which is an appliance-specific configuration. To change the Deployment Mode for an appliance, navigate to **Configuration > Networking > Deployment**.

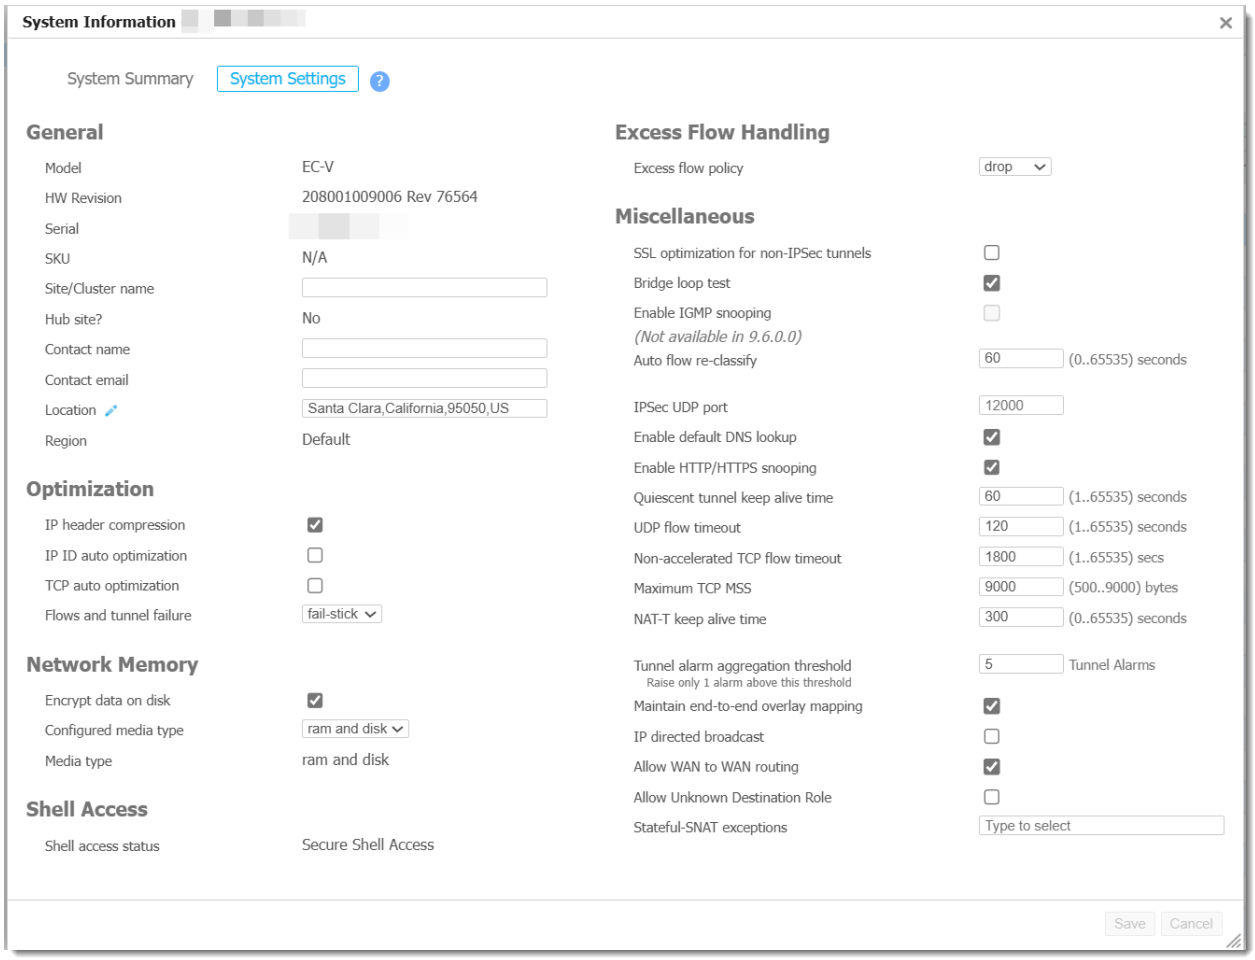
When you click the edit icon next to a specific appliance, the System Information dialog box opens. The System Information dialog box features two subtabs: System Summary and System Settings, which are described in the following information.

System Summary

Property Key	Description
Appliance Key	Orchestrator assigns and uses this key to identify the appliance.
Platform	The underlying cloud platform on which the EdgeConnect appliance runs, such as Amazon EC2, Azure, Google Cloud, or VMware.
Uptime	The time elapsed since the appliance became operational and available.
Active Release	Specifies the software release the appliance is running.
Appliance ID	The unique identifier for the appliance.

Property Key	Description
Discovery Method	Specifies how Orchestrator discovered the appliance: **PORTAL**: Orchestrator discovered the appliance through the portal account. **MANUAL**: The appliance was added manually. **APPLIANCE**: The Orchestrator IP address was added to the appliance. Portal was not involved.
Connection Type	The method that Orchestrator uses to communicate with the appliance. Options are WEBSOCKET, PORTAL, and HTTP.
Appliance Model	Specifies if the appliance is an EC, EC-V, NX, VX, or VRX model.
HW Revision	The hardware number and revision number of the appliance (for example, 208001009006 Rev 76564).
BIOS Version	The version of BIOS firmware that the appliance is using.
Serial Number	The serial number of the appliance.
SKU	The stock keeping unit (SKU) identifier for the appliance.
System Bandwidth	The total outbound bandwidth, determined by appliance model or license.
Mode	Specifies the deployment mode of the appliance: Server, Router, or Bridge.

System Settings



Property Key	Description
Model	Specifies if the appliance is an EC, EC-V, NX, VX, or VRX model.
HW Revision	The hardware number and revision number of the appliance. For example, 208001009006 Rev 76564.
Serial	The serial number of the appliance.
SKU	The stock keeping unit (SKU) identifier for the appliance.
Site/Cluster name	Orchestrator will not build tunnels between appliances with the same user-assigned Site/Cluster name. For more information about Site/Cluster names, see Clusters.
Hub site?	Specifies if the appliance has been assigned the role of Hub in Orchestrator.
Contact name	The name of the person to contact within your organization (optional).

Property Key	Description
Contact email	The email address of the person to contact within your organization (optional).
Location	The physical location of the appliance, optionally specified during appliance setup.
Region	A user-assigned name that is created for segmenting topologies and streamlining the number of tunnels created. When regions contain at least one hub, you can choose to connect regions through hubs only.
IP header compression	IP header compression reduces the size of IP packet headers in order to optimize bandwidth usage, improve data transmission efficiency, and reduce latency in network communication. It is particularly useful in environments with bandwidth constraints, such as wireless networks, low-speed links, and mobile networks.
IP ID auto optimization	Enables any IP flow to automatically identify the outbound tunnel and gain optimization benefits. Enabling this option reduces the number of required static routing rules (route map policies).
TCP auto optimization	Enables any TCP flow to automatically identify the outbound tunnel and gain optimization benefits. Enabling this option reduces the number of required static routing rules (route map policies).
Flows and tunnel failure	If there are parallel tunnels and one fails, Dynamic Path Control determines where to send the flows. There are three options: **fail-stick**: When the failed tunnel comes back up, the flows do not return to the original tunnel. They stay where they are. **fail-back**: When the failed tunnel comes back up, the flows return to the original tunnel. **disable**: When the original tunnel fails, the flows are not routed to another tunnel.
Encrypt data on disk	Enables encryption of all the cached data on the disks. Disabling this option is not recommended.
Configured media type	There are two options, **ram and disk** (VX) or **ram only** (VRX). This can be changed for special circumstances if support recommends it.
Media type	Specifies the actual type of media being used.

Property Key	Description
Shell access status	Specifies the current shell access policy for EdgeConnect appliances. **Open Shell Access**: Full access is granted to the underlying Linux operating system shell. **Secure Shell Access**: Access is denied to the shell, but Support can grant access. Contact Support for assistance. You cannot change this setting to Open Shell Access. **Disabled Shell Access**: Access permanently denied to the shell. You cannot change this setting to Open Shell Access or Secure Shell Access. This setting is managed on the Advanced Security Settings page (**Configuration > Overlays & Security > Security > Advanced Security Settings**). Changes to this setting affect all appliances in your network.
Excess flow policy	Specifies what happens to flows when the appliance reaches its maximum capacity for optimizing flows. The default is to **bypass** flows. Or you can choose to **drop** the packets.
SSL optimization for non-IPSec tunnels	Specifies whether the appliance should perform SSL optimization when the outbound tunnel for SSL packets is not encrypted. For example, a GRE or UDP tunnel. To enable Network Memory for encrypted SSL-based applications, you must provision server certificates in Orchestrator. This activity can apply to the entire distributed network of EdgeConnect appliances or to a specified group of appliances.
Bridge loop test	This is only valid for virtual appliances. When enabled, the appliance can detect bridge loops. If it detects a loop, the appliance stops forwarding traffic and raises an alarm. Appliance alarms include recommended actions.
Enable IGMP snooping	IGMP snooping is a common Layer 2 LAN optimization that filters the transmit of multicast frames only to ports where multicast streams have been detected. Disabling this feature floods multicast packets to all ports. IGMP snooping is recommended and enabled by default.
Auto flow re-classify	Specifies how often to do a policy lookup. The default is 60 seconds. Shortening this interval can increase the workload on an EdgeConnect to complete tasks such as VRRP convergence.

Property Key	Description
IPSec UDP port	Specifies the port that Orchestrator uses to build IPSec UDP tunnels. If the field is blank, Orchestrator uses the default.
Enable/disable default DNS lookup	Cloud portal name lookup through default DNS (8.8.8.8) can be enabled or disabled through the System Information tab for appliances.
Enable HTTP/HTTPS snooping	Enables a more granular application classification of HTTP/HTTPS traffic by inspecting the HTTP/HTTPS header, specifically the Host field. Enable is the default setting.
Quiescent tunnel keep alive time	Specifies the rate at which to send keep alive packets after a tunnel has become idle (quiescent mode). The default is 60 seconds.
UDP flow timeout	Specifies how long to keep the UDP session open after traffic stops flowing. The default is 120 seconds (2 minutes).
Non-accelerated TCP flow timeout	Specifies how long to keep the TCP session open after traffic stops flowing. The default is 1800 seconds (30 minutes).
Maximum TCP MSS	Maximum Segment Size. The default value is 1328 bytes. This ensures that packets are not dropped for being too large. You can adjust the value (500 to 9000) to lower the MSS of of a packet, if your environment requires a lower size.
NAT-T keep alive time	If a device is behind a NAT, this specifies the rate at which to send keep alive packets between hosts to keep the mappings in the NAT device intact.
Tunnel alarm aggregation threshold	Specifies the number of alarms to allow before alerting the tunnel alarm.
Maintain end-to-end overlay mapping	Enforces the same overlay to be used end-to-end when traffic is forwarded on multiple nodes.
IP directed broadcast	Allows an entire network to receive data that only the target subnet initially receives.
Allow WAN to WAN routing	Redirects inbound WAN traffic back to the WAN.
Allow Unknown Destination Role	Indicates whether to allow unknown destination roles.
Stateful-SNAT exceptions	The name of the address group configured for Stateful-SNAT exceptions (for example, Stateful-SNAT-Exceptions). To set up this address group, see Disable Stateful+SNAT Processing for Selected LAN-side Subnets below.

Disable Stateful+SNAT Processing for Selected LAN-side Subnets

Most internet providers require that flows originate from the WAN-side IP address assigned to the appliance. When Stateful+SNAT is configured on a WAN-side interface, all traffic that leaves the interface will be Source NATed to the IP address of the WAN-side interface.

In certain situations, you want the original LAN-side IP address to be seen by the upstream network. You can use the Stateful-SNAT exceptions feature to avoid Source NATing for specific IP addresses or subnets.

Considerations:

- Stateful-SNAT exceptions apply only to appliances with firewall mode set to “Stateful+SNAT”.
- Exceptions apply only to outbound flows destined to external addresses.
- Inbound flows initiated from the WAN side toward IP addresses within the address group rely on existing inbound port-forwarding functionality.
- SNAT exceptions apply to the default segment only, not VRF SNAT.
- This feature does not support IPv6 because the address groups feature does not support IPv6.
- The Stateful-SNAT-exceptions feature works for directly connected WAN-side labels. Labels configured on the other side of an EdgeHA link are not supported.

You can use the System template to set up Stateful-SNAT exceptions for all appliances or the System Information dialog box for individual appliances. To set up exceptions for all appliances, see System Template.

To set up Stateful-SNAT exceptions for individual appliances:

1. Create an address group for all public IP space (subnets) used by your network across all branches, as follows:
    1. Navigate to **Configuration > Templates & Policies > ACLs > Address Groups**.

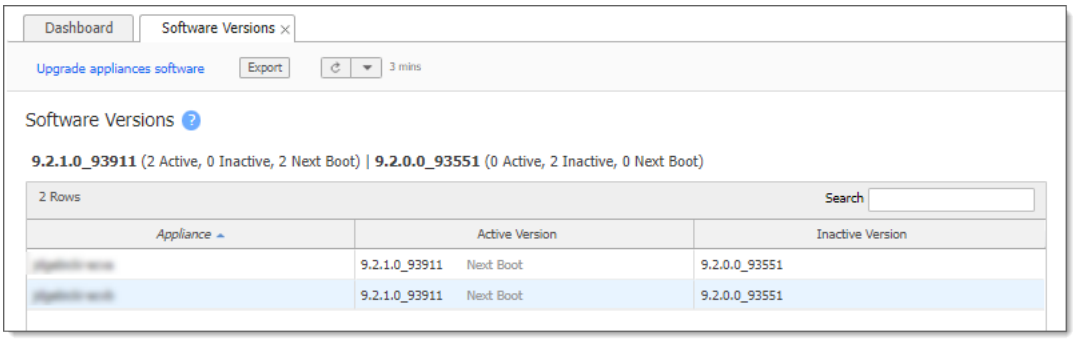
       The Address Groups tab opens.
    2. Click **Add Group**.

       The Add Address Group dialog box opens.
    3. In the **Group name** field, enter an appropriate name for the Stateful-SNAT exceptions (for example, Stateful-SNAT-Exceptions).
    4. In the **IPs to include** and **IPs to exclude** fields, enter IP addresses/masks to include/exclude individually or IP prefixes to include/exclude multiple addresses at once, as appropriate. Use commas to separate entries.
    5. If desired, use the **Comment** field to state the purpose of this address group.
    6. Click **Add**.
2. In the **Stateful-SNAT Exceptions** field on the System Settings page of the System Information dialog box, enter the name of the address group you created for Stateful-SNAT exceptions, and then click **Save**.

Software Versions

Administration > Software > Upgrade > Software Versions

This tab lists the software versions on each appliance.



Upgrade Appliance Software

Administration > Software > Upgrade > Upgrade Appliances

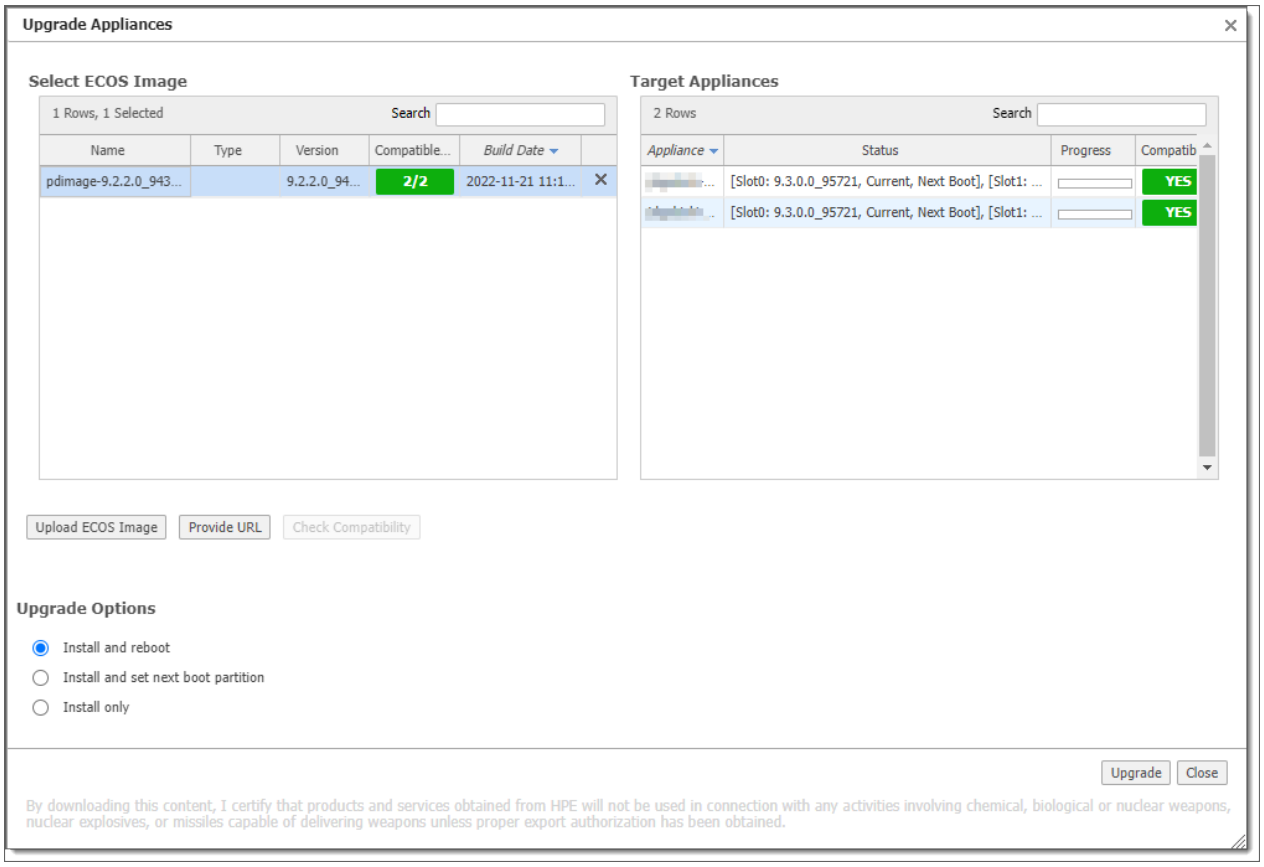
You can download and install new appliance software from your network or computer, or you can download and store (without installing) new appliance software to the Orchestrator server.

NOTE: Before you upgrade the appliance software, navigate to **Administration > Tools > Monitoring > Active Sessions** to verify that no users are logged on.

To upgrade your appliances:

1. Select one or more appliances in the appliance tree, and then navigate to **Administration > Software > Upgrade > Upgrade Appliances**.

   The Upgrade Appliances dialog box opens. The Target Appliances table lists the appliances you selected.



2. Do one of the following:

   - To upload an image file you downloaded to your machine, click **Upload ECOS Image**. File Explorer opens. Navigate to the image file, and then click **Open**. When the upload finishes, the image file appears in the Select ECOS Image table.
   - To provide a URL for the image file, click **Provide URL**. Enter the URL for the image file in the **URL for image file** field. When you click Provide URL, the Check Compatibility button becomes available.

3. Depending on the option you selected in step 2, do one of the following:

   - If you uploaded an ECOS image file, select the new image file in the Select ECOS Image table. The Compatible column in the Target Appliances table indicates compatibility for each appliance.
   - If you provided a URL, click **Check Compatibility** to ensure that the image file is compatible with your current appliance software. The Compatible column in the Target Appliances table indicates compatibility for each appliance.

   NOTE: If the image file is not compatible, do not proceed with the upgrade. Find and select an image file that is compatible with the current versions of the appliances. Refer to the Orchestrator Release Notes for more information.

4. Select one of the following install options:

- **Install and reboot**—Installs the image file into the appliance's inactive partition and reboots the appliance to use the new software.
- **Install and set next boot partition**—Installs the image file into the appliance's inactive partition and points to that partition for the next reboot.
- **Install only**—Downloads the image file into the inactive partition. Use this option to store new appliance software from your network or computer to the Orchestrator server.

5. Click **Upgrade**.

   The upgrade process downloads the image file and then installs it onto the appliances. Monitor the progress of the installation in the Target Appliances table. The Status column will display "Upload in progress," "Installing boot disk image file," "Waiting for reboot to finish," and then "Success."

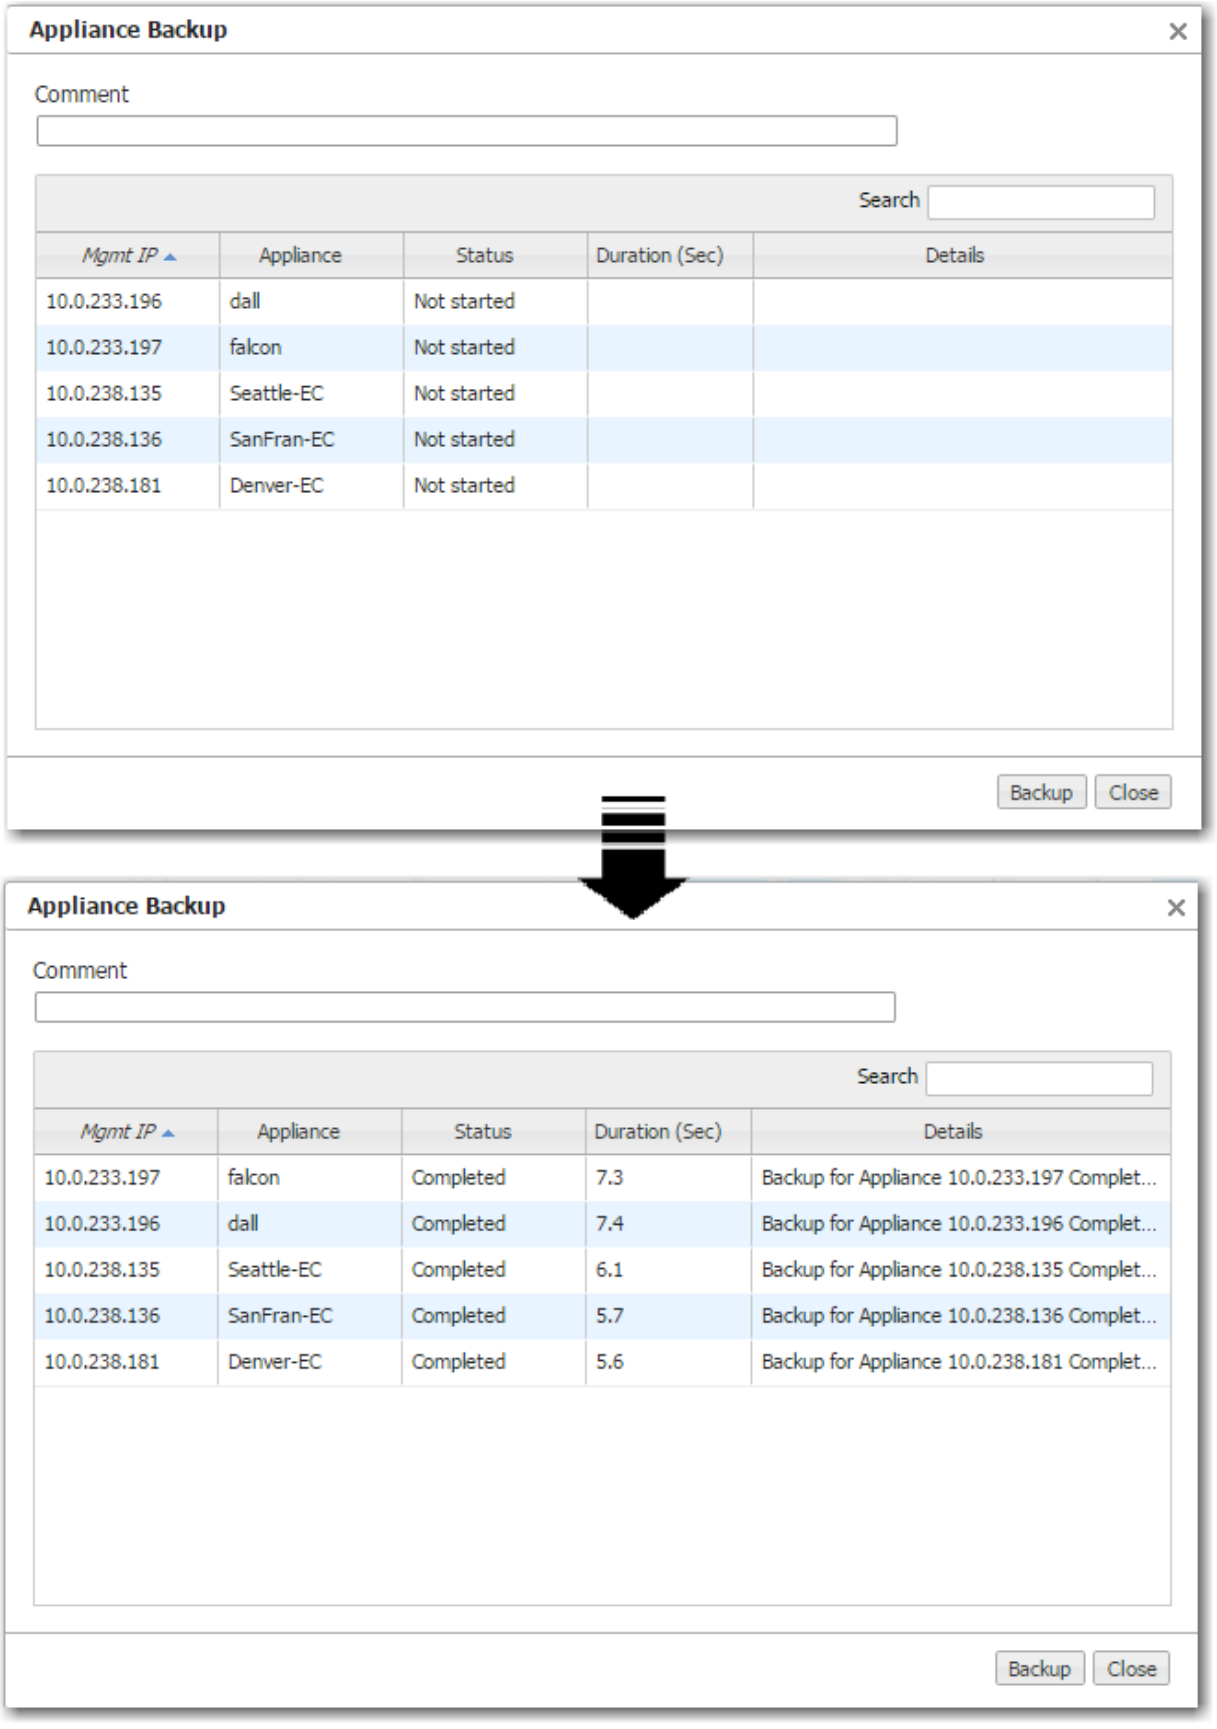
6. Click **Close**.
7. If the upgrade was successful, monitor the appliance tree to verify that each appliance reboots. This might take some time.

Appliance Configuration Backup

Administration> Software > Backup & Restore > Backup Now

Orchestrator automatically creates a weekly backup of each appliance's configuration to the Orchestrator server. Additionally, you can create an immediate backup on demand.

After selecting the appliance(s) in the appliance tree, navigate to **Administration > Software > Backup & Restore > Backup Now**, and then click **Backup**.

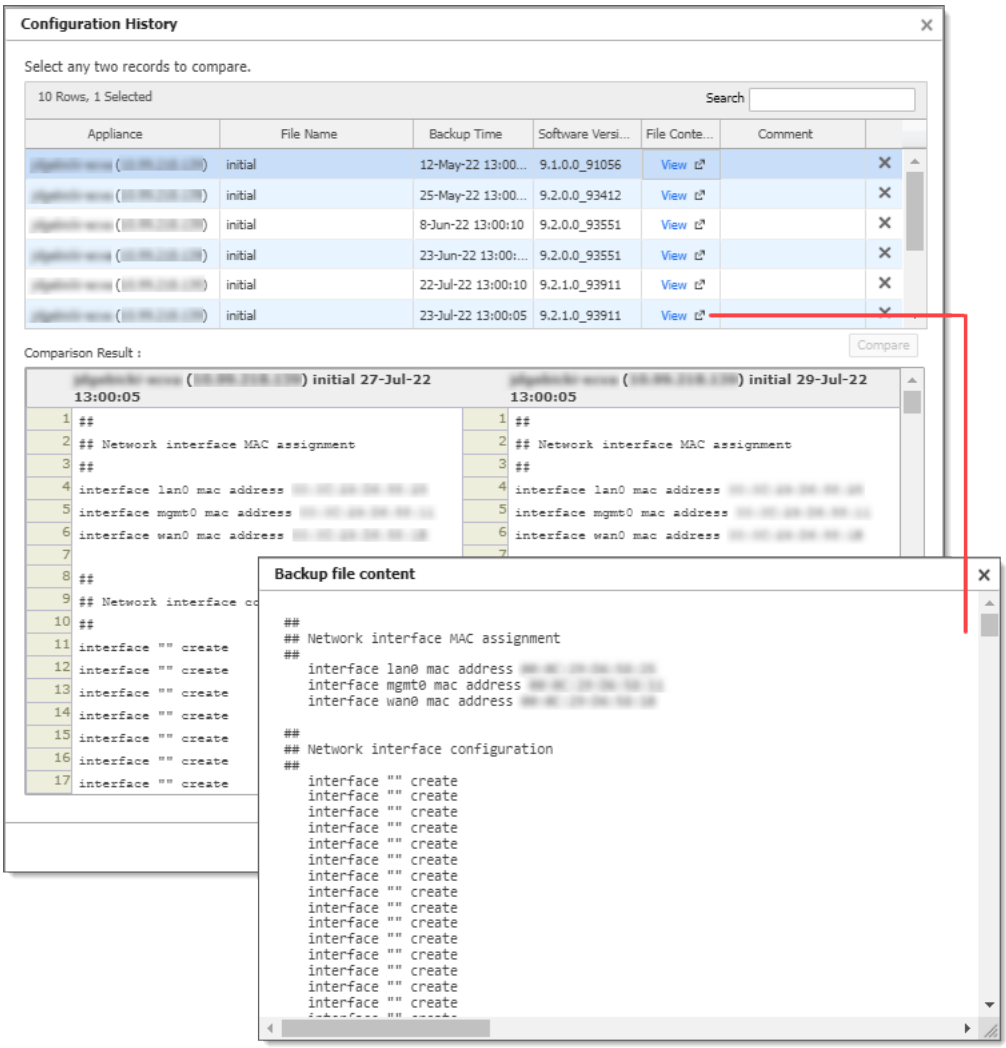

Appliance Backup
Comment
Search
Mgmt IP
Appliance
Status
Duration (Sec)
Details
10.0.233.196 dall Not started
10.0.233.197 falcon Not started
10.0.238.135 Seattle-EC Not started
10.0.238.136 SanFran-EC Not started
10.0.238.181 Denver-EC Not started
Backup
Close
Appliance Backup
Comment
Search
Mgmt IP
Appliance
Status
Duration (Sec)
Details
10.0.233.197 falcon Completed 7.3 Backup for Appliance 10.0.233.197 Complet...
10.0.233.196 dall Completed 7.4 Backup for Appliance 10.0.233.196 Complet...
10.0.238.135 Seattle-EC Completed 6.1 Backup for Appliance 10.0.238.135 Complet...
10.0.238.136 SanFran-EC Completed 5.7 Backup for Appliance 10.0.238.136 Complet...
10.0.238.181 Denver-EC Completed 5.6 Backup for Appliance 10.0.238.181 Complet...
Backup
Close

NOTE: You cannot delete an appliance backup from Orchestrator.

View Configuration History

Administration> Software > Backup & Restore > Configuration History

From the Configuration History tab, you can view an appliance's current or previous configuration, as well as compare any two appliance configuration files.



Appliance Restore and Reboot

Administration > Software > Backup & Restore > Restore

On the Appliance Restore and Reboot tab, you can restore an appliance configuration backup from Orchestrator to one or more EdgeConnect appliances in your network simultaneously. You can also view all backup information for each appliance, as described in the following table.

- To view information for a specific appliance, select it in the appliance tree. The most recent backup for the appliance is shown.
- To view previous backups for an appliance, click the arrow to expand the row for that appliance.
- To view previous backups for all appliances, click **Expand all**.

Column	Description
Appliance	Unique identifier for the appliance.
Appliance name	Name of the appliance.
Backup time	Date and time of the most recent Orchestrator backup of the configuration for this appliance.
File name	The name of the current backup file. The format of the file name is: `{ECOS version}_{appliance hostname}_{backup date in mm_dd_yy format}` Example: 9.6.0_MyAppliance-ECVA_06_01_26
Software version	Version of ECOS software that was running on the appliance at the time the backup occurred.
Comment	Comments entered by a user when they performed a backup of the appliance.
Recent restore time	Date and time the appliance was restored to a previous backup configuration.
Status	When an appliance is being restored to a previous backup configuration, this column displays real-time updates to show the current stage of the restore operation. When the restore operation is complete, it shows *Restored*.

Column	Description
Progress	When an appliance is being restored to a previous backup configuration, this column shows a progress bar that displays real-time updates throughout the restore operation.
Actions	Hover to view available actions for this row (View backup file and Download backup file).

Restore an Appliance Configuration Backup

This section describes how to restore an appliance configuration for one or more appliances. If you restore an appliance from a backup file that is older than the most recent backup file, you could experience:

- A loss of configuration changes that were made between the time the older backup file was created and the most recent backup. Any changes made during that interval are not included in the older backup file and may be lost if an appliance is restored from that backup.
- A loss of appliances that were added between the time the older backup was created and the most recent backup.
- IPSec UDP tunnels may experience a brief outage as the new seed is applied.

CAUTION: Restoring a backup that specifies a static mgmt0 IP address, as opposed to DHCP, can cause an IP conflict if the network configuration has changed since the backup was created. Before restoring, verify that the static IP address in the backup does not conflict with any current network assignments.

1. Navigate to **Administration > Software > Backup & Restore > Restore**.
2. From the appliance tree, select one or more appliances. By default, the check box for the most recent backup is selected. To restore an appliance to the most recent backup, leave it selected.
3. To restore an appliance to an older backup, click the arrow to expand the row for that appliance. Clear the check box for the most recent backup, and then select the check box for the backup to restore.
4. Click **Restore**.

   The appliance receives the configuration from the backup file and reboots. The Progress column indicates the progress of the restore operation.

   NOTE: For EdgeHA appliance pairs, the appliances are restored consecutively. The second appliance begins restoring only after the first appliance has successfully rebooted.

View or Copy an Appliance Configuration Backup File

To view or copy the contents of an appliance configuration backup file:

1. Navigate to **Administration > Software > Backup & Restore > Restore**.
2. From the appliance tree, select one or more appliances.
3. Click the arrow to expand the information for the appliance, if needed.
4. Find the backup file you want to view, hover over the Actions column and click **View backup file**.

   The Backup File Content dialog box opens.
5. To copy the content of the backup file to your clipboard, click the copy button in the upper-right corner of the dialog box.
6. Click the **X** to close the dialog box.

Download an Appliance Configuration Backup File

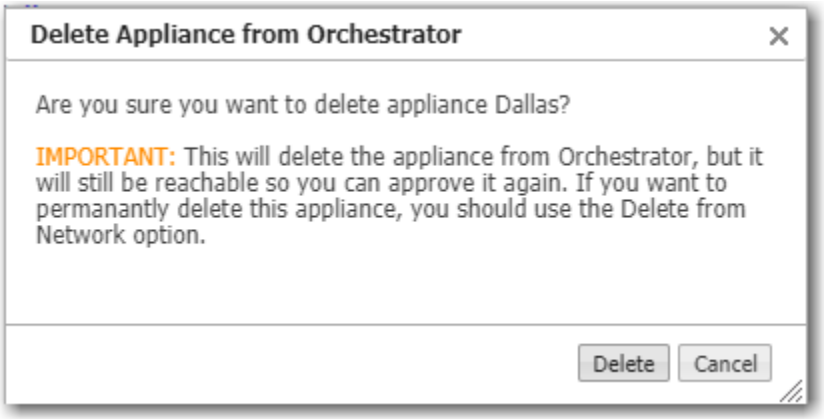
1. Navigate to **Administration > Software > Backup & Restore > Restore**.
2. From the appliance tree, select one or more appliances.
3. Click the arrow to expand the information for the appliance, if needed.
4. Find the backup file you want to download, hover over the Actions column and click **Download backup file**.

   The backup file is saved to your local machine.

Remove Appliance from Orchestrator

Administration > Software > Remove Appliances > Remove from Orchestrator

Removing an appliance with this action returns the appliance to the Discovered Appliances list.

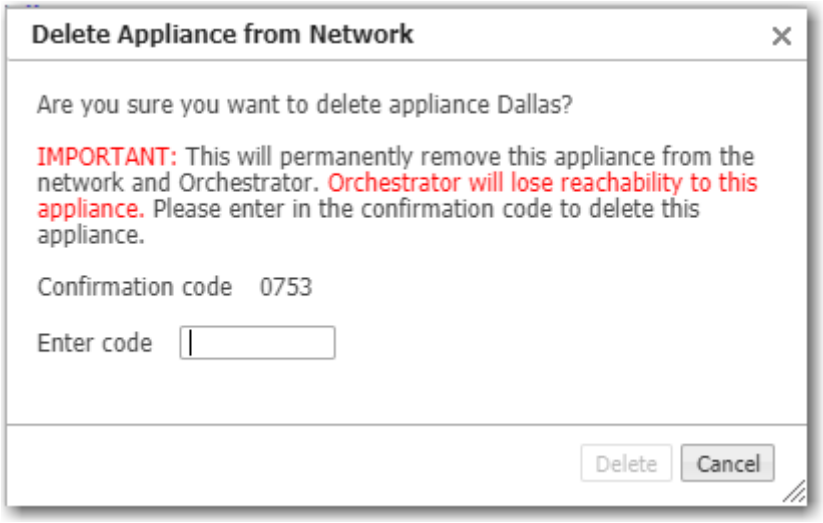


This action deletes the appliance from the navigation tree. In addition, Orchestrator will break all tunnels, overlays, and so forth to this device.

Remove Appliance from Orchestrator and Account

Administration > Software > Remove Appliances > Remove from Orchestrator and Account

Removing an appliance with this action places the appliance on the Denied Devices list, which is located as a link on the Discovered Devices tab.



This action deletes the appliance from the navigation tree. In addition, Orchestrator breaks all tunnels, overlays, and so forth to this device and tells the Portal to “unlicense” the appliance.

Administration > Tools

The options under **Administration > Tools** focus on tools that can help optimize your Orchestrator deployment, including how to synchronize appliances or put them in standby mode, perform a link integrity test, broadcast CLI commands, reboot appliances, and more.

Synchronize Appliance Configuration

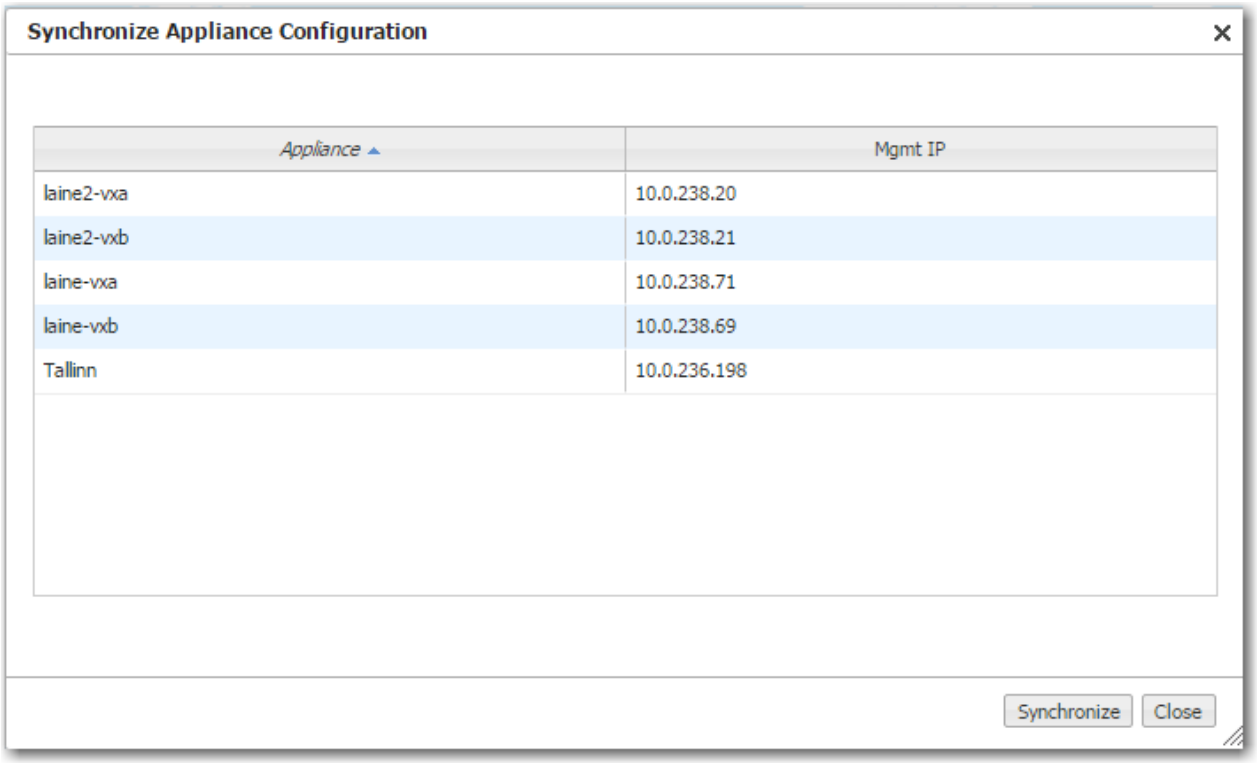
Administration > Tools > Synchronize

Orchestrator keeps its database synchronized with the running configurations for the appliances.

- When you use Orchestrator to make a configuration change to an appliance’s running configuration, the appliance responds by sending an event back to the Orchestrator server to log. This keeps Orchestrator and the appliance in sync.
- Whenever an appliance starts or reboots, Orchestrator automatically inventories the appliances to resync.
- Whenever Orchestrator restarts, it automatically resyncs with the appliances.

- When an appliance is in an OutOfSync management state, the Orchestrator server resyncs with it as it comes back online.

If your overall network experiences problems, you can use this dialog box to manually resync and ensure that Orchestrator has an appliance's current running configuration.



Broadcast CLI Commands

Administration > Tools > Broadcast CLI

You can simultaneously apply Command Line Interface (CLI) commands to multiple selected appliances.

The dialog box automatically provides you with the highest user privilege level.

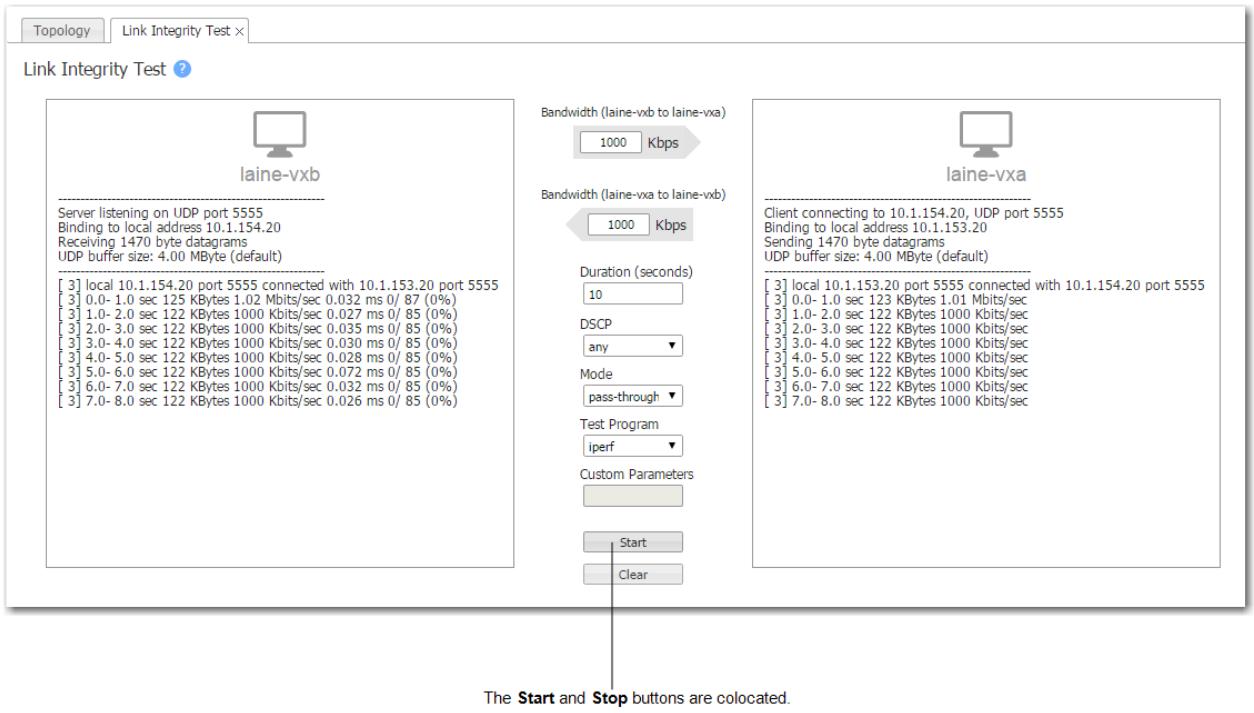


For more information, see the EdgeConnect Command Line Interface (CLI) Reference.

Link Integrity Test

Administration > Tools > Link Integrity Test

Used for debugging, the link integrity test enables you to measure the throughput and integrity (amount of loss) of your WAN link. You can run either **iperf** or **tcpperf** (Version 1.4.8).



- These tests run on the two selected appliances using user-specified parameters for bandwidth, duration, DSCP marking, and type of traffic (tunnelized / pass-through-shaped / pass-through-unshaped).
- Orchestrator runs the selected test twice—once passing traffic from Appliance A to Appliance B, and a second run passing traffic from Appliance B to Appliance A.
- Custom Parameters are available for tcpperf and should be used cautiously by advanced users.

TCPPERF Version 1.4.8

Basic Mode

Option	Description
-h	*help*
-s	*server*: Run tcpperf in server mode (not applicable for file generation). Listens on TCP port 2153 by default. [server_port [server_port [server_port]..]]
-sr	*server range*: <server_port_start:server_port_end>
-c	*client server_IP*: TCPperf Server's IP address (not applicable for file generation). [server_port [server_port [server_port]..]]
-cr	<server_port_start:server_port_end> <server_port_start:server_port_end>
-g	*generate basefilename*: Dump generated data to a file.
-sw	*sgwrite conffilename*

NOTES:

1. The default server ports are 2153 and 2154.
2. You can specify multiple odd-numbered server ports.
3. The next even-numbered server ports will also be assigned automatically.
4. These even numbers are reserved for double connection testing (see -I, *interface IP*).
5. Generate mode generates a local file per flow with the same content that the client would have generated with the specified parameters.
6. SG write mode is like generate mode, except that it writes to an SG device.

General Parameters

Option	Description
-6	*ip6*: Forces tcpperf to use IPv6 addresses only. Default is **IPv4** addresses.
-I	*interface IP*: Specify source interface IP address. Default is **any**.
-o	*outname*: Output filename. Default is **stdout**.
-u	*update* <*secs*>: Frequency of printed updates in seconds. Default is **1**.
-d	*duration* <*secs*>: Set maximum test duration in seconds. Default is **infinite**.
-w	*wait* <*secs*>: Wait until <secs> since 1970 before transmitting data.
-z	*realtime*: Elevate to realtime priority. Requires root privilege.
-cm	*cpu mask*: Specify CPU affinity. Requires root privilege.
-q	*quiet* <*level*>: Suppresses detail based on level: **0**: None. Print results when test is complete. **1**: Default. Periodic packet/byte statistics. **2**: Verbose. Adds connection state changes. **3**: Debug. Prints everything.

TCP Parameters

Option	Description
-tw	*tcpwindow*: TCP window_size. Default is **OS default**.
-tm	*tcpmss*: TCP mss. Default is **OS default**.
-tn	*tcpnodelay*: TCP nodelay option. Default is **nagle enabled**.
-tq	*tcpquickack*: TCP quick ack option. Default is **delayed acks**.
-td	*tcpdscp* <*cp*>: Sets IP DSCP to <cp> (decimal). Default is **0**.

Option	Description
-tr	*tcpretries* <*n*>: Sets number of times to retry TCP connections.
-tp	*tcppace* <*n*> <*mode*>: Pace TCP connection setup rate. Limits number of half-open connections to <n>. Valid <mode> types are: **preestablish**: All connections are established before data transmission. Default. **simultaneous**: Begin data transmission as soon as connection made.
-ta	*tcpabort*: Sends RSTs instead of FINs on close.
-tf	*tcpfindelay* <*secs*>: Time to wait after all data is sent before sending FIN/RST.

Traffic Generation Parameters

Option	Description
-f	*file*: Source filename to load. Default is **10MB** of random data.
-i	*test id* <*i*>: Set test ID. The same test ID produces the same data set. User different test IDs to generate unique data for each test run. Default is **zero**.
-n	*number* <*n*>: Generate <n> flows. Default is **one**.
-b	*begin* <*byte*>: First byte in transmission. Default is **zero**.
-e	*end* <*byte*>: End byte in transmission (number of bytes to transmit). Default is **file size**. Begin and end bytes can be greater than file size. The content is repeated to create extra bytes.
-a	*antipat* <*mode*>: Antipattern mode: default is **mutate**. **none**: Repeats same content verbatim on all flows. Repeats content if end byte exceeds content size. **mutate**: Ensures all flows and data repeats are unique. Preserves short range patterns within flow. Destroys cross flow similarity. Destroys original byte code distribution. **shuffle**: Ensures all flows and data repeats are unique. Preserves short range patterns within flow. Preserves cross flow similarity. Preserves original byte code distribution. **fast**: Ensures all flows and data repeats are unique. Does not preserve short range patterns. Destroys cross flow similarity. Destroys original byte code distribution. Uses less CPU than mutate or shuffle.

Option	Description
-l	*loopback [mode]*: Loopback. Default is **unidirectional**. **uni**: Unidirectional client to server. **rev**: Unidirectional server to client. **bidir**: Bidirectional, client and server independently send data on the same TCP connection. **bidir2**: Bidirectional, client and server independently send data on secondary TCP connections. **loop**: Bidirectional, server loops data back to client on the same TCP connection. **loop2**: Bidirectional, server loops data back to client on a secondary TCP connection. **bidir2**: Bidirectional, transmits one transaction at a time. Client waits for previous transaction to be echoed. Emulates transactional data. **NOTES:** Content source for traffic originating at the server is determined by the server (not client) command line. **loop2** and **bidir2** modes 2 x <n> TCP connections and requires that the server has even-numbered ports available.
-r	*rate* <*bps*>: Limits aggregate transmission rate to . Default is **no rate limit**.
-t	*trans* <*min*> *[max]*: Sets size of each socket transaction. Default is **64000**. If <min> and <max> are specified, client generates transactions with random sizes between <min> and <max>. This feature is often used with -l and -r. Set the minimum transaction size to 100000 to improve single-flow performance.

Option	Description
-v	*verify* <*mode*>: Verify integrity of received data. Default is **global**. **none**: No verification. Fastest/least CPU load. **global**: Single global hash per flow. Fast, but cannot isolate an errored block. **literal**: Literal comparison of data upon reception. Fast, can isolate errors to the byte level. Requires that server has same content as client. Use random data gen or same -f file at server. **embedded**: Embedded hashes every 4096 bytes. Slower, can isolate errors to 4096 byte block.
-p	*repeat* <*n*>: Repeat each content byte n times. Default is **1** (no repeats). Works for both random data and file content.
-k	*corrupt* <*n*> <*m*> <*s*> *[<%change>[<%insert>[<%delete>]]]* : Corrupt 0 to n bytes of data every m bytes using seed s. Delta bytes will require 0.5*n/m percent overhead. Each corrupt can be a change, insert or delete with the probability of each being specifiable. The default is 33.3% changes, 33.3% inserts, and 33/3% deletes.
-x	*excerpts* <*b*> <*e*> <*l*> *[s]*: Send random excerpts of average <l> length bytes from content between <b>egin and <e>nd bytes. The **-b** and **-e** options still specify total bytes to send. Uses random seed **s**.

Option	Description
-y	*defred <s% > <m%> <l%> <sb> <smin> <smax> <mb> <mmin> <mmax> <lb> <lmin lmax>*: Generate content based on defined reduction model. Content is drawn from three data sets: **s**, **m**, and **l**: **s%**: Specifies fraction [50%] of s-type content (short term reducible). **m%**: Specifies fraction [30%] of m-type content (medium term reducible). **l%**: Specifies fraction [20%] of l-type content (long term reducible). Short-term content comes from data set of sb Mbytes [100MB] with excerpts uniformly distributed between smin and smax bytes [10K-1M]. Medium-term content comes from data set of mb Mbytes [100GB] with excerpts uniformly distributed between lmin and lmax bytes [10K-1M]. Long-term content comes from data set of lb Mbytes [100TB] with excerpts uniformly distributed between smin and smax bytes [10K-1M]. The **-b** and **-e** options still specify total bytes to send. Performance is best if **-b** is **0**. Uses random seed **s**.

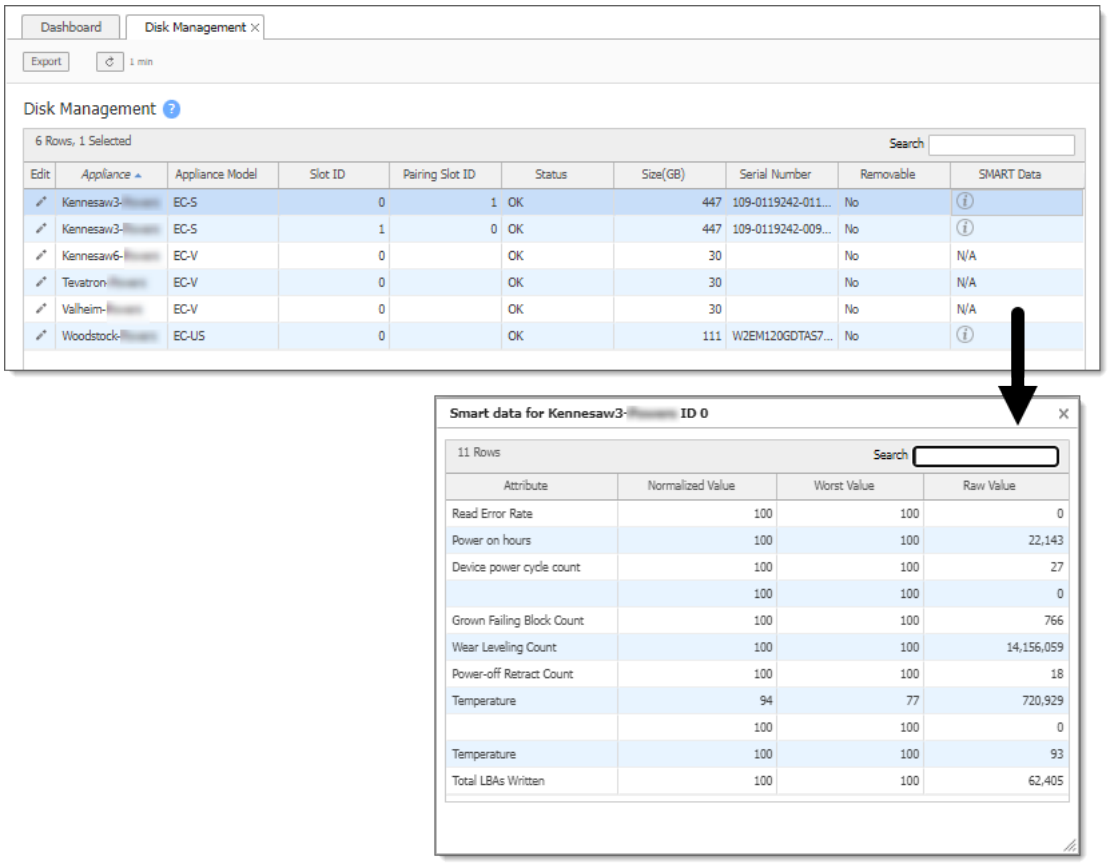
Option	Description
-ssl [param=value . . .]	Enable SSL on connection with optional parameters. **version=2|3|t10|t11|t12**: Set the protocol version. **cipher=OPENSSL-CIPHER-DESC**: Set the choice of ciphers. **ticket=yes|no**: Enable/disable session ticket extension. **cert=FILENAME**: Use this certificate file. **key=FILENAME**: Use this private keyfile. **compression=none|any|deflate|zlib|rle**: Set the compression method. **sslcert**: Print the SSL certificate in PEM format. **sslkey**: Print the SSL key in PEM format.

Disk Management

Administration > Tools > Disk Management

The Disk Management tab lists information about physical and virtual appliance disks.

- The progress bar shows what percentage of the polling is complete.
- Disk failure results in a critical alarm.
- If a row indicates that a disk has failed, click the **Edit** icon to access the appliance, and then follow the directions in the local help to replace the failed disk.
- You can view the SMART (Self-Monitoring Analysis and Reporting Technology) data from physical appliance disks only.



To replace a failed disk:

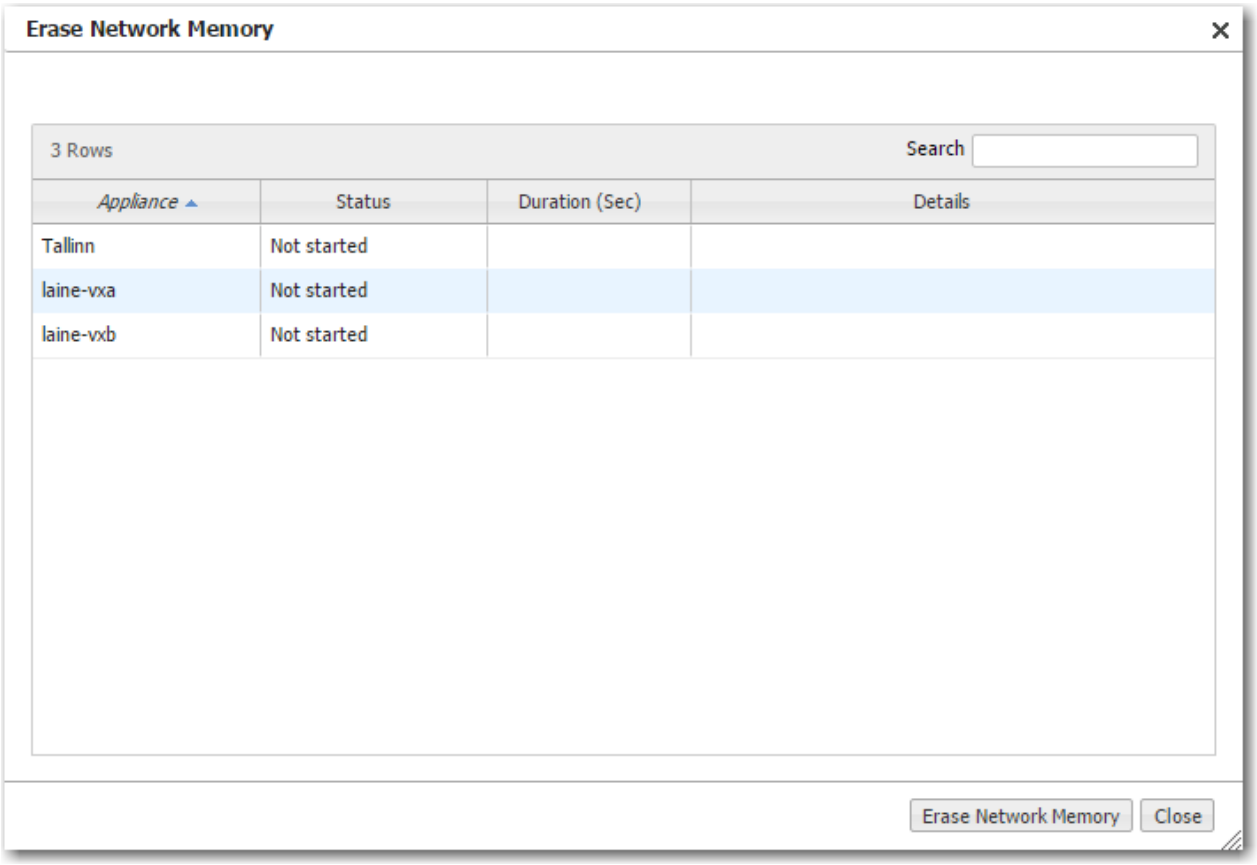
1. Log in to your Support portal account, and then click **Open a Self Service RMA** for disk replacement.
2. Complete the wizard. Use the serial number of the appliance (not the disk).
3. Receive the new disk.
4. Install the new disk by following the steps in HPE Aruba Networking EdgeConnect Hardware Reference Guide under “Physically Replacing a Disk.”
5. To access the Appliance Manager, navigate to **Administration > Tools > Disk Management**, and click the edit icon next to the newly installed appliance.
6. Follow the instructions in the Appliance Manager online help.

Erase Network Memory

Administration > Tools > Erase Network Memory

Erasing network memory removes all stored local instances of data.

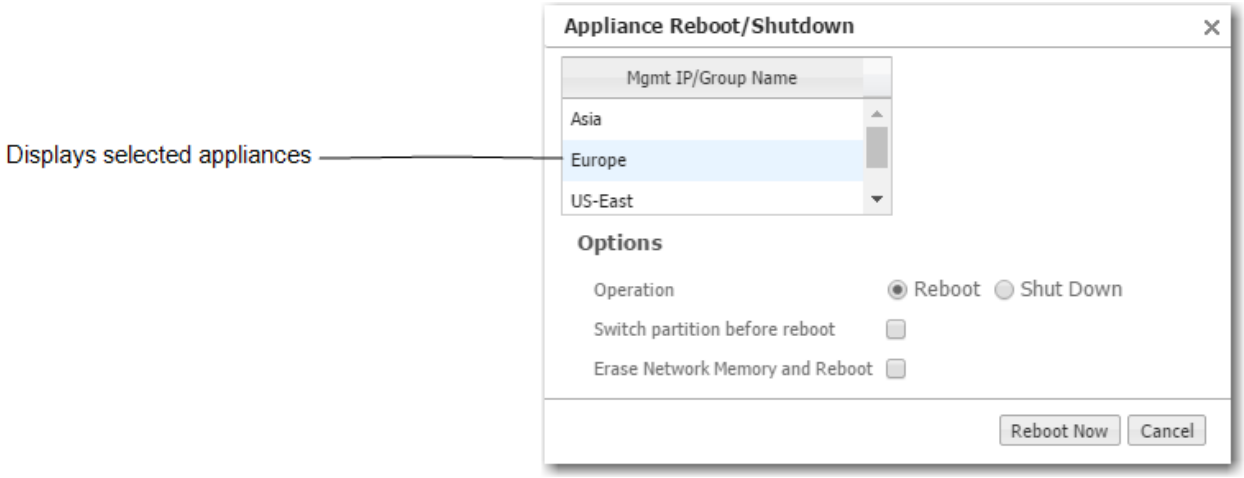
No reboot is required to complete this task.



Reboot or Shut Down an Appliance

Administration > *Tools* > *Reboot* > *Appliance Reboot / Shutdown*

The appliance supports three types of reboot:



- **Reboot**: Reboots the appliance gracefully. This is your typical “vanilla” restart.

Use case: You are changing the deployment mode or other configuration parameters that require a reboot.

- **Erase Network Memory and Reboot**: Erases the Network Memory cache and reboots the appliance.

  Use case: You need to restart the appliance with an empty Network Memory cache.

- **Shutdown**: Shuts down the appliance and turns the power off. To restart, go to the appliance and physically turn the power on with the Power switch.

  Use cases:

  - You are decommissioning the appliance.
  - You need to physically move the appliance to another location.
  - You need to re-cable the appliance for another type of deployment.

Behavior During Reboot

A physical appliance enters into one of the following states:

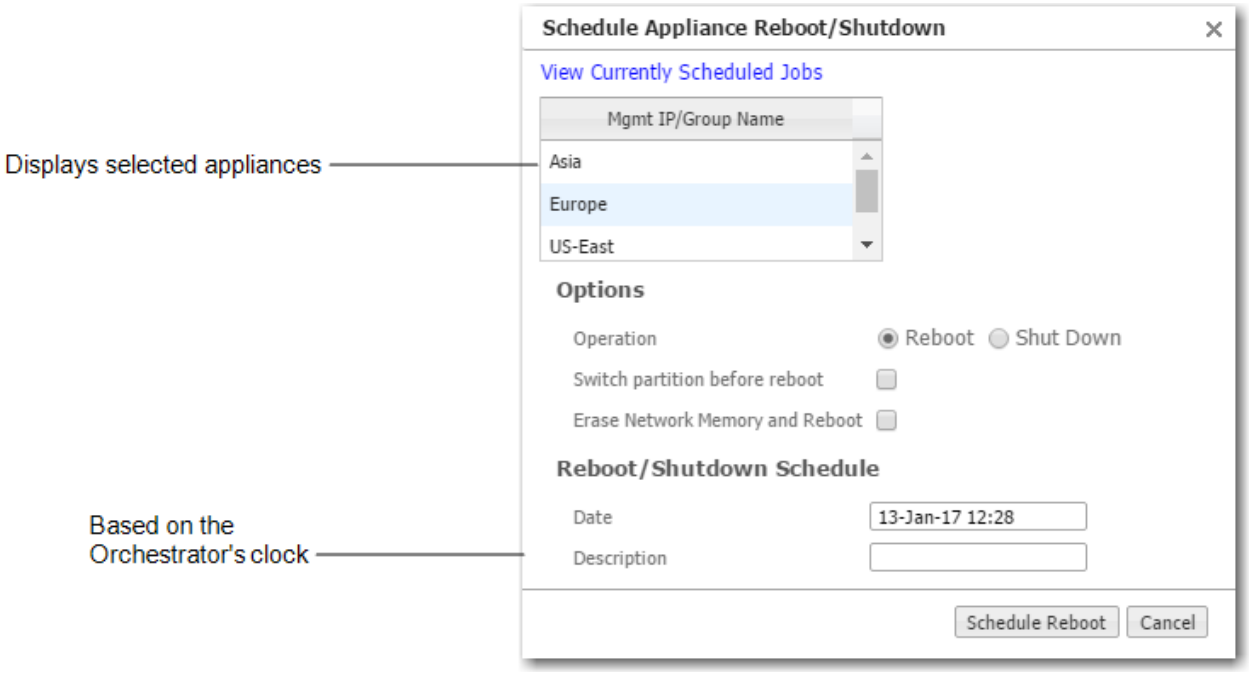
- hardware bypass, if deployed in-line (Bridge mode)
- open-port state, if deployed out-of-path (Router mode or Server mode)

Unless a virtual appliance is configured for a high availability deployment, all flows are discontinued during reboot.

Schedule an Appliance Reboot

Administration > Tools > Reboot > Schedule Appliance Reboot

You can schedule an appliance for any of three types of reboot:



- **Reboot**: Reboots the appliance gracefully. This is your typical “vanilla” restart.

  Use case: You are changing the deployment mode or other configuration parameters that require a reboot.

- **Erase Network Memory and Reboot**: Erases the Network Memory cache and reboots the appliance.

  Use case: You need to restart the appliance with an empty Network Memory cache.

- **Shutdown**: Shuts down the appliance and turns the power off. To restart, go to the appliance and physically turn the power on with the Power switch.

  Use cases:

  - You are decommissioning the appliance.
  - You need to physically move the appliance to another location.
  - You need to re-cable the appliance for another type of deployment.

Behavior During Reboot

A physical appliance enters into one of the following states:

- hardware bypass, if deployed in-line (Bridge mode)
- open-port state, if deployed out-of-path (Router/Server mode)

Unless a virtual appliance is configured for a high availability deployment, all flows are discontinued during reboot.

TIP: To specify the time zone for scheduled jobs and reports, navigate to **Orchestrator > Software & Setup > Setup > Timezone for Scheduled Jobs**.

Reachability Tab

Administration > Tools > Monitoring > Reachability Status

This tab displays status details about communications with EdgeConnect appliances. It includes the following two subtabs:

- Appliances/Orchestrator – Displays status details on communications between your appliances and Orchestrator. This is the default view on the Reachability tab. See Appliances/Orchestrator for information about this subtab.
- Appliances/HPE ANW Central – Displays status details on communications between appliances and HPE Aruba Networking Central. Information is only available on this subtab if you have deployed the Unified Fabric solution and have enabled HPE Aruba Networking Central. See Appliances/HPE ANW Central for information about this subtab.

This tab also provides the following links:

- **Cloud Portal** – Click to view current configuration and registration information for the Cloud Portal connection. You can also view HTTPS and WebSocket reachability statuses.
- **Orchestrator Reachability** – Click to view how your appliances are connecting to Orchestrator, including the default Orchestrator IP address or domain name, whether the Orchestrator management IP is being used, and any labels used to connect to Orchestrator.
- **Software Versions** – Click to view active and inactive software versions for all appliances in your network or those selected in the appliance tree.

Appliances/Orchestrator

This subtab displays status details on communications between your appliances and Orchestrator, which helps in early detection of network or communication failures that might occur over direct or Portal WebSockets.

Orchestrator provides automatic WebSocket failover from direct WebSocket to HPE Aruba Networking Cloud Portal WebSocket if communications over the direct WebSocket fail or become stale. You can control how quickly failover occurs by selecting one of three failover modes: Aggressive, Normal (default), or Slow. For details, see Select Fast WebSocket Failover Mode below.

The listed appliances reflect all appliances in your network or those selected in the appliance tree.

The Ping Interval (Sec) and Max Idle Time (Sec) columns reflect the current WebSocket failover mode configuration. The other columns show live data. Click the refresh button to update this data. You can also click **Export** to export the data to a CSV file.

The following table describes the fields on the Appliances/Orchestrator tab.

Field	Description
Appliance	Name of the appliance.

Field	Description
Direct WebSocket	Status of the appliance connection to Orchestrator by direct WebSocket (Connected or Not Connected).
Appliance Portal WebSocket	Status of the appliance connection to Orchestrator by HPE Aruba Networking Cloud Portal WebSocket (Connected, Connected Secure, or Not Connected). “Connected Secure” indicates that the end-to-end encryption feature is enabled (the endToEndEncryptionEnabled property on the Orchestrator Advanced Properties dialog box is set to “true”) and the appliance is running on an ECOS version that supports the feature. **NOTE:** If Air-Gap is enabled, this status is always “Not Connected” because the HPE Aruba Networking Cloud Portal is not being accessed. Appliance activation, license management, and other supporting functions are performed through the Air-Gap Portal. Users manually exchange data between HPE Aruba Networking Cloud Portal and the EdgeConnect SD-WAN network. Air-Gap is available only for self-hosted Orchestrator deployments. For details, see Air-Gap.
Active Channel	Currently active communication channel (Via Direct WebSocket or Via Portal WebSocket) that Orchestrator is using to communicate with the appliance.
Last Direct WebSocket Message	When the last message (ping or any other message) was received from this appliance. This is the actual time that Orchestrator tracks.
Ping Interval (Sec)	Ping interval in seconds. The appliance polls the Orchestrator at the passing of every ping interval to confirm reachability over the direct WebSocket. This value reflects the current WebSocket failover mode selection. To change the mode, see Select Fast WebSocket Failover Mode below. If Orchestrator does not receive a message on the direct WebSocket within the Redirection Time as configured for the failover mode, Orchestrator redirects communications to the Portal WebSocket if it is available.

Field	Description
Max Idle Time (Sec)	Maximum idle time in seconds. Orchestrator closes the direct WebSocket with the appliance if a message is not received from the appliance within the Max Idle Time. This value reflects the current WebSocket failover mode selection. To change the mode, see Select Fast WebSocket Failover Mode below. The appliance will attempt to re-open the direct WebSocket to Orchestrator. If successful, Orchestrator will redirect traffic to the direct WebSocket.
Sync State	Synchronization state of the appliance and Orchestrator databases. Possible synchronization states are: Normal – The Orchestrator and appliance databases are synchronized, and communications are running normally across the active channel. Not synchronized – Databases for Orchestrator and the appliance are not synchronized. Synchronization in progress – Synchronization is occurring between the Orchestrator and appliance databases. Unknown – The appliance's transitional state when initially added to the network. Unreachable – A problem exists in your network. Check your ports, firewalls, and deployment configuration. Unsupported – The appliance is running on an unsupported ECOS version.

Select Fast WebSocket Failover Mode

You can use the Fast WebSocket Failover Mode dialog box to select a mode for WebSocket failover to the Portal WebSocket when the direct WebSocket fails or becomes stale.

Before you change this mode, consider the following:

- Changes to failover mode affect all appliances on your network. You cannot set this mode for individual appliances.
- The WebSocket failover mode feature requires appliance version 9.4.2.0 or later. Appliances on earlier versions remain in Legacy mode. Changing the failover mode on this dialog box does not affect Legacy mode on these appliances.

- Faster WebSocket failover requires greater resource utilization (such as with CPU or bandwidth).

To change the mode:

1. Click **Change WebSocket Failover Mode** on the Reachability tab.

   The Fast WebSocket Failover Mode dialog box opens.

2. Select a mode.

   The following table describes the WebSocket failover modes. You cannot change the configurations for these modes.

Field	Description
Aggressive	The most optimal mode for most network configurations. Mode configurations: 30 seconds for Redirection Time, 60 seconds for Max Idle Time, and 10 seconds for Ping Interval.
Normal	The default mode for appliances running on ECOS version 9.4.2.0 or later. **NOTE:** If Orchestrator is upgraded to version 9.4.2, existing appliances running on ECOS versions earlier than 9.4.2.0 will operate in Legacy mode, even though the setting displayed on this dialog box for those appliances will be Normal (by default). These appliances ignore changes to Aggressive, Normal, or Slow mode and continue to operate in Legacy mode. Mode configurations: 60 seconds for Redirection Time, 120 seconds for Max Idle Time, and 20 seconds for Ping Interval.
Slow	Select this mode if you are experiencing resource issues (such as with CPU or bandwidth). Mode configurations: 90 seconds for Redirection Time, 180 seconds for Max Idle Time, and 30 seconds for Ping Interval.
Legacy	Appliances running on ECOS versions earlier than 9.4.2.0 operate in this mode regardless of the mode selected on this dialog box. You can also set this mode for all appliances in your network. This mode is not recommended because of latency concerns. Mode configurations: 10 minutes for Redirection Time, 10 minutes for Max Idle Time, and 120 seconds for Ping Interval.

3. Click **Save**.

   Changes to failover mode are orchestrated. Updates can take a while to propagate across the network. Appliances will rebuild direct WebSocket connections.

Appliances/HPE ANW Central

This subtab displays status details on communications between appliances and HPE Aruba Networking Central. The listed appliances reflect only the appliances in your network that are part of the Unified Fabric solution.

IMPORTANT: Information is only available on this subtab if you have deployed the Unified Fabric solution and have enabled HPE Aruba Networking Central.

For more information about the Unified Fabric solution, see HPE Aruba Networking Central.

The following table describes the fields on the Appliances/HPE ANW Central subtab.

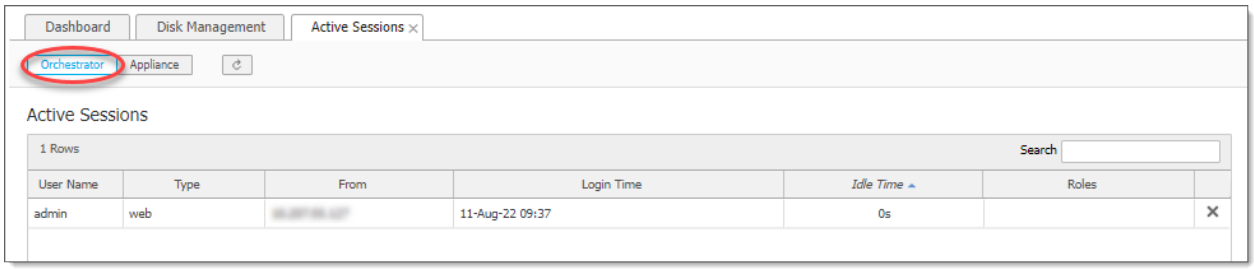
Field	Description
Appliances	Name of the appliance.
HPE ANW Central Hostname	The hostname of the HPE Aruba Networking Central instance.
HPE ANW Central IP	The IP address of the HPE Aruba Networking Central instance.
HPE ANW Central WebSocket	Status of the appliance connection to HPE Aruba Networking Central by WebSocket (Connected or Not Connected).
ORO (gRPC)	Status of the Overlay Route Orchestrator (Connected or Not Connected).
OTO (gRPC)	Status of the OTO (Overlay Tunnel Orchestrator) tunnels (Connected or Not Connected).

Active Sessions Tab

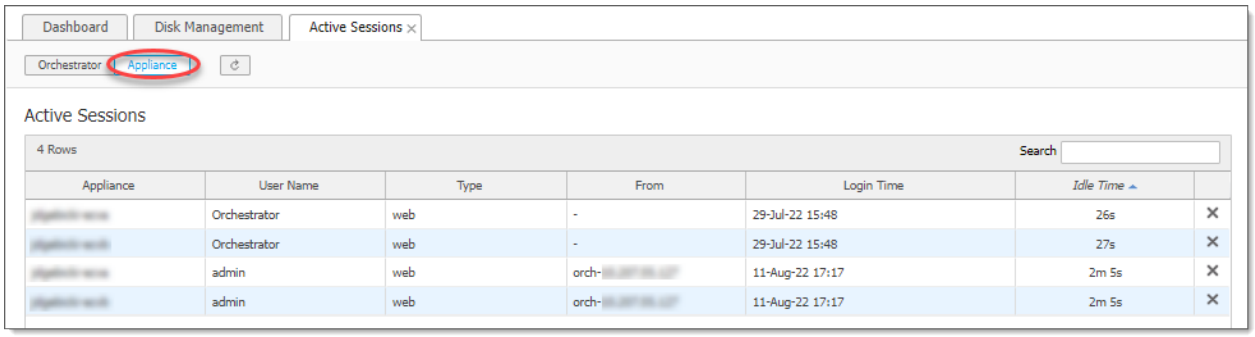
Administration > Tools > Monitoring > Active Sessions

This tab lists users who are logged in to Orchestrator and the appliances that Orchestrator is currently managing.

To list active user sessions, click **Orchestrator**.



To list active appliance sessions, click **Appliance**.



Orchestrator

The menus under **Orchestrator** are used to manage Orchestrator itself, and are not related to managing appliances. The menus under this section are organized as follows:

- Orchestrator Server
- Software & Setup
- HPE Integration Services

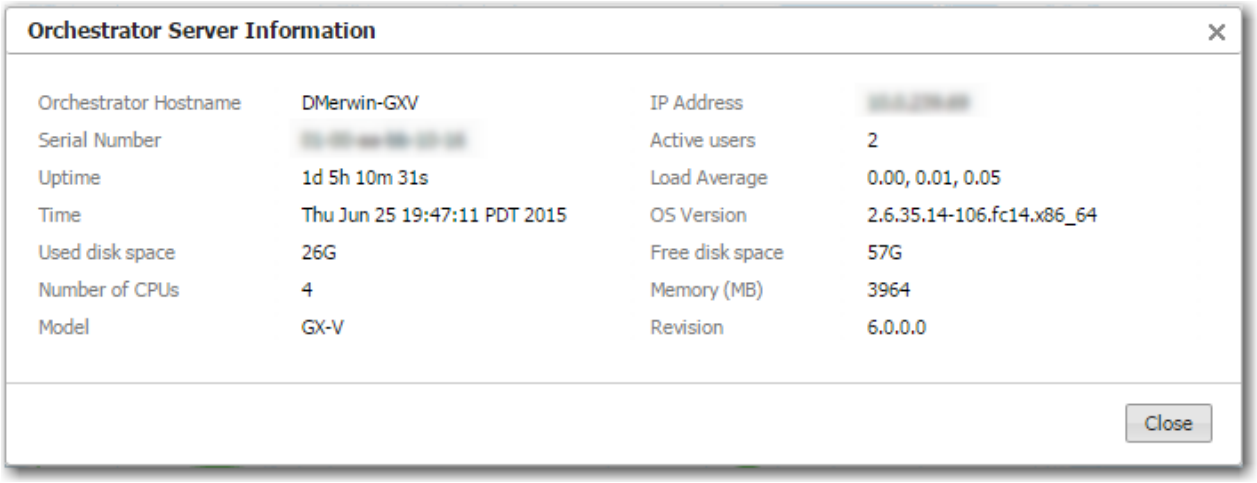
Orchestrator > Orchestrator Server

The options under **Orchestrator > Orchestrator Server** focus on settings and configuration changes you can make to Orchestrator server deployments, including managing users, RBAC options, audit logging, tunnel settings, and more.

View Orchestrator Server Information

Orchestrator > Orchestrator Server > Server Management > Server Information

This dialog box provides data specific to this Orchestrator server.



Restart, Reboot, or Shutdown

Orchestrator > Orchestrator Server > Server Management > Reboot Orchestrator

Orchestrator > Orchestrator Server > Server Management > Shutdown Orchestrator

Orchestrator provides these two convenient actions in the Orchestrator menu:

- **Reboot Orchestrator** reboots the Orchestrator server.
- **Shutdown Orchestrator** results in the server being unreachable. To restart, you must manually power on the server.

Orchestrator High Availability

Orchestrator > Orchestrator Server > Server Management > Orchestrator High Availability

Interruption of Orchestrator functionality can occur for various reasons, such as when the Orchestrator database or application fails, the Cloud Portal becomes unreachable, or routine maintenance occurs. Orchestrator high availability (HA) is designed to minimize interruptions by enabling you to fail over from a primary Orchestrator to a backup (stand-by) Orchestrator. An Orchestrator HA cluster consists of one primary Orchestrator and one or more backups.

In non-HA configurations, a single Orchestrator has its own database on the same machine. With Orchestrator HA, the primary and backup Orchestrators share the same database, which is on its own dedicated machine.

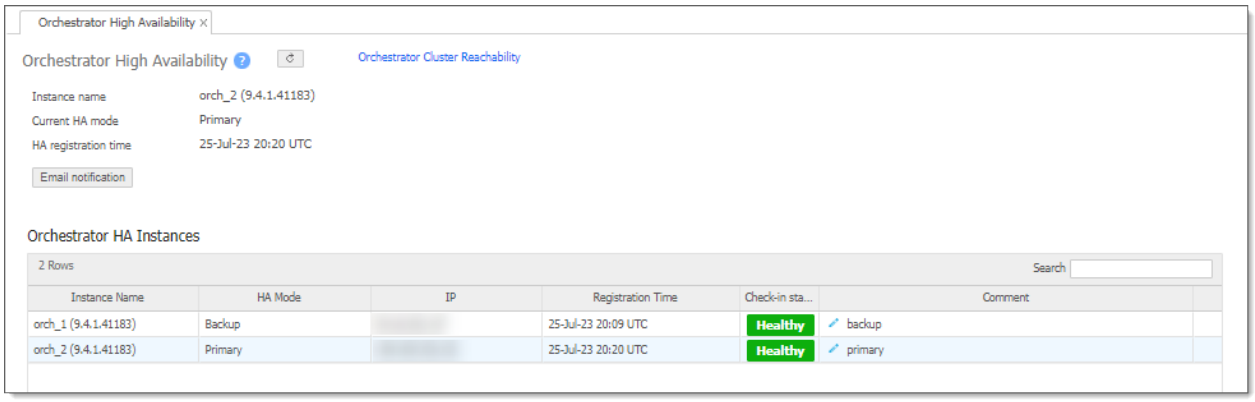
Before you can use Orchestrator HA, you must:

- Install a MySQL database cluster on a dedicated machine.
- Prepare the database cluster for Orchestrator HA.
- Install two or more Orchestrators in HA mode or upgrade existing Orchestrators to HA mode. Connect the Orchestrators to the database cluster.

For details, refer to the HPE Aruba Networking Orchestrator High Availability Installation and Deployment Guide.

Use the Orchestrator High Availability tab from primary or backup Orchestrator instances to view the HA modes of all Orchestrator instances in the HA cluster and their check-in statuses. You can also use the tab to:

- Set up email notifications to be sent when Orchestrator instances do not make health status updates to the database *(from the primary instance only)*.
- Configure Orchestrator HA cluster reachability *(from the primary instance only)*.
- Promote a backup Orchestrator to primary when the current primary Orchestrator is down *(from backup instances only)*.
- Decommission an Orchestrator instance *(from the primary instance only)*.



Field	Description
Instance name	Name of the Orchestrator instance you are currently accessing. The current build version is indicated.
Current HA mode	HA mode of the Orchestrator instance you are currently accessing (Primary or Backup).
HA registration time	Timestamp indicating when the Orchestrator instance you are currently accessing initially connected to the database.
Instance Name	Name of the Orchestrator instance. The current build version is indicated.
HA Mode	HA mode of the Orchestrator instance (Primary or Backup).
IP	IP address for the Orchestrator instance.
Registration Time	Timestamp indicating when the Orchestrator instance initially connected to the database.

Field	Description
Check-in status	Indicates the health status of the Orchestrator instance (whether the instance can check in to the database). Each Orchestrator checks in to the database every ten seconds. Statuses are: Healthy – Orchestrator instance made a successful health status update in the past 20 seconds. Failed – Orchestrator instance failed to make a health status update for longer than 20 seconds.
Comment	Comment you want to include for the Orchestrator instance. Click the edit icon, enter a comment, and then click **Save**.

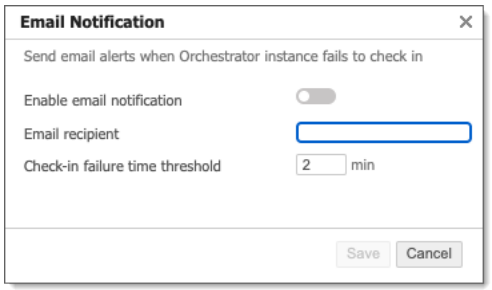
Set Up Email Notifications

If an Orchestrator instance goes down, an alarm is raised and an alarm email notification is sent. You can also set up an additional email notification to be sent when an Orchestrator instance in the HA cluster fails to make a health status update to the database.

To set up the additional email notification:

1. Ensure that you are logged in to the primary Orchestrator instance. You cannot set up email notifications from a backup instance.
2. On the Orchestrator High Availability tab, click **Email Notification**.

   The Email Notification dialog box opens.



3. Move the **Enable email notification** toggle to the right.
4. Complete the following fields.

Field	Description
Email recipient	Email address of the person to receive notifications.

Field	Description
Check-in failure time threshold	Number of minutes that healthy Orchestrators wait to send an email notification after an Orchestrator instance fails to check in. The notification includes the statuses of all registered instances. The default is 2. You can specify a number from 1 to 30.

5. Click **Save**.

Configure Orchestrator HA Cluster Reachability

You need to establish how appliances will communicate with Orchestrator instances in the HA cluster. You can configure reachability in either of the following ways:

- Assign a default IP address or domain name to each Orchestrator instance. If the Orchestrator IP address is private, you must provide a reachable IP address so that appliances can communicate with it. Appliances will use this reachable IP address to reach Orchestrator by opening a direct web socket to the Orchestrator.
- Assign IP addresses or domain names to each Orchestrator instance for specific appliance interface labels. You can set a priority to these label-specific connections. Use this option if the appliance needs to communicate to Orchestrator using a WAN interface over a private network.

To configure Orchestrator HA cluster reachability:

1. Ensure that you are logged in to the primary Orchestrator instance. You cannot configure reachability from a backup instance.
2. On the Orchestrator High Availability tab, click **Orchestrator Cluster Reachability**.

   The Orchestrator Cluster Reachability dialog box opens.



3. For default reachability, specify a default IP address or domain name for each Orchestrator instance listed in the Default Orchestrator IPs or Domain Names area.
4. For reachability by interface label:
    1. Click **Add** in the Reachability for Labels area.

       A row is added to the table.
    2. From the **Label** drop-down list, select the label for which you want to assign reachability.

       The IP or Domain Name field for each Orchestrator instance automatically populates with an IP address and mask. You must change the IP address if a specific label needs to reach the Orchestrator using a different IP address or domain.
    3. Change the value in the **Priority** field as appropriate.
5. Click **Save**.

Promote Backup Orchestrator to Primary

To promote a backup Orchestrator to primary, the current primary Orchestrator must be unreachable because it is down or has been properly shut down so that the VM is not running.

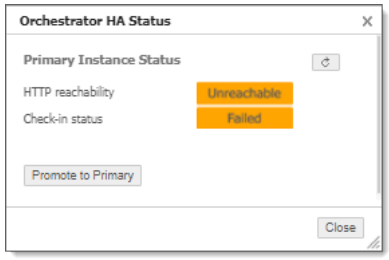
NOTE: The Orchestrator High Availability tab for the backup Orchestrator displays the same way as it does for the primary Orchestrator, except that it includes the Promote to Primary button.

To promote a backup Orchestrator instance to primary:

1. Ensure that the current primary Orchestrator is down or has been properly shut down so that the VM is not running.

2. Log in to the backup Orchestrator instance you want to promote to primary.
3. Navigate to the Orchestrator High Availability tab, and then click **Promote to Primary**.

   The Orchestrator HA Status dialog box opens.



   This dialog box provides status information for the current primary Orchestrator.

   NOTE: If the current primary instance has a Reachable status rather than an Unreachable status as shown above, you cannot proceed to promote the backup instance to primary. Only one Orchestrator can be working as primary at any given time. Ensure that the previous primary is completely down before promoting a backup.
4. To promote the backup Orchestrator to primary, click **Promote to Primary**.

   A confirmation dialog box opens.
5. To proceed, click **Promote to Primary**.

   The backup Orchestrator will restart in ten seconds and will be in Primary mode. This new primary Orchestrator will function fully, the same as the previous one did.

Decommission an Orchestrator Instance

To decommission an Orchestrator instance:

1. Ensure that the Orchestrator instance you want to decommission is completely shut down.
2. Log in to the primary Orchestrator instance. You cannot decommission Orchestrators from a backup instance.
3. On the Orchestrator High Availability tab, click the **X** in the far-right column of the row associated with the Orchestrator instance you want to decommission.

   A confirmation dialog box opens.
4. Confirm that you want to decommission the instance.

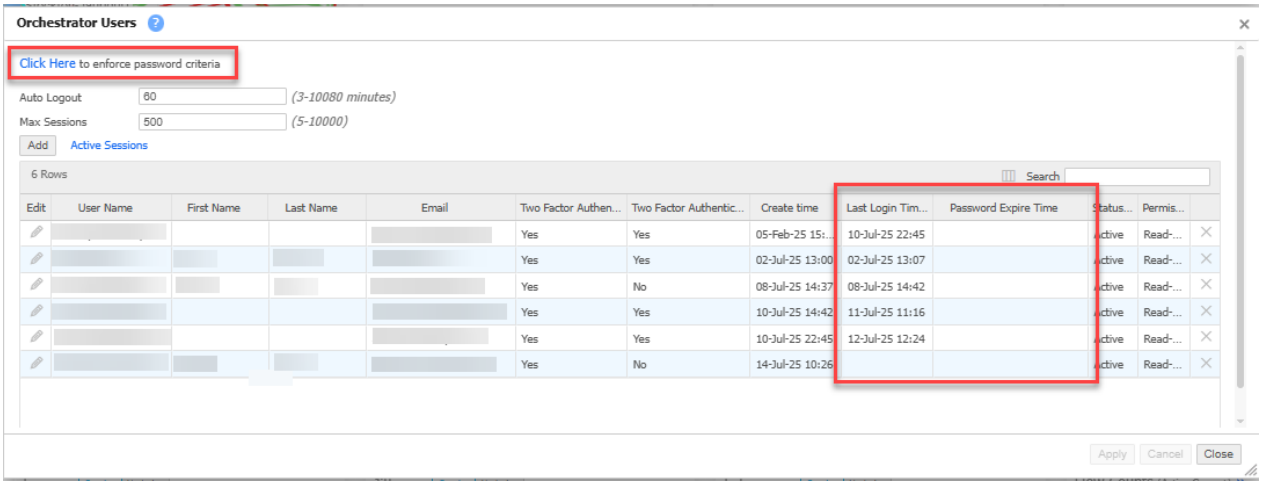
Orchestrator Users

Orchestrator > Orchestrator Server > Users & Authentication > Orchestrator Users

Use the Orchestrator Users dialog box to:

- Manage who has Read-Write or Read-Only access to Orchestrator.
- Select whether to enforce enhanced password criteria for all users.

  You can enforce enhanced password criteria for all users and select the password properties you want to enforce. New password criteria were added, and existing criteria were enhanced to comply with enterprise security standards. The following graphic highlights new features including the link to enforce enhanced password requirements and the Last Login Time and Password Expire Time columns.



Enforce Password Criteria and Set Password Requirements

To enforce the enhanced password criteria and set password requirements, you must modify the Orchestrator Advanced Properties. The ability to modify Orchestrator Advanced Properties depends on your Orchestrator deployment.

- Self-hosted/on-prem Orchestrator users can edit the Orchestrator Advanced Properties using the following instructions. You must restart Orchestrator to activate certain Orchestrator Advanced Properties changes.
- Orchestrator-as-a-Service (OaaS) users must open a TAC case to request changes to the Orchestrator Advanced Properties. TAC must also restart OaaS to activate Orchestrator Advanced Properties changes.

To enforce enhanced password criteria for all users and set password requirements:

1. Click **Click Here** (to enforce password security) at the top of the Orchestrator Users dialog box.

   The Orchestrator Advanced Properties dialog box opens.

   WARNING: Use caution when modifying Orchestrator Advanced Properties. Verify your changes before applying them.

2. Set the following password property values as necessary:

   TIP: Enter “password” in the search field and press Enter to display the password properties.

Password Property	Description	enforcePasswordCriteria required
user.passwordMinLen	Minimum number of characters for the password. Allowable range: 8-64. The default setting is 8.	No
useLowercaseLetterInPassword	Indicates whether a lowercase letter is required in the password. Enter **true** or **false**. The default setting is true.	Yes*
enforcePasswordCriteria	Indicates whether to enforce password criteria for Orchestrator users. Enter **true** or **false**. The default setting is false for both upgraded and new Orchestrators. **IMPORTANT:** If enforcePasswordCriteria is set to true, you cannot revert to false.	NA**
useSpecialSymbolInPassword	Indicates whether a special character is required in the password. Enter **true** or **false**. The default setting is true.	Yes*
user.passwordFailedThreshold	Number of failed login attempts allowed from a given source IP address. The default is 4 attempts. When this threshold is exceeded, Orchestrator locks out the source IP address from further login attempts for the lockout duration, which is 5 minutes by default. You can modify the lockout duration by changing the resetTimerForFailedLogi-nAttempts(mins) login property in step 4 below.	No

Password Property	Description	enforcePasswordCriteria required
passwordHistoryCount	Number of previous passwords to store for the user to ensure passwords are not repeated. Allowable range: 4-8. The default setting is 8.	Yes
passwordExpireDays	Number of days before the password expires. Allowable range: 30-360. The default setting is 90. If enforcePasswordCriteria is set to true, the user's password will expire based on the value configured for passwordExpireDays, otherwise the password will not expire.	Yes
notifyPasswordExpiryDays	Number of days before a password expires to notify the user of the upcoming expiration. Allowable range: 5-30. The default setting is 7. Orchestrator will display a popup when the user logs in within the configured number of days.	Yes
useUppercaseLetterInPassword	Indicates whether an uppercase letter is required in the password. Enter **true** or **false**. The default setting is true.	Yes*
useDigitInPassword	Indicates whether a number is required in the password. Enter **true** or **false**. The default setting is true.	Yes*

Password Property	Description	enforcePasswordCriteria required
userInactiveDays	Indicates whether to delete the user after a specified number of inactive days. Enter a value of 0 or a number in the range of 30-360. A value of 0 indicates that the user deletion setting is not activated. The default value is 0.	Yes

* If enforcePasswordCriteria is set to false, passwords require at least one uppercase letter, one lowercase letter, a number, and a special character.

** Enforcement of the Password Settings template is dependent on the enforcePasswordCriteria property. If the enforcePasswordCriteria field on the Orchestrator Advanced Properties is set to false, the Password Settings template will not be applied. See Password Settings Template for more information.

3. Click **Apply**.
4. Set the following login property values as necessary.

   TIP: Enter “login” in the search field and press Enter to display the login properties.

Login Property	Description	enforcePasswordCriteria required
failedLoginAttemptThreshold	Orchestrator issues an alarm when the number of unsuccessful login attempts from any IP address reaches this threshold within the number of minutes set by the resetTimerForFailedLoginAttempts(mins) login property. The threshold counter resets when a successful login occurs or when no login attempts occur within the lockout duration, which is set using resetTimerForFailedLoginAttempts(mins).	No

Login Property	Description	enforcePasswordCriteria required
resetTimerForFailedLoginAttempts(r	Lockout duration in minutes for a given source IP address. The default is 5 minutes.	No

5. Click **Apply**.
6. Click **Close**.
7. If you are using on-prem Orchestrator, you can activate Orchestrator Advanced Properties changes. Open an SSH session and use the following CLI command to restart Orchestrator:

   ```
   systemctl restart gms
   ```

Add a User

- Users can have either Read-Write or Read-Only privileges.
- Multi-Factor Authentication (MFA) is mandatory for OaaS Orchestrator users. For on-premises Orchestrator users, MFA is optional but strongly recommended.
- A username cannot be more than 512 characters long.

  NOTE: You cannot modify a Username. You must delete it and create a new user.

1. Navigate to **Orchestrator > Orchestrator Server > Users & Authentication > Orchestrator Users**.
2. Click **Add**.

   The Add User dialog box opens.
3. Complete the fields and click **Add**.
4. Navigate to **Orchestrator > Orchestrator and Server > Users & Authentication > Role Based Access Control** and grant the user the proper authorization. See Role Based Access Control.

Multi-Factor Authentication

Orchestrators support Multi-Factor Authentication (MFA) on all platforms, including cloud and on-premise versions. For cloud versions of Orchestrator, MFA is required. For on-premise deployments, MFA is available but not required.

The first step in authentication is always username/password. For added security, users can choose between application- or email-based authentication, as described below.

NOTE: Only users whose role is assigned Read-Write privilege for Orchestrator Users can enable or disable MFA for any user.

Configuring Multi-Factor Authentication Through an Application

Orchestrator supports applications that provide time-based keys for two-factor authentication and are compliant with RFC 4226 / RFC 6238. Google Authenticator is one such app. The example below uses Google Authenticator on a mobile phone. You can also use a desktop version.

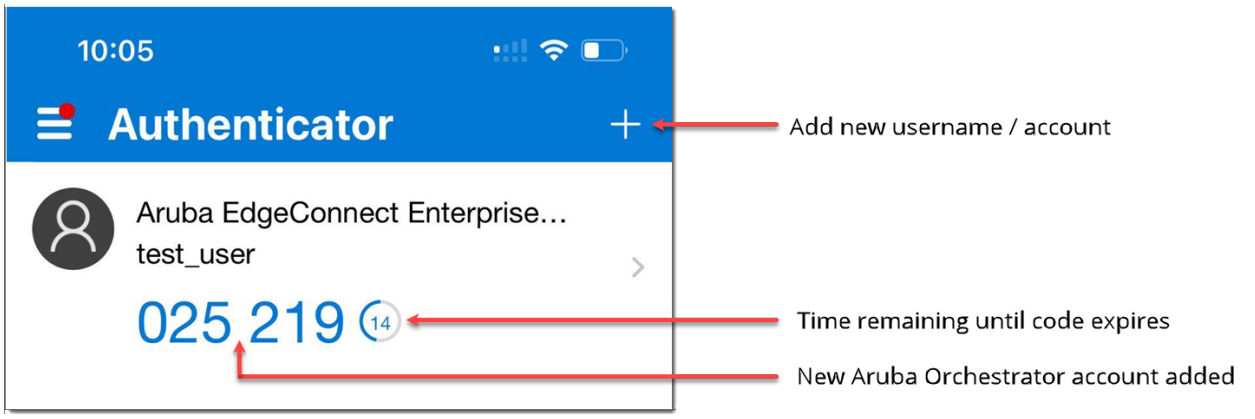
To enable MFA through an application:

1. Navigate to **Orchestrator > Orchestrator Server > Users & Authentication > Orchestrator Users**, and then click your username.
2. In the Two Factor field, select **Application**. Orchestrator generates a time-limited QR code.



3. In the Google Authenticator app, use the **Scan barcode** function to read the QR code. You will be prompted to enter your Orchestrator username and password.

   Here you can see Google Authenticator with the new account added for the Orchestrator.



Configuring Multi-Factor Authentication Through Email

To enable MFA through email:

1. Navigate to **Orchestrator > Orchestrator Server > Users & Authentication > Orchestrator Users**, and then click your username.
2. In the Two Factor field, select **Email**, and then enter your email address.

   If an invalid email address is entered, the account could be locked out and would require password reset procedures.
3. Click **Add**. Orchestrator sends a time-limited authentication code to your email address. To verify your email address, click that link.

   Orchestrator then opens a browser window telling you that your email address has been verified.

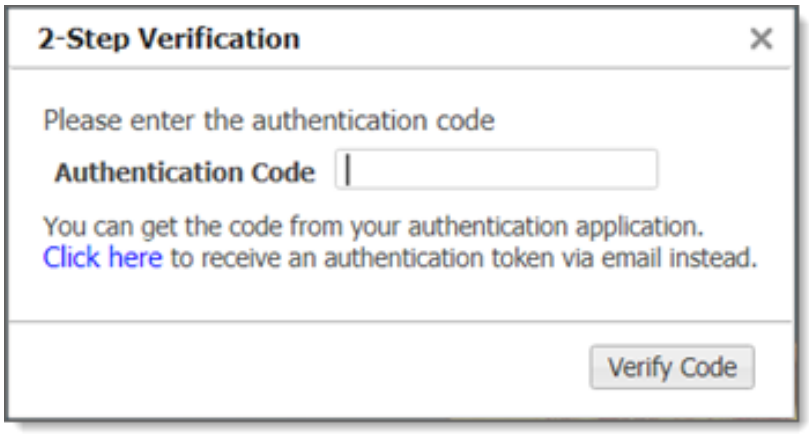
Using Multi-Factor Authentication

After MFA is configured, every login requires two steps: entering the username/password and entering the current token.

Based on the authentication method you choose, do one of the following:

- Use the current token from the Google Authenticator (or other) app.
- Use the code you receive in email.

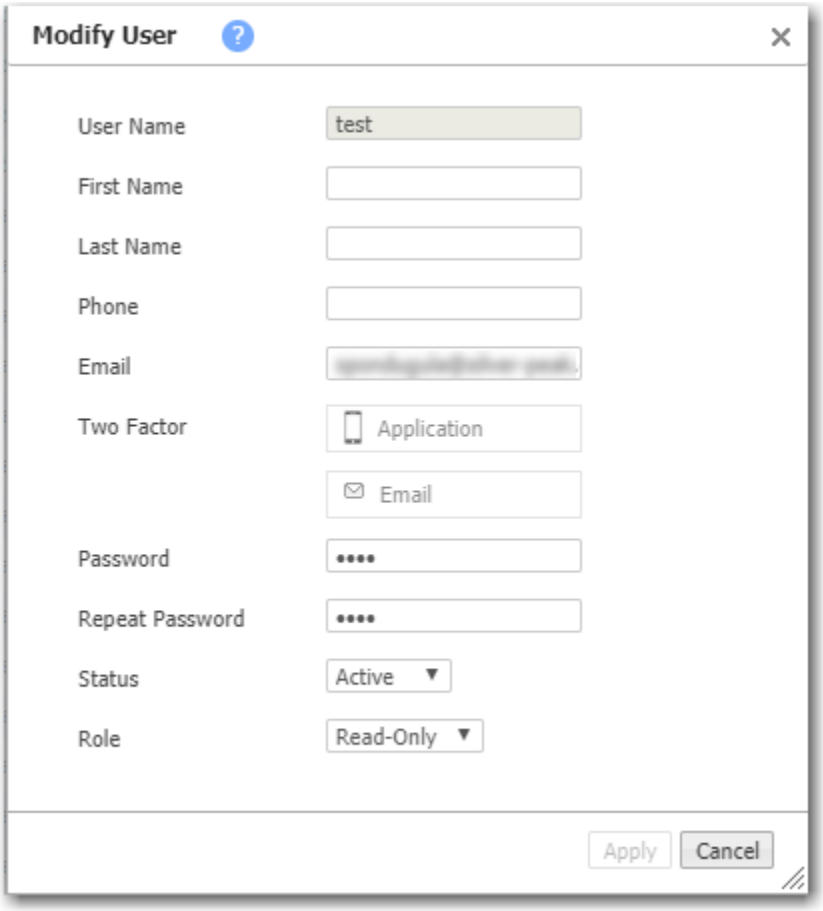
In both cases, the codes have a specific expiration time.



Modify a User

1. Navigate to **Orchestrator > Orchestrator Server > Users & Authentication > Orchestrator Users**.
2. Click the edit icon for the user you want to modify.

   The Modify User dialog box opens.



3. You can modify the following user fields:

   - **User Name** is the identifier the user uses to log in, and it cannot be more than 512 characters long.
   - **First Name**, **Last Name**, and **Phone Number** are optional information.
   - **Email** is required if two-factor authentication is enabled.
   - **Two-Factor Authentication** is a second step in the login process that requires an authentication code. The code can be obtained in two ways:
     - Using an authentication application that generates time-based authentication codes. If this is activated, Orchestrator generates a barcode that can be scanned to set up an authentication app like Google Authenticator for your mobile device.
     - Using your email to receive authentication codes every time you log in. This requires access to your email every time you log in.
   - **Password** is used at login.
   - **Status** determines whether the user can log in.
   - **Role** determines the user's permissions.

Role Based Access Control

Orchestrator > Orchestrator Server > Users & Authentication > Role Based Access Control

Role Based Access Control (RBAC) provides a customized Orchestrator experience. On a per-user basis, you can assign roles that specify access levels for a user, control the menu options available in the Orchestrator UI, and grant or deny access to appliance groups. Starting with Orchestrator 9.3.0, RBAC affects both Orchestrator UI users and Orchestrator REST API users.

NOTE: In Orchestrator 9.3.0, endpoint definitions changed for SD-WAN Orchestrator REST APIs. This required users to update endpoint definitions in their Orchestrator REST API scripts. To reduce the magnitude of change required, Orchestrator 9.3.1 provides the ability to enable support for Interop (before Orchestrator 9.3.0 release) API endpoints so you can continue using your existing Orchestrator REST API scripts for a specific list of frequently used commands (see Pre 9.3 API Endpoints).

You should be aware of the following:

- To enable Interop API support, navigate to **Orchestrator > Software & Setup > Setup > Advanced Properties** and set the *enableLegacyApisSupport* property to *true*.

  NOTE: It is recommended that you restart Orchestrator during a maintenance window.
- RBAC settings affect users accessing the Orchestrator UI as well as the Orchestrator REST APIs whether authentication is via login/password (role associated with the user) or via an API key (role associated with the API key).

- For RBAC users, the RBAC filter will be applied to any Interop Orchestrator REST API calls (for releases before 9.3.0).

  NOTE: For non-RBAC users, the RBAC filter is disabled or will not be applied.
- An API endpoint (rbac/legacyApi) has been introduced to add an Interop API pattern that is not already in the Interop list (see Pre 9.3 API Endpoints. You must test and verify that the pattern in the database does not create issues.
- REST request performance will be impacted due to increased latency in request filtering and routing of Interop REST APIs. It is recommended that all scripts be modified to adapt to the Orchestrator 9.3 REST API endpoints.
- You can specify an Interface Role for an interface when other identity pipelines, such as Radius snooping, NACD, or VXLAN/GPID are not present. For example, a VLAN for unauthenticated IoT devices could have the Interface Role set to "IoT." This Interface Role is transported across the fabric and is used to reconstruct the VXLAN GPID role on the other end of the fabric.

For information about enabling Interop API support, see Orchestrator Advanced Properties. For information about using Interop APIs, see Pre 9.3 API Endpoints.

The following table describes the information displayed on the RBAC tab for each user who has an RBAC role assigned to them.

Field	Description
User	Name of the user.
Appliance Access Group	Appliance access groups assigned to the user.
RBAC Roles	Role Based Access Control (RBAC) roles assigned to the user.

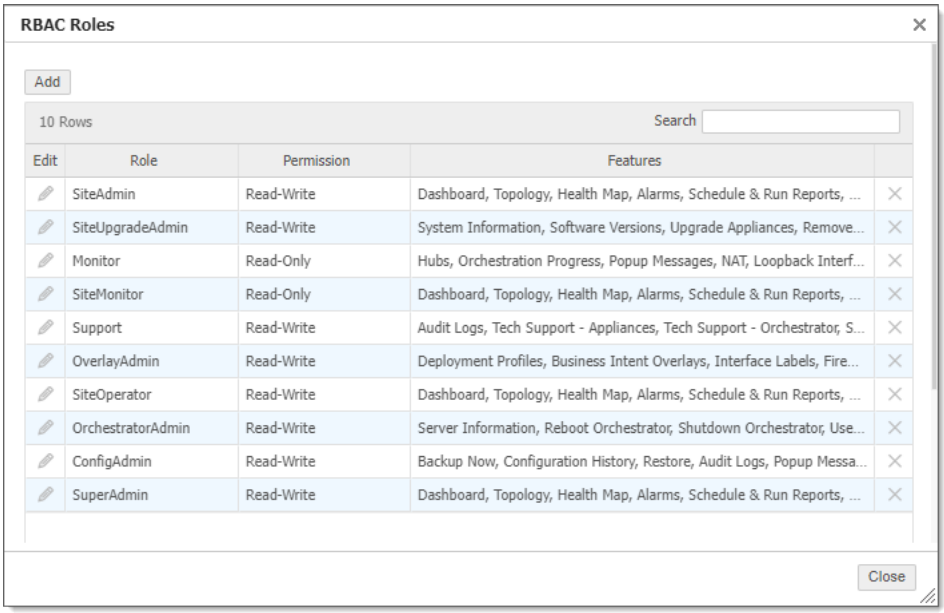
Roles

Orchestrator provides a set of default roles. You can create new roles or modify existing roles. The following table describes the information displayed on the RBAC Roles dialog box for each role.

Field	Description
Role	Name of the role.
Permission	Overall access level assigned to the selected role (Read-Only or Read & Write).
Features	Orchestrator features available to the selected role.

To add a role:

1. Click **Manage RBAC Roles**. The RBAC Roles dialog box opens.



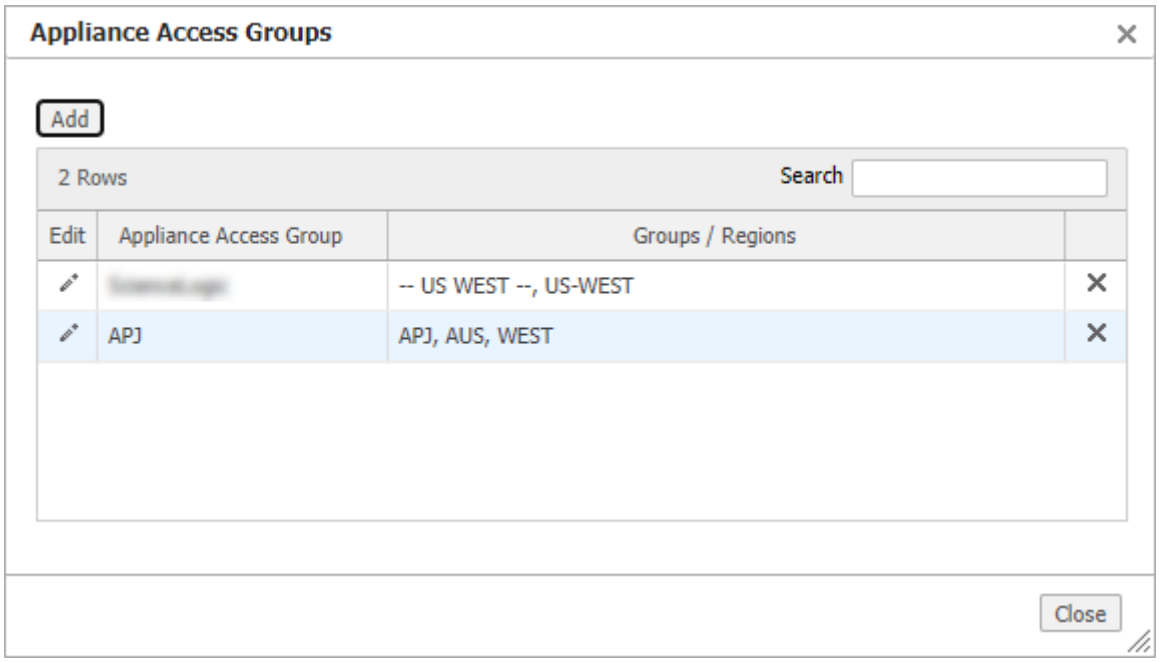
2. Click **Add** to create a new role, or click the edit icon for the role.
3. Enter or modify the role name.
4. Select a category you want to assign to your user from the following tabs: **Monitoring**, **Configuration**, **Administration**, **Orchestrator**, **Support**, or **Miscellaneous**.
5. To assign the overall access level for the role, select **Read Only** or **Read & Write**.
6. Select the check boxes for the Orchestrator menu options you want to make accessible to the role.

   NOTE: You can **Select All** or **Clear All**.
7. Click **Save**.

Appliance Access

With appliance access groups, you can restrict appliance access to one or more groups or regions. Complete the following steps to manage appliance access.

1. On the Role Based Access Control tab, click **Manage Appliance Access Groups**. The Appliance Access Groups dialog box opens.



2. Click **Add** to create a new group, or click the edit icon for an existing group. The Appliance Access Group dialog box opens.
3. Add or modify the name of the appliance access group.
4. Choose how you want to add appliances: **Select Groups** or **Select Regions**. You can manually select groups or regions to include, or use the buttons to select or clear all options.
5. Click **Save**.

WARNING: If an RBAC User does not have appliance access, they will be denied access to Appliance Manager, CLI Session, and Broadcast CLI.

Appliance access and roles determine whether a user is denied or allowed access to the Appliance Manager, CLI Session, and Broadcast CLI. To avoid access issues, ensure that each user has the appropriate appliance access and user roles assigned, as outlined in the following table.

User	Appliance Access	Roles	Access to Appliance Manager, CLI Session, Broadcast CLI
Non-RBAC User	N/A	N/A	Yes
RBAC User	Yes	None assigned	Yes
RBAC User	No	Any	No

Assign Roles and Appliance Access

Complete the following steps to assign roles and appliance access.

1. On the Role Based Access Control tab, click **Assign RBAC Roles & Appliance Access Groups**.
2. In the User field, enter the username of an existing Orchestrator user.
3. In the Appliance Access Group field, select the name of an existing Appliance Access Group.
4. Select the check boxes for one or more roles you want to assign to the user.
5. Click **Save**.

The following table defines the roles provided by default in Orchestrator (roles are listed alphabetically).

Role	Description
ConfigAdmin	Allows user to perform appliance configuration backups, restore appliance configurations, and view the configuration history for an appliance.
Monitor	Provides user read-only access to all menu items.
OrchestratorAdmin	Enables user to perform Orchestrator operations only, such as settings, tools, user management, and Orchestrator upgrades. Appliance operations are not allowed.
OverlayAdmin	Enables user to manage SD-WAN overlays. Overlay management cannot be specific to a site or region. This is a global role.
SiteAdmin	Enables user to manage appliance or site-specific operations, such as configuring appliance-specific policies, ACLs, TCAs, SSL certificates, and upgrades. This role does not allow a user to remove an appliance from the network or perform global SD-WAN functions, such as overlay management or Zscaler orchestration.
SiteMonitor	Grants user read-only permissions equivalent to SiteAdmin.
SiteOperator	Enables user to manage appliance or site-specific operations such as configuring appliance-specific policies, ACLs, TCAs, and SSL certificates. This role does not allow a user to upgrade or remove an appliance from the network or perform global SD-WAN functions, such as overlay management or Zscaler orchestration.
SiteUpgradeAdmin	Allows user to upgrade appliances and remove them from the network.
SuperAdmin	Grants user full read-write access to all menu items.
Support	Grants user access to all support operations.

Authentication

Orchestrator > Orchestrator Server > Users & Authentication > Authentication

Use the Authentication dialog box to manage different remote authentication methods for Orchestrator users.

- To add a new remote authentication method, click **+Add New Server**.
- To view or modify the settings for an existing remote authentication method, click the edit icon in the row of the existing method.

Orchestrator supports the following for remote authentication:

- RADIUS
- TACACS+
- OAuth
- JWT
- SAML

Configure a RADIUS or TACACS+ Server

You need to configure the following when adding or modifying a RADIUS or TACACS+ server.

Field	Description
Attributes	Specify the privilege level/role by entering the attributes on the RADIUS or TACACS+ server: **RADIUS:** In the attributes section, select the **cisco-av-pair** and enter *LOGIN:priv-lvl=[#]*. Level 7 and above equates to the admin role; level 6 and below is the monitor role. Optionally, for RBAC roles, enter *LOGIN:rbac-roles=<RBAC role>*, and for RBAC appliance access groups, enter *LOGIN:rbac-aag=<RBAC aag>*. **TACACS+:** In the custom attributes text box, enter *role=<authorization level>*. Valid authorization values are admin, monitor, and manager. Optionally, for RBAC roles, enter *rbac-roles=<RBAC role>*, and for RBAC appliance access groups, enter *rbac-aag=<RBAC aag>*.
Read-Write Privilege	**RADIUS only:** Lowest value at which a user has Read-Write privileges. This value must be the same as the value configured on the RADIUS server.
Read-Only Privilege	**RADIUS only:** Lowest value at which a user has Read-Only privileges. This value must be the same as the value configured on the RADIUS server.

Field	Description
Authentication Type	Select the authentication type that matches what is configured on the RADIUS or TACACS+ server.
Default Role	If RBAC is enabled, you must specify a default role.
Primary/Secondary Server	For each server in use, enter the IP address or hostname, port, and secret key of the RADIUS or TACACS+ server.

Authenticate Using RADIUS or TACACS+

1. Select the access control protocol you want to use.
2. Under Servers, enter the information for a Primary server of that type. Entering a Secondary server is optional.

Field	Description
Authentication Order	Whether to use the remote map or the local map first. The default is Remote first.
Primary/Secondary Server	IP address or hostname of the RADIUS or TACACS+ server.
Secret Key	String defined as the shared secret on the server.
Read-Write Privilege	Lowest value at which a user has Read-Write privileges. This value must be the same as the value configured on the RADIUS server.
Read-Only Privilege	Lowest value at which a user has Read-Only privileges. This value must be the same as the value configured on the RADIUS server.
Authentication Type	When configuring to use the TACACS+ server, from the drop-down list, select the type that matches what is configured on the TACACS+ server.

Configure an OAuth Server

Orchestrator supports remote authentication via the OAuth 2.0 framework. Before configuring an OAuth server in Orchestrator, you must register Orchestrator as an application with your OAuth provider.

Prerequisites

- The OAuth server must support OAuth 2.0 authorization codes, ID tokens, and (optionally) refresh tokens.
- The ID token is used to get username, RBAC roles, and RBAC appliance access groups.
- The refresh token can be checked periodically to ensure that the user is still authorized.

- Depending on the OAuth server configuration, refresh tokens can be permanent or they can expire. If a token is revoked or expires, the user is forced to authenticate again.

Register Orchestrator as an App

Before adding an OAuth server in Orchestrator, register a new app on your OAuth server for Orchestrator. Provide the following details when registering the app.

Needed Information	Description
Application Type	Register Orchestrator as a Web App.
Allowed Grant Types	Authorization code (required). Refresh token (optional).
Redirect URL	Orchestrator endpoint to which the user is redirected after successful authentication, which should be https://<*Orchestrator_domain_or_IP_address*> /gms/rest/authentication/oauth/redirect.

Configure OAuth Server Properties in Orchestrator

When adding a new OAuth server or modifying an existing server, configure the following fields in the Remote Authentication Server dialog box.

Field	Description
Name	Name to identify the server. This name is displayed on a button on the Orchestrator login page as an alternative method of authentication.
Client ID	Client ID for the Orchestrator application you created in your OAuth provider.
Client Secret	Client secret for the Orchestrator application you created in your OAuth provider.
Scopes	OAuth 2.0 uses scope values, as defined in RFC 6749, to specify which access privileges are being requested for in Access Tokens. The default scopes for Orchestrator are openid, offline_access, and email.
Authentication URL	The Issuer Identifier URL with the authentication request path appended. For example: https://<*your-oauth-domain*> /oauth2/v1/authorize.
Token URL	The Issuer Identifier URL with the token path appended. For example: https://*your-oauth-domain*/oauth2/v1/token.

Field	Description
Username key	The OAuth attribute to be sent as the username. If the username is an email address, use **email**. If any other key is used, ensure that it is mapped to the correct scope on the OAuth server.
(Optional) Roles key	Specify a key name, such as ec-roles. This key name must match what is configured in your OAuth provider. This is a user claim sent in the ID token that maps to Orchestrator roles defined in Role Based Access Control (RBAC). For example, the OAuth server attribute userType maps to ec-roles, and the OAuth user in Orchestrator has userType = OverlayAdmin. **NOTE:** You must specify either **Roles key** or **Default role**. You can specify both if you want to. During Single Sign-On (SSO), if roles and appliance access group keys are not provided, Orchestrator inspects its own configuration to determine the role and appliance access group for the user. If it does not find that information, Orchestrator uses the role specified in the Default role field. If a default role is not specified, the user is not allowed to log in.
(Optional) Appliance Access Group key	Specify a key name, such as ec-aag. This key name must match what is configured in your OAuth provider. This is a user claim sent in the ID token that maps to Orchestrator Appliance Access Groups defined in RBAC. For example, the OAuth server attribute department maps to ec-aag, and the OAuth user in Orchestrator has department = Asia-Admin. **NOTE:** During SSO, if an appliance access group key is not provided or the OAuth server does not return an appliance access group for the user, and if an appliance access group is not assigned to the user in Orchestrator, the user is granted access to all appliances.

Field	Description
(Optional) Default role	If you have not provided a key name in the Roles key field, you must specify a default role. The default role can be any RBAC role, even if RBAC is not enabled. **NOTE:** You must specify either **Roles key** or **Default role**. You can specify both if you want to. During SSO, if roles and appliance access group keys are not provided, Orchestrator inspects its own configuration to determine the role and appliance access group for the user. If it does not find that information, Orchestrator uses the role specified in the Default role field. If a default role is not specified, the user is not allowed to log in.



Configure a JWT Server

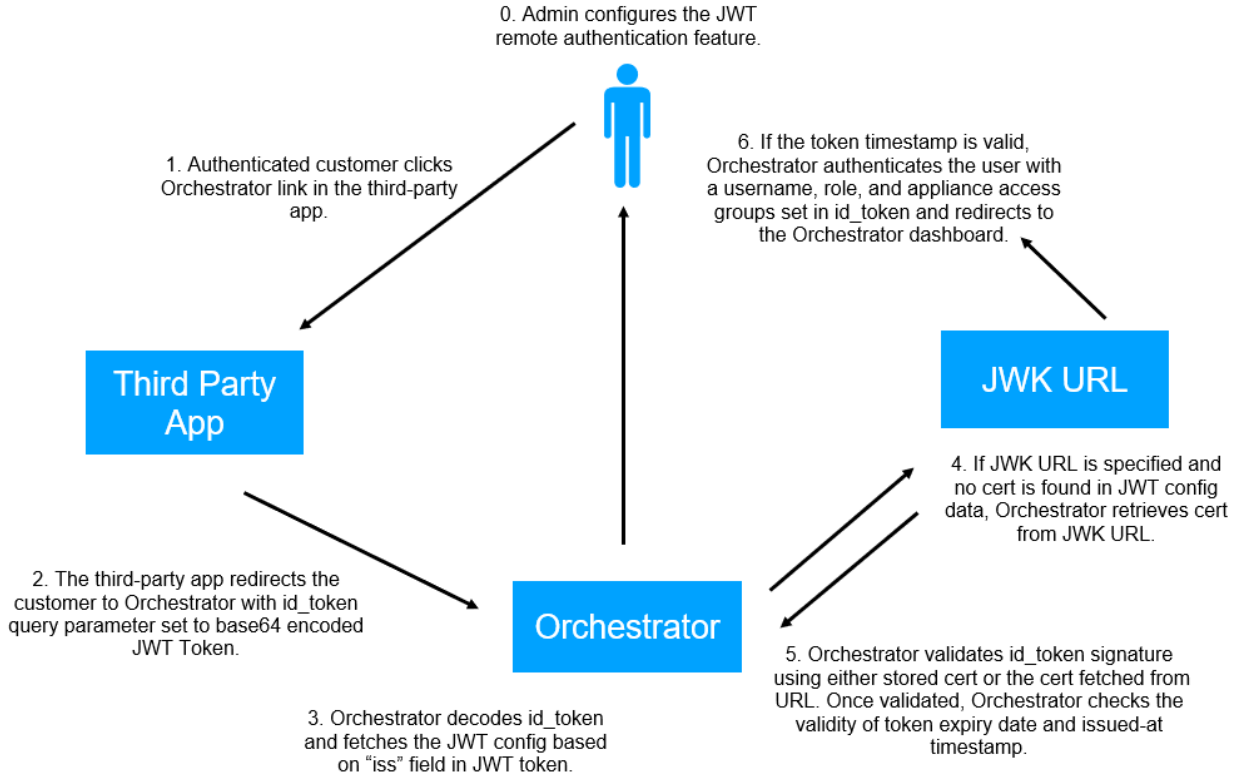
To begin JWT server configuration, the assigned admin must specify the following JWT configuration parameters:

- Issuer ‘iss’
- Auditor ‘aud’
- expiration ‘exp’
- signature
- user, role, and AAG

NOTE: See the following descriptions in the table below.

- Redirect URL based on successful authentication: https://<orchestrator_domainName>?access_token=<token>&id_token=<token>&state=<state>&token_type=Bearer&expires_in=3596

Review the following diagram for more details about the workflow of JWT authentication.



Then, complete the following steps in Orchestrator:

1. Navigate to the Authentication dialog box in Orchestrator (Orchestrator > Users & Authentication > Authentication).
2. Click **+Add New Server**.

   The Remote Authentication Server dialog box opens.
3. From the **Type** drop-down menu, select **JWT**, and then complete the following fields.

Field	Description
Name	Name of your JWT provider.
Cert/Signing Key	HMAC or RSA public key used to verify the id_token.
JWK URL	URL that hosts the public certification.

Field	Description
Validation Window	Maximum amount of time (in minutes) that the expiration is found for the id_token before a new id_token is created.
Issuer	Issuer claim found in the id_token.
Auditor	Auditor claim found in the id_token.
Username key	This attribute is sent as the username. If the username is an email address, use **email**. If any other key is used, ensure that it is mapped to the correct scope on the JWT server.
(Optional) Roles key	Specify a key name, such as ec-roles. This key name must match what is configured in your JWT provider. This is a user claim sent in the ID token that maps to Orchestrator roles defined in Role Based Access Control (RBAC). For example, the JWT server attribute userType maps to ec-roles, and the JWT user in Orchestrator has userType = OverlayAdmin. **NOTE:** You must specify either **Roles key** or **Default role**. You can specify both if you want to. During Single Sign-On (SSO), if roles and appliance access group keys are not provided, Orchestrator inspects its own configuration to determine the role and appliance access group for the user. If it does not find that information, Orchestrator uses the role specified in the Default role field. If a default role is not specified, the user is not allowed to log in.
(Optional) Appliance Access Group key	Specify a key name, such as ec-aag. This key name must match what is configured in your JWT provider. This is a user claim sent in the ID token that maps to Orchestrator Appliance Access Groups defined in RBAC. For example, the JWT server attribute department maps to ec-aag, and the JWT user in Orchestrator has department = Asia-Admin. **NOTE:** During SSO, if an appliance access group key is not provided or the JWT server does not return an appliance access group for the user, and if an appliance access group is not assigned to the user in Orchestrator, the user is granted access to all appliances.

Field	Description
(Optional) Default role	If you have not provided a key name in the Roles key field, you must specify a default role. The default role can be any RBAC role, even if RBAC is not enabled. **NOTE:** You must specify either **Roles key** or **Default role**. You can specify both if you want to. During SSO, if roles and appliance access group keys are not provided, Orchestrator inspects its own configuration to determine the role and appliance access group for the user. If it does not find that information, Orchestrator uses the role specified in the Default role field. If a default role is not specified, the user is not allowed to log in.
JWT token consuming URL	URL of Orchestrator that remains the same.

Configure a SAML Server

Orchestrator supports SAML 2.0 integration, providing authentication and authorization of your credentials through an IdP (Identity Provider), SP (Service Provider), and a Principal. In this example, these are defined as follows:

- **IdP**: Okta
- **SP**: Orchestrator
- **Principal**: A principal end user

SAML and Orchestrator Configuration

Use the following instructions to complete SAML and Orchestrator integration.

TIP: It is recommended to have Orchestrator open next to your Okta window while completing these instructions.

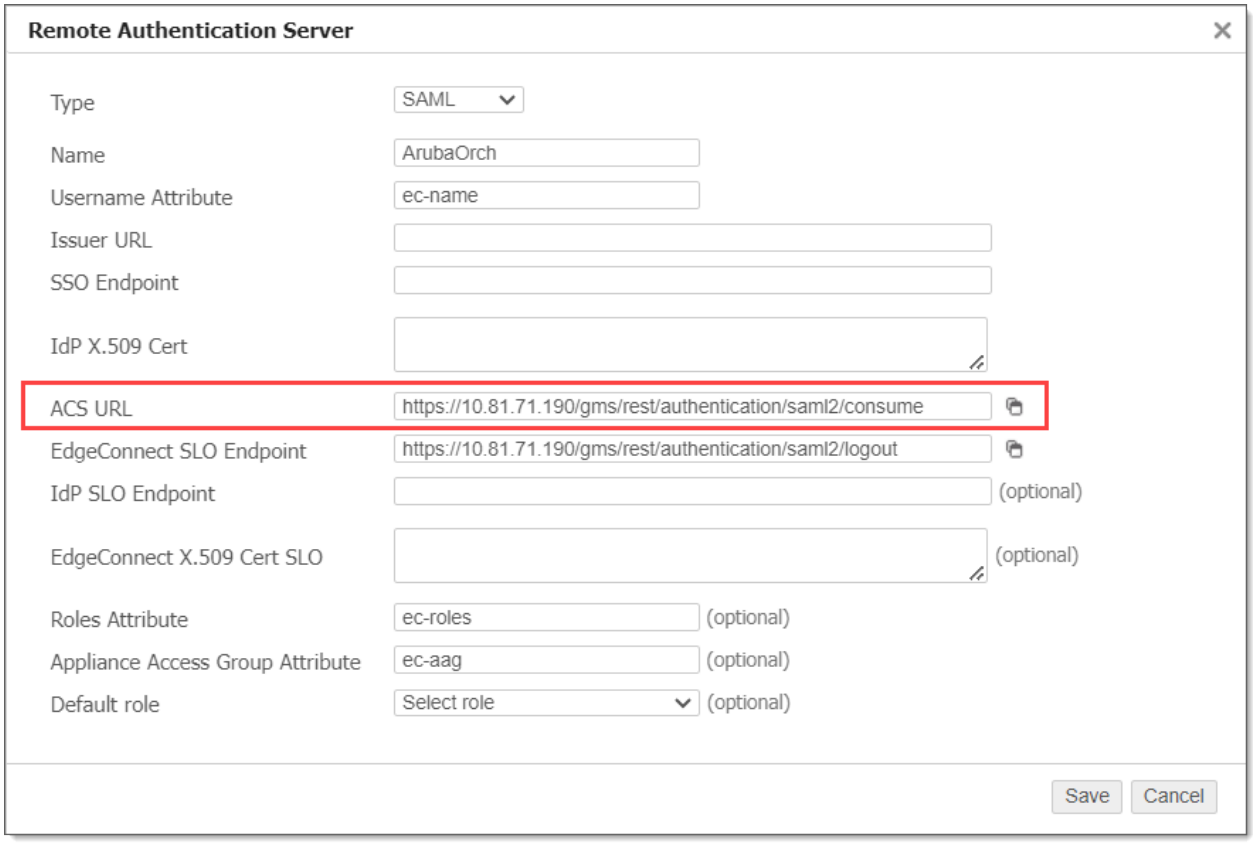
1. Sign in to your Okta account.
2. From the Okta Dashboard menu, click to expand **Applications**, and then click **Applications**.
3. Click **Create App Integration**, and then select **SAML 2.0**.
4. Click **Next**.

   The General Settings dialog box opens.
5. Enter a name in the App name field.
6. Click **Next**.

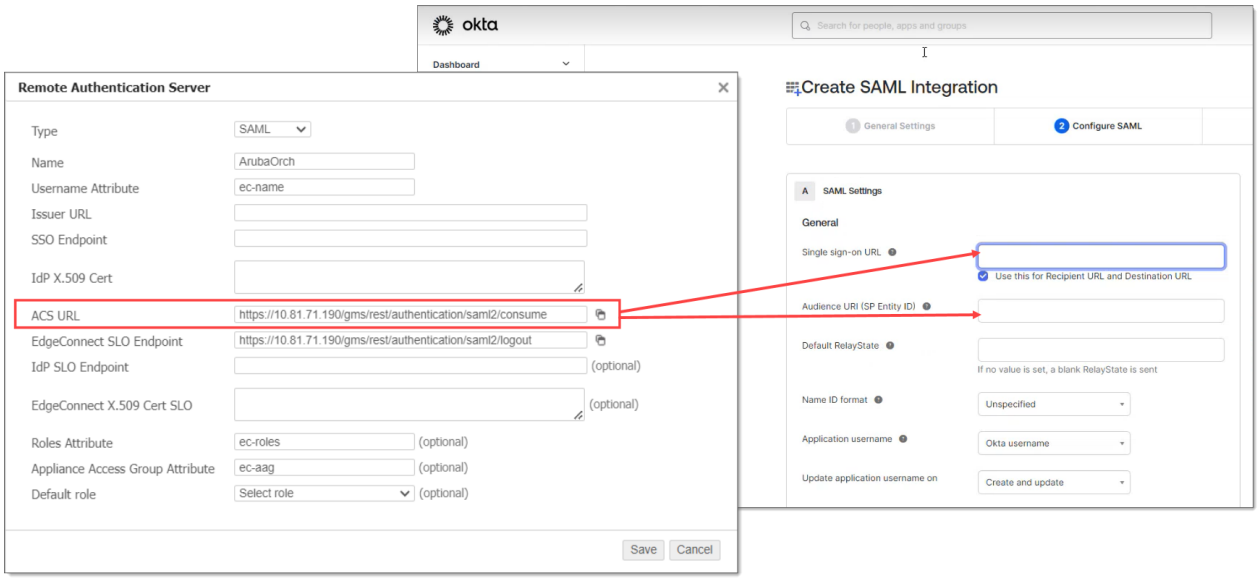
   The SAML Settings dialog box opens.
7. Sign in to Orchestrator and navigate to the Authentication dialog box (Orchestrator > Users & Authentication > Authentication).

8. Click **+Add New Server**.
9. Select **SAML** from the Type field.
10. Enter a **Name** for the server.
11. In Orchestrator, click the copy icon next to the ACS URL field.

    A message at the bottom of the screen notifies you that text is copied to the clipboard.



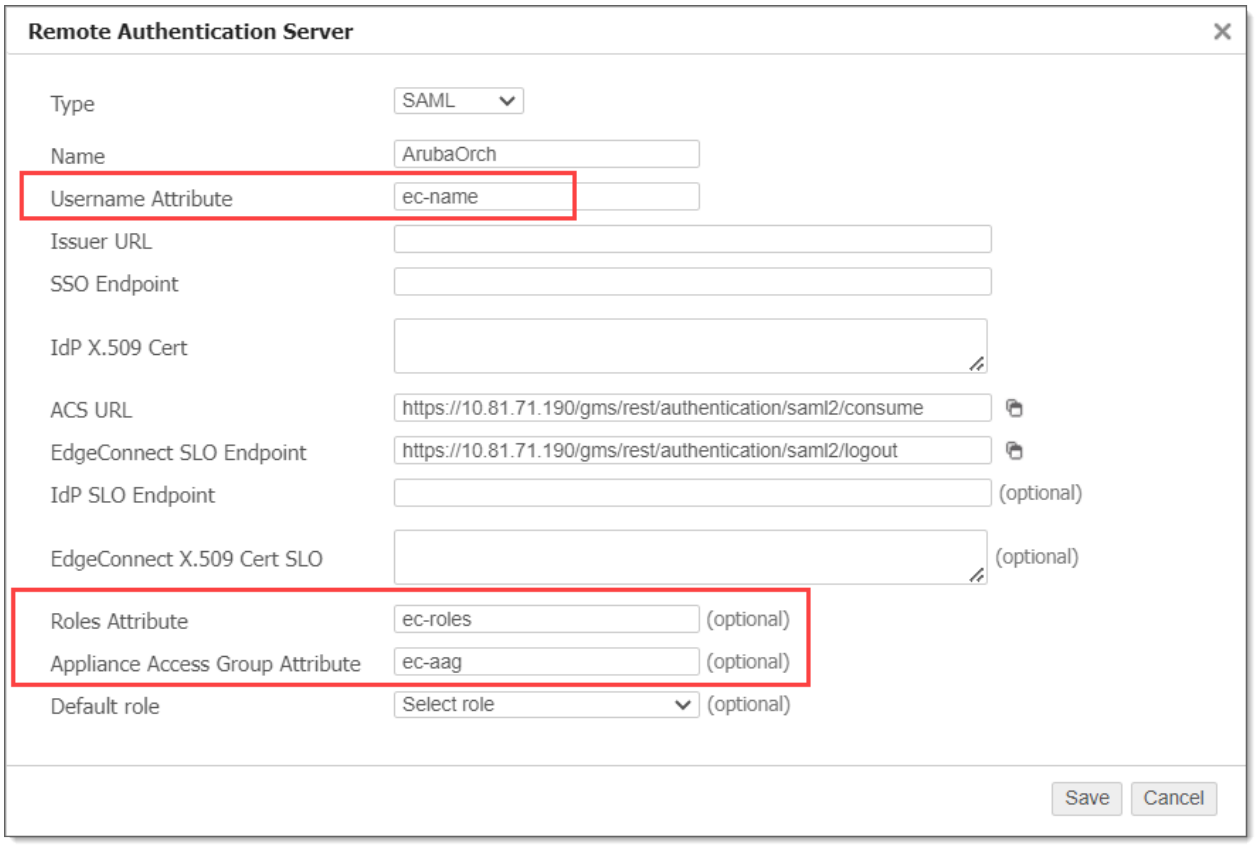
12. Navigate to the Okta SAML Settings dialog box.
13. Paste the ACS URL into both the Single sign-on URL field and the Audience URI (SP Entity ID) field.



14. In Orchestrator, configure the following attributes on the Remote Authentication Server dialog box.

    NOTE: For each attribute, you must click in the field and enter the text for the value. If you click **Save** without entering any text, no value will be saved and the field remains empty.

Attribute Field Name	Example Value
Username Attribute	ec-name
Roles Attribute	ec-roles
Appliance Access Group Attribute	ec-aag



15. Navigate to the Okta SAML Settings dialog box, and in the Attribute Statements (optional) section, create the following attributes and corresponding values.

Name	Example Value
ec-name	user.email
ec-roles	user.usertype
ec-aag	user.department

NOTE: The roles that these claims map to in Orchestrator are configured and assigned to users on the RBAC tab in Orchestrator (Orchestrator > Orchestrator Server > Users & Authentication > Role Based Access Control (RBAC)).

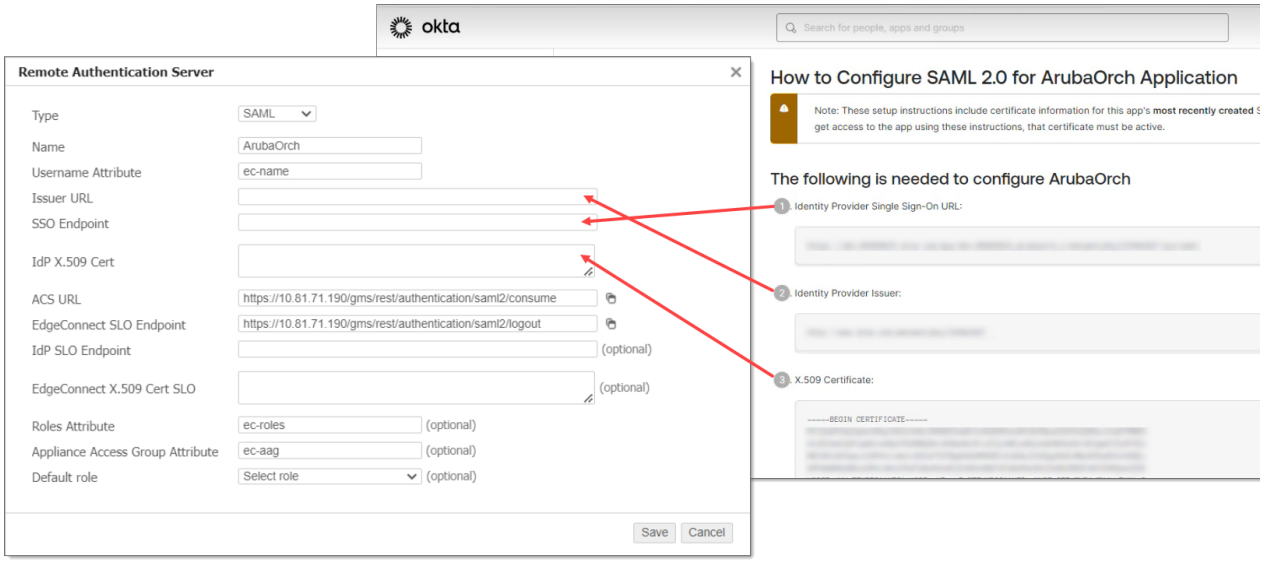
16. Make sure the attribute names match in both applications.

NOTE: The attribute names you enter in the SAML provider (Okta) must exactly match the attribute names you created in Orchestrator, including character case.

17. Click **Next**.
18. Click **Finish**.
19. On the completed SAML settings page, click **View SAML setup instructions**.

20. Copy and paste the URLs from the following Okta fields into the corresponding Orchestrator fields.

Okta Field	Orchestrator Field
Identity Provider Single Sign-On URL	SSO Endpoint
Identity Provider Issuer	Issuer URL
X.509 Certificate	IdP X.509 Cert



21. In Orchestrator, click **Save** to save the changes you made on the Remote Authentication Server dialog box.

    A message at the bottom of the screen notifies you that you have successfully created SAML server configuration.

The following table provides more details about the fields in Orchestrator.

NOTE: Okta supports assigning the attributes at the group level. Refer to either Okta documentation or your Okta administrator for more information.

Field	Description
Name	Any text value for your SAML account for identification purposes.
Username Attribute	Retrieves the username from the SAML XML response.
Issuer URL	Unique identifier of the issuer (for example: Okta, OneLogin).

Field	Description
SSO Endpoint	Unique endpoint for the SAML application created on the IdP server.
IdP X.509 cert	Certificate issued by IdP to verify and validate the response received from the IdP (Okta) server.
ACS URL	Orchestrator endpoint needed for configuration on the IdP server. This is provided as a redirect URL after you are authenticated on the IdP server.
(Optional) EdgeConnect SLO Endpoint	Endpoint used by IdP to initiate the logout request from Orchestrator to the IdP server.
(Optional) IdP SLO Endpoint	Endpoint used by IdP to initiate the logout request from Orchestrator to the IdP server. Endpoint used by Orchestrator to initiate the logout request to IdP.
(Optional) EdgeConnect X.509 Cert SLO	Certificate used by IdP to verify the Single Logout request initiated by Orchestrator to log out the IdP.
(Optional) Roles Attribute	Specify a key name, such as ec-roles. This key name must match what is configured in your SAML provider. This is a claim sent to Orchestrator that maps to roles defined in Role Based Access Control (RBAC). **NOTE:** You must specify either **Roles Attribute** or **Default role**. You can specify both if you want to. During Single Sign-On (SSO), if roles and appliance access group keys are not provided, Orchestrator inspects its own configuration to determine the role and appliance access group for the user. If it does not find that information, Orchestrator uses the role specified in the Default role field. If a default role is not specified, the user is not allowed to log in.

Field	Description
(Optional) Appliance Access Group Attribute	Specify a key name, such as ec-aag. This key name must match what is configured in your SAML provider. This is a claim sent to Orchestrator that maps to Orchestrator Appliance Access Groups defined in RBAC. **NOTE:** During SSO, if an appliance access group key is not provided or the SAML server does not return an appliance access group for the user, and if an appliance access group is not assigned to the user in Orchestrator, the user is granted access to all appliances.
(Optional) Default role	If you have not provided a key name in the Roles Attribute field, you must specify a default role. The default role can be any RBAC role, even if RBAC is not enabled. **NOTE:** You must specify either **Roles Attribute** or **Default role**. You can specify both if you want to. During SSO, if roles and appliance access group keys are not provided, Orchestrator inspects its own configuration to determine the role and appliance access group for the user. If it does not find that information, Orchestrator uses the role specified in the Default role field. If a default role is not specified, the user is not allowed to log in.

API Key

Orchestrator > Orchestrator Server > Users & Authentication > API Keys

Use this page to allow your applications to utilize REST APIs without session authentication and management. You can specify permissions, status, name, and IP allow list for your API keys.

An API key can be passed either in the HTTP request header field *X-Auth-Token* or as a header parameter *apiKey*.

NOTE: It is recommended to use different keys for different applications and users.

To add and define a new API key, click the **Edit** icon and configure the fields below.

Field	Description
Key Name	Name of the key you are creating.
Key	Text you cut, paste, and insert into your client code.
Permission	Read-Only or Read-Write.
Description	Enter details in this field that describe the purpose of the key you are configuring.
Expiration	Set the expiration date if you want a certain application or script to access the key for a fixed amount of time.
Active	To display if the key is active or inactive, select **Yes** or **No**.
IP Allow List	Filters traffic to your private resources through this specified IP range. Traffic is able to pass through with the IP addresses defined in this field.

Create, Change, or Reset Password

Orchestrator > Orchestrator Server > Users & Authentication > Change Password

The Change Password menu is only applicable if you are logged in as a local user. If so, use the Create Password dialog box to create or change your password.

NOTE: If you are logged in through an identify provider via OAuth or SAML, the change password feature is not applicable.

To reset your password, see Reset Password. Both local and OaaS users can reset their passwords.

Create or Change Password

To create or change your password:

1. In the **New Password** field, enter your new password.

   The password must adhere to the following requirements:

   - Must be 8 to 64 characters. The default minimum length is 8. Your administrator might have changed the minimum password length requirement.
   - At least one uppercase letter
   - At least one lowercase letter
   - At least one digit
   - At least one special character

   If your Orchestrator administrator has enabled password criteria for all users, your password must comply with the following additional requirements:

   - Your password cannot be repeated. Your password cannot match your previous history of passwords. (The default history setting is 8 previous passwords.)

- Your password cannot contain your username.
- Your password cannot use well known common phrases.

2. In the **Confirm New Password** field, enter the password again.
3. Click **Save Password**.

   Orchestrator logs you out. Log in again with your new password.

Reset Password

To reset your password:

1. If you are logged in to Orchestrator, log out.
2. Click **Reset Password / QR Code** on the login screen.

   The Reset Password / QR Code dialog box opens.
3. Enter the username or email address associated with the account.
4. Click **Request Password / QR Code Reset**.

   An email is sent to the associated account with instructions on how to reset your password.
5. Click **Close**.
6. Follow the steps in the email to reset your password.
7. Log in to Orchestrator with your new password.

Cloud Portal

Configuration > Overlays & Security > Licensing > Cloud Portal

Orchestrator > Orchestrator Server > Licensing > Cloud Portal

The Cloud Portal dialog box is used to register cloud-based features and services, such as SaaS optimization and EdgeConnect.

NOTE: Orchestrator 9.5.2 and later support IPv6 and IPv4. To support IPv6, the Cloud Portal URL changed to portal2.silverpeak.cloud. All EdgeConnect appliances on the network must be at least on ECOS 9.5.2.x or a later release to use this option. If you have EdgeConnect appliances running on releases earlier than ECOS 9.5.2.x or you have an environment with a mix of releases (some appliances are on ECOS 9.5.2.x and later releases and some are on releases earlier than ECOS 9.5.2.x), you must use portal.silverpeak.cloud. See the Orchestrator 9.5.2 release notes for more details.

NOTE: If you enabled Air-Gap mode, you cannot access the Cloud Portal. You will log in to the Air-Gap Portal to generate a new account key, which you will provide on this dialog box. Follow the instructions for enabling Air-Gap and accessing the Air-Gap Portal here: HPE Aruba Networking Orchestrator Air-Gap User Guide.

- When you purchase one of these services, an Account Name and instructions to obtain your Account Key are sent to you. You will use these to register your appliances.

- Use of these services requires that your appliances can access the Cloud Portal via the Internet.
- You can require hardware appliances to be provisioned with the Account Name and Account Key in order to be discovered.
- You can register a secondary account using the **Secondary Accounts** button.
- If you subscribed to an AAS License, the license type is displayed. The AAS license subscription determines which features are available in Orchestrator. The license token is refreshed from Cloud Portal every 24 hours.
- The Orchestrator UUID (universally unique identifier) is available under the Registration heading on the Cloud Portal dialog box.

Secondary Accounts

Secondary accounts allow you to manage multiple license end dates across a single SD-WAN, and can be used in the following situations:

- You want to add an account that has a different end date from the existing primary account in your Orchestrator.
- You want to merge accounts from one Orchestrator to another Orchestrator and have one SD-WAN fabric to manage.

NOTES:
- Hardware assets must co-reside in the same secondary account where the license resides.
- It is recommended to co-terminate licenses when possible for ease of management.

The primary account is where Orchestrator resides, and secondary accounts are associated to Orchestrator through the registration process. Hardware must reside in the same account as the software licenses associated with that hardware. You can register up to 11 secondary accounts.

As with a primary account, an Account Name and instructions to obtain your Account Key are sent to you. You will use these to register your appliances to the secondary account. To add a secondary account click **Secondary Accounts**, then click **+Add**, enter the account information, and click **Add**. To register the added secondary account, click **Save**.

Any secondary account changes, such as adding, updating, or deleting accounts, will trigger the registration process for all secondary accounts, even for those not modified directly in Orchestrator. This will be evident when a spinning icon is displayed next to “Registered” after saving any changes. After the registration process is complete, the spinning icon will be replaced with registration information that Cloud Portal returns to Orchestrator.

Air-Gap Tab

Orchestrator > Orchestrator Server > Licensing > Air-Gap

Use the Air-Gap tab to enable Air-Gap mode and complete Air-Gap registration.

The HPE SD-WAN Orchestrator Air-Gap solution provides a deployment option that allows you to deploy HPE Aruba Networking EdgeConnect SD-WAN in an isolated network environment with no internet connectivity. With the EdgeConnect SD-WAN product running in Air-Gap mode, appliance activation, license management, and other supporting functions are performed through the Air-Gap Portal where authorized customers can manually exchange information between HPE Aruba Networking Cloud Portal and the customer's EdgeConnect SD-WAN deployment, removing the requirement of an internet connection.

The HPE Aruba Networking Orchestrator Air-Gap solution is only available for self-hosted Orchestrator deployments.

IMPORTANT: Once Air-Gap mode is enabled, it cannot be disabled. Only Silver Peak Support can disable Air-Gap mode.

Before You Begin

You must purchase the Air-Gap mode deployment option before you can deploy EdgeConnect SD-WAN in Air-Gap mode. Once purchased, HPE Aruba Networking Operations will establish an account for you on the Air-Gap Portal. Your licenses and hardware serial numbers will be updated in the HPE Aruba Networking Cloud Portal. HPE Aruba Networking Operations will send an email invitation to the designated contact(s) to set up a user account for the Air-Gap Portal. This email will include a link for resetting passwords and accessing the Air-Gap Portal, allowing the designated contacts to reset their passwords and gain access to the portal.

After the user obtains access to the Air-Gap Portal, they will physically move ("sneakernet") the data using removable media from the Air-Gap Portal to their isolated SD-WAN network. The user must enable Air-Gap mode on the Orchestrator and then register the Orchestrator to the Air-Gap Portal. You will use this tab to enable Air-Gap mode on Orchestrator, complete Air-Gap registration, and upload the license file and supporting file to the Orchestrator within your isolated network.

Complete instructions for enabling Air-Gap and accessing the Air-Gap Portal are here: HPE Aruba Networking Orchestrator Air-Gap User Guide.

Enable Air-Gap Mode

To enable Air-Gap mode:

1. Navigate to **Orchestrator > Orchestrator Server > Licensing > Air-Gap**.

   The Air-Gap tab opens.

2. Select **Enable Air-Gap mode**.

   The Enable/Disable Air-Gap Mode dialog box displays.

3. Click **Confirm**.

   The Air-Gap tab refreshes.

Register Air-Gap to Orchestrator

To register Air-Gap to Orchestrator:

1. Click **Air-Gap Registration**.

   The Air-Gap Registration dialog box opens.
2. Click **Show Registration Key**.
3. Click **Copy Registration Key**. You will need this key to register your account in the Air-Gap Portal.
4. Leave this dialog box open and navigate to the Air-Gap Portal to obtain a response.
5. Paste the response from the Air-Gap Portal in the Portal response field, and then click **Save Portal Response** to complete your registration.
6. Click **Close**.

   Air-Gap registration is complete. You can now assign licenses to all appliances registered to your account. Navigate to the Air-Gap Portal to complete this task.

Upload Air-Gap Files

To upload the license file and supporting files:

1. Click **Air-Gap File Upload**.

   The Air-Gap File Upload dialog box opens.
2. Click **Upload license file**.

   The file Open dialog box opens.
3. Navigate to and select the license file you downloaded in the Air-Gap Portal, and then click **Open**.
4. Click **Upload supporting file**.

   The file Open dialog box opens.
5. Navigate to and select the supporting file you downloaded in the Air-Gap Portal, and then click **Open**.
6. Click **Close**.

Audit Logs

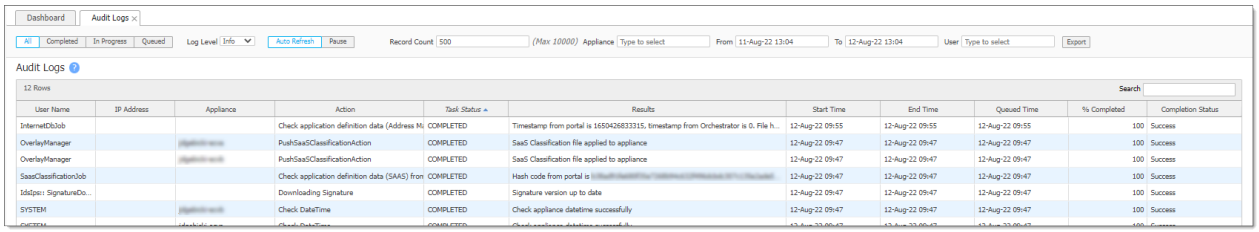
Orchestrator > Orchestrator Server > Tools > Audit Logs

The Audit Logs tab lists actions from a user or the system itself, initiated by Orchestrator.

You can apply the following filters to your audit logs:

- To determine which actions you want to display in the table, select **Completed**, **In Progress**, or **Queued** filters.
- Select the following different log levels to apply to your filter: **Debug**, **Info**, or **Error**.
- To refresh or pause the table, select either **Auto Refresh** or **Pause**. By default, the table refreshes automatically.

- Enter the **Record Count** to limit the filtering criteria. The default value is **500**, and the maximum value is **10,000**.
- Select the name of the **Appliance** from the lists to apply as a filter.
- You can search a wild card character (*) as a username to display all user logs. If you enter any value in the user field, no filter is applied to the search. The following are true for audit log wild cards:
  - x*= anything that starts with the entered value
  - *x= anything that ends with the entered value



Field	Description
User Name	Filter/search for an audit log by the username of the appliance.
IP Address	IP address of the selected appliance.
Appliance	Name of the appliance the audit log comes from.
Action	The action that was taken by the user or the system, initiated by Orchestrator.
Task Status	Status of the audit log task.
Results	Contains a brief description of the audit log including any actions taken. If the audit log refers to template changes or segment (VRF) firewall zone policy changes, any comments entered in the **Audit Log Comment** field will be included in the description. Click the cell to view the full description.
Start Time	Time when the search of the audit log started.
End Time	Time when the search of the audit log ended.
Queued Time	Time when the process/task was requested or scheduled in the queue.
% Completed	Percent completed of the audit log task.
Completion Status	Whether the task has been completed.

Orchestration Settings

Orchestrator > Orchestrator Server > Tools > Orchestration Settings

The Orchestration Settings dialog box manages Business Intent Overlays (BIOs) and the properties that control them. It builds new tunnels and fixes existing ones.

Field	Description
Orchestrate appliances by applying and updating overlays	Selected by default, this updates all associated appliances when overlay changes are saved. **NOTE:** Disabling orchestration stops overlays from being applied and updated on all appliances. IPSec UDP tunnels will not be created nor updated while orchestration is disabled. Tunnels are rebuilt only if this field is enabled.
Reset all flows	When selected, Orchestrator automatically resets all flows whenever you edit overlays or change policies or priorities. When deselected, the flows can only be reset manually.
Autosave appliance changes	Selected by default, this automatically saves any changes made to an appliance. If you need a time delay for troubleshooting or testing, deselect this option to suspend automatic saving of configuration changes.
Apply templates	When selected, updates all associated appliances when template changes are saved.
Idle time	Amount of time Orchestrator sleeps or is idle between checking for any configuration changes. For normal-sized networks, the recommended idle time is **60** seconds. For smaller networks, the recommended idle time is **30** seconds.
Auto flow re-classify	Specifies how the Overlay Manager waits before surveying the network when configuration changes are not being made.

NOTE: Click the first **Reset to Defaults** button to reset the above settings to the defaults.

IPSec UDP Settings

Field	Description
Default port	By default, BIOs create IPSec UDP tunnels. The default port is **10002**. If necessary, you can configure this for an individual appliance on its System Information page, under System Settings. This is accessible from the appliance's context-sensitive menu in the Orchestrator navigation pane.
Increment port by	Referenced when configuring an EdgeHA pair. When the value is 1000, the second appliance's default port becomes 11002.

NOTE: Click the second **Reset to Defaults** button to reset the IPSec UDP settings to the defaults.

Maintenance Mode

Orchestrator > Orchestrator Server > Tools > Maintenance Mode

You can set maintenance mode on an appliance in two ways. You can:

- Use the menu available from the appliance tree. This method automatically suppresses alarms and pauses orchestration.
- Use the Orchestrator menu to select appliances and specify settings. This method allows you to specify whether to pause orchestration or suppress alarms.

Set Maintenance Mode Using the Menu Available from the Appliance Tree

1. Right-click on one or more appliances in the appliance tree, and then select **Maintenance Mode**.
2. In the Maintenance Mode dialog box, click **OK**.

   Alarms are automatically suppressed, and orchestration is automatically paused for the selected appliances.

Set Maintenance Mode Using the Orchestrator Menu

1. Navigate to **Orchestrator > Orchestrator Server > Tools > Maintenance Mode**.
2. In the Maintenance Mode dialog box, click **Add**.

   The Configure Maintenance Mode dialog box opens.
3. In the **Appliance** field, enter the name of the appliance you want to put into maintenance mode.
4. To pause orchestration, select **Pause Orchestration**.
5. To suppress alarms associated with this appliance while in maintenance mode, select **Suppress Alarms**.
6. Click **OK**.

7. Click **Save**.

The following table describes the fields on the Maintenance Mode dialog box.

Field	Description
Appliance	Name of the appliance you put into maintenance mode.
Alarms	Indicates whether to suppress alarms while the appliance is in maintenance mode.
Orchestration	If paused, all orchestration is paused on the selected appliance, except IPSec UDP Tunnel Key material.
IP	IP address of the appliance in maintenance mode.
Version	Current version of the appliance.

Tunnel Settings Tab

Orchestrator > Orchestrator Server > Tools > Tunnels Settings

This tab enables you to manage properties for tunnels created by Orchestrator. Tunnel settings are controlled on a per-label basis, such as MPLS, Internet, or LTE.

IPSec Suite B Presets

As of version 9.2, Orchestrator provides you with four IPSec Suite B presets, as follows:

- GCM-128
- GCM-256
- GMAC-128
- GMAC-256

Each preset includes a predetermined set of IKE and ESP (IPSec) cryptographic algorithms. By selecting an IPSec Suite B preset, you can streamline the algorithm aspect of your tunnel setup rather than selecting individual algorithms. However, you can select individual algorithms if you want to. To select a preset, use the **IPSec Suite B Preset** drop-down field on the General tab.

The following tables show the IPSec Suite B presets in the header row and provide the associated algorithm setups for the IKEv2 and ESP (IPSec) stages.

IKEv2 Stage

	GCM-128	GCM-256	GMAC-128	GMAC-256
Encryption (Note)	AES-128-CBC	AES-256-CBC	AES-128-CBC	AES-256-CBC
Pseudo Random Function	HMAC-SHA-256	HMAC-SHA-384	HMAC-SHA-256	HMAC-SHA-384

	GCM-128	GCM-256	GMAC-128	GMAC-256
Integrity (IKE Data Authentication)	HMAC-SHA-256-128	HMAC-SHA-384-192	HMAC-SHA-256-128	HMAC-SHA-384-192
Key Exchange (NIST Elliptic Curve Groups)	DH-19 256-bit Prime Size	DH-20 384-bit Prime Size	DH-19 256-bit Prime Size	DH-20 384-bit Prime Size

ESP (IPSec) Stage

	GCM-128	GCM-256	GMAC-128	GMAC-256
Encryption	AES-128-GCM with 16 octet ICV	AES-256-GCM with 16 octet ICV	NULL	NULL
Integrity (Data Authentication)	NULL	NULL	AES-128-GMAC	AES-256-GMAC

Notice in the second table that the encryption and data authentication is done in one step for GCM. For GMAC, there is no encryption.

General Tab

Access the following fields by clicking the General Tab.

General

Field	Description
Mode	Indicates whether the tunnel protocol is **UDP**, **GRE**, **IPSec**, or **IPSec UDP**. The default setting is IPSec UDP. If you select IPSec, you can specify the IKE version on the IKE tab. **NOTES:** - If this field is set to IPSec UDP, it is recommended that you use the AES_256_GCM_16 algorithm, which performs both encryption and authentication, resulting in better performance. - Due to external firewall rules, some users may have issues with SD-WAN fabric tunnels configured in GRE mode. In these cases, it is recommended to use IPSec UDP.

Field	Description
IPSec Suite B Preset	This field is available only if the Mode field is set to IPSec. Select an IPSec Suite B preset if required by the security service (**GCM-128**, **GCM-256**, **GMAC-128**, or **GMAC-256**). The default setting is None. If IPSec Suite B Preset is set to None, no preset is selected, but GCM and GMAC algorithms are available to set independently. If an IPSec Suite B preset is selected, various settings on the IKE and IPSec tabs are configured automatically based on the selected preset.
Auto max BW enabled	When enabled, allows the appliances to auto-negotiate the maximum tunnel bandwidth. Enabled by default.
Auto discover MTU enabled	When enabled, allows the appliances to auto-negotiate the maximum tunnel bandwidth. Enabled by default.
MTU	Maximum Transmission Unit (MTU) is the largest possible unit of data that can be sent on a given physical medium. For example, the MTU of Ethernet is 1500 bytes. MTUs up to 9000 bytes are supported. Auto allows the tunnel MTU to be discovered automatically, and it overrides the MTU setting. This field is not available if the Auto discover MTU enabled check box is selected.
UDP destination port	Used in UDP mode. Accept the default value unless the port is blocked by a firewall.
UDP flows	Used in UDP mode. Number of flows over which to distribute tunnel data.

Packet

NOTE: FEC settings do not apply when overlays are used. FEC settings only apply when routing directly to an underlay via Route Policy.

Field	Description
Reorder wait	Maximum time (in milliseconds) the appliance holds an out-of-order packet when attempting to reorder. **100** ms is the default value and should be adequate for most situations. FEC can introduce out-of-order packets if the reorder wait time is not set high enough.
FEC	Forward Error Correction (FEC) can be set to **enable**, **disable**, or **auto**.
FEC ratio	When FEC is set to **auto**, FEC will range dynamically from off to 1:10 based on detected loss. The options are **1:1**, **1:2**, **1:5**, **1:10**, or **1:20**. This field is available only if FEC is set to enable.

Tunnel Health

Field	Description
Retry count	Number of failed keep-alive messages allowed before the appliance brings the tunnel down.
DSCP	Determines the DSCP marking that the keep-alive messages should use.

FastFail Thresholds

NOTE: FastFail thresholds do not apply when overlays are used. FastFail only applies when routing directly to an underlay via Route Policy.

Field	Description

Fastfail enabled	When multiple tunnels are carrying data between two appliances, this feature determines how quickly to disqualify a tunnel from carrying data.
	
	The Fastfail connectivity detection algorithm for the wait time from receipt of last packet before declaring a __*brownout__* is:
	
	`Twait = Base + N * RTTavg`
	
	where `Base` is a value in milliseconds, and `N` is the multiplier of the average Round Trip Time over the past minute.
	
	For example, if:
	
	`Base = 200mS`
	`N = 2`
	
	then,
	
	`RTTavg = 50mS`
	
	The appliance declares a tunnel to be in brownout if it does not see a reply packet from the remote end within 300 mS of receiving the most recent packet.
	
	In the Tunnel Advanced Options, `Base` is expressed as **Fastfail wait-time base offset** (ms), and `N` is expressed as **Fastfail RTT multiplication factor**.
	
	Fastfail enabled - This option is triggered when a tunnel's keep-alive signal does not receive a reply. The options are **disable**, **enable**, and **continuous**. If the disqualified tunnel subsequently receives a keep-alive reply, its recovery is instantaneous.
	
	For disable, keep-alives are sent every second, and 30 seconds elapse before failover. In that time, all transmitted data is lost.
	
	For enable, keep-alives are sent every second, and a missed reply increases the rate at which keep-alives are sent from one per second to ten per second. Failover occurs after one second.
	
	For continuous, keep-alives are continuously sent at ten per second. Therefore, failover occurs after one tenth of a second.

Field	Description
Latency	Amount of latency in milliseconds. Thresholds for Latency, Loss, or Jitter are checked once every second. Receiving three successive measurements in a row that exceed the threshold puts the tunnel into a brownout situation and flows will attempt to fail over to another tunnel within the next 100 ms. Receiving three successive measurements in a row that drop below the threshold will drop the tunnel out of brownout.
Loss	Amount of data lost as a percentage.
Jitter	Amount of jitter in milliseconds.
Fastfail wait-time base offset	Fastfail basic timeout time in milliseconds.
Fastfail RTT multiplication factor	Amount of RTT (Round Trip Time) added to the basic timeout.

IKE Tab

Access the following fields by clicking the IKE tab. This tab is displayed only if the Mode field on the General tab is set to IPSec.

IKE

Field	Description
Peer Authentication	There are two options for IKE authentication, End entity certificate or Pre-shared key, choose one of the options. End entity certificate is the recommended option. **End entity certificate** – If selected, select a profile from the End entity certificate profile drop-down menu. **NOTE:** To select an orchestrated end entity certificate profile, you must first add an EST server profile and create an appliance end entity profile with a Purpose of “SD-WAN”. To do this, see End Entity Certificates Tab(./Mediate protection against future “harvest now, decrypt later” attacks by quantum computers. For orchestrated tunnels, Orchestrator automatically generates the PPK ID and PPK Value for the PPK. **Post quantum pre-shared key mandatory** – If selected, RFC 8784 must be used.

Field	Description
Authentication algorithm	Authentication algorithm used for IKE security association (SA). If the IPSec Suite B Preset field on the General tab is set to None, you can select **SHA1**, **SHA2-256**, **SHA2-384**, or **SHA2-512**. The default setting is SHA1. If the IPSec Suite B Preset field is set to any other setting, this field is automatically set to the appropriate algorithm. **NOTE:** With IKEv2 and the Encryption algorithm field set to auto, AES-GCM will probably be negotiated, which includes encryption and authentication. In this case, this field might show a SHA setting that is not actually used. If the Encryption algorithm field is set to AES-GCM-128 or AES-GCM-256, this field will show as NA because the authentication algorithm is already included.
Encryption algorithm	Encryption algorithm used for IKE security association (SA). If the IPSec Suite B Preset field on the General tab is set to None, and the IKE Version field is set to IKE v1, you can select **AES-CBC-128**, **AES-CBC-256**, or **auto**. The default setting is auto. If the IPSec Suite B Preset field is set to None, and the IKE Version field is set to IKE v2, you can select **AES-CBC-128**, **AES-CBC-256**, **AES-GCM-128**, **AES-GCM-256**, or **auto**. The default setting is auto. If the IPSec Suite B Preset field is set to any other setting, this field is automatically set to the appropriate algorithm.
Pseudo Random Function	This field is displayed only if the IKE Encryption Algorithm field is set to AES-GCM-128 or AES-GCM-256. For AES-GCM-128, you can select **SHA2-256**, **SHA2-384**, or **SHA2-512**. For AES-GCM-256, you can select **SHA-384** or **SHA-512**.

Field	Description
Diffie-Hellman group	Diffie-Hellman Group used for IKE security association (SA) negotiation. If the IPSec Suite B Preset field on the General tab is set to None, you can select the appropriate group. Available groups are 14 through 21, 26, and 31. If the IPSec Suite B Preset field is set to any other setting, this field is automatically set to the appropriate group.
Rekey interval/lifetime	Rekey interval/lifetime of IKE security association (SA) in minutes. The default is 360 minutes.
Dead peer detection	**Delay time**: Interval (in seconds) to check the lifetime of the IKE peer. **Retry count**: Number of times to retry the connection before determining that the connection is dead. This field is not editable.
Phase 1 mode	Exchange mode for the IKE security association (SA) negotiation. If the IKE Version field is set to IKE v1, you can select **Main** or **Aggressive**. If the IKE Version field is set to IKE v2, this field is automatically set to Aggressive.
IKE version	If the IPSec Suite B Preset field on the General tab is set to None, you can select **IKE v1** or **IKE v2**. If the IPSec Suite B Preset field is set to any other setting, this field is automatically set to IKE v2.

IPSec Tab

Access the following fields by clicking the IPSec tab. This tab is displayed only if the Mode field on the General tab is set to **IPSec** or **IPSec UDP**.

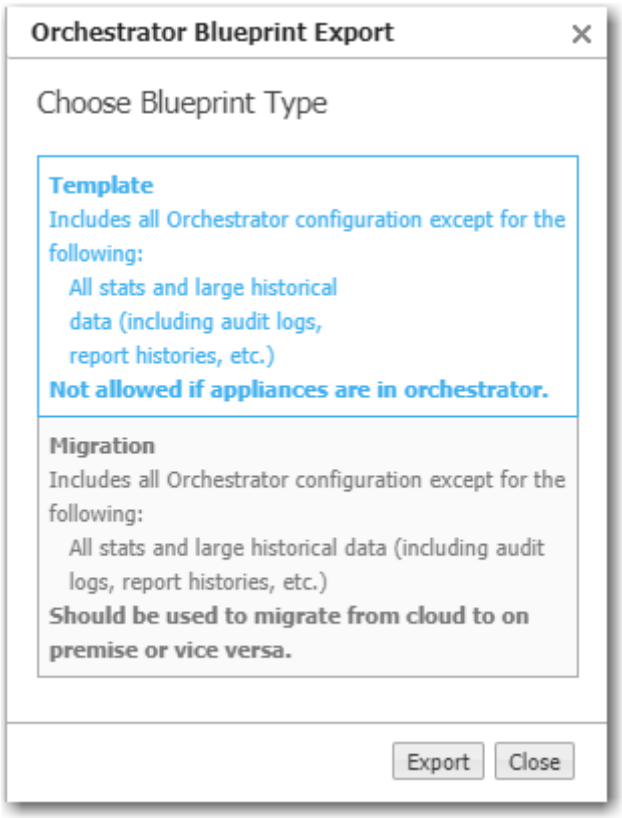
IPSec

Field	Description
Authentication algorithm	Authentication algorithm used for the IPSec security association (SA). If the IPSec Suite B Preset field on the General tab is set to None, you can select **SHA1**, **SHA2-256**, **SHA2-384**, **SHA2-512**, **AES-GMAC-128**, or **AES-GMAC-256**. The default setting is SHA1. If the IPSec Suite B Preset field is set to GMAC-128 or GMAC-256, this field is automatically set to the appropriate algorithm. If the IPSec Suite B Preset field is set to GCM-128 or GCM-256, this field is not applicable.
Encryption algorithm	Encryption algorithm used for the IPSec security association (SA). If the IPSec Suite B Preset field on the General tab is set to None, and the IPSec Authentication algorithm field is set to SHA1, SHA2-256, SHA2-384, or SHA2-512, you can select **AES-CBC-128**, **AEC-CBC-256**, **AES-GCM-128**, **AES-GCM-256**, **NULL**, or **Auto**. The default setting is Auto. If the IPSec Suite B Preset field is set to None, and the IPSec Authentication algorithm field is set to AES-GMAC-128 or AES-GMAC-256, this field is automatically set to NULL.
IPSec anti-replay window	Select a size from the drop-down list or **Disable** to disable the IPSec anti-replay window. If a size is selected, protection is provided against an attacker duplicating encrypted packets by assigning a unique sequence number to each encrypted packet.
Relay interval/lifetime	Rekey interval/lifetime of the IPSec security association (SA) in minutes. The default is 360 minutes.
Perfect forward secrecy group	Diffie-Hellman group used for IPSec security association (SA) negotiation. Based on the setting of the IPSec Suite B Preset field on the General tab, this field is set to the following Diffie-Hellman group: For None: 14 (by default) For GCM-128 or GMAC-128: 19 For GCM-256 or GMAC-256: 20

Orchestrator Blueprint Export

Orchestrator > Orchestrator Server > Tools > Orchestrator Blueprint Export

Use this dialog box to export the current Orchestrator configuration to a blueprint that you can apply to another Orchestrator instance.



You can use a blueprint when creating a new Orchestrator or when migrating an existing Orchestrator to on-prem or cloud.

- Blueprints can only be created from Orchestrators that have no appliances associated with them. If the source Orchestrator manages any appliances, blueprint creation fails.
- You can create and store multiple blueprints with the same Orchestrator.
- After creating as many blueprints as you need, you can add appliances to the source Orchestrator.
- Blueprints automatically exclude all statistics, large historical data files (including audit logs, report histories, and so forth), and account information.

To export an Orchestrator blueprint:

1. In the Orchestrator Blueprint Export dialog box, select the blueprint type: **Template** or **Migration**
    - Blueprint Template mode - In this mode, IPSec UDP key is not included in the exporting. MTO uses Blueprint Template to create new orchestrators.

- Blueprint Migration mode - In this mode, IPSec UDP key is included in the exporting. This works as making clones on an orchestrator.

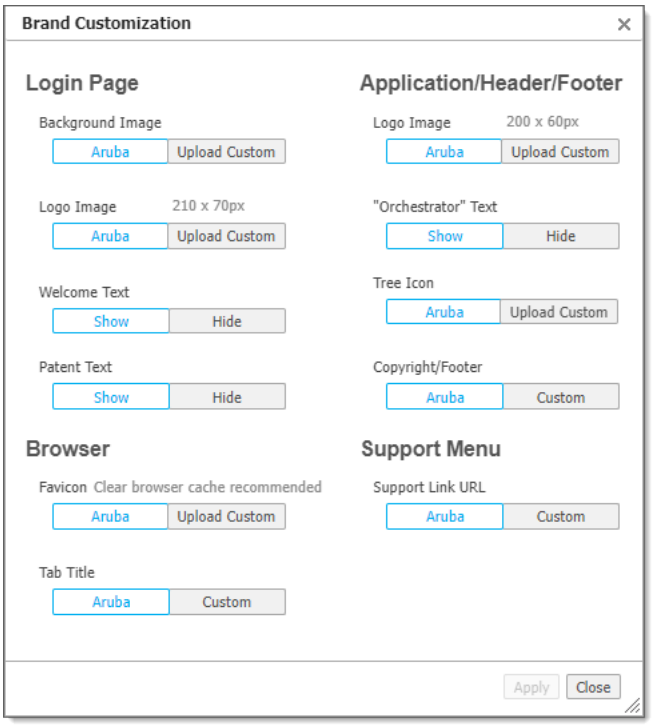
2. Click **Export**. Export downloads an SQL file to your local desktop.

WARNING: This completely replaces the configuration of the existing Orchestrator.

Brand Customization

Orchestrator > Orchestrator Server > Tools > Brand Customization

Use this dialog box to customize the branding aspects of the Orchestrator user interface.



Orchestrator > Software & Setup

The options under **Orchestrator > Software & Setup** focus on configuring the software elements of Orchestrator, including SMTP settings, creating banners, updating/upgrading Orchestrator, and more.

Upgrade Orchestrator Software

You can upgrade Orchestrator software for an on-prem installation or an installation in the cloud.

Upgrade On-Prem Orchestrator

If you are using on-prem Orchestrator 8.6.0 or later and want to upgrade to a newer version, complete the following steps.

WARNING: An upgrade that fails can cause Orchestrator to be in a corrupt state. Back up Orchestrator before you start the upgrade process.

1. Create an SSH shell to Orchestrator.
2. Open an SSH session to Orchestrator.
3. Log in as admin or a user with administrator privileges.
4. Switch to root: `su - root`
5. Enter the root password when prompted. If you do not know your root password, contact Support.
6. Enter `cd /home`

Depending on your environment, you can upgrade Orchestrator in either of the following ways:

NOTE: Upgrading via SCP is only available for versions 8.10.20 and later.

- Upgrade Orchestrator Software
  - Upgrade On-Prem Orchestrator
    - Upgrade via HTTP
    - Upgrade via SCP
  - Upgrade Orchestrator in the Cloud

Upgrade via HTTP

If you have an HTTP URL to the Orchestrator installation file, enter the following in the existing SSH console to run the install script and point it to the hosted installation file:

`/home/gms/gms/setup/install_orchestrator.sh <HTTP URL of the Orchestrator Installation File>`

NOTE: The upgrade process can take several hours to complete.

Upgrade via SCP

If you do not have an HTTP server, copy the Orchestrator installation file to Orchestrator by using SCP, run the install script, and point it to the local installation file.

NOTE: This procedure assumes that the scp programs on both ends are patched for CVE-2020-15778 and/or you trust the remote server from which you will scp the installation file. From the Orchestrator SSH console, enter the following as root:

1. From the Orchestrator SSH console, enter the following: `mv /bin/scp-local /bin/scp #`
2. From your local PC console, do one of the following:

- If you are running Orchestrator release 9.1.9, 9.2.10, 9.3.3, 9.4.x, or later, do the following:
  1. Enter *scp <Orchestrator Installation file> admin@<orchestrator_ip_address>:/home/admin*
  2. From the Orchestrator SSH shell console, enter *mv /home/admin/<Orchestrator Installation file> /home/gms/*
  3. From the Orchestrator SSH shell console, enter *chown gms.gms /home/gms/<Orchestrator Installation file>*
- If you are running an Orchestrator release earlier than those listed above, enter *scp <Orchestrator Installation file> admin@<orchestrator_ip_address>:/home/gms*

3. From the Orchestrator SSH shell console, enter the following: */home/gms/gms/setup/install_orchestrator.sh <full_path_to_Orchestrator_Installation_file>*

NOTE: The upgrade process can take several hours to complete.

Upgrade Orchestrator in the Cloud

Complete the following steps to upgrade Orchestrator in the cloud.

1. Using Orchestrator in the cloud, navigate to **Orchestrator > Software & Setup > Upgrade > Upgrade Orchestrator**.

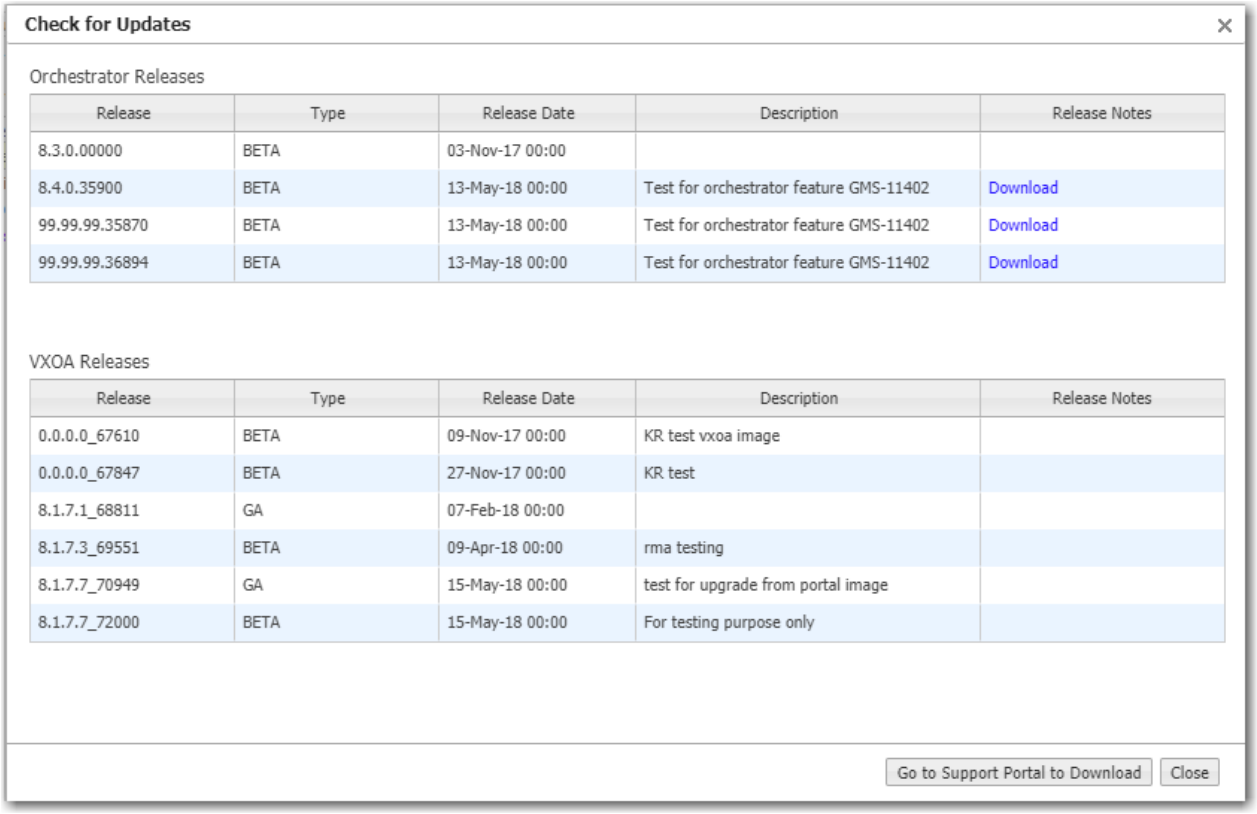
   The Upgrade Orchestrator dialog box opens.
2. Click in the **Select Version** field to display a drop-down list of Orchestrator versions you can upgrade to, and then select the appropriate version.
3. Click **Upgrade**.

   You will receive an email message indicating that the Orchestrator version upgrade has started. Orchestrator service is not available during the upgrade process. When the process completes, you will receive another message indicating that the upgrade was successful. If a failure occurs, however, you will receive a failure message with instructions to contact HPE Networking support.

Check for Orchestrator and Appliance Software Updates

Orchestrator > Software & Setup > Upgrade > Check for Updates

These pages show what appliance and Orchestrator server software is available for download.



Back Up on Demand

Orchestrator > Software & Setup > Backup > Backup Now

Use this dialog box to back up the Orchestrator database on demand.



Schedule Orchestrator Backup

Orchestrator > Software & Setup > Backup > Schedule Backup

Use this dialog box to schedule backups of the Orchestrator database and the Orchestrator Stats Collector, and specify the details for the appliance backup behavior in Orchestrator.

User-Configured Orchestrator Backup (On-Premises Only)

For on-premises Orchestrators, you can configure a scheduled backup job.

- This scheduled backup performs both an Orchestrator backup and an EdgeConnect appliance backup according to the schedule you set.
- When backing up EdgeConnect, the Orchestrator instructs EdgeConnect to generate and send the backup file, which is then stored in the Orchestrator database.

NOTE: Cloud Orchestrator does not support scheduled backup.

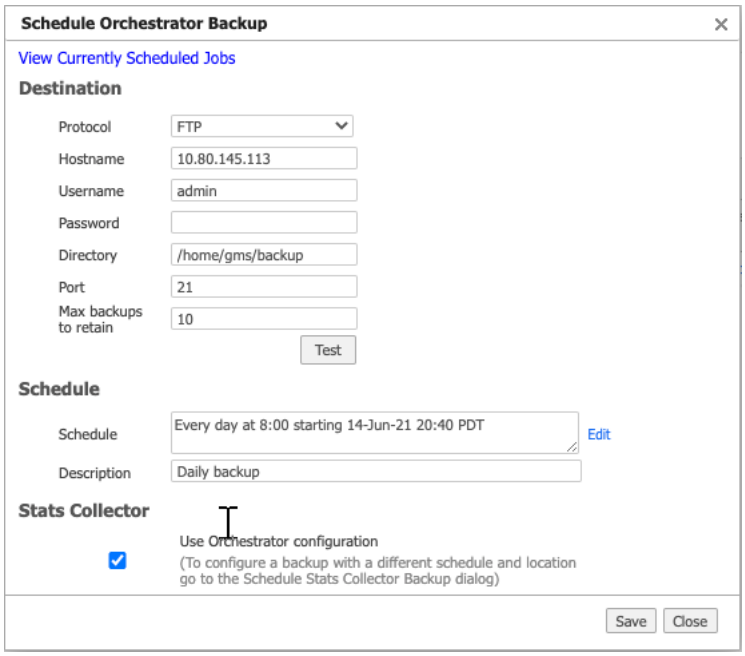
Daily EdgeConnect Backup (All Orchestrators)

The Orchestrator automatically backs up EdgeConnect appliances every day at 12:00 pm. This schedule is not editable.

- This daily backup runs on both on-premises and cloud Orchestrators.
- For each backup, the Orchestrator compares the new backup file with the one in the database and saves the new file only if there is a change.

Scheduling Details

You can view and edit scheduled jobs, set the destination protocol (for example, FTP), and configure backup retention and timing. To configure a backup with a different schedule and location for the Stats Collector, use the **Schedule Stats Collector Backup** dialog box.



Field	Description
View Currently Scheduled Jobs	Click to open the Scheduled Jobs tab.
Protocol	Protocol to apply: **FTP**, **SCP**, **HTTP**, **HTTPS**, or **SFTP**.
Hostname	Host name of the backup server.
Username	Username that the Orchestrator server uses to log in to the backup server.
Password	Password for the username.
Directory	Directory name of the backup server.
Port	Port number of the backup server.
Max backups to retain	Maximum number of backups to retain.
Test	To verify that Orchestrator can reach the destination, click **Test**.

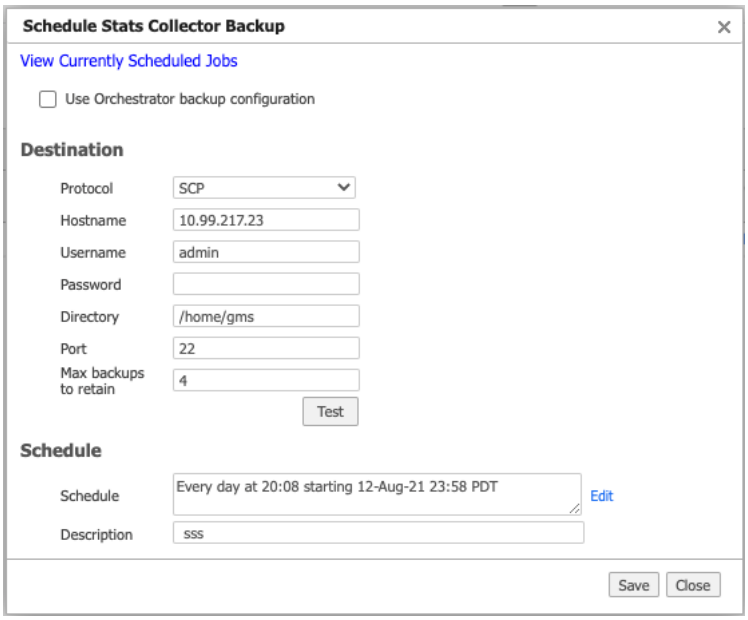
Field	Description
Schedule	To create a schedule, click **Add**. To modify a schedule, click **Edit**. In the Schedule dialog box, select **Daily**, **Weekly**, **Monthly**, or **Yearly**. Complete the remaining fields, and then click **OK**. **TIP:** To specify the timezone for scheduled jobs and reports, navigate to **Orchestrator > Software & Setup > Setup > Timezone for Scheduled Jobs**.
Description	*(Optional)* Description for the backup schedule.
Stats Collector	Do one of the following: Select the **Use Orchestrator configuration** check box to back up the Orchestrator Stats Collector on the same schedule and to the same destination. Clear the **Use Orchestrator configuration** check box to specify a different backup destination and set a different schedule for the Orchestrator Stats Collector. **CAUTION:** If you clear the **Use Orchestrator configuration** check box and you do not complete the Schedule Stats Collector Backup dialog box, the Stats Collector will not be backed up. For more information, see Schedule Stats Collector Backup.

Schedule Stats Collector Backup

Orchestrator > Software & Setup > Backup > Schedule Stats Collector Backup

Use this dialog box to schedule backups of the Orchestrator Stats Collector.

NOTE: The backslash (\) character is not allowed in any field.



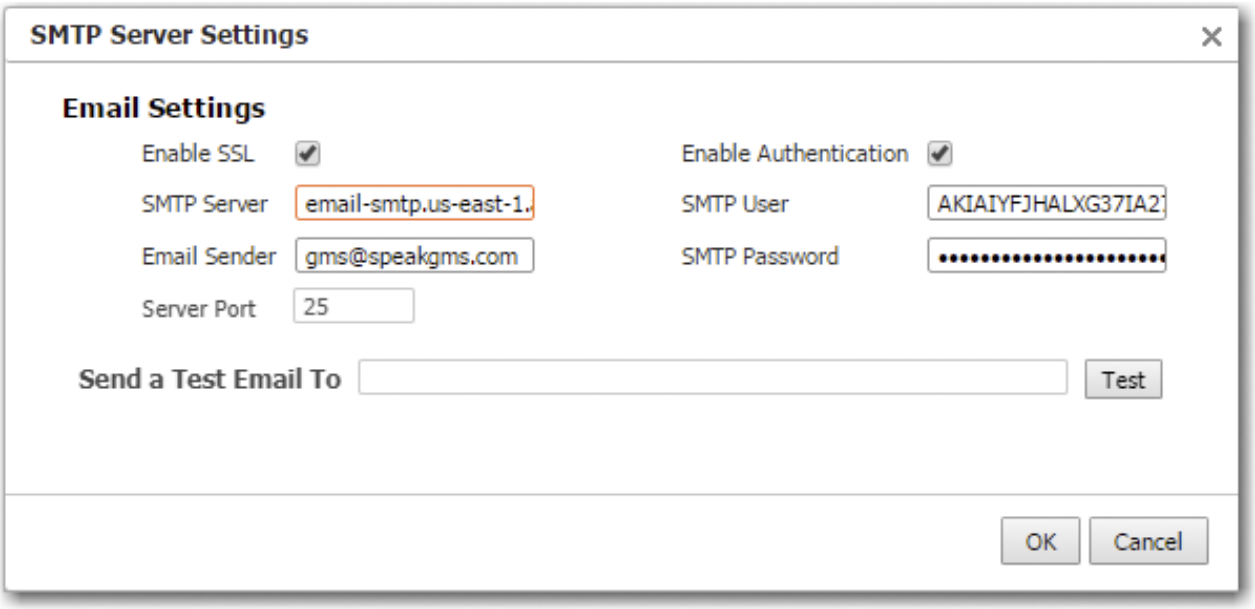
Field	Description
View Currently Scheduled Jobs	Click to open the Scheduled Jobs tab.
Use Orchestrator backup configuration	Select this check box to back up the Stats Collector using the same destination and schedule set in the Schedule Orchestrator Backup dialog box. For more information, see Schedule Orchestrator Backup.
Protocol	Protocol to apply: **FTP**, **SCP**, **HTTP**, **HTTPS**, or **SFTP**.
Hostname	Host name of the backup server.
Username	Username that the Orchestrator server uses to log in to the backup server.
Password	Password for the username.
Directory	Directory name of the backup server.
Port	Port number of the backup server.
Max backups to retain	Maximum number of backups to retain.
Test	To verify that Orchestrator can reach the destination, click **Test**.

Field	Description
Schedule	To create a schedule, click **Add**. To modify a schedule, click **Edit**. In the Schedule dialog box, select **Daily**, **Weekly**, **Monthly**, or **Yearly**. Complete the remaining fields, and then click **OK**. **TIP:** To specify the timezone for scheduled jobs and reports, navigate to **Orchestrator > Software & Setup > Setup > Timezone for Scheduled Jobs**.
Description	*(Optional)* Description for the backup schedule.

SMTP Server Settings

Orchestrator > Software & Setup > Setup > SMTP Server Settings

For permanent and private email delivery, change the SMTP (Simple Mail Transfer Protocol) server and settings to your company's SMTP settings.



- If a test email does not arrive within minutes, check your firewall.
- After configuring the SMTP settings, you can specify email recipients for the following:
  - **alarms** (Monitoring > Alarms > Alarm Email Recipients)
  - **reports** (Monitoring > Reporting > Schedule & Run Reports)

Proxy Configuration

Orchestrator > Software & Setup > Setup > Proxy Configuration

The following two proxy configuration scenarios are discussed below:

- Web proxy
- Orchestrator as a proxy for EdgeConnect onboarding

Web Proxy

Configure EdgeConnect appliances to use a Web Proxy. This configuration addresses security requirements and enables graceful scaling of such traffic.

Configure Orchestrator as a Proxy for EdgeConnect Onboarding

For EdgeConnect onboarding, configure Orchestrator as a proxy to enable Cloud Portal EdgeConnect traffic.

NOTE: Orchestrator as a proxy is not a substitute for using a web proxy. Orchestrator as a proxy is not supported if you have more than 200 appliances.

Go to *Orchestrator > Software & Setup > Setup > Proxy Configuration.*

The following table describes the fields on the Proxy Configuration dialog box.

Field	Description
Use Proxy	Select the Use Proxy check box to enable Orchestrator as a proxy. If you enable Orchestrator as a proxy, see Set Up the Appliance to Use Orchestrator as a Proxy.
Proxy Host	The URL of the proxy host.
Proxy Port	The port of the proxy host.
Proxy User Name (optional)	The user name for the proxy host. This field is optional.
Use Password (optional)	The user password for the proxy host. This field is optional.

Onboarding Appliances Using Orchestrator as a Proxy

If you enable Orchestrator as a proxy for onboarding EdgeConnect appliances, complete the following steps.

1. Boot up the EdgeConnect appliance and retrieve the IP address on your management network.
2. Navigate to the appliance directly using the IP address.
3. Log in to the appliance.

4. In the appliance, navigate to **Administration > Basic Settings > HPE Aruba Networking Cloud Services**.
5. In the Orchestrator field, enter the IP address of the Orchestrator.

   NOTE: Do not update the host or port fields and do not configure the DNS settings on the appliance.
6. Navigate to **Administration > Basic Settings > License & Registration**.
7. Copy and paste the account name and key into the appropriate fields.
8. Click **Apply** and navigate back to the HPE Aruba Networking Cloud Services tab.

   NOTE: If the Management Interface of the EdgeConnect is not on the same subnet as the Orchestrator, navigate to the Management Routes tab and add a route to the Management Route table.
9. Refresh the screen after a few minutes. The following alarm should display: “Appliance Registered but not Approved.”
10. In Orchestrator, navigate to **Configuration > Discovery > Discovered Appliances**. You should see the appliance you just configured.
11. Click **Approve**.

Orchestrator HTTPS Certificate

Orchestrator > Software & Setup > Setup > HTTPS Certificate

By default, Orchestrator presents a self-signed server certificate to any client opening a TLS connection. This includes web browsers and EdgeConnect appliances. To ensure secure communications, TLS clients will cryptographically verify that a trusted Certificate Authority (CA) issued the Orchestrator certificate. For self-hosted Orchestrators, enterprises must set up an HTTPS server certificate for their Orchestrator.

Orchestrator provides two methods to set up an HTTPS server certificate. The first is new to release 9.4 where Orchestrator builds end entity certificates. This is the preferred method. With this method the user builds a Certificate Signing Request (CSR) in Orchestrator. As part of this process, Orchestrator creates the public key private key pair. The user downloads and submits the CSR for signing by a Certificate Authority (CA). The signed end entity certificate is then uploaded in Orchestrator for use in one of several applications. The end entity certificate contains a label, which is significant to Orchestrator and allows this certificate to be used by referring to its label.

NOTE: The Orchestrator HTTPS certificate cannot be added using EST.

To use an end entity certificate as Orchestrator HTTPS server certificate:

NOTE: To use an end entity certificate, you must first create an end entity certificate for use. To do this, see End Entity Certificates Tab.

1. After the certificate is uploaded on the End Entity Certificates tab, navigate to **Orchestrator > Software & Setup > Setup > HTTPS Certificate**.

   The HTTPS Certificate dialog box opens.

2. Click **Use End Entity Certificate**.
3. Select the label for the certificate you uploaded from the **End Entity Certificate** drop-down menu.
4. Click **Save**.

NOTE: After saving, you must manually restart Orchestrator for the web server to pick up the new certificate.

NOTE: To have the EdgeConnect appliance verify the Orchestrator certificate, you must click the **Verify Orchestrator Certificate** check box on the Advanced Security Settings dialog box. To do this, navigate to **Configuration > Overlays & Security > Advanced Security Settings**.

The other method (legacy method) requires everything to be done externally including creating the public key private key pair and creating the CSR. This legacy method is not recommended.

To use the legacy method with Orchestrator:

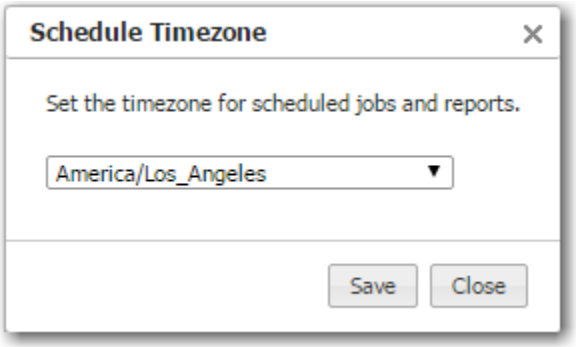
1. Consult with your IT security team to generate a certificate signing request (CSR), and submit it to your organization's chosen SSL Certificate Authority (CA).
   - Examples of Certificate Authorities include GoDaddy, Verisign, Comodo, Symantec, Microsoft Entrust, and GeoTrust.
   - All certificate and key files must be in PEM format.
2. After the Certificate Authority provides a CA-verified certificate:
   - Navigate to **Orchestrator > Software & Setup > Setup > HTTPS Certificate**.
   - If your IT security team advises the use of an Intermediate CA, use an Intermediate Certificate File. Otherwise, skip this file.
   - Load the Certificate File from the CA.
   - Upload the Private Key File that was generated as part of the CSR.
3. To associate the CA verified certificate for use with Orchestrator, click **Upload**.

NOTE: To have the EdgeConnect appliance verify the Orchestrator certificate, you must click the **Verify Orchestrator Certificate** check box on the Advanced Security Settings dialog box. To do this, navigate to **Configuration > Overlays & Security > Advanced Security Settings**.

Timezone for Scheduled Jobs

Orchestrator > Software & Setup > Setup > Timezone for Scheduled Jobs

Use this dialog box to set the timezone for scheduled jobs and reports.



Orchestrator Advanced Properties

Orchestrator > Software & Setup > Setup > Advanced Properties

Orchestrator Advanced Properties provide low-level configuration settings for the Orchestrator server.

WARNING: Modifying Orchestrator Advanced Properties can adversely affect your network. Changing the default settings is not recommended without consulting Support.

The ability to modify Orchestrator Advanced Properties depends on your Orchestrator deployment.

- Self-hosted/on-prem Orchestrator users can edit the Orchestrator Advanced Properties. You must restart Orchestrator to activate certain Orchestrator Advanced Properties changes.
- Orchestrator-as-a-Service (OaaS) users must open a TAC case to request changes to the Orchestrator Advanced Properties. TAC must also restart OaaS to activate Orchestrator Advanced Properties changes.

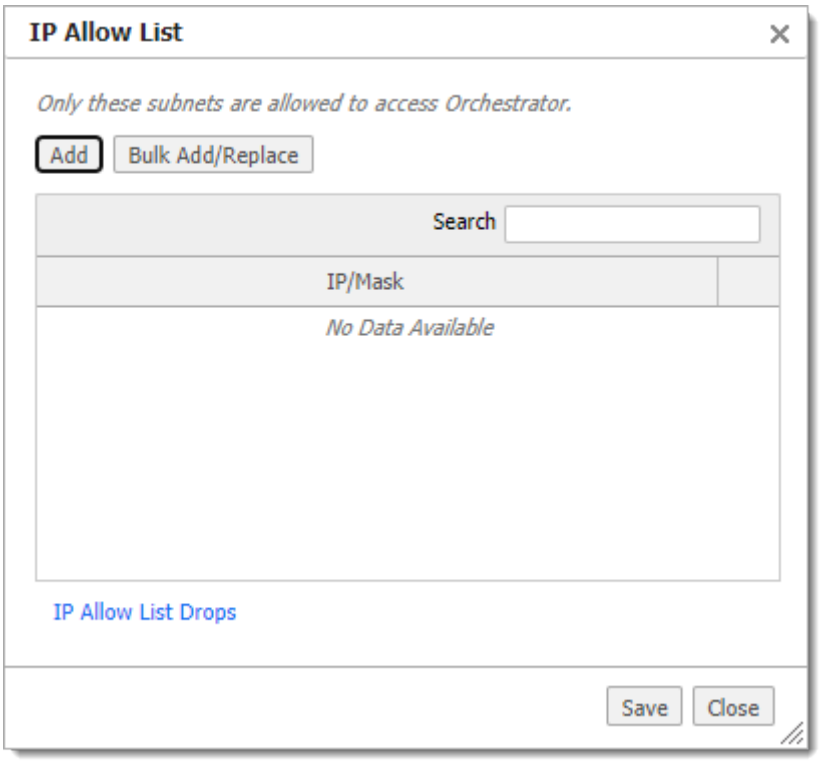
To enforce enhanced password criteria and set password requirements for Orchestrator, you must modify certain Orchestrator Advanced Properties. For more information, see Enforce Password Criteria and Set Password Requirements

IP Allow List

Orchestrator > Software & Setup > Setup > IP Allow List

IP Allow List is a feature that restricts access to Orchestrator to a specified list of source subnets.

If a source IP address changes (for example, with NAT IP), users can get locked out of Orchestrator.

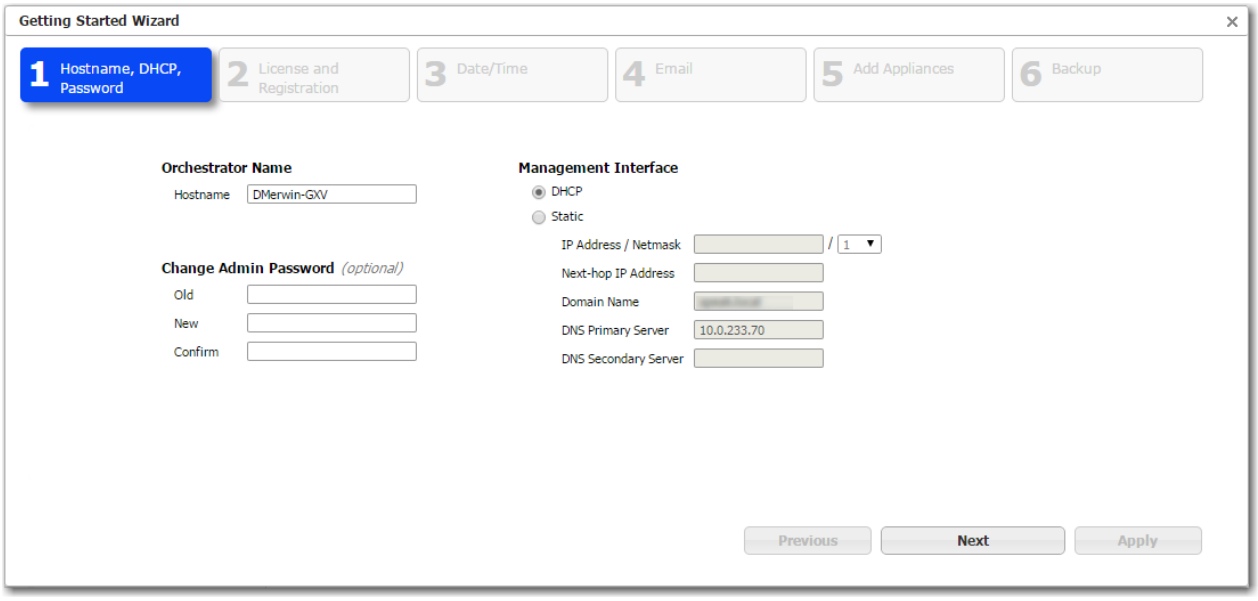


To view a list of traffic that has been dropped because of these restrictions, click **IP Allow List Drops**.

Orchestrator Setup

Orchestrator > Software & Setup > Setup > Orchestrator Setup

When you first install Orchestrator and use a web browser to access the IP address you have assigned it, the Getting Started Orchestrator Setup wizard opens.



The wizard guides you through the basics of configuring the following:

Setting	Description
Orchestrator Name, management IP address, and password	The default for username and password is **admin**.
License and Registration	EdgeConnect registration is required for Cloud-based features and products, including CPX and SaaS. The associated Account Name and Account Key enable Orchestrator to discover EdgeConnect appliances via the Cloud Portal, as they are added to your network.
Date/Time	Using an NTP server is strongly recommended so that data is synchronized across Orchestrator and the appliances.
Email	Change the default settings to your Company's SMTP server, and then test. Separate fields are provided for Global Report recipients and Alarm recipients.
Add Appliances	*(Optional)* You can use this to add NX, VX, and VRX appliances that are already up and running in your network. You can also add them later.
Backup	Specifies the database backup destination, transfer protocol, and backup schedule.

If you do not click **Apply** after you complete the last page, the wizard reappears at your next login.

To access the wizard again after initial configuration, navigate to **Orchestrator > Software & Setup > Setup > Orchestrator Setup**.

Availability Time Settings

Orchestrator > Software & Setup > Setup > Availability Time Settings

Use the Availability Time Settings tab to configure business hours and time zones for your appliances. Orchestrator collects availability data based on each appliance's time setting. The default business hour value is 24 hours, but you can customize the number of business hours in a day. For more information see Availability.

NOTES

- The Date/Time template sets the time and time zone for the appliance. Make sure the proper Date/Time template is applied to the appliance. If the appliance does not have the proper template applied, business hours set on the Availability Time Settings tab will be inaccurate.
- Reports are tied to calendar months (not the last 30 days) and calendar day boundaries.
- Availability trends charts display data with the Coordinated Universal Time (UTC) timestamp.

To set Availability Time Settings for an appliance:

1. Select an appliance in the table, and then complete the following fields as appropriate.

Field	Description
Time Zone	Availability time zone for the appliance. Select one of the following: **Use appliance timezone** – This setting causes Orchestrator to use the appliance's uniquely configured time zone. It is the default setting for appliances newly added to your network and for existing appliances when Orchestrator is upgraded to 9.5.x, except when Orchestrator is upgraded from 9.3.x to a later version. In this specific case, if the availability time zone is UTC, the appliance's time zone is set to availability time zone. If it is not UTC, the appliance's time zone remains unchanged. A listed time zone – If the time zone you select is different than the appliance's uniquely configured time zone, this setting will override the appliance's time zone.

Field	Description
Business Hours	Availability business hours for the appliance. Select one of the following: **24Hours** – This is the default setting for appliances newly added to your network and for existing appliances when Orchestrator is upgraded to 9.5.x, except when Orchestrator is upgraded from 9.3.x to a later version. In this specific case, the appliance's business hours remain unchanged. **Custom** – Select start and end times.
Changes	Indicates changes made to the time zone's business hours, which will be applied when you save.

2. Click **Save**.

Statistics Retention

Orchestrator > Software & Setup > Setup > Statistics Retention

This tab displays all the statistics Orchestrator collects from appliances. Orchestrator saves the statistics data in a database with the retention policies defined on this tab.

To begin, complete the following steps:

1. Click the **Edit** icon in the table next to the statistic you want Orchestrator to collect.
2. To enable or disable statistics collection, select the **Collect this statistic in Orchestrator** check box.
3. Enter how long you want Orchestrator to retain the statics for **Minute Granularity**, **Hourly Granularity**, and **Daily Granularity** before it collects data and stores in the partition.

   TIP: If you click **More Options**, you can enter values for the Database Duration.
4. Click **Apply**.

For more detail, refer to the following table:

Field	Description
Statistic	The selected statistic of which you want Orchestrator to collect data.
Enabled	If you have enabled or disabled statistics retention.
Minute Granularity (hours)	Amount of times in one minute Orchestrator stores data.
Hourly Granularity (days)	Amount of times in one hour Orchestrator stores data.

Field	Description
Daily Granularity (months)	Amount of time in one day Orchestrator stores data.
Estimated Disk Space	Estimated amount of disk space the selected statistic uses. At the bottom of the screen, you can get an estimated disk space required for a number of appliances, overlays, and tunnels.

DoS Stats

You can monitor your network behavior based on the Firewall Protection Profile (FPP) settings and the thresholds you set. The denial of service (DoS) statistics described in the following table provide visibility into the statistics for all Firewall Protection Profiles that are configured and mapped to zone/segment pairs. DoS statistics help you tune FPP settings with the right thresholds, identify violating sources, enable response actions, and so on.

Orchestrator retains statistics related to DoS services on appliances for 30 days. If you require longer retention periods, you can purchase a Storage subscription. See your HPE Aruba Networking representative for details.

DoS Statistic	Description
Ddostotalstats	Includes the following: **Source/Host Total Drop Count:** The total number of flows dropped from all hosts in the zone because the source-level threshold was exceeded. For real-time statistics, stats are collected every second. For historical statistics, stats are collected every 5 minutes. **Zone Level - Total Drop Count:** The total number of flows dropped at the zone level. For real-time statistics, stats are collected every second. For historical statistics, stats are collected every 5 minutes. **Source/Host - Total Deny Flow Count:** The total number of flows dropped/denied from all hosts that are in a deny list because the source level threshold was exceeded. For real-time statistics, stats are collected every second. For historical statistics, stats are collected every 5 minutes. **Source/Host Level - Total Deny IP Count:** The total number of hosts denied in the zone because a deny action was executed due to the source level threshold being exceeded. For real-time statistics, stats are collected every second. For historical statistics, stats are collected every 5 minutes. **Source/Host - Max Threshold Exceed Count:** The total number of times the Max Threshold was exceeded from all hosts in the zone. For real-time statistics, stats are collected every second. For historical statistics, stats are collected every 5 minutes. **Source/Host – Min Threshold Exceed Count:** The total number of times the Min Threshold was exceeded from all hosts in the zone. For real-time statistics, stats are collected every second. For historical statistics, stats are collected every 5 minutes. **Zone Level - Min Threshold Exceed Count:** The total number of times the Min Threshold was exceeded at the zone level. For real-time statistics, stats are collected every second. For historical statistics, stats are collected every 5 minutes. **Zone Level - Max Threshold Exceed Count:** The total number of times the Max Threshold was exceeded at the zone level. For real-time statistics, stats are collected every second. For historical statistics, stats are collected every 5 minutes. **Zone Level - Total Error Drop Count:** Total number of error drop counts collected at the zone level. For real-time statistics, stats are collected every second. For historical statistics, stats are collected every 5 minutes.

DoS Statistic	Description
Ddostoptalkers	The top 50 IPs that used the most number of concurrent flows, embryonic flows, and flows per second (FPS) and the top 50 IPs that violated the minimum and maximum thresholds. This statistic is common for all zones and is refreshed every hour. **NOTE** If there are more IPs than the list can support, IPs in this list will be replaced resulting in some IPs getting dropped from the hourly top talkers list.
Ddossrcipsample10	A list of 10 sample source IPs that exceed the minimum threshold and maximum threshold set in Firewall Protection Profile for concurrent, embryonic, and FPS metrics respectively. These are historical stats (not real-time). These stats are not zone-specific. Sample IPs are from zone/segment pairs where FPP is configured. These stats are updated every 5 minutes.
Ddospeakandpea	Includes the following: Peak stats for IPs and Counts are collected for the FPS, concurrent flows, and embryonic flows metrics every one second for all protocols. These stats report the highest number for Ddostotalstats, Ddostoptalkers, and Ddossrcipsample10. The metrics are each recorded in 5 minute intervals. The Ddospeakandpeakdroprate stats are reported even if no thresholds are configured in the FPP. For real-time statistics, stats are collected every second. For historical statistics, stats are collected every 5 minutes. The peak drop rate stats report the highest flow drop rate per second that occurred within the 5 minute interval. For real-time statistics, stats are collected every second. For historical statistics, stats are collected every 5 minutes.

Stats Configuration

The Stats Configuration dialog box displays the parameters for the Statistics Retention.

WARNING: Changing the default values of these settings is not recommended without consulting Support.

Click **Advanced Properties** to display the Stats Configuration dialog box. For more detail, refer to the following table.

NOTE: When changing the stats retention value for different granularities (for example, minute, hour, day), the estimated disk space displayed on the page only reflects the impact of the new retention setting on statistics generated after the change. Existing stats generated before the retention change are not included in the updated estimate.

Field	Description
minuteRetention	The number of minutes of stats the EdgeConnect will retain. 1440 equals 1 day (EdgeConnect generates a zip file of CSVs every minute.)
verticalRetention	The number of days EdgeConnect will retain data in a format that allows EdgeConnect's UI to display historical charts. Note: If you are using Orchestrator to see historical stats, set verticalRetention to zero to reduce disk space use.
app max_items	EdgeConnect calculates the "Top Applications." These stats are displayed hourly but computed every minute. Increasing this value beyond 100 will impact EdgeConnect performance because EdgeConnect computes a larger number of Top Applications every minute by traversing every flow in the system.
port max_items	EdgeConnect calculates "Top Ports." These stats are displayed hourly but computed every minute. Increasing this value beyond 100 will impact EdgeConnect performance because EdgeConnect computes a larger number of Top Ports every minute by traversing every flow in the system.
dns max_items	EdgeConnect calculates "Top Domains." These stats are displayed hourly but computed every minute. Increasing this value beyond 100 will impact EdgeConnect performance because EdgeConnect computes a larger number of Top Domains every minute by traversing every flow in the system.
ip max_items	EdgeConnect calculates "Top Talkers." These stats are displayed hourly but computed every minute. Increasing this value beyond 100 will impact EdgeConnect performance because EdgeConnect computes a larger number of Top Talkers every minute by traversing every flow in the system.

Field	Description
behavioral max_items	EdgeConnect calculates “Top Traffic Behaviors.” These stats are displayed hourly but computed every minute. Increasing this value beyond 100 will impact EdgeConnect performance because EdgeConnect computes a larger number of Top Traffic Behaviors every minute by traversing every flow in the system.
flows_csv_enable	Do not enable this field. Enabling this field will disable EdgeConnect appliances in most networks, which might require you to RMA the EdgeConnect. If you need a NetFlow log to view historical data, enable Netflow/IPFIX and use an external NetFlow/IPFIX collector.

Stats Collector Configuration

Orchestrator > Software & Setup > Setup > Stats Collector Configuration

IMPORTANT: Starting with Orchestrator 9.6.0, statistics are collected on a single database and do not require configuration. If you are running Orchestrator 9.6.0 or later, you do not need to perform the procedures in this topic.

If you upgrade an on-prem Orchestrator to Orchestrator 9.6.0 or later, your statistics data will not be migrated automatically from the Distributed Stats Collector to the new Stats Collector that runs on a single virtual instance. You must manually migrate the stats from your local or distributed stats collectors to the new Stats Collector. For details, see Migrate Stats Collectors to 9.6.0 On-Prem Orchestrator in the Install, Upgrade, Move, and Restore SD-WAN Orchestrator Guide.

Orchestrator collects statistical data from your appliances to monitor performance, network traffic, and appliance status. Before Orchestrator release 9.1.0, the process of collecting, storing, and retrieving this data impacted performance due to the amount of data stored on and requested from the database.

To improve Orchestrator performance, Orchestrator 9.1.0 includes a new method called the Distributed Stats Collector (simply referred to as Stats Collector) that eliminates the use of Orchestrator resources for monitoring your appliances. This new architecture enables you to scale your network with greater performance.

There are two variations of stats collection:

- **Local Stats Collector**: Orchestrator and Stats Collector in a single VM. This is ideal for deployments with less than 100 appliances.

- **Remote Stats Collector**: Orchestrator on a separate VM and Stats Collector on a different VM (Orchestrator VM deployed in Stats Collector mode only). HPE Aruba Networking recommends one Remote Stats Collector per 150 appliances.

The Distributed Stats Collector feature collects statistics from appliances and provides the information to Orchestrator. When enabled, the Distributed Stats Collector runs in parallel with the Local Stats Collector to collect the necessary historical statistical data. After collecting that data, you will discontinue local stats collection. You will not experience performance improvement until you discontinue local stats collection.

If you are running ECOS 9.1 or later and Orchestrator 9.1 or later, HPE Aruba Networking recommends that you set up the Distributed Stats Collector so that you will be able to take advantage of new stats that are introduced in future releases.

Features that Require Distributed Stats Collector

The Distributed Stats Collector feature allows Orchestrator to collect the following statistics.

- Protection Profile Top Talkers (9.5.x)
- Protection Profile Trends (9.5.x)
- Flow Baseline (9.5.x)
- Flow Baseline Trends (9.5.x)
- Availability (ECOS 9.3.0)
- AppExpress (ECOS 9.4.1) (Reporting and monitoring functions will be limited if Distributed Stats Collector is not enabled.)
- IP SLA Summary (ECOS 9.3.0)
- Internet Breakout (ECOS 9.3.0)
- Application Summary (ECOS 9.3.0)
- Application Trends (ECOS 9.3.0)
- User Trends/Bandwidth (ECOS 9.3.0)

Prerequisites

- Both the Orchestrator and Stats Collector must be on the same release.
- Upgrade all appliances to version 9.1.0 before enabling the Distributed Stats Collector feature.
- By default, when you install Orchestrator for the first time, Orchestrator automatically creates a Local Stats Collector. The Local Stats Collector can accommodate a maximum of 100 appliances. If you need to scale beyond 100 appliances, you must use the Remote Stats Collector. (There might be some dependencies based on your deployed topology.) Refer to the following table to determine the number of appliances per Remote Stats Collector you will need for your topology.

Topology	Appliances per Stats Collector
Mesh	Up to 150
Hub and Spoke	Up to 300

- Each Stats Collector must meet the following virtual machine minimum requirements:
  - vCPU: 6 (greater than 2 GHz with hyperthreading enabled)
  - RAM: 16 GB
  - Disk usage will vary based on retention policies. Disk estimates are based on 30 days of alarm and audit log retention. See the appropriate Orchestrator Host System Requirements document on the System Requirements page for detailed information on recommended resources.

  NOTE: If you have less than 100 appliances, you can use the predefined Local Stats Collector. You do not need to perform the steps in Before You Begin.

Enable Distributed Stats Collector

To enable Distributed Stats Collector on a self-hosted Orchestrator (On-Prem), navigate to **Orchestrator > Software & Setup > Setup > Stats Collector Configuration** and complete the configuration as explained below.

To enable Distributed Stats Collector on an Orchestrator as a Service (Cloud Orchestrator managed by HPE Aruba Networking), contact Technical Support as this can only be done through the backend system.

Before You Begin

For deployments larger than 100 appliances, you must do the following before you configure the Distributed Stats Collector in Orchestrator:

1. Create a VM for the Stats Collector.
2. Upgrade the Stats Collector to the Orchestrator release. See Upgrade Orchestrator or Stats Collectors in the Install, Upgrade, Move, and Restore SD-WAN Orchestrator Guide.
3. Configure the VM as a Stats Collector.
4. Configure the Stats Collector with a Public Key.
5. Create and Install an End Entity Certificate. Skip this step if you are not using a custom HTTPS certificate on Orchestrator or you are using the Local Stats Collector.

Create, configure, and encrypt as many Stats Collectors as needed.

Create a VM for the Stats Collector

To create and set up a Stats Collector, do one of the following:

- If you have an on-prem Orchestrator deployment, bring up an Orchestrator VM by following the steps explained in "On-Prem Orchestrator - Download, Deploy, and Install" in Install, Upgrade, Move, and Restore SD-WAN Orchestrator .
- If you have a cloud Orchestrator deployment, bring up an Orchestrator in your selected cloud platform by following the steps explained in "Orchestrator in IaaS - Deploy" in Install, Upgrade, Move, and Restore SD-WAN Orchestrator .

Configure the VM as a Stats Collector

1. Open an SSH session to the Orchestrator you want to use as a Stats Collector.
   - For on-prem Orchestrator deployment: Enter `$ su`
   - For cloud Orchestrator deployment: Enter `$ sudo su - root`
2. If prompted, enter the root password. If you do not know your root password, contact Support.
3. Change to the /home/gms/gms directory:`cd /home/gms/gms`
4. To run the Orchestrator setup script, enter `orch-setup -m`, and then press **Enter**.
5. To select the stats collector only mode, at the prompt, enter `s`.
6. To proceed, enter `y`.

   This VM is now a Stats Collector.

   NOTE: Orchestrator and EdgeConnect appliances communicate with the Stats Collector over HTTPS (port 443). Orchestrator and EdgeConnect appliances will raise alarm if the Stats Collector is not reachable.

Configure the Stats Collector with a Public Key

After you create and configure a Stats Collector, you must copy the Orchestrator public key and paste it into the same folder on the Stats Collector, as follows. This will establish an HTTPS connection with Orchestrator and the data from the Stats Collector to Orchestrator will be encrypted.

Copy the Public Key File from Orchestrator to the Stats Collector

1. Open an SSH session to the *Orchestrator VM*.
   - For on-prem Orchestrator deployment: If you are running release 9.1.9, 9.2.10, 9.3.3, 9.4.x, or later, enter the following:

     `whoami`

     If the response is not `gms`, enter the following, and then provide the admin password.

     `sudo su - gms`
   - For cloud Orchestrator deployment: Enter `sudo su - gms`
2. Go to: `cd /home/gms/sc/publickeys`
3. To list the file that contains the public key, enter `ls`

4. Enter the followinig command to determine if you are running Rocky Linux:

   ```
   cat /etc/redhat-release
   ```

5. Do one of the following to copy the public key file to your new Stats Collector.

   NOTE: Depending on your underlying OS and security settings, the following SCP steps might not work. If SCP does not work, copy the public key file using another file transfer method, such as FTP, or you can manually copy the file (make sure to copy escape characters) to the new Stats Collector in the /home/gms/sc/publickeys/ directory.

   - If you are running Rocky Linux, enter the following command (you might need to omit `-O` if your operating system does not support this parameter).

     ```
     sudo scp -O <public_key_file_name.pub> admin@<new-Orch-IP>:/home/admin/
     ```

   - If you are *not* running Rocky Linux, enter `scp <public_key_file_name.pub> admin@<remote_stats_collector_ip>:/home/gms/sc/publickeys/`

     where:

     - *public_key_file_name.pub* is the name of the file listed in step 3. For example, `d1ab581df8c745b59eec548ef5a2f011.pub`. The public key file name will be different for each case.
     - *admin* is the user name

Confirm the New Public Key File

1. Open an SSH session to the *Stats Collector VM*.

   - For on-prem Orchestrator deployment: If you are running release 9.1.9, 9.2.10, 9.3.3, 9.4.x, or later, enter the following:

     ```
     whoami
     ```

     If the response is not `gms`, enter the following, and then provide the admin password.

     ```
     sudo su - gms
     ```

   - For cloud Orchestrator deployment: Enter `sudo su - gms`

2. Enter the following command to determine if you are running Rocky Linux:

   ```
   cat /etc/redhat-release
   ```

3. If you are running Rocky Linux, enter the following commands.

   NOTE: If you used FTP or manually copied the public key file in step 5 above, skip this step.

   ```
   sudo mv /home/admin/<public_key_file_name.pub> /home/gms/sc/publickeys/
   ```

   ```
   sudo chown gms.gms /home/gms/sc/publickeys/<public_key_file_name.pub>
   ```

4. Go to: `cd /home/gms/sc/publickeys` and execute `ls -l`.

   Ensure that the new public key file is on the Stats Collector and has the following privileges and ownership.

   [gms@silverpeak-gxv:~/sc/publickeys] $ ls -l -rw-r–r– 1 gms gms 451 Oct 7 09:28 d1ab581df8c745b59eec548ef5a2f011.pub

Create and Install an End Entity Certificate

Complete the following tasks to create and install an end entity certificate for each Remote Stats Collector.

NOTE: Skip this procedure if you did not install a custom HTTPS certificate on Orchestrator or you are using the Local Stats Collector.

1. Add the Root CA Certificate for the Certificate Authority (CA).
2. Create and send the CSR in Orchestrator and upload the signed certificate in Orchestrator.
3. Create and send the CSR in Stats Collector and upload the signed certificate in Stats Collector.

Add the Root CA Certificate for the Certificate Authority

In order for your Remote Stats Collectors to establish connectivity with the appliances, you must add the certificate to the Custom CA Certificate Trust Store.

To add the certificate to the Custom CA Certificate Trust Store:

1. In Orchestrator, navigate to **Configuration > Overlay & Security > Security > Custom CA Certificate Trust Store**.
2. Select **Use Custom Certificate Store**.
3. Click **Add Default Certificates**.
4. Click **Add Certificate to Custom Trust Store**.

   The Add/Edit Custom Certificates dialog box displays.
5. Enter a meaningful Alias for the certificate in the Alias field. For example, “ClearPass_CA_Root”.
6. Paste the root certificate into the Certificate field.
7. Click **Save**.
8. Click **Apply Changes**.
9. Click **Close**.

   IMPORTANT: After you add a root CA certificate to the Custom Trust Store, you must restart Orchestrator from the CLI.
10. Enter the following commands from the CLI to restart Orchestrator.

```
ssh admin@xx.xx.xx.xx

su

service gms status

service gms stop

service gms start
```

where *xx.xx.xx.xx* is the IP address of Orchestrator.

Proceed to Create and Send the CSR in Orchestrator and Upload the Signed Certificate in Orchestrator.

Create and Send the CSR in Orchestrator and Upload the Signed Certificate in Orchestrator

1. Create the CSR in Orchestrator. Follow the steps in Create a Certificate Signing Request (CSR).
2. Send the CSR to your Certificate Authority to receive a signed certificate.
3. Obtain the signed certificate from the CA. Follow the steps in Obtain the Signed Certificate From the CA.
4. After you receive the signed certificate, follow the steps in Upload the Signed Certificate to the End Entity Certificate Tab to upload the signed certificate.

Proceed to Create and send the CSR in Stats Collector and upload the signed certificate in Stats Collector.

Create and send the CSR in Stats Collector and upload the signed certificate in Stats Collector

You must create and send the CSR in Stats Collector and upload the signed certificate in Stats Collector.

1. Enter the following commands to log in to the Remote Stats Collector:

   ```
   ssh admin@xx.xx.xx.xx

   cd /home
   ```

   where *xx.xx.xx.xx* is the IP address of the Remote Stats Collector.
2. Create a file called *openssl conf.cnf*.
3. Copy the following content and paste it into the file.

   ```
   [ req ]
   default_bits = 1024
   distinguished_name = req_distinguished_name
   req_extensions = SAN
   extensions = SAN
   [ req_distinguished_name ]
   commonName = xx.xx.xx.xx
   countryName = US
   stateOrProvinceName = CA
   localityName = San Jose
   organizationName = HPE
   organizationalUnitName = Aruba
   ```

```
[SAN]

#authorityKeyIdentifier=keyid,issuer

#basicConstraints=CA:FALSE

#keyUsage = digitalSignature, nonRepudiation, keyEncipherment, dataEncipherment

subjectAltName = IP:xx.xx.xx.xx
```

where *xx.xx.xx.xx* in both commonName and subjectAltName is the IP address of the Remote Stats Collector.

NOTE: commonName represents the name of the server. However, commonName can also be a FQDN name if there is a DNS entry in your DNS server.

4. Enter the following command to create a new private key and CSR with the config file you just created:

   ```
   openssl req -new -newkey rsa:2048 -keyout newkey.key -config opensslconf.cnf -out newSc_serverCrt.csr -nodes
   ```

5. Use SCP to copy the new CSR from the Remote Stats Collector to your laptop.

   NOTE: If you cannot copy the new CSR via SCP, you can create a new file with the same name on your laptop and copy the contents into the new file.

6. Send the CSR file that you copied in step 5 to your Certificate Authority for signing.
7. Obtain the signed certificate from the CA.

   When you receive the signed certificate from the CA, if there are multiple files you need to combine all the files into a single file, which includes the end entity certificate, all intermediate CA certificates, and the root CA certificates. This is necessary because you must upload the entire certificate chain in Orchestrator as a single file. The sequence of certificates in the single-file chain is important and should be as follows:

   1. End entity certificate (top of file)
   2. One or more certificates of the intermediate CA(s)
   3. Self-signed root CA certificate

   You now have a signed end entity certificate for the Remote Stats Collector.

8. Rename the downloaded and signed end entity certificate *scSignedCrt.pem*.
9. Use SCP or Filezilla to transfer *scSignedCrt.pem* for the Remote Stats Collector from your laptop to the */home* directory on the Remote Stats Collector.
10. Log in to the Remote Stats Collector.

    ```
    ssh admin@xx.xx.xx.xx


    cd /home
    ```

11. Back up the existing *server.crt* and *server.key* files in */home/gms/gms/properties*.

```
cd /home/gms/gms/properties

mv server.crt server_backup.crt

mv server.key server_backup.key
```

12. Enter the following command to move the private key you created while generating the CSR from */home* to */home/gms/gms/properties*.

    ```
    mv /home/newkey.key /home/gms/gms/properties/server.key
    ```

13. Enter the following command to move the signed end entity certificate for the Remote Stats Collector from */home* directory to */home/gms/gms/properties*

    ```
    mv /home/scSignedCrt.pem /home/gms/gms/properties/server.crt
    ```

14. Restart the Remote Stats Collector.

    ```
    service sc restart
    ```

15. After you restart the Remote Stats Collector, the new end entity certificate for the Remote Stats Collector will be installed and the private key will be loaded. Run the following command to verify the status of the Remote Stats Collector.

    ```
    service sc status
    ```

    The status should be “up and active”.

Configure the Stats Collector Feature

After the Stats Collectors are created, configured, and authenticated, configure the Distributed Stats Collector feature in Orchestrator. Complete the following tasks:

1. Back up Orchestrator. For more information about backing up Orchestrator, see Back Up on Demand.

   Before you enable the Distributed Stats Collector feature and discontinue legacy stats collection, it is recommended that you back up the Orchestrator database. Discontinuing legacy stats collection is permanent. To return to your previous configuration, you must restore the Orchestrator configuration backup.

2. Add a Stats Collector. If your network contains less than 200 appliances, you can use the predefined Local Stats Collector.
3. Associate Appliances with a Stats Collector or Associate Appliances with the Predefined Local Stats Collector
4. When the necessary historical data has been collected, Discontinue Legacy Stats Collection.

Add a Stats Collector

To add a stats collector:

1. Navigate to **Software & Setup > Setup > Stats Collector Configuration**.

   The Stats Collector Configuration tab opens.
2. Click **Edit Stats Collectors**.

   The Edit Stats Collectors dialog box opens.
3. Click **Add Stats Collector**.

   The Stats Collector dialog box opens.
4. Configure the following elements as needed:

Field	Description
Name	Name of the stats collector.
DNS Name	DNS name or IP address of this Stats Collector.
Port	Port number the Stats Collector is running on.
Protocol	HTTPS

   NOTE: If your organization uses URL allowlisting, you must also allow the DNS name of the Stats Collector. The Stats Collector URL is unique for each customer and follows this format:
   >
   > *<customer name>-orch-<region>-sc-<string>.silverpeak.cloud*
   >
   > Ensure this DNS entry is included in your allowlist to allow proper communication with the Stats Collector.
5. Click **Save**.

Delete a Stats Collector

To delete an existing Remote Stats Collector, click the delete icon (**X**) in the last column of the entry in the table.

Associate Appliances with a Stats Collector

To associate appliances with a Stats Collector:

1. Navigate to **Software & Setup > Setup > Stats Collector Configuration**.

   The Stats Collector Configuration tab opens.
2. In the Orchestrator appliance tree, select one or more appliances to associate with a specific Stats Collector.

   WARNING: The statistics for an appliance are tied to the Distributed Stats Collector it is associated with. If you associate an appliance with a different Distributed Stats Collector, you lose all statistical data associated with that appliance.

3. Select the **Add** check box next to the Stats Collector you want to associate the selected appliance(s) with.
4. Click **Apply**.

   The Apply Changes dialog box opens.
5. Click **Apply Changes**.

Associate Appliances with the Predefined Local Stats Collector

If you are installing Orchestrator version 9.1.0 or upgrading to version 9.1.0 or later, Orchestrator provides a default Stats Collector called *local*. You cannot edit or delete the Local Stats Collector. You can associate up to 200 appliances with the Local Stats Collector.

NOTE: If you are upgrading to Orchestrator 9.1.0, all appliances will be automatically associated with the Local Stats Collector.

NOTE: If you run Orchestrator in Orchestrator Only mode (`orch-setup -m o`), the Local Stats Collector will be disconnected.

To associate appliances with the Local Stats Collector:

1. Navigate to **Software & Setup > Setup > Stats Collector Configuration**.

   The Stats Collector Configuration tab opens. This tab displays the Stats Collector configuration for all appliances selected in the appliance tree to the left.
2. In the Orchestrator appliance tree, select one or more appliances to associate with the Local Stats Collector.
3. Select the **Add** check box next to the Local Stats Collector.
4. Click **Apply**.

   The selected appliances are associated with the Local Stats Collector. The Changes column indicates the Stats Collectors that were added and removed.

Enable the Distributed Stats Collector

After you associate appliances with either the Local Stats Collector or the newly added Stats Collectors, you must enable the Distributed Stats Collector feature to begin collecting data.

NOTE: The legacy Stats Collector continues to collect statistics in parallel with the Distributed Stats Collector feature until you discontinue legacy stats collection. For more information, see Discontinue Legacy Stats Collection.

NOTE: The backslash (\) character is not allowed in any field in the Orchestrator > Software & Setup > Backup > Schedule Stats Collector menu or the Orchestrator > Software & Setup > Backup > Schedule Backup menu.

WARNING: You cannot disable the Distributed Stats Collector after you enable it. It is recommended that you back up Orchestrator before you enable the Distributed Stats Collector. For more information about backing up Orchestrator, see Back Up on Demand.

To enable the Stats Collector:

1. Navigate to **Software & Setup > Setup > Stats Collector Configuration**.

   The Stats Collector Configuration tab opens.

2. Click **Enable New Stats Collection**.

   The Enable New Stats Collection dialog box opens.

   Before you can enable the Distributed Stats Collector feature, you must upgrade all appliances to version 9.1.0. The Enable New Stats Collection dialog box lists appliances that must be upgraded to support the distributed stats collection.

3. Click **Enable New Stats Collection Now**.

Discontinue Legacy Stats Collection

WARNING: Do not discontinue legacy stats collection until you have collected sufficient historical data with the Distributed Stats Collector feature. For example, if you need 30 days of statistical data, enable the Distributed Stats Collector, wait 30 days, and then disable the legacy stats collection.

To verify that data has been collected:

1. Navigate to **Support > Technical Assistance > Partition Management**.
2. Verify that the Stats Collector table contains sufficient data.

To discontinue legacy stats collection:

1. Navigate to **Software & Setup > Setup > Stats Collector Configuration**.

   The Stats Collector Configuration tab opens.

2. Click **Discontinue Legacy Stats Collection**.

   The Discontinue Legacy Stats Collection dialog box opens.

   WARNING: This step permanently disables legacy Stats Collection and deletes all legacy statistics.

3. Click **Discontinue Legacy Stats Collection**.

Back Up and Restore Stats Collector

Scheduling a Stats Collector backup is a prerequisite to accomplishing Stats Collector recovery/redundancy. If Stats Collector failure occurs, you need the latest backup of that Stats Collector to restore it on a Distributed Stats Collector.

Back Up the Stats Collector

You can schedule a backup or use the CLI to back up Stats Collector on-demand.

- To schedule a backup of the Stats Collector, see Schedule Stats Collector Backup.
- To back up the Stats Collector with the CLI, log in to the CLI and run the following commands:

```
cd /home/gms/sc
```

```
./sc_backup.sh
```

The script creates a `sc.zip` file in the `/home/gms/` directory.

Restore the Stats Collector from the CLI

1. Log in to the Stats Collector virtual machine.
2. Copy and paste the backup zip file in the `/home/gms/sc_backup/` directory for Orchestrator 9.6.0 and later or `/home/gms/` directory for Orchestrator releases earlier than 9.6.0, and then rename it to `sc.zip`
3. Log in as root user and run `service sc stop`
4. Do one of the following:
   - If you are running Orchestrator release 9.1.9, 9.2.10, 9.3.3, 9.4.x, or later, enter `sudo su - gms`
   - If you are running an Orchestrator release earlier than those listed above, enter `su - gms` to log in as gms user.
5. Enter the following commands:

   ```
   bash /home/gms/sc/sc_restore.sh 2>&1 | tee /tmp/restorelog
   ```
6. To confirm that the restore process completed successfully, open the `/tmp/sc_restorelog` file and verify the “Restore successful!” message is listed.
7. Log out and log in as root user.
8. Enter `service sc start`

Notification Banner

Orchestrator > Software & Setup > Setup > Notification Banner

If you are conducting downtime or for maintenance reasons, you can add a notification in the header of your Orchestrator UI. To add a notification, complete the following steps.

1. Navigate to **Orchestrator > Software & Setup > Setup > Notification Banner** in Orchestrator.

   The Notification dialog box opens.
2. Enter the message you want to display in the Orchestrator header.
3. Click **Save**.

Change Stats Collector Log Level

Orchestrator > Software & Setup > Setup > Change Stats Collector Log Level

This topic is for future content.

Orchestrator > HPE Integration Services

The options under **Orchestrator > HPE Integration Services** focus on integrating Orchestrator with HPE Aruba Networking Central, including creating an HPE Aruba Networking Central account, mapping EdgeConnect appliances, and integrating with ClearPass Policy Manager and Mist.

HPE ANW Central Site Mapping

Orchestrator > HPE Integration Services > HPE ANW Central Site Mapping

Use this tab to create an HPE Aruba Networking Central account in Orchestrator. After you create an HPE Aruba Networking Central account, Orchestrator maps EdgeConnect appliances to HPE Aruba Networking Central sites. When mapped, EdgeConnect appliances display in the Network Health tab in HPE Aruba Networking Central and provide real-time site health updates.

NOTE: Single Sign-On (SSO) to HPE Aruba Networking Central from Orchestrator is not supported. If your account is SSO-enabled, or if two-factor user verification is enabled, you will not be able to use the account for HPE Aruba Networking Central Site Mapping integration.

Prerequisites

Before you can integrate Unity EdgeConnect devices with HPE Aruba Networking Central, you must do the following:

1. Create an HPE Aruba Networking Central account. For more information on creating an HPE Aruba Networking Central account, see HPE Aruba Networking Central Online Help and search for “Unity EdgeConnect Integration.”
2. Generate an API token for Orchestrator in HPE Aruba Networking Central. For more information on generating an API token for Orchestrator, see HPE Aruba Networking Central Online Help and search for “Unity EdgeConnect Integration.”
3. Have existing HPE Aruba Networking Central sites to map EdgeConnect appliances to. If you do not have any existing HPE Aruba Networking Central sites, you can export the location details for EdgeConnect appliances and create HPE Aruba Networking Central sites in bulk from that exported list. For more information on creating HPE Aruba Networking Central sites in bulk, see Create HPE Aruba Networking Central Sites in Bulk.

You need the following details from your HPE Aruba Networking Central account.

Field	Steps
Customer ID	Navigate to **Account Home**, and then click the **User** icon in the upper-right corner.
Email	Navigate to **Account Home > API Gateway > System Apps & Tokens**. The email is listed in the **Name** column.

Field	Steps
Password	Navigate to **Account Home > API Gateway > System Apps & Tokens**, and then click **View Tokens**. **NOTE:** If the HPE Aruba Networking Central password changes, you must update this password whether authentication is configured as a system user or as a federated user. **NOTE:** If you do not remember the password, you must reset it from HPE Aruba Networking Central. For more information on resetting your HPE Aruba Networking Central password, see HPE Aruba Networking Central Online Help.
Client ID	Navigate to **Accounts Home > API Gateway > APIs > System Apps & Tokens**.
Client Secret	Navigate to **Accounts Home > API Gateway > APIs > System Apps & Tokens**.
API Gateway domain	Navigate to **Account Home > API Gateway**. The URL is listed in the Documentation column. **NOTE:** Copy the URL without the protocol (for example, *internal-apigw.central.arubanetworks.com*).

Create HPE Aruba Networking Central Sites in Bulk

1. In Orchestrator, navigate to **Administration > Software > Upgrade > System Information**.
2. In the appliance tree, select the appliances you want to create HPE Aruba Networking Central sites for, and then click **Export**.

   Orchestrator creates and downloads a .csv file.
3. Open the .csv file, and then delete the three header rows.

   TIP: Refer to the sample import file provided by HPE Aruba Networking Central for proper formatting. To view the sample import file, in HPE Aruba Networking Central, navigate to **Launch > Network Operations > Organization > Sites > Bulk Upload**, and then click **Download a sample file** on the Bulk Import dialog.
4. Save and close the file.
5. In HPE Aruba Networking Central, navigate to **Launch > Network Operations > Organization > Sites**.
6. Scroll to the bottom of the page, click **Bulk Upload**, and then follow the prompts.

Create an HPE Aruba Networking Central Account in Orchestrator

To create an HPE Aruba Networking Central account in Orchestrator:

1. On the HPE ANW Central Site Mapping tab, click **HPE ANW Central Account**.

   The HPE ANW Central Account dialog box opens.

2. Configure the following elements as needed:

Field	Description
HPE ANW Central	Status of the connection.
Customer ID	Customer ID generated from HPE Aruba Networking Central.
Email	Email provided by HPE Aruba Networking Central.
Password	HPE Aruba Networking Central password. **NOTE:** If the HPE Aruba Networking Central password changes, you must update this password whether authentication is configured as a system user or as a federated user. **NOTE:** If you do not remember the password you must reset the HPE Aruba Networking Central password from HPE Aruba Networking Central. For more information on resetting your HPE Aruba Networking Central password, see HPE Aruba Networking Central Online Help.
Client ID	Client ID generated from HPE Aruba Networking Central.
Client Secret	Client Secret generated from HPE Aruba Networking Central.
API Gateway domain	API Gateway URL without protocol (for example, *internal-apigw.central.arubanetworks.com*).

3. (Optional) To test the connection, click **Test**. To save without testing, see the **IMPORTANT** notice below, and then skip to step 4.

   IMPORTANT: Because of how HPE Aruba Networking Central processes account information, if you click **Test** or **Save**, you must wait 30 minutes before you click **Test** or **Save** again. If you click **Test** or **Save** a second time before 30 minutes have passed, you will receive an error that the connection failed even if you successfully connected to HPE Aruba Networking Central. To resolve this issue, wait 30 minutes before clicking **Test** or **Save** again.

4. To save the connection, click **Save**.

   Orchestrator maps EdgeConnect appliances to HPE Aruba Networking Central sites based on geolocation. (Addresses assigned to EdgeConnect appliances are converted to geolocations.)

5. If “None” displays in the HPE Aruba Networking Central Site column of an appliance, Orchestrator did not locate an HPE Aruba Networking Central site within range of the appliance (within 0.2 degrees of the latitudeDelta and the longitudeDelta combined). Do one of the following:

- Edit the appliance and manually map it to any HPE Aruba Networking Central site. For more information on mapping an EdgeConnect appliance to an HPE Aruba Networking Central site, see Edit EdgeConnect to HPE Aruba Networking Central Site Mapping.
- Add an HPE Aruba Networking Central Site within range of the EdgeConnect appliance, and then check for site list updates. For more information on checking for site list updates, see Check for Site List Updates.

Edit EdgeConnect to HPE Aruba Networking Central Site Mapping

Orchestrator maps EdgeConnect appliances to HPE Aruba Networking Central sites based on geolocation. Orchestrator maps EdgeConnect appliances to HPE Aruba Networking Central sites that are within 0.2 degrees of the latitudeDelta and the longitudeDelta combined.

You can edit an EdgeConnect appliance to map it to a different HPE Aruba Networking Central site. You can also edit an EdgeConnect appliance to map it to an HPE Aruba Networking Central site if Orchestrator did not locate an HPE Aruba Networking Central site within range of the EdgeConnect appliance.

To map an EdgeConnect appliance to an HPE Aruba Networking Central site:

1. Click the **Edit** icon next to an EdgeConnect appliance.

   The Edit EdgeConnect to HPE ANW Central Site Mapping dialog box opens.
2. Configure the following elements as needed:

Field	Description
EdgeConnect Appliance	Selected EdgeConnect appliance.
HPE ANW Central Site	Available sites to map the EdgeConnect appliance to.
Geolocation Suggested Site	HPE Aruba Networking Central site that Orchestrator mapped by geolocation to the EdgeConnect appliance. **NOTE:** If you map the EdgeConnect appliance to any other site, the site that Orchestrator suggests based on geolocation will display next to that site in parentheses.

3. Click **Save**.

   Orchestrator maps the appliance to the HPE Aruba Networking Central site you selected.

Check for Site List Updates

To refresh the HPE Aruba Networking Central site list in Orchestrator to check for HPE Aruba Networking Central site list updates, click **Check for Site List Updates**.

If new HPE Aruba Networking Central sites are detected within range of unmapped EdgeConnect appliances (within 0.2 degrees of the latitudeDelta and the longitudeDelta combined), Orchestrator maps the EdgeConnect appliances to the new HPE Aruba Networking Central sites.

ClearPass Policy Manager

Orchestrator > HPE Integration Services > ClearPass Policy Manager

NOTE: For detailed instructions on how to configure ClearPass Policy Manager to work with Orchestrator, see the SD-WAN Orchestrator and ClearPass Policy Manager Integration Guide.

Orchestrator supports association with ClearPass Policy Manager, which provides role-based and secure network access for devices. This integration provides user and role information for an IP address, which you can view on the Flows and Top Talkers tabs of Orchestrator.

The ClearPass Policy Manager tab displays information about users and devices provisioned to access your network via ClearPass. The searchable information on this tab includes details such as username, IP address, and role.

You can apply the following filters to your ClearPass logs:

- To determine which actions you want to display in the table, select the **All**, **Active**, or **Historical** filters.
- To refresh or pause the table, select **Auto Refresh** or **Pause**. By default, the table refreshes automatically.
- To limit the filtering criteria, enter a value in the **Record Count** field. The default value is **500**, and the maximum value is 10,000.
- To filter by date and time, enter values in the **From** and **To** fields.
- To search for a specific username, enter a value in the **User** field. You can search a wild card character (*) as a username using the following schema:
  - x* = anything that starts with the entered value
  - *x = anything that ends with the entered value
- To search for a specific IP address, enter a value in the IP field.

To export a .csv file of your table, click **Export**.

Field	Definition
Start Time	Time when the device began its network session.
End Time	Time when the device ended its network session.
CPPM	ClearPass Policy Manager server used to authenticate.
IP Address	IP address authenticated to the network.
Username	Username authenticated to the network.
Role	Role assigned to the user that authenticated to the network.
Device Type	Device type used to connect to the network.
MAC Address	MAC address of the system connecting to the network.
Posture	Security health posture of the connected device.
Location ID	Location ID of the user connecting to the network.

Field	Definition
Protocol	Type of authentication server used to connect to the network.
Details	All user information sent from CPPM but not required by Orchestrator. Values are in JSON format.

Manage ClearPass Policy Manager Accounts

To view and manage ClearPass accounts that are associated with Orchestrator, click **Accounts** on the ClearPass Policy Manager tab.

NOTE: Before you begin the ClearPass Policy Manager (CPPM) configuration in Orchestrator, you must have a ClearPass account to authenticate and authorize Orchestrator. If you do not have these credentials, contact your system administrator.

View ClearPass Policy Manager Accounts

The ClearPass Policy Manager Accounts dialog box displays the following information about ClearPass accounts that are already associated with Orchestrator:

Field	Definition
Edit	Click the icon to edit your CPPM instance.
Name	Name of your CPPM instance.
Domain/IP	Domain or URL of your CPPM instance.
Connectivity	Status of the connection between Orchestrator and your CPPM instance. The status may appear as **Connected**, **Connecting**, **Auth Failed**, and **Unreachable**.
Service Status	Status of your CPPM instance. A status other than Connected could indicate a problem with your CPPM configuration. To troubleshoot, click the **Info** icon, and then reset any service that is not currently connected.
Pause	To pause the connection for your CPPM instance, click this toggle.

Add a ClearPass Policy Manager Server

Follow the steps below to add a new ClearPass Policy Manager account.

1. If not already opened, click **Accounts** to open the ClearPass Policy Manager Accounts dialog box.
2. Click **+Add New Server**.

   The ClearPass Policy Manager Server Configuration dialog box opens.
3. Enter the following information:

Field	Definition
Name	Name of your CPPM instance.
Domain/IP	Domain or URL of your CPPM instance.
Client ID	Client ID generated from your CPPM account.
Secret Key	Secret key generated from your CPPM account.
Verify server certificate	If you are using cloud instances of both CPPM and Orchestrator, or if you are using an on-premise instance of CPPM with a valid certificate, select this check box. If you are using an on-premise instance of Orchestrator or an on-premise instance of CPPM without a valid certificate, clear this check box.

4. Click **Save**.

Your CPPM instance now appears in the ClearPass Policy Manager Accounts dialog box. The Connectivity and Service Status fields should both appear as **Connected**.

Edit a ClearPass Policy Manager Server

1. If not already opened, click **Accounts** to open the ClearPass Policy Manager Accounts dialog box.
2. Click the **Edit** icon next to the instance you want to edit.

   The ClearPass Policy Manager Server Configuration dialog box opens.
3. Edit the information in the dialog box, and then click **Save**.

Pause ClearPass Policy Manager Integration

To pause the integration between CPPM and Orchestrator, click **Pause Orchestration** from the ClearPass Policy Manager tab.

NOTE: Clicking **Pause Orchestration** pauses the connection between all instances of CPPM configured in Orchestrator. To pause an individual instance, click **Accounts**, and then click the toggle under Pause for the instance you want to pause.

Mist

Orchestrator > HPE Integration Services > Mist

Orchestrator release 9.7.2 integrates EdgeConnect SD-WAN appliances into the Mist cloud as a supported WAN Edge device. This initial integration phase includes the following functionality:

- Onboard EdgeConnect appliances into Mist using the Adopt WAN Edge method.

- Monitor EdgeConnect appliances within your Mist-managed infrastructure without alternating between Orchestrator and Mist.
- Use the Mist capabilities applicable to EdgeConnect, including Insights, Service Level Experiences (SLEs), Security Events, Marvis AI, and Premium Analytics.

NOTE: EdgeConnect appliances are currently managed in Orchestrator and can be monitored in both Mist and Orchestrator.

Prerequisites

Before you onboard an EdgeConnect appliance in Mist, ensure the following conditions are met:

- EdgeConnect appliances have EdgeConnect OS (ECOS) 9.7.2.0 or later installed.
- Orchestrator has release 9.7.2 or later installed.
- A WAN Assurance subscription has been activated in the Mist Organization. See Activate a Subscription. To understand how your EdgeConnect appliance aligns to its WAN Assurance Subscription class, see Subscriptions for EdgeConnect Gateways.
- Configure a Site in Mist for your EdgeConnect appliance(s), if one doesn't already exist. You should create one site for each EdgeConnect appliance, or one site for each EdgeConnect HA pair.

Onboard EdgeConnect in Mist

Onboarding EdgeConnect appliances in Mist is a three-step process:

1. Add the registration code to Orchestrator using one of the following methods:
   - Use Automation to Add the Registration Code to Orchestrator
   - Manually Add the Registration Code to Orchestrator
2. Enable WAN Assurance in Orchestrator
3. Assign the EdgeConnect Appliance to a Site in Mist

Add the Registration Code to Orchestrator

To add EdgeConnect appliances to Mist, you must have a Mist registration code saved in Orchestrator. The registration code enables WAN Assurance for the appliance. There are two methods for adding the registration code in Orchestrator and onboarding EdgeConnect appliances in Mist:

- Use Automation to Add the Registration Code to Orchestrator

  The automated process uses Mist Authentication to add EdgeConnect appliances to Mist.

  NOTE: This method only supports username/password-based authentication, not single-sign on authentication.
- Manually Add the Registration Code to Orchestrator

  For the manual process, you copy the registration code in Mist and then add it to Orchestrator.

Use Automation to Add the Registration Code to Orchestrator

1. Log in to Orchestrator as an admin or a user with administrator privileges.
2. Navigate to **Orchestrator > HPE Integration Services > Mist**.
3. Click **Add EdgeConnect to Mist** and select **Mist Organization** from the drop-down list.

   The Sign in to Mist Cloud dialog box opens.
4. Enter the email address you use to sign in to your Mist cloud account and click **Next**.
5. Select the Mist cloud environment to sign into and click **Next**.
6. Enter the password for your Mist cloud account and click **Sign in**.
7. If you have two-factor authentication enabled, enter the authentication code you received from your authenticator application and click **Verify**.
8. Select the organization within your Mist cloud account that you want to add the EdgeConnect appliance to.
9. Click **Save**.
10. Click **Close**. Continue to Enable WAN Assurance in Orchestrator.

Manually Add the Registration Code to Orchestrator

Complete the following steps to obtain the registration code in Mist:

1. On the Mist portal, from the left menu, navigate to **Organization > Inventory > WAN Edges**.
2. Select the EdgeConnect appliance you want to onboard by selecting the check box to the left.
3. Click the **Adopt WAN Edges** button on the Inventory page. This is where you obtain the code to enter into Orchestrator.

   NOTE: The “Claim Code or QR Code” method of claiming process in Mist does not exist for EdgeConnect today.
4. In the WAN Edge Adoption pop-up window, select the **EdgeConnect** radio button.
5. Copy the code (this will be pasted into Orchestrator in the next steps).

Complete the following steps to add the registration code in Orchestrator:

1. Log in to Orchestrator as an admin or a user with administrator privileges.
2. Navigate to **Orchestrator > HPE Integration Services > Mist**.
3. Click **Add EdgeConnect to Mist** and select **Registration code** from the drop-down list.

   The Registration dialog box opens.
4. Enter the registration code you obtained from Mist.
5. Click **Save**. Continue to Enable WAN Assurance in Orchestrator.

Enable WAN Assurance in Orchestrator

1. From the appliance tree, select the appliances to onboard in Mist.
2. Click **Enable WAN Assurance** to begin the onboarding process in Mist.
3. To view details and progress about the onboarding process for an appliance, click the info icon (**i**) in the Details column.

   The Device details dialog box opens. It outlines the tasks in the onboarding process. When a task in the onboarding process is completed, a check mark appears next to it. When all onboarding tasks are completed, the message *Enabled* appears at the top of the dialog box.
4. Click **X** to close the Device details dialog box.

   The Status column for the appliance now shows *Ready for Site Assignment in Mist*.

Assign the EdgeConnect Appliance to a Site in Mist

1. Return to the Mist portal. From the left menu, navigate to **Organization > Inventory**, and you should see the "unassigned" EdgeConnect appliance listed.
2. Assign the EdgeConnect appliance to a site by selecting the check box to the left of it, then select **More > Assign to Site**.

   NOTE: You should assign one EdgeConnect appliance per site, or one EdgeConnect HA pair per site.
3. In the pop-up window, select the site to assign the EdgeConnect appliance to, then click **Assign to Site**.

   The appliance should now display a site in the Site column in the Mist inventory.

NOTE: After assigning an EdgeConnect appliance to a site in Mist, verify the status of the EdgeConnect appliance in Orchestrator by navigating to **Orchestrator > HPE Integration Services > Mist**. The appliance status should show *Connected*.

Disable WAN Assurance in Orchestrator

You can disable an EdgeConnect appliance in Orchestrator in order to stop the connection to Mist, which stops the telemetry flow.

1. Navigate to **Orchestrator > HPE Integration Services > Mist**.
2. Select the check box for the EdgeConnect appliance you want to disable.
3. Click **Disable WAN Assurance**.

NOTE: This step does not "release" the appliance from the Mist inventory, but is a required first step if you want to release an EdgeConnect appliance from the Mist organization, which releases the subscription from the appliance. See the Release a Device from a Mist Organization documentation.

Release a WAN Edge Device from a Mist Organization

If you need to fully remove an EdgeConnect appliance from Mist, you can do so by using the Release workflow.

NOTE: For EdgeConnect appliances, you must Disable WAN Assurance in Orchestrator before you can release the appliance from the Mist organization.

1. On the Mist portal, from the left menu, navigate to **Organization > Inventory > WAN Edges**.
2. Select the check box to the left of the WAN Edge device you want to release.
3. Click **More > Release**.
4. Follow the on-screen prompts to complete the action.
5. After you've released the appliance from Mist, in Orchestrator, navigate to **Orchestrator > HPE Integration Services > Mist** to verify that the appliance was successfully released.
6. In the Details column for the appliance, click the info icon (**i**) to open the Device details dialog box.

   If the appliance is fully released from the Mist organization, no information about the appliance appears in the Device details dialog box (for example, you should not see an appliance ID).

   If information about the appliance appears in the dialog box, it is still assigned to a Mist organization. Use the EdgeConnect CLI to fully release the appliance:

   1. Access the CLI by either right-clicking an appliance in the Orchestrator appliance tree, and then selecting **CLI Session**, or by opening an SSH session to the appliance.
   2. Enter *en* to enable the CLI.
   3. Enter *config terminal*
   4. Enter *hpe-integration mist release*
   5. To verify that the appliance was successfully released, click the info icon (**i**) to open the Device details dialog box. Confirm that the appliance information is no longer present.

   This removes (deletes) the appliance from the Mist organization's inventory, and as a result, appliance information is no longer displayed. This action also releases the subscription from the appliance.

Monitor Mist Connections in Orchestrator

After you successfully onboard EdgeConnect appliances in Mist, you can monitor appliances in Mist or use the Mist tab in Orchestrator to:

- Disable WAN Assurance for an appliance that has been onboarded to Mist.
- Verify that an appliance remains connected to Mist.

The following table describes the columns shown on the Mist tab.

Column	Description
Hostname	EdgeConnect appliance hostname.
WAN Assurance	Indicates if WAN Assurance is enabled or disabled for the appliance.
Status	Indicates if the appliance is onboarded and connected to Mist (Connected, Connection failed, or Ready for Site Assignment in Mist).
Device ID	Unique identifier for the EdgeConnect appliance assigned by Mist.
Telemetry	Click the info icon (**i**) to open the Last telemetry dialog box. In the dialog box, the date and timestamp of the last telemetry to Mist is noted for the following data about the appliance: OC-Stats, Heartbeat, App-Summary, OC-Event, Security-Event. This information is helpful for HPE Networking Technical Assistance Center (TAC) troubleshooting.
Details	Click the info icon (**i**) to open the Device details dialog box and view details about the EdgeConnect appliance and its connection to Mist.

HPE Aruba Networking Central

Orchestrator > HPE Integration Services > HPE Aruba Networking Central

Use the HPE Aruba Networking Central dialog box to enable HPE Aruba Networking Central. Enabling HPE Aruba Networking Central in SD-WAN Orchestrator is one step of the process to create a Unified Fabric by integrating EdgeConnect SD-WAN gateways into HPE Aruba Networking Central and HPE GreenLake.

Complete instructions for integrating EdgeConnect gateways with HPE GreenLake are here: Unified Fabric Deployment Guide

Overview of the Unified Fabric

When HPE Aruba Networking Central is enabled, appliances designated as a hub (**Configuration > Hub**) will attempt to open WebSocket connections to HPE Aruba Networking Central. If the appliances have been claimed, licensed, and assigned to an HPE Aruba Networking Central region in HPE GreenLake, they will be eligible to be configured as a VPN concentrator (VPNC)

in HPE Aruba Networking Central. This integration allows EdgeConnect gateways to act as a hub with the EdgeConnect SD-WAN branch gateways and Microbranch access points. It also enables Orchestrator to share the routes and tunnel information from Central into the EdgeConnect SD-WAN fabric.

The EdgeConnect SD-WAN receives routes to the tunnels from the overlay tunnel orchestrator (OTO) and the overlay route orchestrator (ORO) from Central. However, the configuration and deployment of the EdgeConnect SD-WAN gateways still come from the SD-WAN Orchestrator. Central only provides OTO and ORO.

Prerequisites

The following prerequisites are required to deploy the Unified Fabric solution.

- You must have an HPE GreenLake Cloud Platform (GLCP) account.
- You must have a GLCP Workspace.
- You must have Central version 2.5.8 installed with an active account.
- You must have an EdgeConnect SDWAN Foundation subscription license.
- You must have Orchestrator deployed and running version 9.5 or later and gateways running version ECOS 9.5 or later.

Enable HPE Aruba Networking Central

To enable HPE Aruba Networking Central:

1. Navigate to **Orchestrator > HPE Integration Services > HPE Aruba Networking Central**.

   The Enable HPE Aruba Networking Central dialog box opens.
2. Click **Enable HPE Aruba Networking Central**.
3. Complete the following elements:

Field	Description
Segment	Automatically set to “default” and cannot be changed.
Unified fabric zone	Select the zone for all OTO traffic. This allows you to map Unified Fabric traffic to your preferred network zones.
Unified fabric fallback role	Select the role for all OTO traffic. This ensures that Unified Fabric traffic has a consistent default role assigned to it.

4. Click **Save**.
5. Click **OK**.
6. Follow the steps in the Unified Fabric Deployment Guide to complete the Unified Fabric integration.

Support

The menus under **Support** provide troubleshooting tools and different options for working with Support, including opening a support case. You can use these menus to gather information to help Support troubleshoot issues. These menus are organized as follows:

- Technical Assistance
- User Documentation
- Reporting

Support > Technical Assistance

The options under **Support > Technical Assistance** provide resources for working with Support, such as logging in to the Support Portal, creating support cases, uploading files, capturing packets from appliances, and running the RMA Wizard, which automates exchanging and replacing appliances.

Tech Support - Appliances

Support > Technical Assistance > Tech Support - Appliances

Use this tab to create a new case, generate a system dump, upload files to an existing case, or download selected files to Orchestrator.

By default, the table displays all files available on the selected appliances. Click the appropriate button to filter files by type (Logs, Sys Dump, Snapshot, TCP Dump). The table includes the following details for each file:

Field	Description
Appliance	Name of the appliance on which the file is available.
File type	Specific file type (log, sys dump, snapshot, or TCP dump).
File Name	Name of the file.
Last Modified	Date when the file was last modified.
File Size	Size of the file.

Download to Orchestrator

Complete the following steps if you want to download one or more files to Orchestrator.

1. Select one or more files in the table (use Ctrl or Shift to select multiple files).
2. Click the **Download to Orchestrator** button above the table.

3. When prompted, click **Download** to confirm or click **Close** to cancel.

   The Monitor Transfer Progress window appears, showing the status of current and previous downloads.

4. To stop any downloads that are not yet finished, click **Cancel**.

   NOTE: To access any files that have been downloaded, open the **Tech Support - Orchestrator** tab under the Support menu. After selecting one or more files, you can create a new case, upload files to an existing case, or download files to your local machine.

Tech Support - Orchestrator

Support > Technical Assistance > Tech Support - Orchestrator

This tab displays a list of Orchestrator log files and system dump files, as well as support files that have been downloaded from appliances. You can use these files to create or update support cases, or you can download files to your local machine from Orchestrator.

By default, the table displays all files available on Orchestrator. Click the appropriate button to filter files by type (logs, system dumps, or appliance files). The table includes the following details for each file:

Field	Description
Source	Source of the selected file (Orchestrator or a specific appliance).
File Type	Specific file type (log, sys dump, snapshot, or TCP dump).
File Name	Name of the file.
Last Modified	Date when the file was last modified.
File Size	Size of the file.

Take Action with Files

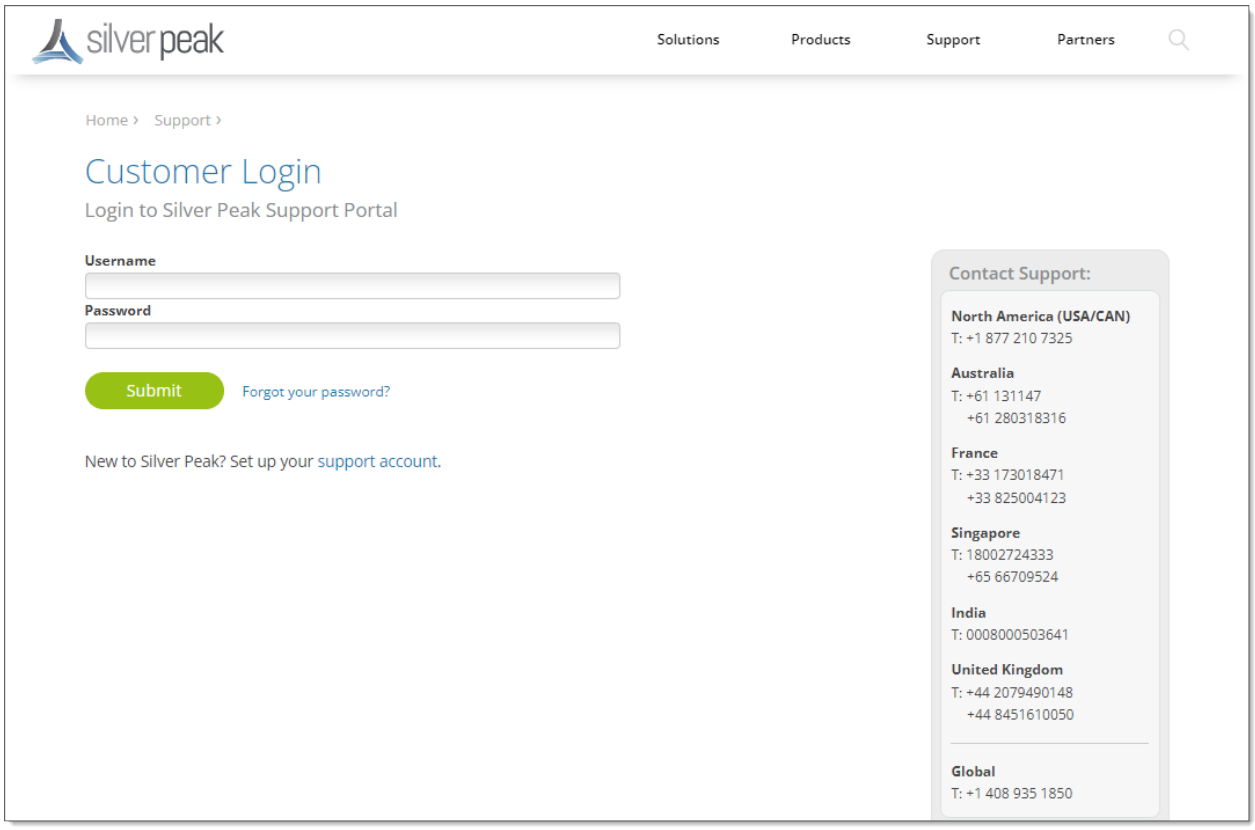
With one or more files selected, you can create a new support case, add files to an existing case, or download files to your local machine.

- Click **Create Case** to open a new support case. Fill in a few additional details and the selected files will be attached to a new support case.
- Click **Upload Selected Files** to attach files to an existing support case. You will need to know the case number when using this option.
- Click **Download selected Files** to download files to your local machine. Confirm the download and select a location where you want to save the files.

Log In to the Support Portal

Support > Technical Assistance > Support Portal Log-in

If you have a Silver Peak account and need technical assistance or customer support, select **Support > Technical Assistance > Support Portal Log-in**. The following page opens in a separate browser tab.

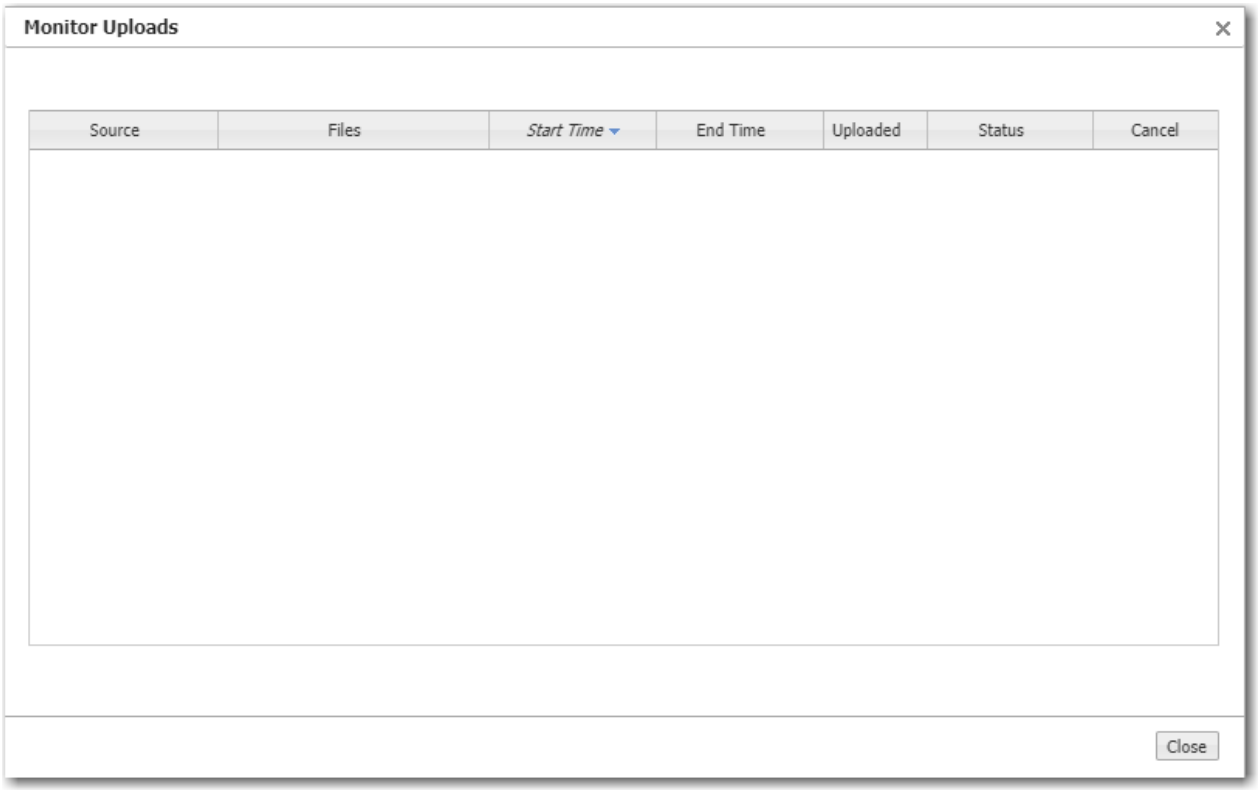


You can also access this page by navigating to Silver Peak's web page and selecting **Support > Customer Login** from the menu bar.

Monitor Transfer Progress

Support > Technical Assistance > Monitor Transfer Progress

This table displays the current status of any files being uploaded to Support.



Packet Capture

Support > *Technical Assistance* > *Packet Capture*

When requested by Support, use this tab to capture packets for appliances that are selected in the appliance tree.

The following table describes each field on this tab.

Field	Description
Maximum number of packets	Enter the maximum number of packets to capture.
Filter by IP	(Optional) Enter the host or IP address to capture from.
Filter by port number	(Optional) Enter the port to capture from. For example, to capture DNS traffic, enter *53*.

Field	Description
Filter by interface or tunnel	(Optional) Specify the tunnels or interfaces to capture from. To filter by tunnel, click **Tunnels**, and then click the appropriate tunnel. To filter by interface, click **Interface**, and then click the appropriate interface. You can use the Ctrl or Shift key to select multiple tunnels or interfaces. **NOTE:** This field is not displayed if more than one appliance is selected in the appliance tree.
Bytes to capture from each packet (snap length)	Enter the number of bytes (the amount of data) for each frame to capture. For example, enter *96* to capture headers only or *1500* to capture full frames. **NOTE:** Configuring a large snap length will result in larger packet capture file sizes.
Additional filter options	(Optional) Enter other options to filter the capture. For example, *host 1.1.1.1 and host 10.0.0.1*
Enable circular storage	Select this check box to limit the amount of data to store by setting a maximum number of files and maximum file size for the capture. For example, set the **Number of files** to *5* and **Max size per file** to *100* (MB). When the size limit is reached for a file, a new file will be written. Once the maximum number of files is reached, the oldest file will be overwritten.
Number of files	If you enabled circular storage, enter the maximum number of files that can be stored for this packet capture.
Max size per file	If you enabled circular storage, enter the maximum file size that can be stored for this packet capture.
Command preview	Displays the progress of the packet capture.

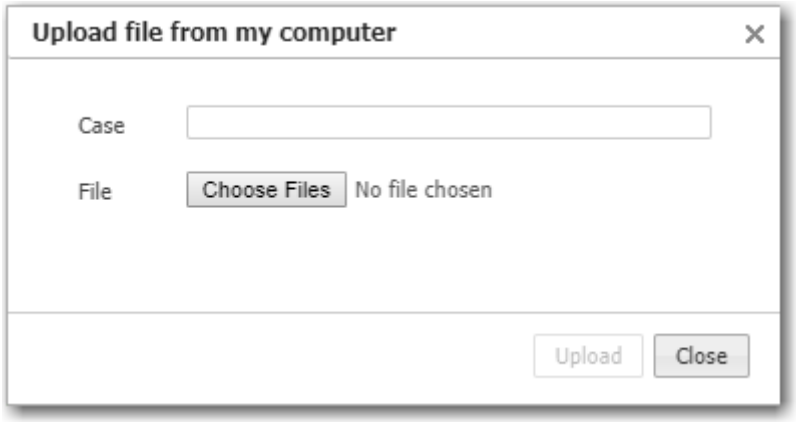
Click **Run** to initiate the packet capture.

Click **Stop** to cancel the packet capture.

Upload Local Files

Support > Technical Assistance > Upload Local Files

Use this dialog box to upload files related to your Support case from your computer.

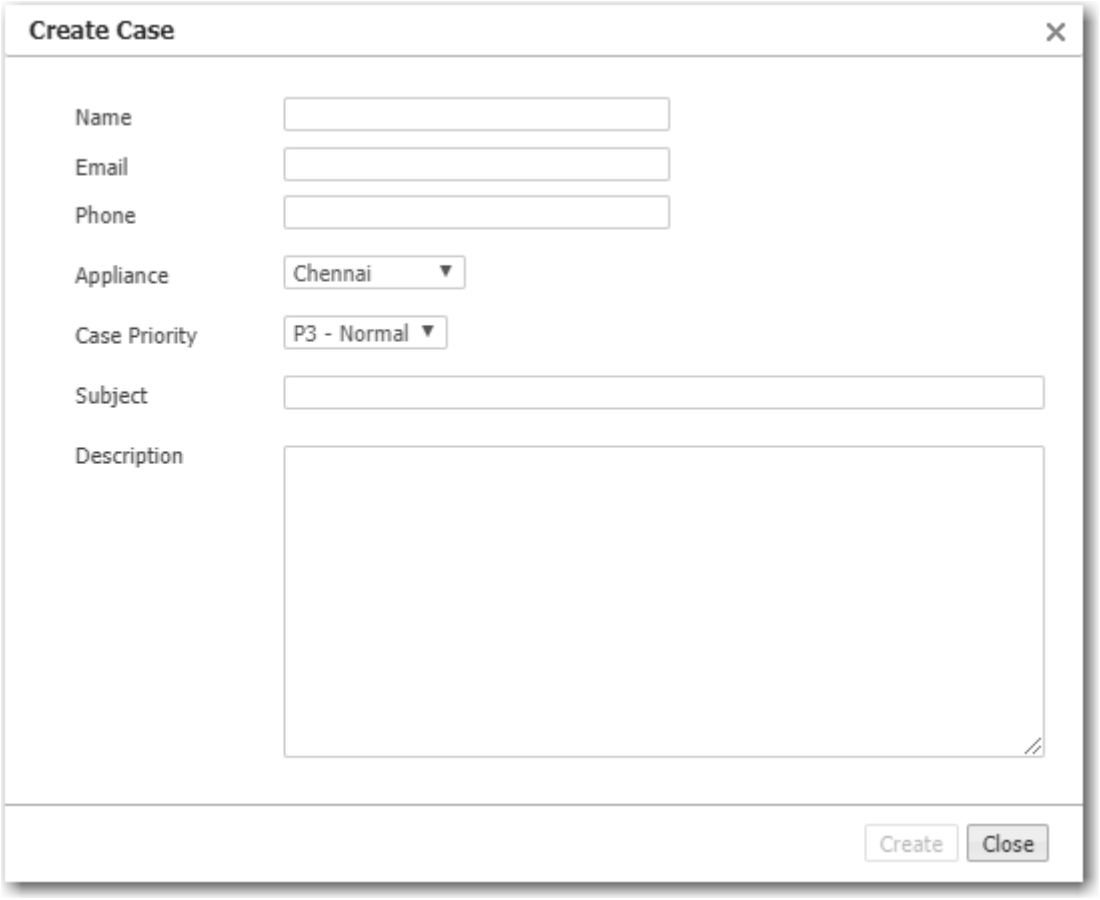


Create a Support Case

Support > *Technical Assistance* > *Create Case*

Use the Create Case dialog box to create a Support case.

You will receive a case number and instructions for what to do next.



Partition Management

Support > Technical Assistance > Partition Management

Use this tab to regain Orchestrator disk space by selectively eliminating statistics no longer needed.

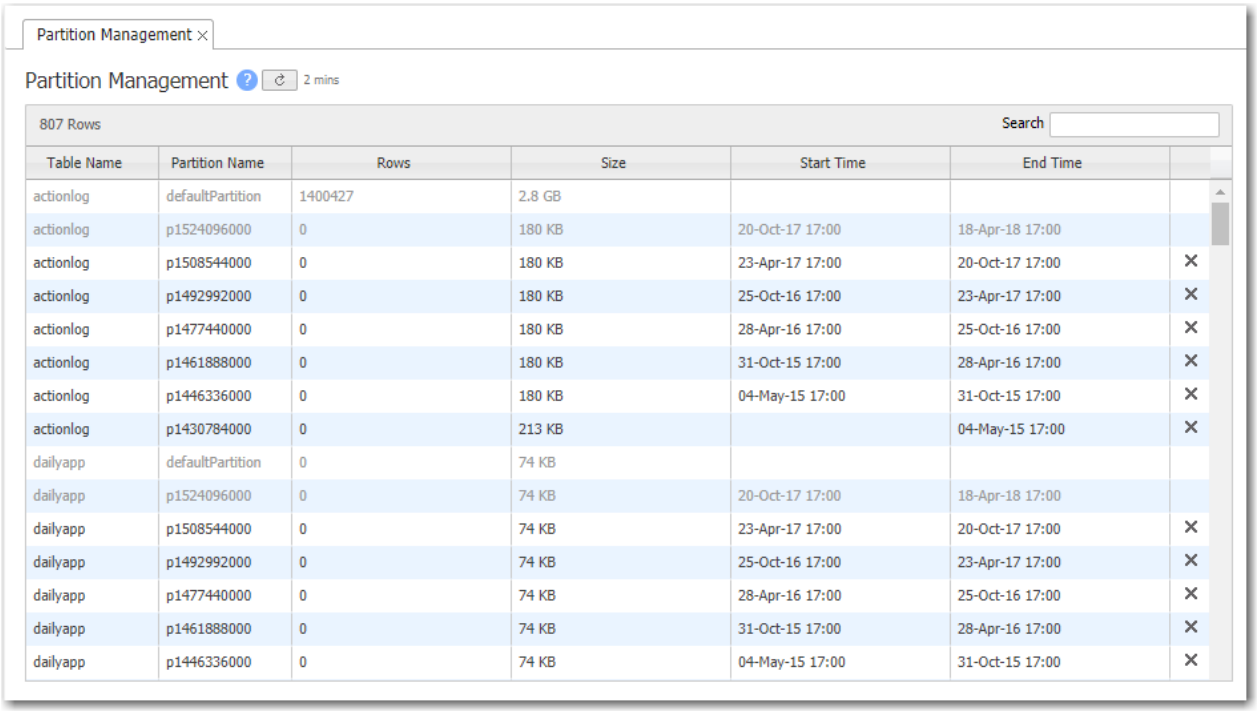


Table Name	Partition Name	Rows	Size	Start Time	End Time
actionlog	defaultPartition	1400427	2.8 GB		
actionlog	p1524096000	0	180 KB	20-Oct-17 17:00	18-Apr-18 17:00
actionlog	p1508544000	0	180 KB	23-Apr-17 17:00	20-Oct-17 17:00
actionlog	p1492992000	0	180 KB	25-Oct-16 17:00	23-Apr-17 17:00
actionlog	p1477440000	0	180 KB	28-Apr-16 17:00	25-Oct-16 17:00
actionlog	p1461888000	0	180 KB	31-Oct-15 17:00	28-Apr-16 17:00
actionlog	p1446336000	0	180 KB	04-May-15 17:00	31-Oct-15 17:00
actionlog	p1430784000	0	213 KB		04-May-15 17:00
dailyapp	defaultPartition	0	74 KB		
dailyapp	p1524096000	0	74 KB	20-Oct-17 17:00	18-Apr-18 17:00
dailyapp	p1508544000	0	74 KB	23-Apr-17 17:00	20-Oct-17 17:00
dailyapp	p1492992000	0	74 KB	25-Oct-16 17:00	23-Apr-17 17:00
dailyapp	p1477440000	0	74 KB	28-Apr-16 17:00	25-Oct-16 17:00
dailyapp	p1461888000	0	74 KB	31-Oct-15 17:00	28-Apr-16 17:00
dailyapp	p1446336000	0	74 KB	04-May-15 17:00	31-Oct-15 17:00

Orchestrator Remote Log Receivers

Support > Technical Assistance > Orchestrator Remote Log Receiver

This page describes how to configure remote log receivers to forward Orchestrator alarm and audit log events to third-party collection endpoints. When an event occurs, Orchestrator sends the event message to all enabled receivers registered on this page. For information about the data contained in remote log messages, see Remote Log Messages.

Supported Receiver Types

Receiver Type	Protocol	Message Format
HTTP	HTTP POST	JSON
HTTPS	HTTP POST	JSON
KAFKA	Kafka producer	JSON
SYSLOG	TCP, UDP, or TCP over SSL	RFC 5424
WEBSOCKET	WSS	JSON

Receiver Type	Protocol	Message Format

Add a Remote Log Receiver

You can add HTTP/HTTPS, KAFKA, SYSLOG, and WEBSOCKET servers as remote log receivers.

Use one of the following options:

- To add an HTTPS receiver, click **HTTPS**. HTTPS is the default receiver type.
- To add a different receiver type, click the **more (...)** menu next to the list of receiver types, and then select the receiver type.

Enter the appropriate information for the receiver, and then click **Save**.

Each receiver entry in the table includes an Actions column on the right with controls to edit or delete the receiver.

Configure an HTTP Receiver

When you add or edit an HTTP receiver, complete the following fields in this order:

Field	Description
Enable receiver	Click this slider to toggle between enabled and disabled state.
Name	Name of the receiver the logs are sent to.
Log type	Select **Audit Log** or **Alarm**.
URL	URL served by the HTTP log server that Orchestrator uses to send log data with POST REST calls.
User name (optional)	User name used in basic authentication when making REST calls.
Password (optional)	Password used in basic authentication when making REST calls.
Repeat password (optional)	Re-enter the password used in basic authentication.
Header name (optional)	Optional custom HTTP header name required by the third-party receiver.
Header value (optional)	Optional value for the custom HTTP header.
Data wrapper key (optional)	Optional key used to wrap the JSON event payload, as required by the third-party receiver.

Configure an HTTPS Receiver

When you add or edit an HTTPS receiver, complete the following fields in this order:

Field	Description
Enable receiver	Click this slider to toggle between enabled and disabled state.
Name	Name of the receiver the logs are sent to.
Log type	Select **Audit Log** or **Alarm**.
URL	The endpoint URL of the remote HTTPS server.
User name (optional)	User name used in basic authentication when making REST calls.
Password (optional)	Password used in basic authentication when making REST calls.
Repeat password (optional)	Re-enter the password used in basic authentication.
Header name (optional)	Optional custom HTTP header name required by the third-party receiver.
Header value (optional)	Optional value for the custom HTTP header.
Data wrapper key (optional)	Optional key used to wrap the JSON event payload, as required by the third-party receiver.

NOTE: The values for Header Name, Header Value, and Data Wrapper Key vary depending on the third-party receiver you are integrating with (for example, Splunk). Refer to your third-party receiver's documentation for the correct values to enter in these fields.

Configure a KAFKA Receiver

When you add or edit a KAFKA receiver, complete the following fields in this order:

Field	Description
Enable receiver	Click this slider to toggle between enabled and disabled state.
Name	Name of the receiver the logs are sent to.
Log type	Select **Audit Log** or **Alarm**.
Topic	Topic name on the KAFKA receiver.
Bootstrap servers	Comma-separated list of KAFKA broker addresses in `<ip>:<port>` format.

Field	Description
Acks (optional)	Defines the number of KAFKA servers that must acknowledge a message before it is considered delivered. Values include `acks=0`, `acks=1`, or `acks=all`.
Retries (optional)	Number of retry attempts KAFKA makes before returning an error.
Batch size (optional)	Number of bytes produced in a batch before KAFKA sends the batch.
Buffer size (optional)	Maximum memory size available for buffering messages before send operations are blocked.
Linger time (optional)	Time KAFKA waits before sending the next message batch.

Configure a SYSLOG Receiver

When you add or edit a SYSLOG receiver, complete the following fields in this order.

General settings:

Field	Description
Enable receiver	Click this slider to toggle between enabled and disabled state.
Name	Name of the receiver the logs are sent to.
Log type	Select **Audit Log** or **Alarm**.
Protocol	Select **TCP**, **UDP**, or **TCP over SSL**.
Hostname/URL	Hostname or URL of the remote SYSLOG server.
Port	Port number on the SYSLOG server that accepts incoming events.
Custom data (optional)	Optional token embedded in the SYSLOG message for third-party integrations.

Facility settings:

Field	Description
Audit facility	SYSLOG facility used for audit log messages.
Alarm facility	SYSLOG facility used for alarm messages.

Audit severity settings:

Field	Description
Audit Error severity	Maps Orchestrator Error audit events to an external SYSLOG severity.
Audit Info severity	Maps Orchestrator Info audit events to an external SYSLOG severity.
Audit Debug severity	Maps Orchestrator Debug audit events to an external SYSLOG severity.

Alarm severity settings:

Field	Description
Alarm Critical severity	Maps Orchestrator Critical alarms to an external SYSLOG severity.
Alarm Major severity	Maps Orchestrator Major alarms to an external SYSLOG severity.
Alarm Minor severity	Maps Orchestrator Minor alarms to an external SYSLOG severity.
Alarm Warning severity	Maps Orchestrator Warning alarms to an external SYSLOG severity.

Orchestrator sends SYSLOG messages in RFC 5424 format.

Configure a WEBSOCKET Receiver

When you add or edit a WEBSOCKET receiver, complete the following fields in this order:

Field	Description
Enable receiver	Click this slider to toggle between enabled and disabled state.
Name	Name of the WEBSOCKET receiver.
Log type	Select **Audit Log** or **Alarm**.
IP allow list	List of source IP addresses that are allowed WEBSOCKET access to Orchestrator.

After completing the fields, click **Generate key** to create the WEBSOCKET key.

To connect to the Orchestrator WEBSOCKET, do one of the following:

- Include the key in the HTTP header *X-Auth-Token* when connecting to the WEBSOCKET.

- Append the key as the query parameter *key* to the WEBSOCKET URL.

The generated WEBSOCKET key can be copied from the dialog and provided to the connecting client.

Event Message Retention

For an on-premises deployment, remote log resubmission availability depends on the current Statistics Retention settings for the non-statistics tables *alarmnotification2* and *actionlog*. To configure these settings, go to Orchestrator > Software & Setup > Setup > Statistics Retention, and then select **Others**.

For orchestrator as a service, .all stats retention are HPE and license governed.

Each message includes a consecutive unique sequence ID that the receiver can use to detect and catch up on missed events.

Remote Log Messages

If you register a remote log receiver, Orchestrator electronically sends a log message to the receiver when certain events take place, such as alarms. The remote log messages that are sent by Orchestrator provide details about the events and are sent in either JSON format or RFC5424 syslog format depending on the type of server you register as a receiver. For information on how to register a remote log receiver, see Remote Log Receivers.

JSON Format

For HTTP, HTTPS, KAFKA, and WEBSOCKET servers, remote log messages are sent using REST POST requests, and the messages are sent in JSON format.

There are two JSON message formats, one for alarm messages and one for audit log messages. The following tables describe the data found in each type of JSON message.

JSON Alarm Message Format

Key	Type	Description
sequenceID	number	The unique ID for the alarm event.
timestamp	UTC	The time at which Orchestrator sent the log message.
hostname	string	The hostname of the Orchestrator.
appName	string	ALARM. Indicates that the message is for an alarm.
severity	string	Indicates the severity of the alarm. You can customize alarm severity in Orchestrator.
msgId	number	An HPE Aruba Networking defined alarm type ID that can be used for parsing the product and alarm category information.

Key	Type	Description
data	JSON	Provides detailed alarm information. **Note:** See JSON Data Key for Orchestrator and Appliance Alarms for detailed information about the data fields that appear in the data key.
message	string	A pipe delimited message that provides brief details about the event.

JSON Audit Log Message Format

Key	Type	Description
sequenceID	number	The unique ID for the audit log event.
timestamp	UTC	The time at which Orchestrator sent the log message.
hostname	string	The hostname of the Orchestrator.
appName	string	AUDIT_LOG. Indicates that the message is for an audit log.
severity	string	Indicates the severity of the audit log event. All audit log events have a severity of INFO'.
msgId	number	-1
data	JSON	Provides detailed audit log information. **Note:** See JSON Data Key for Audit Logs for detailed information about the data fields that appear in the data key.
message	string	A pipe delimited message that provides brief details about the event.

The following tables describe the detailed information contained in the data key for each type of JSON message. The data key contains the bulk of the message data and has detailed information about the alarm or the audit log.

JSON Data Key for Orchestrator and Appliance Alarms

Field	Type	Description
clearable	boolean	Indicates if the alarm can be cleared by the user.
acknowledged	boolear	Indicates if the alarm was acknowledged by the user.
severity	string	Indicates the severity of the alarm. WARNING(1), MINOR(2), MAJOR(3), or CRITICAL(4).
alarmCategory	string	This field is not used and does not contain any data.
source	string	Indicates the module or component that generated the alarm. Some appliance alarms do not have a source defined.

Field	Type	Description
systemId	string	For Orchestrator alarms, this is the Orchestrator hostname. For appliance alarms, this is the unique ID of the appliance (for example, “0.NE” or “1.NE”).
systemHostname	string	The hostname of the Orchestrator or appliance.
alarmId	number	The unique ID of the Orchestrator or appliance alarm.
raisedTime	epoch mil-lisec-onds	The time in UTC at which the alarm was raised.
clearedTime	epoch mil-lisec-onds	The time in UTC at which the the user cleared the alarm. A value of “0” indicates the alarm is still active.
description	string	A description of the alarm.
recommendedA	string	Recommended actions the user can take to clear the alarm.
closed	boolean	Indicates if the alarm has been cleared.

JSON Data Key for Audit Logs

Field	Type	Description
id	number	The unique ID for the audit log event.
user	string	Either the person who or the system that originated the action.
ipAddress	string	The IP address of the Orchestrator.
nepk	string	If the audit log event was an action performed on an appliance, this field shows the unique ID of the appliance.
name	string	The name of the action performed in the audit log event.
description	string	A description of the action performed in the audit log event.
taskStatus	enum	Indicates the status of the action. NOT_STARTED (0), IN_PROGRESS (1), or COMPLETED (2).
startTime	epoch milliseconds	The time in UTC when the action started.
endTime	epoch milliseconds	The time in UTC when the action ended.
logLevel	enum	Indicates the type of log. DEBUG (1), INFO (2), or ERROR (3).

Field	Type	Description
result	string	Indicates the result of the action and usually contains detailed data about the action that was taken.
queuedTime	epoch milliseconds	The time when the action is enqueued for execution.
percentComplete	%	Indicates the percentage of the action that has been completed.
completion status	boolean	Indicates if the action is completed.

RFC5424 Syslog Format

For syslog servers, remote log messages are sent using TCP/UDP, and the messages are sent in RFC5424 syslog format. All RFC5424 syslog remote messages contain the general details described in the following table, as well as a structured data section, which is described in Structured Data.

RFC5424 Syslog Message Details

Field	Type	Description
PRI	number	Indicates the syslog priority.
version	numbe	Indicates the syslog version.
timestamp	UTC	The time at which Orchestrator sent the log message.
hostname	string	The hostname of the originator; the originator will be either an Orchestrator instance or an appliance. Minimally this should be the IPv4/6 address, but ideally it should be a string name, such as “ec-SF-123.”
appName	string	AUDIT_LOG or ALARM. Indicates whether the message is for an alarm or an audit log.
facility	string	Indicates the syslog facility level, as set by the user.
severity	string	Indicates the syslog severity level. For audit log messages, either “Info” or “Debug” appears. For alarm messages, the user can map Orchestrator severity to syslog severity during configuration.
structured data	RFC542	This field contains the sequenceId and detailed alarm or audit log data. Data wrapped in brackets [...], can have one or more structured data elements. **Note:** See Structured Data for detailed information.
msgId	number	For alarms, use the HPE Aruba Networking defined alarm type ID. For audit logs, set to -1.

Field	Type	Description
message	string	A pipe delimited message that provides brief details about the event. Format: “<severity> | <source> | <description> | <userId> | <seqId> | <timestamp> | <src object id> | <target object id> | <result> | <hostname> | <customData>”

Structured Data

Each syslog message contains two structured data elements. The first element contains information about the alarm or audit log, and there are three formats for the first element; one for Orchestrator alarm messages, one for appliance alarm messages, and one for audit log messages. The second structured data element contains the metadata. The following tables describe the information contained in the structured data elements.

RFC5424 Structured Data Element for Orchestrator Alarms

For an Orchestrator alarm, the first structured data element in the syslog message contains detailed Orchestrator alarm information.

Field	Type	Description
structured data id	string	HPE Aruba Networking enterprise number “SP@23867.”
clearable	boolean	Indicates if the alarm can be cleared by the user.
acknowledged	boolean	Indicates if the alarm was acknowledged by the user.
severity	string enum	Indicates the severity of the alarm. WARNING(1), MINOR(2), MAJOR(3), or CRITICAL(4).
alarmCategory	string	This field is not used and does not contain any data.
source	string	Indicates the module or component that generated the alarm (for example, “/orchestration”, “/email/smtp”, “/system/backup”).
systemId	string	The Orchestrator hostname.
systemHostname	string	The Orchestrator hostname.
alarmId	number	The unique ID of the Orchestrator alarm.
raisedTime	epoch milliseconds	The time in UTC at which the alarm was raised.
clearedTime	epoch milliseconds	The time in UTC at which the alarm was cleared by the user. A value of “0” indicates the alarm is still active.
description	string	A description of the alarm.
recommendedAction	string	Recommended actions the user can take to clear the alarm.

Field	Type	Description
closed	boolean	Indicates if the alarm has been cleared.

RFC5424 Structured Data Element for Appliance Alarms

For an appliance alarm, the first structured data element in the syslog message contains detailed appliance alarm information.

Field	Type	Description
structured data id	string	HPE Aruba Networking enterprise number “SP@23867.”
clearable	boolean	Indicates if the alarm can be cleared by the user.
acknowledged	boolean	Indicates if the alarm was acknowledged by the user.
severity	string enum	Indicates the severity of the alarm. WARNING (1), MINOR (2), MAJOR (3), or CRITICAL (4).
alarmCategory	string	This field is not used and does not contain any data.
source	string	Indicates the module or component that generated the alarm (for example, tunnel name “tunnel1” is used for a tunnel down alarm.). Some appliance alarms do not have a source defined.
systemId	string	The unique ID of the appliance.
systemHostname	string	The appliance hostname.
alarmId	number	The unique ID for the appliance alarm.
raisedTime	epoch milliseconds	The time in UTC at which the alarm was raised.
clearedTime	epoch milliseconds	The time in UTC at which the user cleared the alarm. A value of “0” indicates the alarm is still active.
description	string	A description of the alarm.
recommendedAction	string	Recommended actions the user can take to clear the alarm.
closed	boolean	Indicates if the alarm has been cleared.

RFC5424 Structured Data Element for Audit Logs

For audit logs, the first structured data element in the syslog message contains detailed audit log information.

Field	Type	Description
structured data id	string	HPE Aruba Networking enterprise number “SP@23867.”
id	number	The unique ID for the audit log event.
user	string	Either the person who or the system that originated the action.
nepk	string	If the audit log event was an action performed on an appliance, this field shows the unique ID of the appliance.
name	string	The name of the action performed in the audit log event.
description	string	A description of the action performed in the audit log event.
taskStatus	enum	Indicates the status of the action. NOT_STARTED (0), IN_PROGRESS (1), or COMPLETED (2).
startTime	epoch milliseconds	The time in UTC when the action started.
endTime	epoch milliseconds	The time in UTC when the action ended.
logLevel	enum	Indicates the type of log. DEBUG (1), INFO (2), or ERROR (3).
result	string	Indicates the result of the action and usually contains detailed data about the action that was taken.

RFC5424 Meta Structured Data Element

The second structured data element is the metadata.

Field	Type	Description
structured data id	string	The metadata for the event.
sequenceId	number	The unique sequence ID for each event. Alarms and audit log events use a different sequence ID.

Routing Peer Table

Support > Technical Assistance > Routing Peer Table

You can use the Routing Peer Table tab to track the communication between multiple peers within a network and for troubleshooting purposes. This tab also reflects the details of the subnet information being shared between each set of peers.

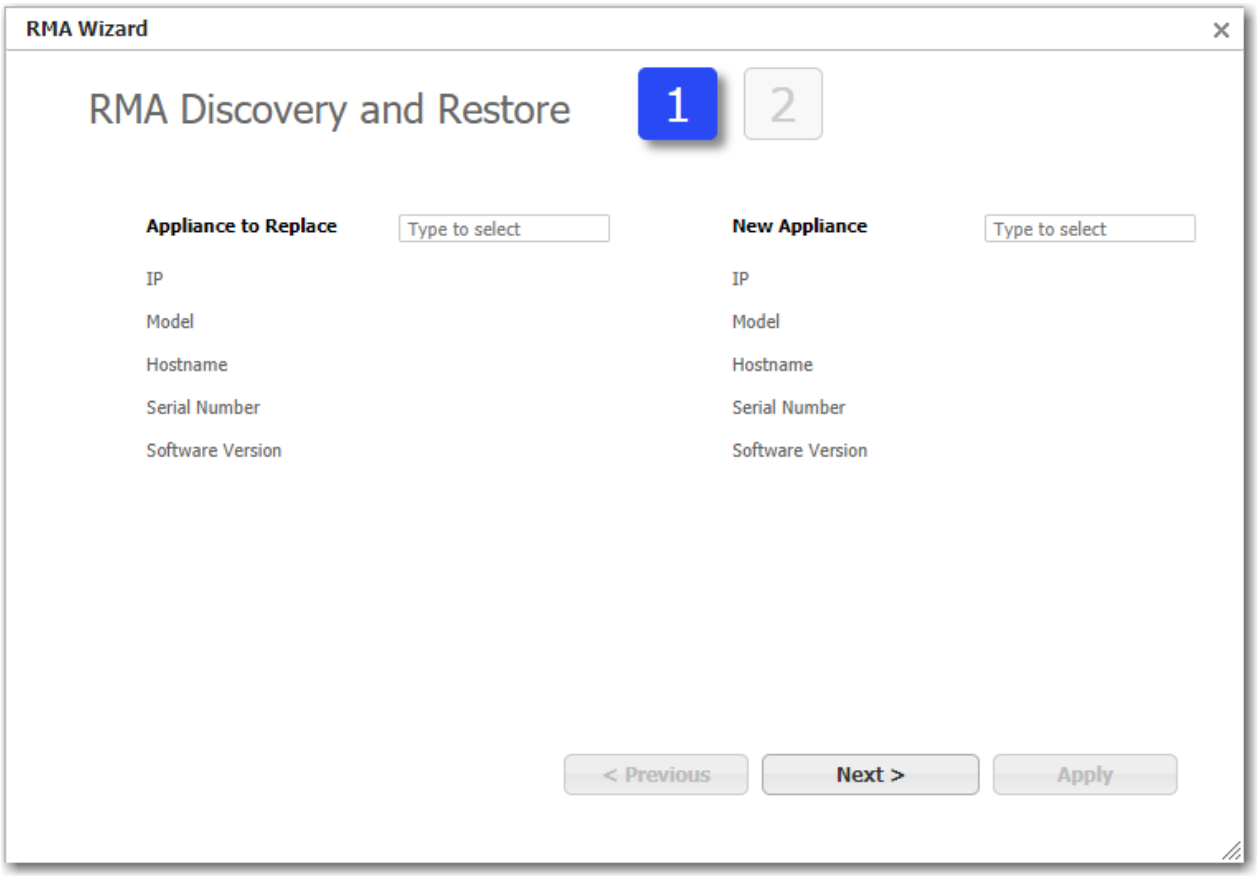
The following table describes the fields on the Routing Peer Table tab.

Field	Description
Appliance	Name of the appliance.
Peer Id	ID of the peer.
Peer Name	Name of the peer.
Role	Whether the hub or spoke topology is being used for the specified peer.
Last transmission count	Last transaction count the peer was sent.
Time since last transmission	How many seconds have elapsed since the last subnet update was sent to the peer.
Last received count	Last transaction count from the peer that was received.
Time since last received	Amount of time since the last received update.
MainVer and Region	Main version and the region of the designated peer.
Message	Peer information to assist Support with troubleshooting.

RMA Wizard

Support > *Technical Assistance* > *RMA*

The RMA (Return Merchandise Authorization) Wizard automates the RMA process for an exchange or replacement of your appliance, if needed. It includes appliance discovery, the version of the appliance, and a backup selection. Use this screen as instructed by Support to prepare an RMA.



Note the following before you begin the RMA process.

- Upgrade or downgrade the new appliance to the same software version before shipping to the site. This will save time.
- Perform a backup of the Orchestrator and EdgeConnect appliances.
- Install the new EdgeConnect appliance onsite.
- When Orchestrator discovers the new device, do not approve it. Start the RMA process to move the license to the new EdgeConnect appliance.

Run the RMA Wizard

Complete the following steps to RMA your appliance.

1. Navigate to the **RMA** tab in Orchestrator.
2. Select the appliance you want to replace from the menu.

   NOTE: The IP address, appliance model, hostname, serial number and software version will auto-populate after you select the appliance.
3. Select the newly discovered appliance that will replace the current appliance.

   NOTE: The IP address, appliance model, hostname, serial number and software version will auto-populate after you select the appliance.

4. Click **Next >**.
5. If you are adding a backup appliance, proceed to the next section. Otherwise, click **Apply**.

   The Applying Configuration dialog box opens and displays the status of the upgrade and restore.

Add a Backup Appliance

If you choose to add a backup appliance from the table, complete the following steps.

1. Select the backup appliance from the table.
2. Select the version you want the backup appliance to have from the drop down menu.

   NOTE: If your selection results in a software downgrade, a backup must be provided.

Upgrade and Downgrade

If the software version you selected for your backup appliance is **higher** than that of the discovered appliance, you will need to do the following:

- Upgrade to the new version using Orchestrator.
- Back up the appliance from a restore, if applicable.

If the software version you selected for your backup appliance is **lower** than that of the discovered appliance, you will need to do the following:

- Install the desired version as a next boot on the appliance.
- Restore from backup.

Support > User Documentation

The options under **Support > User Documentation** provide resources and documentation that support your use of Orchestrator.

Alarm Descriptions

Support > User Documentation > Alarm Descriptions

Orchestrator enables you to export to a CSV file a full list of alarms you could potentially receive. To automatically export the CSV file, navigate to **Support > User Documentation > Alarm Descriptions**.

The CSV file includes the following information:

- Type ID: Unique ID assigned to the alarm.
- Severity: Severity level of the alarm, as follows:

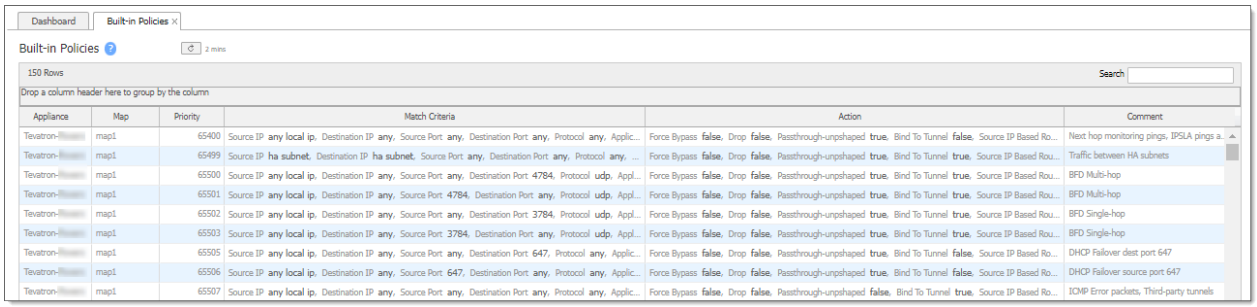
  - Critical: Critical alarms are service-affecting and require immediate attention. They reflect conditions that adversely affect an appliance or indicate the loss of a broad category of service.
  - Major: While service-affecting, major alarms are less severe than critical alarms. They reflect conditions that should be addressed in the next 24 hours. An example would be an alarm caused by an unexpected traffic class error.
  - Minor: Minor alarms are not service-affecting and can be addressed at any time. Examples include alarms caused by a user who has not changed their account's default password, a degraded disk, or a software version mismatch.
  - Warning: Warning alarms are not service-affecting. They warn of conditions that could become problems over time—for example, an alarm caused by IP SLA being down.

- Description: Brief description of the alarm.
- Recommended Action: Recommended actions to take to resolve the alarm.
- Service Affecting: Indicates whether the alarm is service affecting.
- Source: Indicates where the alarm originated.
- System Type: Identifies the type of system the alarm originated from, as follows:
  - 0: EdgeConnect appliance
  - 100: Orchestrator
  - 200: Orchestrator-SP or Orchestrator Global Enterprise
- Source Type: Identifies the alarm category, as follows:
  - 1: Tunnel (applicable to both Orchestrator and appliance alarms)
  - 2: Traffic Class (applicable to appliance alarms only)
  - 3: Equipment
  - 4: Software
  - 5: Threshold Crossing (applicable to appliance alarms only)
- Alarm Type: Indicates an index into the specific alarm category. For example, within the Tunnel alarm category, there is an alarm type associated with index 0 (INTERFACES_WITH_DUPLICATE_IP_EXIST), another with index 1 (INTERFACES_WITH_NO_PUBLIC_IP_EXIST), and so forth. Each alarm type within an alarm category has a unique ID.
- Clearable: Indicates whether you can clear the alarm.

Built-in Policies

Support > User Documentation > Built-in Policies

This table displays read-only built-in policies, which are executed before any other policies.



Support > Reporting

The options under **Support > Reporting** focus on reports that can assist you with troubleshooting.

Dropped Packet Trends

Support > Reporting > Dropped Packet Trends

The Dropped Packet Trends tab displays charts that show dropped packet trends for a single selected appliance. The charts vary based on the cause of the drop.

Charts are available in real time or for a specific time period. Real-time charts show drops over the last five minutes and refresh every five seconds.

You can customize chart settings using the controls at the top of the tab, as follows:

Option	Description
Time period	Click **Realtime** to enable live statistics for all available interfaces. Click a predefined time period (**1h**, **4h**, **1d**, **7d**) to display statistics over the last hour, four hours, day, or seven days. Click **Custom** and set your own custom time range to display statistics for that time period.
Show in UTC	Toggles chart times between local appliance time or UTC.
Lock scale	By default, each chart uses its own scale that is relative to the data displayed. This option applies and locks the same scale to each chart.

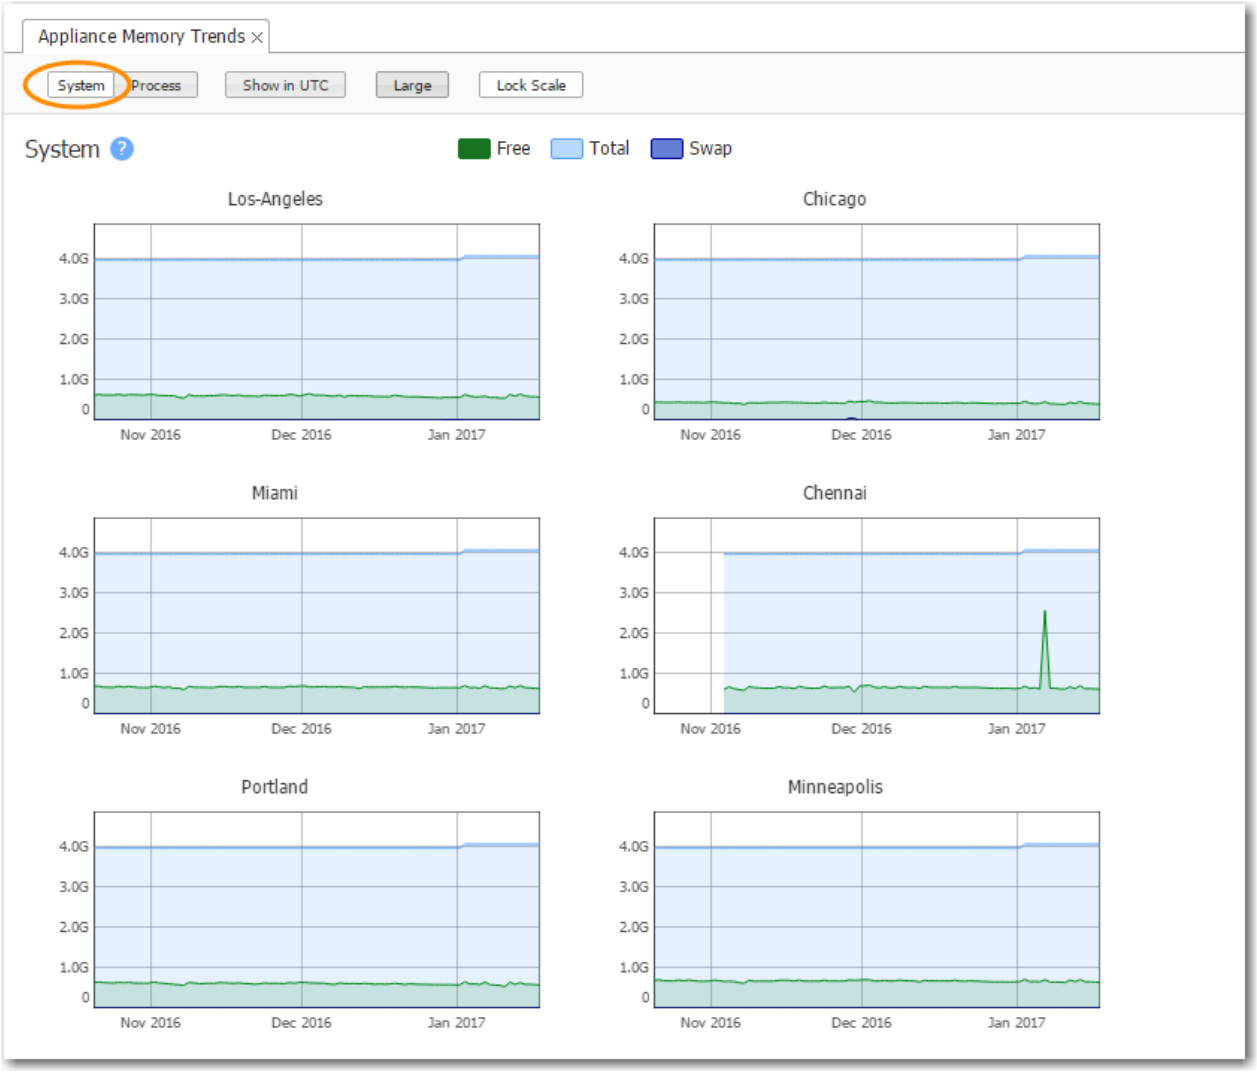
To display a small view of the charts, click the tile view icon. To display a large view, click the large icon.

To refresh data on the tab for the selected time period, click the refresh button.

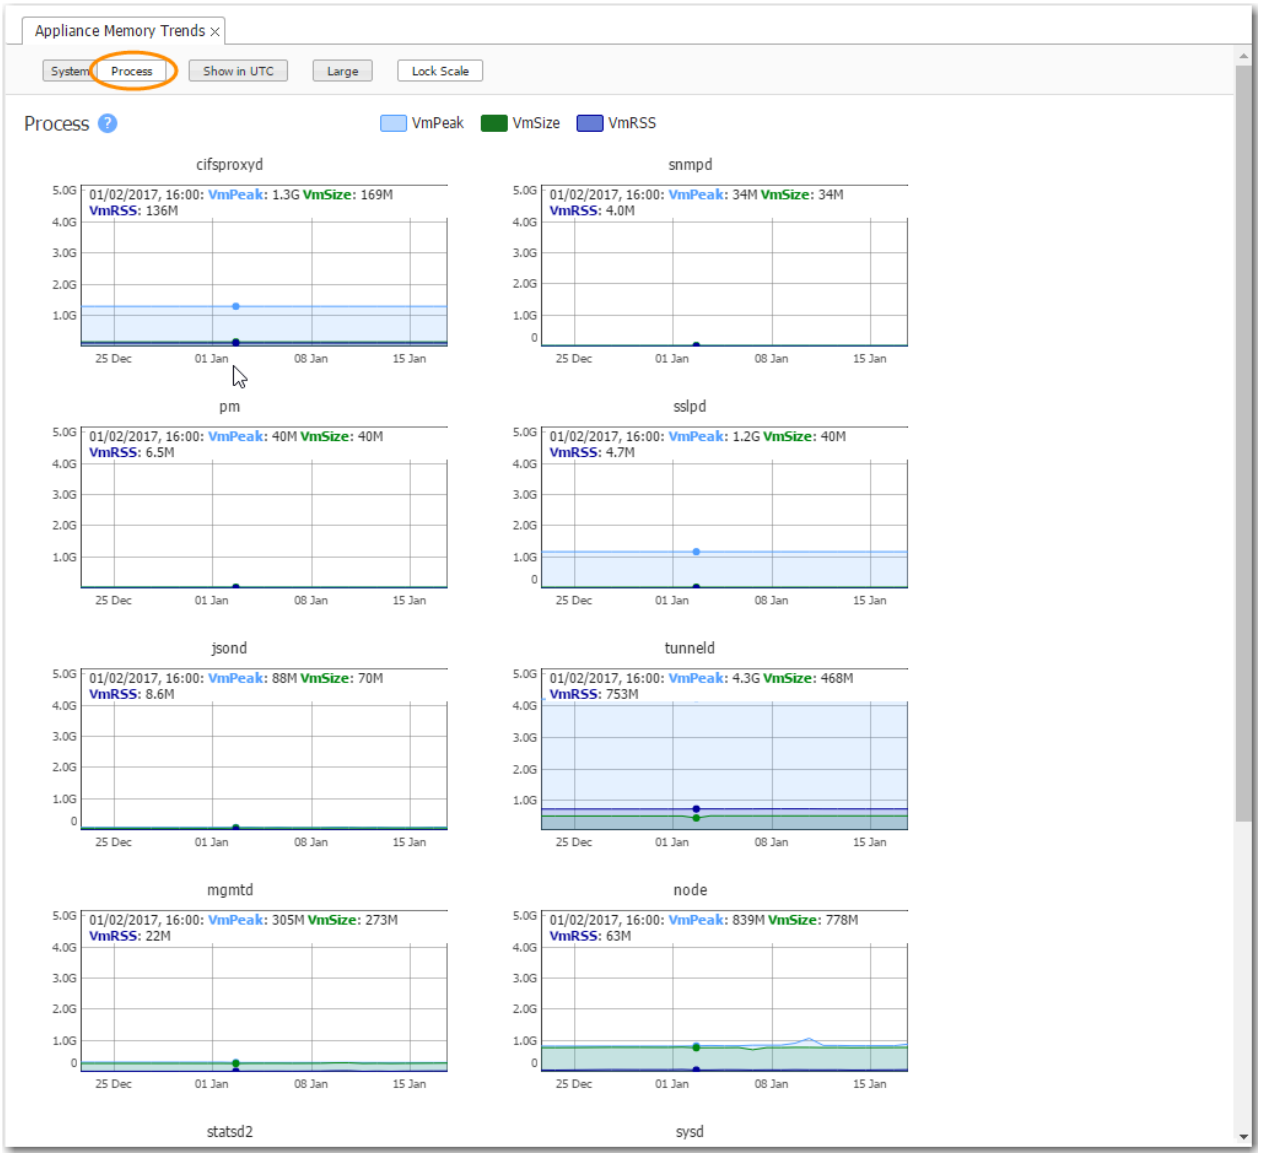
Appliance Memory Trends

Support > Reporting > Appliance Memory Trends

The **System** view shows appliance daily memory usage.



The **Process** view is for individual appliances.

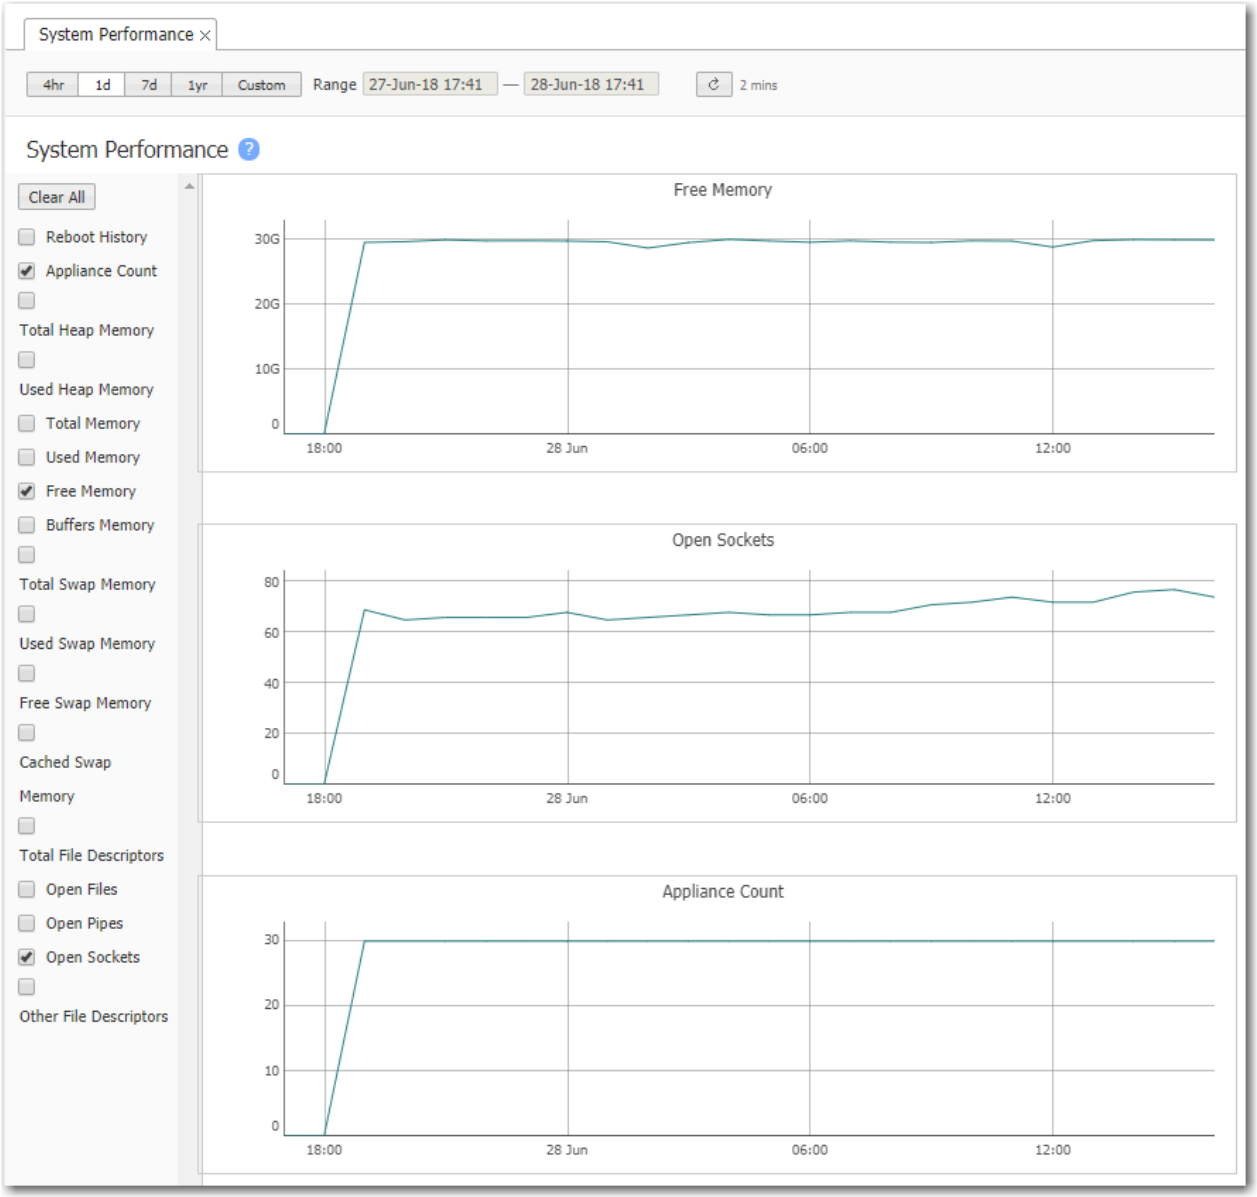


System Performance

Support > Reporting > System Performance

This tab shows Orchestrator metrics.

Orchestrators located in the cloud cannot display useful information about host memory, file descriptors, sockets, or pipes.



System PoE State

Support > *Reporting* > *System PoE State*

The System PoE State tab displays the PoE resources for the appliances selected in the appliance tree. The **Show PoE appliances only** checkbox filters the table to display only appliances that support PoE. This checkbox is selected by default.

Descriptions of the fields on this tab follow:

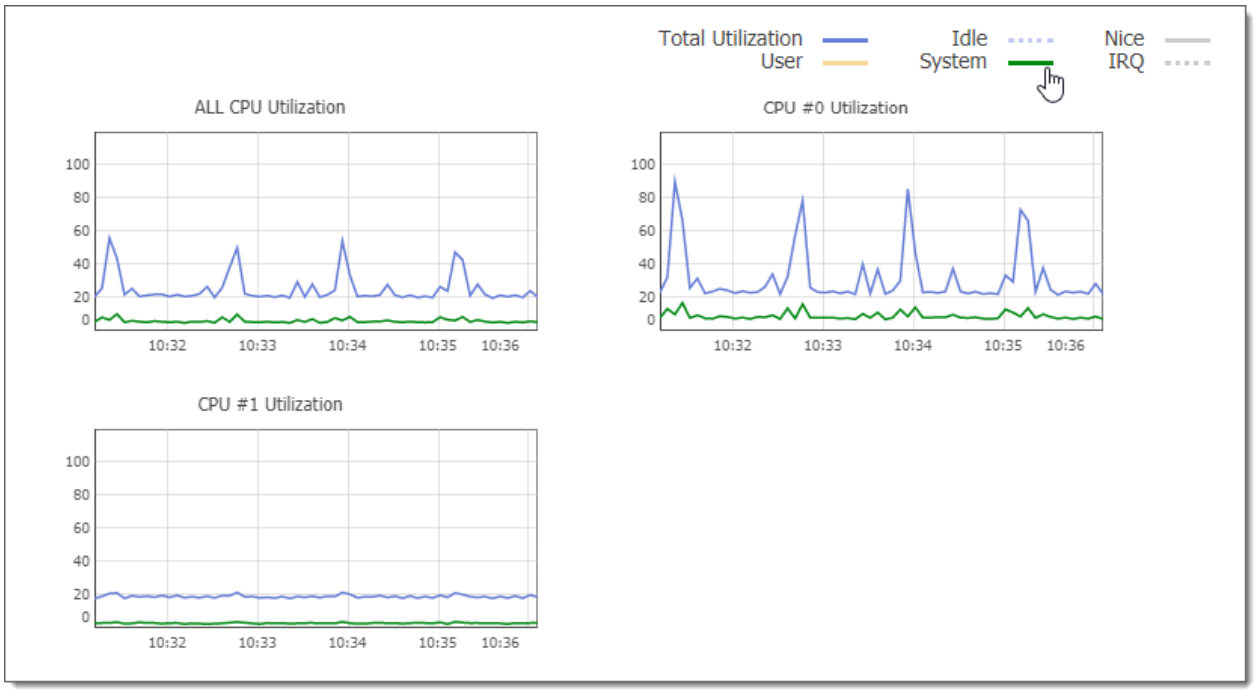
Field	Description
Appliance	Name of the appliance.
Available Power (W)	Maximum power available for provisioning active PoE ports. This is the usable power available to devices connected to the appliance.
Drawn Power (W)	Power (watts) drawn by connected devices.
Remaining Power (W)	PoE power (watts) still available.
Input Voltage (C)	Total voltage available on the appliance.
Temperature (°C)	Temperature of appliance PoE interfaces.

Appliance CPU Usage

Support > Reporting > Appliance CPU Usage

The charts on this page provide real-time views of combined and individual CPU usage statistics for a single selected appliance. Charts show the past five minutes of usage and refresh every five seconds. By default, only total utilization is displayed on the charts. You can toggle the available statistics on or off by clicking the sample indicator line next to each statistic name.

NOTE: On appliances with WAN Optimization enabled, it is common for non-CPU0 cores to run at or close to 100%. CPU0 will show occasional spikes of high usage when statistics are rolled up and archived.



Appliance Hardware Debug

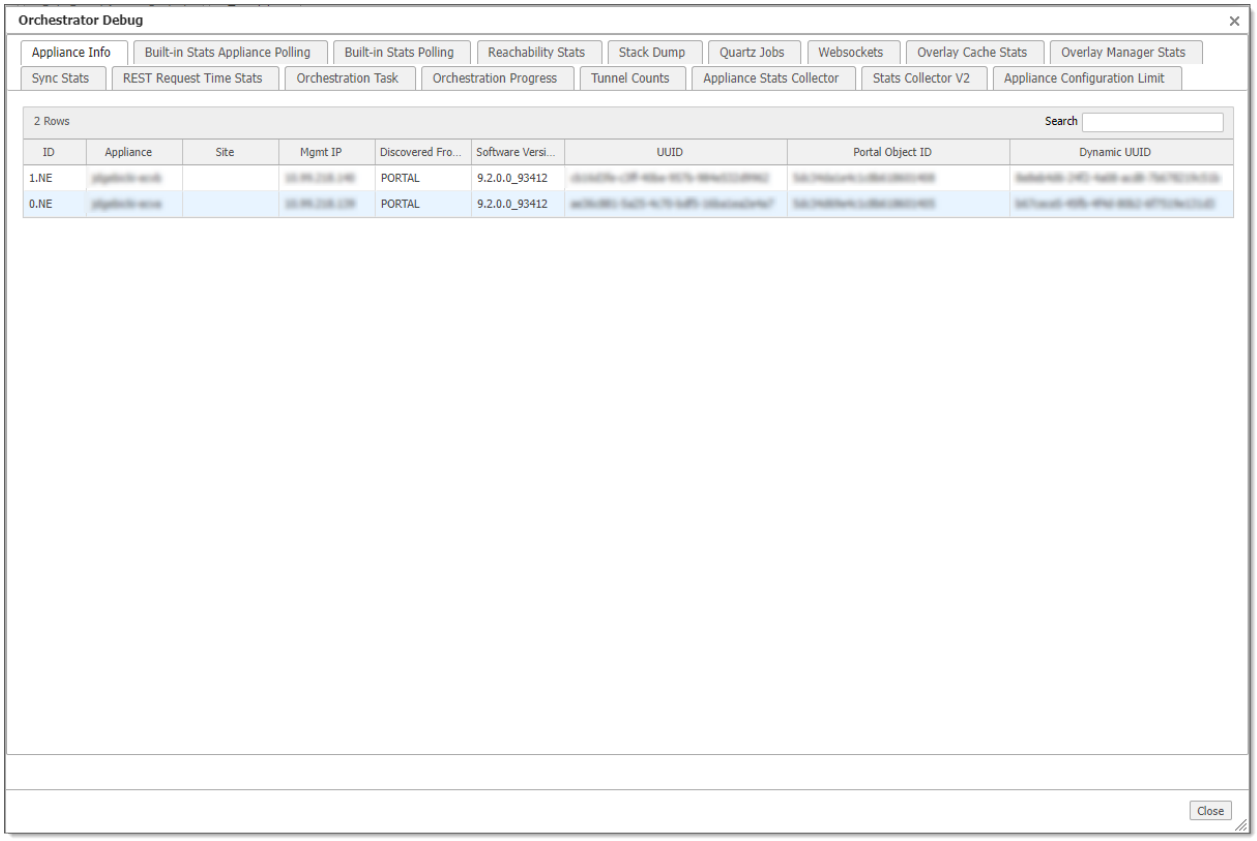
Support > Reporting > Appliance Hardware Debug

This topic is for future content.

Orchestrator Debug

Support > Reporting > Orchestrator Debug

This dialog box provides the various debugging tools available to Support for troubleshooting and debugging issues with Orchestrator.



IPSec UDP Status

Support > Reporting > IPSec UDP Status

Use this tab to review and monitor the IPSec UDP key material status for all appliances in your network.

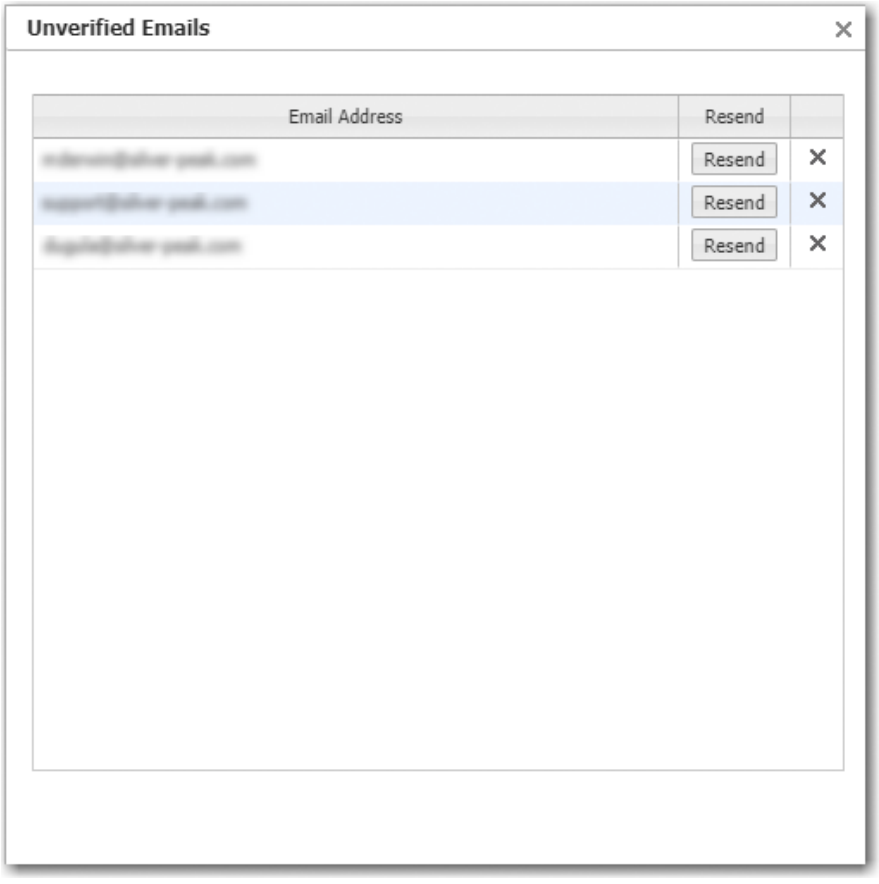
Field	Description
Appliance	Name of the appliance.
Active Key	Indicates whether the appliance is using the active IPSec UDP key.
Active Key Pushed Time	Time when the active key was pushed to the appliance.
Active Key Activation Time	Time when the key was activated on the appliance.
Reachability	Indicates whether the appliance is reachable.
Detail	Additional details about reachability or key material status.

Unverified Emails

Support > Reporting > Unverified Emails

When you add an email address to either the Alarms or the Reports email distribution list, Orchestrator sends the recipient an email that contains a link, asking them to click to provide verification.

If Orchestrator does not receive a verification, either the recipient has not responded or the email address is invalid.



- An unverified email address remains inactive and does not generate an alarm.
- You can retest an address with **Resend**.
- You can only correct an email address in the Alarm or Reports email distribution list.

Orchestrator Debug Log Level

Support > Reporting > Orchestrator Debug Log Level

Use the Orchestrator Debug Log Level dialog box to change what level of server-side Orchestrator logs are retained. You can set a global log level or customized log levels based on your need.

Global Log Level

Select a severity level from the Global log level drop-down menu.

Severity Level	Description
Debug	Used by Support for debugging.
Info	Informational. Also used by Support for debugging. This setting is the default selection.
Error	An error. This is a non-urgent failure.
Warn	A warning condition. This indicates that an error will occur if action is not taken.

- The Severity Level displays in Orchestrator logs.
- If you select **Info** (default), the log records any event with a severity of Info, Warn, and Error.
- These designations are related to event logging levels, not alarm severities, although some naming conventions overlap. Events and alarms have different sources. When alarms clear, they list as the ALERT level in the event log.

Customized Log Levels

Using customized log levels provide more functionality and more granular control over logging, allowing you to add specific loggers with specific log levels. To set customized loggers:

1. Under the Customize log level header, click the **Select Logger** field and either scroll to your selection or begin typing in the field (the logger options will narrow to meet your search terms).
2. Select the log level from the drop-down menu (Debug, Info, Error, Warn).
3. Click **Add**.

4. Click **Save**.

   The customized loggers appear in the table. You can edit or delete customized loggers from this table at any time.

NOTE: Setting customized loggers with their own log levels takes precedence over the settings in Global log level. For example, if the Global log level is set to Info and the OrchestrationTask logger is set to the Debug level, logs are created from OrchestrationTask at the Debug level and higher, while all other server logs remain at the Info level (and higher).

Notable Customized Loggers

There are numerous customized loggers available through the UI, with new loggers planned for future versions of Orchestrator. Here is some context around 10 loggers that might be most useful in your deployment:

Logger	Description
OrchestrationTask	Provides information about orchestration-related tasks.
WatchDog	Provides details about the system (heap, CPU usage, disk, and so on).
SynchronizationManager	Provides details about appliance synchronization activity.
NEReachabilityMana	Provides appliance reachability information.
OverlayApplier	Provides details about overlay changes.
TemplatesManager	Provides details about template changes.
ApplianceConfiguration	Provides information about changes to appliance configuration.
IkelessSeedApplier	Provides details about applying IKE-less seeds.
AuthenticationManager	Provides details about user authentication.
WebUserSessions	Provides details about active user sessions.

Live Tail Logger

Support > Reporting > Live Tail Logger

The live tail logger is a debug utility that pulls logs from Orchestrator in real-time and displays them on the Live Tail Logger tab. To do this, the live tail logger opens a WebSocket connection between the Orchestrator interface and server. From the Live Tail Logger tab, you can view and export the logs.

To view logs, enter text in the Select Loggers field to search for a specific logger, or click and scroll to select a logger. You must select a logger for the utility to work.

- You have the option to apply keywords to filter the logs. Keyword filtering only takes place in the Orchestrator interface and the filters are not sent back to the server. To do this, enter a keyword in the Keyword Filter field and click **Enter**. If you apply multiple keyword filters, the filters are applied using a logical OR operator.

- You can select the log level from the drop-down menu. Only logs from that level and above are included. The system defaults to the Info log level, which includes Info, Warn, Error, and Fatal logs.
- To run the live tail logger utility, click **Run**. To pause the utility click **Pause**, and to resume the utility click **Run**.
- To clear all logs from the Live Tail Logger tab, click **Clear**.
- To export the logs displayed on the Live Tail Logger tab, click **Export Terminal**. The logs are exported in a .log file.
- The maximum number of lines that you can scroll back through on the Live Tail Logger tab is 10,000. To configure this setting, navigate to **Orchestrator > Software & Setup > Setup > Advanced Properties** and change the value for maxLiveTailLoggerTerminalLinesScrollBack.

  NOTE: Changing this default is not recommended without consulting HPE Aruba Networking.

