Juniper Networks | Engineering Simplicity

Contrail Service Orchestration Customer Portal User Guide

Published
2021-01-18

Release
5.1.2

Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, or registered service marks are the property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

Contrail Service Orchestration Customer Portal User Guide
5.1.2
Copyright © 2021 Juniper Networks, Inc. All rights reserved.

The information in this document is current as of the date on the title page.

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. Junos OS has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement (“EULA”) posted at https://support.juniper.net/support/eula/. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

About the Documentation | xxxi

Documentation and Release Notes | **xxxi**

Documentation Conventions | **xxxi**

Documentation Feedback | **xxxiv**

Requesting Technical Support | **xxxiv**

> Self-Help Online Tools and Resources | **xxxv**
>
> Creating a Service Request with JTAC | **xxxv**

1 **Introduction**

Customer Portal Overview | 2

About the Customer Portal User Guide | **2**

Customer Portal Overview | **3**

Accessing Customer Portal | **7**

Switching the Tenant Scope | **10**

Setting Up Your Network with Customer Portal | **10**

About the Customer Portal Dashboard | **11**

> Tasks You Can Perform | **11**
>
> Field Descriptions | **12**

Changing the Customer Portal Password | **15**

Resetting the Password | **16**

Changing the Password on First Login | **17**

Extending the User Login Session | **19**

View and Edit Tenant Settings | 20

Users and Roles | 27

Role-Based Access Control Overview | 27

About the Users Page in Customer Portal | 28

- Tasks You Can Perform | 29
- Field Descriptions | 29

Adding Tenant and OpCo Tenant Users | 30

Editing and Deleting Tenant and OpCo Tenant Users | 32

- Editing Tenant and OpCo Tenant Users | 32
- Deleting Tenant and OpCo Tenant Users | 33

Resetting the Password for Tenant Users | 33

Roles Overview | 34

- Types of Roles | 34
- Role Scopes | 35
- Access Privileges | 35
- Relationship Between User, Roles, and Access Privileges | 35
- Benefits of role-based access control (RBAC) | 36

About the Tenant Roles Page | 37

- Tasks You Can Perform | 37
- Field Descriptions | 37

Adding User-Defined Roles for Tenant Users | 38

Editing, Cloning, and Deleting User-Defined Roles for Tenant Users | 39

- Editing Roles | 40
- Cloning Roles | 40
- Deleting Roles | 41

Access Privileges for Role Scopes (Tenant and Operating Company) | 42

2 Managing Sites, Site groups, and Site Templates

Managing Sites | 52

About the Sites Page | 53

- Tasks You Can Perform | 53
- Field Descriptions | 53

Multihoming Overview | 55

Enterprise Hubs Overview | **55**

- Benefits of Enterprise Hubs | **57**

Adding and Provisioning Switches to Provide LAN Capability to a Site Overview | **57**

- Standalone Switch Overview | **57**
- Switch Behind a CPE or Next Generation Firewall Overview | **58**
- Monitoring Switches Overview | **61**

Adding Enterprise Hubs with SD-WAN Capability or SD-WAN and LAN Capabilities | **61**

Adding Provider Hub Sites for SD-WAN Deployment | **81**

Adding Cloud Spoke Sites for SD-WAN Deployment | **83**

Provisioning a Cloud Spoke Site in AWS VPC | **91**

- Add a Cloud Spoke Site | **91**
- Download the Cloud Formation Template | **92**
- Provision the Device on AWS Server | **93**
- Activate the Device | **94**

Manually Adding On-Premise Spoke Sites | **95**

Adding an On-Premise Spoke Site with Hybrid WAN Capability | **95**

Adding an On-Premise Spoke Site with SD-WAN Capability | **100**

Adding an On-Premise Spoke Site with SD-WAN and LAN Capabilities | **117**

Add an On-Premise Spoke Site with LAN Capability | **136**

Adding an On-Premise Spoke Site with Next Generation Firewall and LAN Capabilities | **151**

Adding and Provisioning a Next Generation Firewall Overview | **158**

- Overview | **158**
- Topology | **159**
- Workflow | **160**

Add a Switch to an Existing SD-WAN Site Or Next-Generation Firewall Site | **160**

Add Switches to an Existing SD-LAN Site | **166**

Enabling Integration with Mist Access Points | **177**

Adding a Standalone Next Generation Firewall Site | **177**

Managing LAN Segments on a Tenant Site | **182**

- Adding LAN Segments | **182**
- Deploying LAN Segments | **186**
- Reassigning a Department to a LAN Segment | **187**
- Deleting LAN Segments | **188**

Managing a Single Site | **188**

Viewing the Sites History | **190**

- Viewing Jobs Initiated to Add and Configure Sites | **190**
- Viewing Jobs Initiated to Delete Sites | **191**

Edit Site Properties | **193**

Deleting a Site | **194**

Managing Site Groups | 196

About the Site Groups Page | **196**

- Tasks You Can Perform | **196**
- Field Descriptions | **196**

Creating Site Groups | **197**

Managing Site Templates | 198

About the Site Templates Page | **198**

- Tasks You Can Perform | **198**

Adding On-Premise Spoke Sites by Using a Site Template | **199**

Cloning, Editing, and Deleting Site Templates | **201**

- Cloning Site Templates | **201**
- Editing Site Templates | **202**
- Deleting Site Templates | **202**

Adding a Site Template | **203**

Adding and Configuring Sites by Importing a JSON File | **215**

Managing Mesh Tags | 217

Mesh Tags Overview | **217**

About the Mesh Tags Page | **218**

- Tasks You Can Perform | **218**
- Field Descriptions | **218**

Creating User-defined Mesh Tags | **219**

Managing Dynamic Mesh | 220

Dynamic Mesh Tunnels Overview | **221**

Adding On-Demand Mesh Tunnels | **222**

Deleting On-Demand Mesh Tunnels | **223**

3

Managing Devices and Resources

Managing Devices | 226

Device Redundancy Support Overview | **227**

- Prerequisites for SRX Series Devices | **227**
- Supported Connection Plans | **227**
- Create and Configure an SD-WAN Site | **228**
- Dual CPE Devices Logical Topology for NFX Network Services Platform | **228**
- Dual CPE Devices Logical Topology for SRX Series Gateway Devices | **228**

Activate a Device | **229**

- Manually Activate a Device That Supports Phone-Home Client | **230**
- Manually Activate a Device That Does Not Support Phone-Home Client | **231**

Manually Activating a Switch | **232**

Manage an EX Series Switch | **233**

- View the Chassis Information of an EX Series Switch | **234**
- View Information about an EX Series Switch | **237**
- View Information about Ports on an EX Series Switch | **239**
- Deploying an Access Profile on a Switch | **243**
- Dissociating an Access Profile | **244**

Managing Ports on an EX Series Switch | **245**

- View Port Details | **245**
- Enable Ports | **246**
- Disable Ports | **247**
- Deploy or Redeploy a Port Profile | **247**
- Configure EX Series Switch Ports Overview | **248**
  - Configure Switch Ports | **249**
  - Life Cycle of a Port Profile | **250**
- Edit Configuration of Ports | **251**

Dissociate a Profile from a Port | 253

Activating Dual CPE Devices (Device Redundancy) | 254

Viewing the History of Tenant Device Activation Logs | 257

Zero Touch Provisioning Overview | 259

Devices Supported | 259

Benefits | 259

Workflow for Onboarding a Device Using ZTP | 260

Managing Device Images | 263

Device Images Overview | 263

About the Device Images Page | 263

Tasks You Can Perform | 263

Field Descriptions | 264

Deleting Device Images | 264

Managing Resources | 266

Multidepartment CPE Device Support | 267

About the Devices Page | 268

Tasks You Can Perform | 268

Field Descriptions | 269

Perform Return Material Authorization (RMA) for a Single-CPE Device or an EX Series Device | 271

Performing Return Material Authorization (RMA) for Dual-CPE Devices | 274

Performing RMA for an NFX Cluster | 274

Performing RMA for an SRX Cluster | 276

Granting RMA for a Device | 278

Grant RMA for a Single-CPE Device or an EX Series Device | 279

Grant RMA for a Dual-CPE Device | 280

Grant RMA for an SRX Device within an SRX Cluster | 282

Managing a Single CPE Device | 284

Rebooting a CPE Device | 287

Configuring APN Settings on CPE Devices | 288

Configuring APN Settings with SIM Change on CPE Devices | 289

Configuring APN Settings without SIM Change on CPE Devices | 291

Identifying Connectivity Issues by Using Ping | 291

Identifying Connectivity Issues by Using Traceroute | 295

Remotely Accessing a Device CLI | 297

Configuring the Firewall Device | 299

About the Physical Interfaces Page | 300

- Tasks You Can Perform | 300
- Field Descriptions | 301

About the Logical Interfaces Page | 301

- Tasks You Can Perform | 301
- Field Descriptions | 302

Adding a Logical Interface | 302

Editing, Deleting, and Deploying Logical Interfaces | 305

- Editing Logical Interfaces | 305
- Deleting Logical Interfaces | 306
- Deploying Logical Interfaces | 306

Adding a Security Zone | 306

Adding a Routing Instance | 309

About the Static Routes Page | 310

- Tasks You Can Perform | 310
- Field Descriptions | 311

Adding a Static Route | 311

Editing, Deleting, and Deploying Static Routes | 314

- Editing Static Routes | 314
- Deleting Static Routes | 315
- Deploying Static Routes | 315

Managing Configuration Templates | 316

About the Configuration Templates Page | 316

- Tasks You Can Perform | 317
- Field Descriptions | 317

Edit, Clone, and Delete Configuration Templates | 319

- Edit a Configuration Template | 319
- Clone a Configuration Template | 320

Delete a Configuration Template | 321
Deploy Configuration Templates to Devices | 322
Deploy from the Configuration Templates Page | 322
Deploy from the Devices Page | 326
Preview and Render Configuration Templates | 327
Import Configuration Templates | 328
Assign Configuration Templates to Device Templates | 330
Add Configuration Templates | 332
View the Configuration Deployed on Devices | 339

Managing Licenses | 342

About the Device Licenses Page | 342
Tasks You Can Perform | 342
Field Descriptions | 342
About the CSO Licenses Page | 343
Tasks You Can Perform | 344
Field Descriptions | 344

Managing Signature Database and Certificates | 346

Signature Database Overview | 346
About the Signature Database Page | 347
Tasks You Can Perform | 347
Field Descriptions | 347
Installing Signatures | 348
Certificates Overview | 349
About the Certificates Page | 350
Tasks You Can Perform | 350
Field Descriptions | 350
Importing a Certificate | 352
Installing and Uninstalling Certificates | 354
Installing a Certificate | 354
Uninstalling a Certificate | 355

About the VPN Authentication Page | 355

Tasks You Can Perform | 356

Field Descriptions | 356

Modify PKI Settings for All Sites | 358

Modify PKI Settings for Selected Sites | 361

Managing Juniper Identity Management Service | 363

Juniper Identity Management Service Overview | 363

Access Token Query | 364

Batch or Periodic Query | 364

IP Address Query | 365

User Mapping Query | 365

About the Identity Management Page | 366

Tasks You Can Perform | 366

Configuring CSO and JIMS Connection | 367

Configuring JIMS for an SRX Device | 369

4 **Managing Policies, Profiles, and Proxies**

Managing Firewall Policies | 373

Firewall Policy Overview | 374

About the Firewall Policy List Page | 376

Tasks You Can Perform | 376

Field Descriptions | 376

About the Firewall Policy Name Page | 377

Tasks You Can Perform | 377

Field Descriptions | 378

Adding a Firewall Policy | 378

Editing and Deleting Firewall Policies | 380

Editing Firewall Policies | 380

Deleting Firewall Policies | 380

Adding Firewall Policy Intents | 381

Editing, Cloning, and Deleting Firewall Policy Intents | 387

Editing Firewall Policy Intents | 388

Cloning Firewall Policy Intents | 388

Deleting Firewall Policy Intents | **389**

Selecting Firewall Source | **389**

Adding an End Point as Firewall Source | **390**

Selecting Firewall Source Using Abbreviations | **391**

Selecting a Firewall Source from the End Points Panel | **391**

Creating and Selecting a Firewall Source from the End Points Panel | **392**

Creating Addresses from Source | **392**

Selecting Firewall Destination | **393**

Adding an End Point as Firewall Destination | **393**

Selecting Firewall Destination Using Abbreviations | **394**

Selecting a Firewall Destination from the End Points Panel | **394**

Creating and Selecting a Firewall Destination from the End Points Panel | **395**

Creating Addresses from Destination | **395**

Firewall Policy Examples | **396**

Example 1: Firewall Policy that Permits Traffic from Departments in Site A to the Departments in Site B | **398**

Example 2: Firewall Policy that Permits Internet Access for all Departments in Site A and Site B | **400**

Example 3: Firewall Policy that Permits Any Public Internet Address to Access the Sales Department in Site B | **403**

Example 4: Firewall Policy that Permits Social Media Access to all Departments in Site A | **404**

Example 5: Firewall Policy that Controls Access to Specific Applications for Various Departments | **406**

Example 6: Firewall Policy that Denies Access to Social Networking Sites | **414**

Example 7: Firewall Policy that Controls Access to an Address over the Internet (HTTP) | **417**

Example 8: Firewall Policy that Permits or Denies the Use of HTTP or FTP as a Service | **423**

Example 9: Firewall Policy that Denies Access to BitTorrent to the Finance Departments across both Site A and Site B | **425**

Example 10: Firewall Policy that Allows Access to Facebook for Users in User Group A | **428**

Example 11: Firewall Policy that Permits User B in Site A Access to YouTube with UTM Enabled | **432**

Example 12: Firewall Policy that blocks access to Internet and allow access to Google Drive. | **435**

Firewall Policy Schedules Overview | **436**

About the Firewall Policy Schedules Page | **437**

- Tasks You Can Perform | **437**
- Field Descriptions | **437**

Creating Schedules | **438**

Editing, Cloning, and Deleting Schedules | **440**

- Editing Schedules | **440**
- Cloning Schedules | **440**
- Deleting Schedules | **441**

Deploying Firewall Policies | **441**

About the Default Profiles for Unified Firewall Policy Page | **442**

- Tasks You Can Perform | **443**
- Field Descriptions | **443**

Editing Default Settings for the Unified Firewall Policy | **444**

Importing Policies Overview | **446**

Importing Firewall Policies | **448**

Managing UTM Profiles | 450

UTM Overview | **451**

- UTM Licensing | **452**
- UTM Components | **452**

Configuring UTM Settings | **453**

About the UTM Profiles Page | **455**

- Tasks You Can Perform | **455**
- Field Descriptions | **455**

Creating UTM Profiles | **457**

Editing, Cloning, and Deleting UTM Profiles | **460**

- Editing UTM Profiles | **460**
- Cloning UTM Profiles | **461**
- Deleting UTM Profiles | **461**

About the Web Filtering Profiles Page | **462**

- Tasks You Can Perform | **462**
- Field Descriptions | **463**

Creating Web Filtering Profiles | **464**

Editing, Cloning, and Deleting Web Filtering Profiles | **468**

- Editing Web Filtering Profiles | **469**
- Cloning Web Filtering Profiles | **469**
- Deleting Web Filtering Profiles | **470**

About the Antivirus Profiles Page | **470**

- Tasks You Can Perform | **471**
- Field Descriptions | **471**

Creating Antivirus Profiles | **472**

Editing, Cloning, and Deleting Antivirus Profiles | **475**

- Editing Antivirus Profiles | **475**
- Cloning Antivirus Profiles | **475**
- Deleting Antivirus Profiles | **476**

About the Antispam Profiles Page | **477**

- Tasks You Can Perform | **477**
- Field Descriptions | **477**

Creating Antispam Profiles | **478**

Editing, Cloning, and Deleting Antispam Profiles | **480**

- Editing Antispam Profiles | **481**
- Cloning Antispam Profiles | **481**
- Deleting Antispam Profiles | **482**

About the Content Filtering Profiles Page | **482**

- Tasks You Can Perform | **482**
- Field Descriptions | **483**

Creating Content Filtering Profiles | **484**

Editing, Cloning, and Deleting Content Filtering Profiles | **488**

- Editing Content Filtering Profiles | **488**
- Cloning Content Filtering Profiles | **488**
- Deleting Content Filtering Profiles | **489**

About the URL Patterns Page | **490**

Tasks You Can Perform | **490**

Field Descriptions | **490**

Creating URL Patterns | **491**

Editing, Cloning, and Deleting URL Patterns | **492**

Editing URL Patterns | **493**

Cloning URL Patterns | **493**

Deleting URL Patterns | **494**

About the URL Categories Page | **494**

Tasks You Can Perform | **494**

Field Descriptions | **495**

Creating URL Categories | **495**

Editing, Cloning, and Deleting URL Categories | **497**

Editing URL Categories | **497**

Cloning URL Categories | **497**

Deleting URL Categories | **498**

Managing SLA Profiles and SD-WAN Policies | 499

SLA Profiles and SD-WAN Policies Overview | **499**

SLA Profiles | **500**

SD-WAN Policies | **501**

About the SD-WAN Policy Page | **503**

Tasks You Can Perform | **503**

Field Descriptions | **503**

Creating SD-WAN Policy Intents | **505**

Editing and Deleting SD-WAN Policy Intents | **512**

Editing SD-WAN Policy Intents | **512**

Deleting SD-WAN Policy Intents | **513**

About the SLA-Based Steering Profiles Page | 513
Tasks You Can Perform | 513
Field Descriptions | 514
Adding SLA-Based Steering Profiles | 517
Editing and Deleting SLA-Based Steering Profiles | 524
Editing an SLA-Based Steering Profile | 524
Deleting SLA-Based Steering Profiles | 525
About the Path-Based Steering Profiles Page | 525
Tasks You Can Perform | 526
Field Descriptions | 526
Adding Path-Based Steering Profiles | 528
Editing and Deleting Path-Based Steering Profiles | 530
Editing a Path-Based Steering Profile | 531
Deleting a Path-Based Steering Profile | 532
Breakout and Breakout Profiles Overview | 532
Cloud Breakout | 534
Breakout Profiles | 534
SD-WAN Policy Intents for Breakout | 534
Benefits of Breakout Profiles | 535
About the Breakout Profiles Page | 535
Tasks You Can Perform | 535
Breakout Profiles Field Descriptions | 536
Cloud Breakout Settings Field Descriptions | 538
Adding Breakout Profiles | 540
Adding Cloud Breakout Settings | 542
Assigning Cloud Breakout Settings to Sites | 546
Detaching Cloud Breakout Settings from Sites | 547
Editing Breakout Profiles and Cloud Breakout Settings | 548
Editing Breakout Profiles | 549
Editing Cloud Breakout Settings | 549
Deleting Breakout Profiles and Cloud Breakout Settings | 550
Deleting Breakout Profiles | 551
Deleting Cloud Breakout Settings | 551
Configuring Breakout on SD-WAN Sites | 552

Managing NAT Policies | 554

NAT Policies Overview | **555**

About the NAT Policies Page | **558**

- Tasks You Can Perform | **558**
- Field Descriptions | **558**

Creating NAT Policies | **559**

Editing and Deleting NAT Policies | **561**

- Editing NAT Policies | **561**
- Deleting NAT Policies | **561**

About the Single NAT Policy Page | **562**

- Tasks You Can Perform | **562**
- Field Descriptions | **563**

Creating NAT Policy Rules | **564**

Editing, Cloning, and Deleting NAT Policy Rules | **571**

- Editing NAT Policy Rules | **571**
- Cloning NAT Policy Rules | **571**
- Deleting NAT Policy Rules | **572**

Deploying NAT Policy Rules | **573**

Selecting NAT Source | **574**

- Adding an Endpoint as NAT Source | **574**
- Selecting Interfaces when GWR Resides Inside an NFX Box | **574**
- Selecting NAT Source Using Abbreviations | **575**
- Selecting a NAT Source from the End Points Panel | **576**
- Creating and Selecting a NAT Source from the End Points Panel | **576**
- Creating Addresses from Source Field | **577**

Selecting NAT Destination | **578**

- Adding an Endpoint as NAT Destination | **578**
- Selecting Interfaces when GWR Resides Inside an NFX Box | **578**
- Selecting NAT Destination Using Abbreviations | **579**
- Selecting a NAT Destination from the End Points Panel | **580**
- Creating and Selecting a NAT Destination from the End Points Panel | **580**
- Creating Addresses from Destination Field | **581**
- Creating Services from Destination Field | **581**

NAT Pools Overview | **582**

About the NAT Pools Page | **582**

Tasks You Can Perform | **583**

Creating NAT Pools | **584**

Editing, Cloning, and Deleting NAT Pools | **586**

Editing NAT Pools | **586**

Cloning NAT Pools | **587**

Deleting NAT Pools | **587**

Deploying NAT Policies | **588**

Importing NAT Policies | **588**

Managing IPS Signatures and Profiles | 591

About the IPS Signatures Page | **591**

Tasks You Can Perform | **592**

Field Descriptions | **592**

Create IPS Signatures | **596**

Create IPS Signature Static Groups | **604**

Create IPS Signature Dynamic Groups | **605**

Edit, Clone, and Delete IPS Signatures | **611**

Edit IPS Signatures | **611**

Clone IPS Signatures | **612**

Delete IPS Signatures | **612**

Edit, Clone, and Delete IPS Signature Static Groups | **613**

Edit IPS Signature Static Groups | **613**

Clone IPS Signature Static Groups | **614**

Delete IPS Signature Static Groups | **615**

Edit, Clone, and Delete IPS Signature Dynamic Groups | **616**

Edit IPS Signature Dynamic Groups | **616**

Clone IPS Signature Dynamic Groups | **617**

Delete IPS Signature Dynamic Groups | **618**

About the IPS Profiles Page | **618**

Tasks You Can Perform | **619**

Field Descriptions | **619**

Create IPS Profiles | **620**

Edit, Clone, and Delete IPS Profiles | **621**

- Edit IPS Profiles | **621**
- Clone IPS Profiles | **622**
- Delete IPS Profiles | **622**

About the <IPS-Profile-Name> / Rules Page | **623**

- Tasks You Can Perform | **623**
- Field Descriptions | **624**

Create IPS or Exempt Rules | **625**

- Create IPS Rules | **625**
- Create Exempt Rules | **632**

Edit, Clone, and Delete IPS or Exempt Rules | **633**

- Edit IPS or Exempt Rules | **633**
- Clone IPS or Exempt Rules | **634**
- Delete IPS or Exempt Rules | **634**

Managing SSL Proxies | 636

SSL Forward Proxy Overview | **636**

- Supported Ciphers in Proxy Mode | **638**
- Server Authentication | **639**
- Root CA | **640**
- Trusted CA List | **640**
- Session Resumption | **641**
- SSL Proxy Logs | **641**

About the SSL Proxy Policy Page | **642**

- Tasks You Can Perform | **643**
- Field Descriptions | **643**

Creating SSL Proxy Policy Intents | **644**

Editing, Cloning, and Deleting SSL Proxy Policy Intents | **648**

- Editing SSL Proxy Policy Intents | **649**
- Cloning SSL Proxy Policy Intents | **649**
- Deleting SSL Proxy Policy Intents | **650**

Understanding How SSL Proxy Policy Intents Are Applied | **651**

- Example 1: Firewall Policy Intent and SSL Proxy Policy Intent Match | **651**
- Example 2: Firewall Policy Intent and SSL Proxy Policy Intent Do Not Match | **652**

Example 3: Applying SSL Proxy Policy Intents on Internal (Site-to-Site) Traffic | **652**

About the SSL Proxy Profiles Page | **653**

Tasks You Can Perform | **653**

Widget Descriptions | **654**

Creating SSL Forward Proxy Profiles | **655**

Editing, Cloning, and Deleting SSL Forward Proxy Profiles | **659**

Editing SSL Forward Proxy Profiles | **660**

Cloning SSL Forward Proxy Profiles | **660**

Deleting SSL Forward Proxy Profiles | **661**

Configuring and Deploying an SSL Forward Proxy Policy | **662**

Deploying Policies | 664

Deploying Policies Overview | **664**

About the Deployments Page | **665**

Tasks You Can Perform | **665**

Field Descriptions | **665**

Using the Deployment Icon to Deploy Policies | **667**

Deploying Policies | **668**

Configuring Policies for SD-LAN | 670

NAC Overview | **671**

Types of Authentication | **671**

About the Port Profiles Page | **673**

Tasks You Can Perform | **673**

Field Descriptions | **673**

Add Port Profiles | **674**

Edit, Clone, and Delete Port Profiles | **679**

Edit a Port Profile | **680**

Clone a Port Profile | **680**

Delete a Port Profile | **681**

About the Authentication Profiles Page | 682

Tasks You Can Perform | 682

Field Descriptions | 682

Add Authentication Profiles | 684

Edit, Clone, and Delete an Authentication Profile | 690

Edit Authentication Profiles | 690

Clone an Authentication Profile | 691

Delete Authentication Profiles | 691

About the Access Profiles Page | 692

Tasks You Can Perform | 692

Field Descriptions | 692

Add Access Profiles | 693

Edit, Clone, and Delete Access Profiles | 696

Edit Access Profiles | 696

Clone an Access Profile | 697

Delete Access Profiles | 698

About the RADIUS Server Profiles Page | 699

Tasks You Can Perform | 699

Field Descriptions | 699

Add RADIUS Server Profiles | 700

Edit, Clone, and Delete RADIUS Server Profiles | 702

Edit RADIUS Server Profiles | 703

Clone a RADIUS Server Profile | 703

Delete RADIUS Server Profiles | 704

Firewall Filters Overview | 705

Firewall Filter | 705

Firewall Filter Components | 705

Firewall Filter Processing | 706

Configure a Firewall Filter for an EX Series Switch | 706

About the EX Firewall Filters Page | 706

Tasks You Can Perform | 706

Field Descriptions | 707

Add Firewall Filters | 707

Delete Firewall Filters | 708

About the < Firewall-Filters-Name> / Terms Page | **709**

- Tasks You Can Perform | **709**
- Field Descriptions | **709**

Add Terms to Firewall Filters | **710**

Edit, Clone, and Delete Terms | **713**

- Edit a Firewall Filter Term | **713**
- Clone a Firewall Filter Term | **714**
- Delete a Firewall Filter Term | **715**

Deploy or Redeploy a Port Profile | **715**

Enable Ports | **716**

Disable Ports | **717**

Edit Configuration of Ports | **717**

5 **Managing Network Services and Shared Objects**

Configuring Network Services in a Distributed Deployment | 721

Network Service Overview | **721**

About the Network Services Page | **722**

- Tasks You Can Perform | **722**
- Field Descriptions | **722**

About the Service Overview Page | **724**

- Tasks You Can Perform | **724**
- Field Descriptions | **724**

About the Service Instances Page | 726

Tasks You Can Perform | 726

Field Descriptions | 726

Configuring VNF Properties | 727

vSRX VNF Configuration Settings | 728

Managing Shared Objects | 734

Addresses and Address Groups Overview | 735

About the Addresses Page | 736

Tasks You Can Perform | 736

Field Descriptions | 736

Creating Addresses or Address Groups | 737

Editing, Cloning, and Deleting Addresses and Address Groups | 740

Editing Addresses and Address Groups | 741

Cloning Addresses and Address Groups | 741

Deleting Addresses and Address Groups | 742

Services and Service Groups Overview | 742

About the Services Page | 743

Tasks You Can Perform | 743

Field Descriptions | 743

Creating Services and Service Groups | 744

Creating Protocols | 746

Editing and Deleting Protocols | 749

Editing Protocols | 750

Deleting Protocols | 750

Editing, Cloning, and Deleting Services and Service Groups | 751

Editing Services and Service Groups | 751

Cloning Services or Service Groups | 752

Deleting Services and Service Groups | 752

Application Signatures Overview | 753

About the Application Signatures Page | 754

Tasks You Can Perform | 754

Field Descriptions | 754

Understanding Custom Application Signatures | 755

Adding Application Signatures | **757**

Editing, Cloning, and Deleting Application Signatures | **762**

- Editing Application Signatures | **762**
- Cloning Application Signatures | **763**
- Deleting Application Signatures | **763**

Adding Application Signature Groups | **764**

Editing, Cloning, and Deleting Application Signature Groups | **765**

- Editing Application Signature Groups | **765**
- Cloning Application Signature Groups | **765**
- Deleting Application Signature Groups | **766**

About the Departments Page | **766**

- Tasks You Can Perform | **767**
- Field Descriptions | **767**

Adding a Department | **767**

Deleting a Department | **768**

About the MAC Addresses Page | **769**

- Tasks You Can Perform | **769**
- Field Descriptions | **769**

Add a MAC Address Endpoint | **770**

Edit or Delete MAC Address Endpoint | **771**

- Edit MAC Address | **772**
- Delete MAC Address | **772**

About the Protocols Page | **773**

- Tasks You Can Perform | **773**
- Field Descriptions | **773**

Add a Protocol Endpoint | **774**

Edit or Delete Protocol Endpoint | **775**

- Edit Protocols | **775**
- Delete Protocols | **776**

About the Ports Page | **776**

- Tasks You Can Perform | **776**
- Field Descriptions | **777**

Add a Port Endpoint | **777**

Edit or Delete Port Endpoint | **779**

- Edit Ports | **779**
- Delete Ports | **780**

6 **Monitoring Jobs and Audit Logs**

Managing Jobs | 782

About the Jobs Page | **782**

- Tasks You Can Perform | **782**
- Field Descriptions | **782**
- Field Descriptions | **783**

Editing and Deleting Scheduled Jobs | **784**

- Editing Scheduled Jobs | **785**
- Deleting Scheduled Jobs | **785**

Viewing Job Details | **786**

Retrying a Failed Job on Devices | **787**

Managing Audit Logs | 788

Audit Logs Overview | **788**

About the Audit Logs Page | **789**

- Tasks You Can Perform | **789**

Viewing the Details of an Audit Log | **790**

Exporting Audit Logs | **793**

Purging Audit Logs (After Archiving or Without Archiving) | **794**

7 **Monitoring Alarms, Events, and Threats**

Monitoring Security Alerts and Alarms | 799

About the Monitor Overview Page | **799**

- Tasks You Can Perform | **799**
- Field Descriptions | **800**

Alerts Overview | **801**

About the Generated Alerts Page | **801**

- Tasks You Can Perform | **802**
- Field Descriptions | **802**

About the Alert Definitions/Notifications Page | **803**

Tasks You Can Perform | **803**

Field Descriptions | **803**

Managing Security Alerts Definitions | **804**

Tasks You Can Perform | **804**

Field Descriptions | **804**

Creating Security Alert Definitions | **805**

Editing, Cloning, and Deleting Security Alert Definitions | **807**

Editing Security Alert Definitions | **807**

Cloning Security Alert Definitions | **807**

Deleting Security Alert Definitions | **808**

About the Alarms Page | **809**

Tasks You Can Perform | **809**

Field Descriptions | **810**

Enable E-mail Notifications for SD-LAN and SD-WAN Alarms | **810**

Monitoring Security and Device Events | 813

About the All Security Events Page | **813**

Tasks You Can Perform | **813**

Summary View | **814**

Detail View | **814**

About the Firewall Events Page | **818**

Tasks You Can Perform | **818**

Summary View | **818**

Detail View | **819**

About the Web Filtering Events Page | **821**

Tasks You Can Perform | **821**

Summary View | **821**

Detail View | **822**

About the IPsec VPNs Events Page | **824**

Tasks You Can Perform | **824**

Summary View | **824**

Detail View | **825**

About the Content Filtering Events Page | **826**

- Tasks You Can Perform | **826**
- Summary View | **826**
- Detail View | **827**

About the Antispam Events Page | **828**

- Tasks You Can Perform | **828**
- Summary View | **829**
- Detail View | **829**

About the Antivirus Events Page | **830**

- Tasks You Can Perform | **830**
- Summary View | **831**
- Detail View | **831**

About the IPS Events Page | **833**

- Tasks You Can Perform | **833**
- Summary View | **833**
- Detail View | **834**

About the Device Events Page | **836**

- Tasks You Can Perform | **836**
- Advanced Search | **836**
- Field Descriptions | **837**

About the Screen Events Page | **840**

- Tasks You Can Perform | **840**
- Summary View | **840**
- Detail View | **841**

Monitoring SD-WAN Events | 845

SD-WAN Events Overview | **845**

About the SD-WAN Events Page | **846**

- Tasks You Can Perform | **846**
- Field Descriptions | **846**

Monitoring Applications | 848

About the SLA Performance of a Single Tenant Page | **848**

- Tasks You Can Perform | **848**
- Field Descriptions | **849**

Viewing the SLA Performance of a Site | **851**

- SLA Not Met by SLA Profiles | **851**
- Applications SLA Performance by Throughput | **852**
- SLA Performance for ALL | **854**

Viewing the SLA Performance of an Application or Application Group | **855**

Application Visibility Overview | **857**

About the Application Visibility Page | **857**

- Tasks You Can Perform | **857**
- Chart View | **857**
- Grid View | **859**

About the User Visibility Page | **861**

- Tasks You Can Perform | **861**
- Chart View | **861**
- Grid View | **863**

Viewing Application or User Visibility Data for Specific Sites | **864**

- Viewing Application Visibility Data for Specific Sites | **864**
- Viewing User Visibility Data for Specific Sites | **865**

Monitoring Threats | 866

About the Threats Map (Live) Page | **866**

- Tasks You Can Perform | **867**
- Field Descriptions | **868**
- Threat Types | **869**

8

Managing Reports

Security Reports | 873

Reports Overview | **873**

About the Security Report Definitions Page | **874**

- Tasks You Can Perform | **875**
- Field Descriptions | **875**

Scheduling, Generating, Previewing, and Sharing Security Reports | **877**

- Editing Report Generation Schedule | **877**
- Generating Reports | **878**
- Previewing Reports in PDF | **879**
- Sharing Reports through E-mail | **879**

About the Security Generated Reports Page | **880**

- Field Descriptions | **880**

Creating Log Report Definition | **881**

Creating Bandwidth Report Definition | **885**

Creating ANR Report Definition | **887**

Editing, Deleting, and Cloning Log Report Definitions | **890**

- Editing the Log Report Definition | **890**
- Deleting Log Report Definitions | **890**
- Cloning Log Report Definitions | **891**

Editing, Deleting, and Cloning Bandwidth Report Definitions | **892**

- Editing Bandwidth Report Definitions | **892**
- Deleting Bandwidth Report Definitions | **892**
- Cloning Bandwidth Report Definitions | **893**

Editing, Deleting, and Cloning ANR Report Definitions | **894**

- Editing ANR Report Definitions | **894**
- Deleting ANR Report Definitions | **895**
- Cloning ANR Report Definitions | **895**

SD-WAN Reports | 897

About the SD-WAN Report Definitions Page | **897**

- Tasks You Can Perform | **897**
- Field Descriptions | **898**

Editing, Deleting, and Cloning SD-WAN Report Definitions | **899**

- Editing the SD-WAN Report Definition | **899**
- Deleting SD-WAN Report Definitions | **899**
- Cloning SD-WAN Report Definitions | **900**

Creating SD-WAN Tenant Performance Report Definitions | **901**

Creating SD-WAN Site Performance Report Definitions | **905**

About the SD-WAN Generated Reports Page | **908**

- Field Descriptions | **908**

About the Documentation

IN THIS SECTION

- Documentation and Release Notes | xxxi
- Documentation Conventions | xxxi
- Documentation Feedback | xxxiv
- Requesting Technical Support | xxxiv

Use this guide to understand the features and tasks that you can configure and perform from the Cloud-based Contrail Service Orchestration (CSO) Customer Portal UI . This guide provides, feature overviews, and procedures that help you understand the features and perform CSO configuration tasks.

Documentation and Release Notes

To obtain the most current version of all Juniper Networks® technical documentation, see the product documentation page on the Juniper Networks website at https://www.juniper.net/documentation/.

If the information in the latest release notes differs from the information in the documentation, follow the product Release Notes.

Juniper Networks Books publishes books by Juniper Networks engineers and subject matter experts. These books go beyond the technical documentation to explore the nuances of network architecture, deployment, and administration. The current list can be viewed at https://www.juniper.net/books.

Documentation Conventions

Table 1 on page xxxii defines notice icons used in this guide.

Table 1: Notice Icons

Icon	Meaning	Description
	Informational note	Indicates important features or instructions.
	Caution	Indicates a situation that might result in loss of data or hardware damage.
	Warning	Alerts you to the risk of personal injury or death.
	Laser warning	Alerts you to the risk of personal injury from a laser.
	Tip	Indicates helpful information.
	Best practice	Alerts you to a recommended use or implementation.

Table 2 on page xxxii defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the **configure** command: user@host> **configure**
`Fixed-width text like this`	Represents output that appears on the terminal screen.	`user@host>` **show chassis alarms** `No alarms currently active`
Italic text like this	• Introduces or emphasizes important new terms. • Identifies guide names. • Identifies RFC and Internet draft titles.	• A policy *term* is a named structure that defines match conditions and actions. • *Junos OS CLI User Guide* • RFC 1997, *BGP Communities Attribute*

Table 2: Text and Syntax Conventions *(continued)*

Convention	Description	Examples
Italic text like this	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# **set system domain-name** *domain-name*
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	• To configure a stub area, include the **stub** statement at the **[edit protocols ospf area area-id]** hierarchy level. • The console port is labeled **CONSOLE**.
< > (angle brackets)	Encloses optional keywords or variables.	**stub <default-metric *metric*>;**
| (pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	**broadcast | multicast** **(*string1* | *string2* | *string3*)**
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	**rsvp { # Required for dynamic MPLS only**
[ ] (square brackets)	Encloses a variable for which you can substitute one or more values.	**community name members [*community-ids*]**

Indention and braces ({ })	Identifies a level in the configuration hierarchy.	[edit]
		routing-options {
		  static {
		  route default {
		  nexthop *address*;
		  retain;
		  }
		  }
		}

; (semicolon)	Identifies a leaf statement at a configuration hierarchy level.	
GUI Conventions		

Table 2: Text and Syntax Conventions *(continued)*

Convention	Description	Examples
Bold text like this	Represents graphical user interface (GUI) items you click or select.	• In the Logical Interfaces box, select **All Interfaces**. • To cancel the configuration, click **Cancel**.
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select **Protocols>Ospf**.

Documentation Feedback

We encourage you to provide feedback so that we can improve our documentation. You can use either of the following methods:

- Online feedback system—Click TechLibrary Feedback, on the lower right of any page on the Juniper Networks TechLibrary site, and do one of the following:

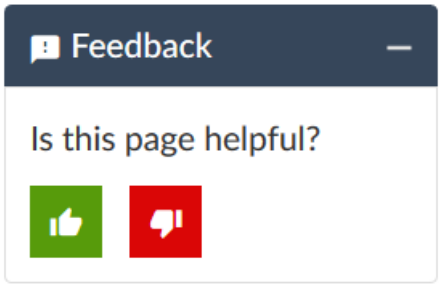

  - Click the thumbs-up icon if the information on the page was helpful to you.
  - Click the thumbs-down icon if the information on the page was not helpful to you or if you have suggestions for improvement, and use the pop-up form to provide feedback.
- E-mail—Send your comments to techpubs-comments@juniper.net. Include the document or topic name, URL or page number, and software version (if applicable).

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active Juniper Care or Partner Support Services support contract, or are

covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the *JTAC User Guide* located at https://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf.
- Product warranties—For product warranty information, visit https://www.juniper.net/support/warranty/.
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: https://www.juniper.net/customers/support/
- Search for known bugs: https://prsearch.juniper.net/
- Find product documentation: https://www.juniper.net/documentation/
- Find solutions and answer questions using our Knowledge Base: https://kb.juniper.net/
- Download the latest versions of software and review release notes: https://www.juniper.net/customers/csc/software/
- Search technical bulletins for relevant hardware and software notifications: https://kb.juniper.net/InfoCenter/
- Join and participate in the Juniper Networks Community Forum: https://www.juniper.net/company/communities/
- Create a service request online: https://myjuniper.juniper.net

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool: https://entitlementsearch.juniper.net/entitlementsearch/

Creating a Service Request with JTAC

You can create a service request with JTAC on the Web or by telephone.

- Visit https://myjuniper.juniper.net.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, see https://support.juniper.net/support/requesting-support/.

1

PART

Introduction

Customer Portal Overview | 2

Users and Roles | 27

CHAPTER 1

Customer Portal Overview

IN THIS CHAPTER

- About the Customer Portal User Guide | 2
- Customer Portal Overview | 3
- Accessing Customer Portal | 7
- Switching the Tenant Scope | 10
- Setting Up Your Network with Customer Portal | 10
- About the Customer Portal Dashboard | 11
- Changing the Customer Portal Password | 15
- Resetting the Password | 16
- Changing the Password on First Login | 17
- Extending the User Login Session | 19
- View and Edit Tenant Settings | 20

About the Customer Portal User Guide

This guide provides an understanding of how to use the Contrail Service Orchestration (CSO) Customer Portal to implement your use cases. This guide is appropriate for tenant administrators and operators who need to know how to use Customer Portal.

Refer to Table 3 on page 2 for additional CSO documentation resources.

Table 3: Additional CSO Documentation Resources

Title	Available At
What is SD-WAN?	https://www.juniper.net/us/en/products-services/what-is/sd-wan/
What is Network Functions Virtualization (NFV)?	https://www.juniper.net/us/en/products-services/what-is/network-functions-virtualization/

Table 3: Additional CSO Documentation Resources *(continued)*

Title	Available At
Learn About NFV	https://www.juniper.net/documentation/en_US/learn-about/LearnAbout_NFV.pdf
Administration Portal User Guide	https://www.juniper.net/documentation/product/en_US/contrail-service-orchestration (User Guides section)
Other Resources	https://www.juniper.net/documentation/product/en_US/contrail-service-orchestration

RELATED DOCUMENTATION

Customer Portal Overview | 3

Customer Portal Overview

The Customer Portal in Contrail Service Orchestration (CSO) provides a Web-based UI that tenants of service providers (SPs) and operating companies (OpCos) can use to manage sites and devices, apply policies (firewall, NAT, SD-WAN, and so on), and perform administrative tasks. In Customer Portal, the objects that you manage and the actions that you perform are done in the context of a single tenant. Therefore, objects belonging to one tenant are isolated from objects belonging to other tenants.

Customer Portal supports role-based access control (RBAC), which means that the roles assigned to users determine their access privileges and the actions that they can perform. The following predefined roles are available in Customer Portal:

- Tenant Admin
- Tenant Operator

Administrator users have read access and write access to the Customer Portal UI and API capabilities, whereas operator users only have read access. Administrators can also create more users with specific roles and access privileges.

When you log in to Customer Portal, the main menu (left sidebar) that is displayed and the actions that you can perform depend on your access privileges. Table 4 on page 4 displays the main menu available in the Customer Portal, a brief description of each menu item, and a link to the relevant topic in the *Contrail Service Orchestration Customer Portal User Guide*.

Table 4: Customer Portal Main Menu

Main Menu	Description
Dashboard	Access a user-configurable dashboard that you can customize with available widgets (also known as dashlets). For more information, see "About the Customer Portal Dashboard" on page 11.
Monitor	Monitor or view the following: • Sites: See "About the Monitor Overview Page" on page 799. • Alerts and alarms: See "About the Generated Alerts Page" on page 801 and "About the Alarms Page" on page 809. • Alert definitions and notifications: See "About the Alert Definitions/Notifications Page" on page 803. • Link switchover events (for SD-WAN sites): See "About the SD-WAN Events Page" on page 846. • Device events: See "About the Device Events Page" on page 836. • Security events: See "About the All Security Events Page" on page 813. • Application service-level agreement (SLA) performance (for SD-WAN sites): See "About the SLA Performance of a Single Tenant Page" on page 848. • Application visibility (for SD-WAN sites): See "About the Application Visibility Page" on page 857. • User visibility: See "About the User Visibility Page" on page 861. • Threat maps: See "About the Threats Map (Live) Page" on page 866. • Jobs (ongoing or completed): See "About the Jobs Page" on page 782.

Table 4: Customer Portal Main Menu *(continued)*

Main Menu	Description
Resources	Manage the following resources: • Sites: See "About the Sites Page" on page 53. • Devices: See "About the Devices Page" on page 268. • Deploy or stage images: See "About the Device Images Page" on page 263. • Site groups: See "About the Site Groups Page" on page 196. • Mesh tags: See "About the Mesh Tags Page" on page 218. • Site templates: See "About the Configuration Templates Page" on page 316. • Device templates: See *About the Device Template Page*. • Configuration templates: See "About the Configuration Templates Page" on page 316.

Table 4: Customer Portal Main Menu *(continued)*

Main Menu	Description

Configuration	• Configure the following:
	  ◦ Intent-based firewall policies: See "About the Firewall Policy List Page" on page 376.
	  ◦ Unified threat management (UTM): See "About the UTM Profiles Page" on page 455.
	  ◦ NAT policies: See "About the NAT Policies Page" on page 558.
	  ◦ Intrusion prevention system (IPS): See "About the IPS Signatures Page" on page 591 and "About the IPS Profiles Page" on page 618.
	  ◦ SSL proxy: See "About the SSL Proxy Policy Page" on page 642.
	  ◦ SD-WAN policies: See "About the SD-WAN Policy Page" on page 503.
	  ◦ SLA-based and path-based steering profiles: See "About the SLA-Based Steering Profiles Page" on page 513 and "Adding Path-Based Steering Profiles" on page 528.
	  ◦ SD-WAN breakout profiles: See "About the Breakout Profiles Page" on page 535.
	  ◦ SD-LAN profiles and firewall filters: See "About the Port Profiles Page" on page 673 and "About the EX Firewall Filters Page" on page 706.
	  ◦ Shared objects (for example, IP addresses): See "About the Addresses Page" on page 736.
	• View and manage deployments: See "About the Deployments Page" on page 665.
	• View network services: See "About the Network Services Page" on page 722.

Reports	• Add and manage security and SD-WAN report definitions, and generate reports: See "About the Security Report Definitions Page" on page 874 and "About the SD-WAN Report Definitions Page" on page 897. • View generated security and SD-WAN reports: See "About the Security Generated Reports Page" on page 880 and "About the SD-WAN Generated Reports Page" on page 908.

Table 4: Customer Portal Main Menu *(continued)*

Main Menu	Description
Administration	Perform various administrative tasks including the following: • Manage users and roles: See "About the Users Page in Customer Portal" on page 28 and "About the Tenant Roles Page" on page 37. • Monitor audit logs: See "About the Audit Logs Page" on page 789. • View device and CSO licenses: See "About the Device Licenses Page" on page 342 and "About the CSO Licenses Page" on page 343. • Modify tenant settings: See "View and Edit Tenant Settings" on page 20. • Install signatures: See "About the Signature Database Page" on page 347. • Manage certificates and VPN authentication: See "About the Certificates Page" on page 350 and "About the VPN Authentication Page" on page 355. • Set up and configure Juniper Identity Management Service (JIMS) See "About the Identity Management Page" on page 366. • Integrate with Mist Access Points (APs): See "Enabling Integration with Mist Access Points" on page 177.

RELATED DOCUMENTATION

Accessing Customer Portal | 7

Changing the Customer Portal Password | 15

Accessing Customer Portal

To access Customer Portal:

1. If you are logging in to Customer Portal for the first time, do the following. If not, skip to 2.

> **NOTE:** When your administrator creates a CSO account for you, an e-mail (with the subject line CSO Account Created) is sent. This e-mail contains a URL that allows you to log in to Customer Portal. The URL is active for only 24 hours and is valid only for the first log in.

a. Click the URL that you have received in the e-mail.

   The Change Password page appears with a message that you must change your password for security purposes.

b. Change your password following the guidelines provided in Table 5 on page 9.

c. (Optional) Click the Terms of Use link to view the Terms of Use document.

d. Click the check box to accept CSO terms of use.

e. Click **Ok**.

The login password is changed and you are logged out of the system. When you log in you must use the changed password.

2. Login to Customer Portal using the link provided in the account activation e-mail.

> **NOTE:** We recommend that you use Google Chrome Version 60 or later to access the Contrail Service Orchestration (CSO) GUI.

3. Enter your username (E-mail ID) and password.

   The Welcome page appears listing the key features of the release.

4. (Optional) If you want to hide the Welcome page on your next login, select the **Hide this on next login** check box.

5. (Optional) If you want to review the tenant setting, select **Review Settings**.

   The Review Settings page appears. For more information see, "View and Edit Tenant Settings" on page 20.

6. Click **Go to Dashboard**. The menu bar on the left-hand side of the every page allows you to access the different tasks easily. The top-level menu items are listed in Table 6 on page 9.

Table 5: Fields on the Change Password Page

Field	Description
New Password	Enter your new password. The password must be between 6 and 21 characters long, and must include at least one lowercase letter, one uppercase letter, one special character, and one number. **NOTE:** The password strength indicator displays the efficiency of the password that you enter. You cannot proceed to the next step if the password strength indicator shows that the password is weak.
Confirm Password	Reenter the password for confirmation. You can select **Show Password** to view the password.

Table 6: Customer Portal Menu

Menu Name	Description
Dashboard	Configurable dashboard that offers you a customized view of network services through its widgets
Monitor	Monitor alerts and alarms, security, device, and software-defined WAN (SD-WAN) events; applications and jobs
Resources	Manage device and software image, sites, site templates, mesh tags, and site groups
Configuration	Configure network services, shared objects, and policies (firewall, NAT, SD-WAN), and view and manage configuration deployments
Reports	Create report definitions and view reports
Administration	Manage users, licenses, audit logs, tenant settings, certificate management, Identity Management, and the signature database

RELATED DOCUMENTATION

Changing the Customer Portal Password | 15

Customer Portal Overview | 3

Switching the Tenant Scope

Administration Portal users can change the tenant scope from all tenants to a specific tenant by using the tenant switcher displayed on the banner.

When you switch scope from all tenants to a specific tenant, the menu and pages displayed are almost the same as those displayed for Customer Portal users, with some additional actions visible to the Administration Portal users. When you switch back to the **All Tenants** scope, the menu and pages for the Administration Portal are displayed.

To switch from one scope to another:

- From the top right corner of the page, select the **All Tenants** scope to access Administration Portal or select a specific tenant (for example, aaa) to access Customer Portal. The menu and pages for Administration Portal or Customer Portal are displayed based on the scope selected from the drop-down list.

RELATED DOCUMENTATION

Role-Based Access Control Overview | 27

Setting Up Your Network with Customer Portal

Your service provider specifies which sites appear in your network and the network services that you can use. When you start working in Customer Portal, you must set up your network using the available sites and network services.

To set up your network with Customer Portal:

1. You can add the following types of sites from the Sites page:
   - Provider hub site: A provider hub site connects to multiple spoke sites using overlay connections. To add a provider hub site, see “Adding Provider Hub Sites for SD-WAN Deployment” on page 81.
   - On-premise spoke sites: An on-premise spoke represents an endpoint that is part of customer premise equipment (CPE) at some physical location such as branch office or point of sale location. Typically, these points are connected using overlay connections to hub sites. You can add on-premise spoke sites manually or by using site templates:
     - To manually add an on-premise site with one WAN capability (SD-WAN, or Hybrid WAN, or Next Gen Firewall), LAN capability, or both WAN and LAN capabilities, see “Manually Adding On-Premise Spoke Sites” on page 95.

    - To add multiple on-premise spoke sites, use site templates. See “Adding On-Premise Spoke Sites by Using a Site Template” on page 199.

  - Cloud spoke site: A cloud spoke site connects to a hub site using overlay connections. To add a cloud spoke site, see “Adding Cloud Spoke Sites for SD-WAN Deployment” on page 83.
  - Enterprise hub site: An enterprise hub site carries site-to-site traffic between on-premise spoke sites and to break out backhaul (central breakout) traffic from on-premise spoke sites. To add an enterprise hub site, see “Adding Enterprise Hubs with SD-WAN Capability or SD-WAN and LAN Capabilities” on page 61.

2. Activate the site.
3. Deploy network services. See “Managing a Single Site” on page 188.
4. View and manage policies.
   - View and manage a firewall policy. See “Adding Firewall Policy Intents” on page 381 and “Deploying Policies” on page 668.
   - View and manage an SD-WAN policy. See “Adding SLA-Based Steering Profiles” on page 517, “Adding Path-Based Steering Profiles” on page 528, “Creating SD-WAN Policy Intents” on page 505, and “Deploying Policies” on page 668.

RELATED DOCUMENTATION

Accessing Customer Portal | 7

About the Customer Portal Dashboard

To access the dashboard, select **Customer Portal > Dashboard**.

The user-configurable dashboard that offers you a customized view of network services through its widgets.

You can drag these widgets from the top of the dashboard to your workspace, where you can add, remove, and rearrange them to meet your needs.

The dashboard automatically adjusts the placement of the widgets to dynamically fit on your browser window without changing their order. You can manually reorder the widgets by using the drag and drop option. In addition, you can press and hold the top portion of the widget to move it to a new location.

Tasks You Can Perform

You can perform the following tasks from this page:

- Customize the dashboard by adding, removing, and rearranging the widgets.
- Update the dashboard or an individual widget by clicking the refresh icon.
- Show or hide widget thumbnails in the carousel by selecting the category of widgets that you want to view from the list at the top left of the carousel; the default is **All Widgets**.
- Add a widget to the dashboard by dragging the widget from the palette or thumbnail container into the dashboard.
- Delete a widget from the dashboard page by clicking delete icon (**X**) in the title bar of the widget and confirming the delete operation.
- Add a dashboard tab by clicking the **+** icon, (optionally) entering a name, and pressing Enter.

  You can then add widgets to the dashboard as needed.
- Rename a dashboard by double-clicking on the title bar of the dashboard, entering a name, and pressing Enter.
- Delete a dashboard by clicking the delete icon (**X** icon) in the title bar of the dashboard and confirming the delete operation.
- Search for a widget by clicking the search icon (magnifying glass) at the top left of the carousel, entering search text, and pressing Enter.

Field Descriptions

You can quickly view important data by using the widgets at the top of your dashboard.

Table 7 on page 12 describes the dashboard widgets.

Table 7: Widgets on the Customer Portal Dashboard

Widget	Description
Tenant Sites: Total Alerts	Displays the total number of alerts grouped by severity level. Click each alert name to view the total number of tenant sites receiving alerts that are critical, major, or minor.
Top 5 Sites with Alerts	Displays the top five sites in the tenant receiving alerts. • **Name**—Name of the tenant site. • **Location**—Location of the tenant site. • **Status**—Type of alerts received: critical, major, or minor.

Table 7: Widgets on the Customer Portal Dashboard *(continued)*

Widget	Description
Top Sites not meeting SLA	Displays a bar chart of the top sites in the tenant that did not meet SLA requirements and the percentage of time that SLA requirements were not met. You can sort the information based on profile and period ranging from the last hour to the last month.
Top Profiles not meeting SLA	Displays a bar chart of the top SLA profiles that did not meet SLA requirements and the percentage of time that SLA requirements were not met. You can sort the information based on location and period ranging from the last hour to the last month.
Top Sites Switching Links	Displays a column chart of the top sites in the tenant that switched WAN links to meet SLA requirements and the number of link-switch events for the sites. You can sort the information based on profile and period ranging from the last hour to the last month.
Top Profiles Switching Links	Displays a column chart of the top SLA profiles that switched WAN links and the number of link-switch events for the SLA profiles. You can sort the information based on location and period ranging from the last hour to the last month.
Top Applications by Throughput	Displays a bar chart of the top sites in the tenant that did not meet SLA requirements and the percentage of time that SLA requirements were not met. You can sort the information based on profile, location, and time period.
Firewall: Top Denials	Displays a column chart of the top requests denied by the firewall based on their source IP addresses, sorted by count. You can sort the information based on time period ranging from 5 minutes to 7 days.
Firewall: Top Events	Displays a bar chart of the top firewall events of the network traffic, sorted by count. You can sort the information based on time period ranging from 5 minutes to 7 days.
IPS: Top Events	Displays the top IPS events of the network traffic, sorted by count. You can sort the information based on time period ranging from 5 minutes to 7 days.

Table 7: Widgets on the Customer Portal Dashboard *(continued)*

Widget	Description
Applications: Top by Sessions	Displays a bar chart of the top applications with a maximum number of sessions, sorted by count. You can sort the information based on time period ranging from 5 minutes to 7 days.
IP: Top Destinations	Displays the top IP destination addresses of the network traffic, sorted by count. You can sort the information based on time period ranging from 5 minutes to 7 days.
IP: Top Sources	Displays the top IP source addresses of the network traffic, sorted by count. You can sort the information based on time period ranging from 5 minutes to 7 days.
IP: Top Spams by Source IPs	Displays the number of spams detected by the source IPs. You can sort the information based on time period ranging from 5 minutes to 7 days.
Virus: Top Blocked	Displays viruses with the maximum number of blocks, sorted by count. You can sort the information based on time period ranging from 5 minutes to 7 days.
Web Filtering: Top Blocked Websites	Displays a bar chart of websites with the maximum number of blocks, sorted by count. You can sort the information based on time period ranging from 5 minutes to 7 days.
IP: Top Source IPs by Volume	Displays the top source IP addresses based on volume of traffic, sorted by count. You can sort the information based on time period ranging from 15 minutes to 7 days.
Application: Top Application by Volume	Displays the applications based on volume of traffic, sorted by count. You can sort the information based on time period ranging from 5 minutes to 7 days and view the information in a bar chart or a bubble chart.
IP: Top Users/IP by Sessions	Displays the top source IP addresses by sessions, sorted by count. You can sort the information based on time period ranging from 15 minutes to 7 days.
Threat Map: Virus	Displays a world map showing total virus event count across countries. You can sort the information based on source, destination, and time period ranging from 5 minutes to 7 days.

Table 7: Widgets on the Customer Portal Dashboard *(continued)*

Widget	Description
Threat Map: IPS	Displays a world map showing total IPS event count across countries. You can sort the information based on source, destination, and time period ranging from 5 minutes to 7 days.

RELATED DOCUMENTATION

Customer Portal Overview | 3

Changing the Customer Portal Password

To change the Customer Portal password:

1. Click the customer username that is located at the right side of the Customer Portal banner.

   The drop-down list appears.

2. Click **Change Password**.

   The Change Password page appears.

3. Specify the current password.

4. In the New Password text box, specify your new password.

   The login password that you set must conform to a particular set of requirements such as minimum length of 6 characters, a maximum length of 21 characters, and that includes at least one lowercase letter, one uppercase letter, an alpha-numeric character, and a numeric character.

   You must change the password periodically (every 90 days) and you cannot reuse the old password. Your user account will be locked after five consecutive unsuccessful login attempts.

5. In the Confirm Password text box, specify your new password again.

   Select the Show Password option to view the password.

6. Click **OK**.

You are logged out of the system. To log in to Customer Portal again, you must use your new password. Other sessions logged in with the same username are unaffected until the next login.

RELATED DOCUMENTATION

Customer Portal Overview | 3
Accessing Customer Portal | 7

Resetting the Password

If you have forgotten your password, you can reset the password from the Contrail Service Orchestration (CSO) login page.

> **NOTE:** If you have entered an incorrect password, your account will be locked after five consecutive unsuccessful login attempts.

To reset the password:

1. On the login page, enter the username , and then press **Enter**.
2. Click the **Forgot Password** link.

   The Forgot Password page appears, with a message that an e-mail notification with a verification code is sent to your e-mail address.

   > **NOTE:** The **Forgot Password** link appears only after you specify the username.

3. In **Verification Code**, specify the verification code that you have received through an e-mail.

   > **NOTE:** The verification code expires after a time duration of 15 minutes.

4. Click **OK**.

   The Reset Password page appears.

5. Change your password following the guidelines provided in Table 8 on page 17.

6. Click **OK**.

Your password is reset.

Table 8: Fields on the Reset Password Page

Field	Description
Username	Enter your username.
New Password	Enter your new password. The login password that you set must be between 6 and 21 characters long, and it must include at least one lowercase letter, one uppercase letter, one special character, and one number. **NOTE:** The password strength indicator displays the efficiency of the password that you enter. You cannot proceed to the next step if the password strength indicator shows that the password is weak.
Confirm Password	Reenter the password for confirmation. You can select **Show Password** to view the password.

RELATED DOCUMENTATION

- Accessing Customer Portal | 7
- Changing the Password on First Login | 17
- Changing the Customer Portal Password | 15

Changing the Password on First Login

To enhance the security related to login credentials, you are prompted to change the password when you login to the portal for the first time.

To change the password when you log in for the first time:

1. Log in to the portal with the default login credentials.

   The Change Password page appears with a message that you must change your password for security purposes.

   > **NOTE:** The Change Password page appears only if you are logging in to the portal for the first time.

2. Change your password following the guidelines provided in Table 5 on page 9.

3. (Optional) Click the Terms of Use link to view the Terms of Use document

4. Click the check box to accept CSO terms of use.

5. Click **Ok**.

   > **NOTE:** It is mandatory to change the password when you log in to the portal for the first time. If you click **Cancel**, you are redirected to the login page.

The login password is changed and you are logged out of the system. When you log in you must use the changed password.

Table 9: Fields on the Change Password Page

Field	Description
New Password	Enter your new password. The password must be between 6 and 21 characters long, and must include at least one lowercase letter, one uppercase letter, one special character, and one number. **NOTE:** The password strength indicator displays the efficiency of the password that you enter. You cannot proceed to the next step if the password strength indicator shows that the password is weak.
Confirm Password	Reenter the password for confirmation. You can select **Show Password** to view the password.

RELATED DOCUMENTATION

Accessing Customer Portal | 7

Changing the Customer Portal Password | 15

Resetting the Password | 16

Extending the User Login Session

In the unified portal, a login session expires in 60 minutes. After 55 minutes, the **Extend Session** page is displayed and, prompting you to enter your password. You must enter your password to extend the session. The **Extend Session** page is displayed when the **Local** authentication method is configured.

If you have logged in to the portal with SSO authentication, the **Extend Session** page is displayed and you can authenticate with the external SSO server. However, the SSO expiration is not under the control of CSO and the following can happen:

- If the external SSO session is expired, you will be authenticated in the **Extend Session** page. After successful authentication, the **Extend Session** page is closed automatically.
- If the external SSO session is not expired, the **Extend Session** page is closed automatically.

To extend the login session:

1. On the **Extend Session** page, enter your password in the **Password** field. If you want to end your session and exit from the portal, click **Cancel** instead and you are redirected to the Login page.

2. Click **OK**.

   The success message **Your Session has been successfully extended** is displayed.

RELATED DOCUMENTATION

Changing the Customer Portal Password | 15

View and Edit Tenant Settings

Users with a tenant administrator role can view and modify the tenant settings that are configured on the Administration Portal, while users with tenant operator role can only view the tenant settings.

> **NOTE:** You cannot add or remove services (configured in Administration Portal) for the tenant.

To modify the settings configured for a tenant:

1. If the Welcome to CSO *Release-Number* page is displayed after you log in, click **Review Settings**. Alternatively, select **Administration > Tenant Settings**.

   The Tenant Settings page appears.

2. (Optional) Click the Expand icon or the Collapse icon on the top-right corner of the page to expand or collapse the different sections displayed.

3. Modify the tenant settings as explained in Table 10 on page 21.

4. Click **Save** to save the changes.

   A tenant edit job is triggered and a confirmation message, indicating that a tenant edit job is created successfully, appears on the Tenant Settings page.

5. (Optional) You can click the job name in the message to view details of the job (including job status, start date and time, and end date and time) on the **Update tenant settings Details** page. Alternatively, you can view the status of the job on the Jobs (**Monitor > Jobs**) page.

   If the job is completed successfully, a confirmation message appears on top of the Tenant Settings page.

Table 10: Fields on the Tenant Settings Page

Field	Description	Tenant Capabilities (Services)
Services	Displays the services supported for the tenant You cannot modify this setting.	SD-WAN Hybrid WAN Next Gen Firewall LAN
Password Policy		SD-WAN Hybrid WAN Next Gen Firewall LAN
Password Expiration Days	Specify the duration (in days) after which the password expires and must be changed. Range: 1 through 365. Default: 180 days. **NOTE:** The modifications are applicable only to new users and users whose password has expired.	SD-WAN Hybrid WAN Next Gen Firewall LAN
SSL Settings	**NOTE:** You can modify this setting only if you have not added any SD-WAN sites for the tenant.	SD-WAN

Table 10: Fields on the Tenant Settings Page *(continued)*

Field	Description	Tenant Capabilities (Services)
Default SSL Proxy Profile	Click the toggle button to enable or disable a default SSL proxy profile for the tenant. If you enable this option, the following items are created: • A default root certificate with the certificate content specified (in the Root Certificate field) • A default SSL proxy profile • A default SSL proxy profile intent that references the default profile **NOTE:** You use this option to create a tenant-wide default profile; enabling or disabling this option does *not* mean that SSL is enabled or disabled. If you enable this option, you must add a root certificate.	SD-WAN
Root Certificate	**NOTE:** This field is displayed only if you enabled the default SSL proxy profile. You can add a root certificate (X.509 ASCII format) by importing the certificate content from a file or by pasting the certificate content: • To import the certificate content directly from a file: 1. Click **Browse**. The **File Upload** dialog box appears. 2. Select a file and click **Open**. The content of the certificate file is displayed in the Root Certificate field. • Copy the certificate content from a file and paste it in the text box. After the tenant is successfully added, a default root certificate, a default SSL proxy profile, and a default SSL proxy profile intent are created. **NOTE:** • The root certificate must contain both the certificate content and the private key. • For full-fledged certificate operations, such as certificates that need a passphrase, or that have RSA private keys, you must use the Certificates page (**Administration > Certificates**) to import the certificates and install on one or more sites.	SD-WAN

Table 10: Fields on the Tenant Settings Page *(continued)*

Field	Description	Tenant Capabilities (Services)
VPN Authentication		SD-WAN
Authentication Type	**NOTE:** • If PKI Certificate was configured as the authentication type, you can modify the PKI properties (CA Server URL, Password, CRL Server, and Auto Renew) even after you add sites for the tenant. • If Preshared Key was configured as the authentication type, then you can modify the authentication type only if you have not added SD-WAN sites for the tenant. Select the VPN authentication method to establish a secure IPsec tunnel: • **Preshared Key**, which means that CSO establishes IPsec tunnels using keys. • **PKI Certificate**, which means that CSO establishes IPsec tunnels using public key infrastructure (PKI) certificates. If you select this option, you can configure the following: ◦ **CA Server URL**—Specify the Certificate Authority (CA) Server URL. For example, http://*CA-Server-IP-Address*/certsrv/mscep/mscep.dll/pkiclient.exe. ◦ **Password**—Specify the password for the CA server. This field is optional. ◦ **CRL Server URL**—Specify the certificate revocation list (CRL) server URL. For example, http://*Revocation-List-Server-IP-Address*/certservices/abc.crl. CSO retrieves the list of revoked certificates from the CRL server. ◦ **Auto Renew CA Certificates**—Click the toggle button to enable or disable automatic renewal of certificates. If you enable this option, certificates are automatically renewed for all sites in the tenant. If you disable this option, certificates must be manually renewed. **NOTE:** If the certificate expires before the renewal, CSO might not be able to reach the device. ◦ **Renew before expiry**—If you enabled automatic renewal, select the period (3 days, 1 week, 2 weeks, or 1 month) before the expiration date when the certificates get automatically renewed. **NOTE:** You can also change the duration in the VPN Authentication page in Customer Portal (**Administration > Certificate Management > VPN Authentication**) page.	SD-WAN

Table 10: Fields on the Tenant Settings Page *(continued)*

Field	Description	Tenant Capabilities (Services)
Overlay Tunnel Encryption	**NOTE:** You can modify this setting only if you have not added any SD-WAN sites for the tenant.	SD-WAN
Encryption Type	For security reasons, all data that passes through the VPN tunnel must be encrypted. Select the encryption type: • 3DES-CBC—Triple Data Encryption Standard with Cipher-Block Chaining (CBC) algorithm. • AES-128-CBC—128-bit Advanced Encryption Standard with CBC algorithm. • AES-128-GCM—128-bit Advanced Encryption Standard with Galois/Counter Mode (GCM) algorithm. • AES-256-CBC— 256-bit Advanced Encryption Standard with CBC algorithm. • AES-256-GCM—256-bit Advanced Encryption Standard with GCM algorithm. The default encryption type is AES-256-GCM.	SD-WAN
Network Segmentation	**NOTE:** You can modify this setting only if you have not added any SD-WAN sites for the tenant.	SD-WAN
Network Segmentation	Click the toggle button to disable network segmentation on the tenant.	SD-WAN
Dynamic Mesh	**NOTE:** You can modify these settings even after you add sites for the tenant.	SD-WAN
Threshold for Creating a Tunnel		SD-WAN
Number of Sessions	Specify the maximum number of sessions closed (for a time duration of 2 minutes) between two spoke sites. The dynamic mesh tunnel is created between two spoke sites if the number of sessions closed (for a time duration of 2 minutes) is greater than or equal to the value that you specified. The default threshold value (the number of sessions for 2 minutes) is 5.	SD-WAN
Threshold for Deleting a Tunnel		SD-WAN

Table 10: Fields on the Tenant Settings Page *(continued)*

Field	Description	Tenant Capabilities (Services)
Number of Sessions	Specify the minimum number of sessions closed (for a time duration of 15 minutes) between two spoke sites. The dynamic mesh tunnel is deleted between two spoke sites if the number of sessions closed (for a time duration of 15 minutes) is lesser than or equal to the value that you specified. The default threshold value (the number of sessions for 15 minutes) is 2.	SD-WAN
Max Dynamic Mesh Tunnels		SD-WAN
Max tunnels per CSO	Displays the maximum number of dynamic mesh tunnels that can be created in CSO. The total number of dynamic mesh tunnels that can be created by all tenants in CSO is limited to 125000. You cannot modify this field.	SD-WAN
Max tunnels per tenant	Specify the maximum number of dynamic mesh tunnels that the tenant can create. Range: 1 through 50,000.	SD-WAN
Dynamic Mesh	Click the toggle button to disable or enable dynamic meshing between sites in the tenant.	SD-WAN
Cloud Breakout Settings	**NOTE:** You can modify these settings even after you add sites for the tenant.	SD-WAN
Customer Domain Name	Enter the domain name of the tenant. The domain name is used in cloud breakout profiles to generate the fully qualified domain name (FQDN). The cloud security providers use the FQDN to identify the IPsec tunnels.	SD-WAN

Table 10: Fields on the Tenant Settings Page *(continued)*

Field	Description	Tenant Capabilities (Services)
Tenant-Specific Attributes	**NOTE:** You can modify these settings even after you add sites for a tenant. If you have set up a third-party provider edge (PE) device by using software other than CSO, then configure settings on that router by specifying custom parameters and its corresponding values. You can modify existing attributes or add attributes. • To add an attribute: 1. Click the add (+) icon. An editable row appears inline in the table. 2. Specify any information about the site that you want to pass to a third-party router; for example, location. 3. Specify a value for the information about the site that you want to pass to a third-party device; for example, Chicago. 4. Click √ (check mark) to save your changes. The prefix that you entered is displayed in the table. • To modify an attribute, select a row, click the edit (pencil) icon, and modify the name and value.	SD-WAN Hybrid WAN Next Gen Firewall LAN

RELATED DOCUMENTATION

About the Sites Page | 53

CHAPTER 2

Users and Roles

IN THIS CHAPTER

- Role-Based Access Control Overview | 27
- About the Users Page in Customer Portal | 28
- Adding Tenant and OpCo Tenant Users | 30
- Editing and Deleting Tenant and OpCo Tenant Users | 32
- Resetting the Password for Tenant Users | 33
- Roles Overview | 34
- About the Tenant Roles Page | 37
- Adding User-Defined Roles for Tenant Users | 38
- Editing, Cloning, and Deleting User-Defined Roles for Tenant Users | 39
- Access Privileges for Role Scopes (Tenant and Operating Company) | 42

Role-Based Access Control Overview

Contrail Service Orchestration supports the authentication and authorization of users. Both service provider and tenant users access the pages within the unified Administration and Customer Portal based on their role and access permissions.

In addition to predefined roles, CSO enables you to add object-based custom roles. You can create custom roles and assign access privileges (read, create, update, delete, and other actions) to each role.

Table 11 on page 27 shows predefined service provider, tenant, and OpCo roles and their access privileges.

Table 11: Roles and Access Privileges

Role	Role Scope	Access Privileges
Tenant Admin	Tenant	Users with the Tenant Admin role have full access to the Customer Portal UI and APIs. They can add one or more users with the Tenant Administrator or Tenant Operator roles.

Table 11: Roles and Access Privileges *(continued)*

Role	Role Scope	Access Privileges
Tenant Operator	Tenant	Users with the Tenant Operator role have read-only access to the Customer Portal UI and APIs.
OpCo Admin	Operating Company	Users with the OpCo Admin role have full access to the OpCo's Administration Portal UI or API capabilities. They can use the UI or APIs to add one or more users with OpCo Admin, OpCo Operator, and custom roles. They can onboard tenants, and add the first tenant user during the OpCo's tenant onboarding process. They can also add tenant administrators or operators by switching the scope to a specific tenant.
OpCo Operator	Operating Company	Users with the OpCo Operator role have read-only access to the OpCo's Customer Portal UI and APIs.

RELATED DOCUMENTATION

About the Users Page in Customer Portal | 28

About the Users Page in Customer Portal

To access the Users page for a tenant or OpCo tenant, click **Administration** > **Users** in the Customer Portal.

Use this page to manage tenants and OpCo tenant users in the Tenant and OpCo scopes respectively. In the Tenant scope, the SP Admin, SP Operator, Tenant Admin, and Tenant Operator can access the Users page for tenants. The SP Admin and the SP Operator can switch scope from all tenants to a specific tenant.

In the OpCo scope, the SP Admin, SP Operator, OpCo Admin and OpCo Operator can access the OpCo Tenant Users page.

For information about the predefined roles and access permissions of OpCo tenant users and tenant users, see *Role-Based Access Control Overview*.

The information listed on the Users page changes depending on the authentication mode configured:

- **Local Authentication** —The **Users** page lists local users that you can add, edit, and delete.
- **Authentication and Authorization with SSO Server**—The **Users** page is not displayed because users are externally managed in the single sign-on (SSO) server.

Tasks You Can Perform

You can perform the following tasks from this page:

- Add a tenant user or an OpCo tenant user. See "Adding Tenant and OpCo Tenant Users" on page 30.
- Edit or delete a tenant user or an OpCo tenantuser. See "Editing and Deleting Tenant and OpCo Tenant Users" on page 32.
- View details of users in the respective scope. See Table 12 on page 29.
- Show or hide columns about users. Click the **Show Hide columns** icon in the top right corner of the page and select columns that you want to view on the page.
- Search for a user. Click the Search icon in the top right corner of the page to search for a user.

  You can enter partial text or full text of the keyword in the text box and press Enter.

  The search results are displayed on the same page.

Field Descriptions

Table 12 on page 29 describes the fields on the Users page for the Tenant and OpCo scopes in the Customer Portal.

Table 12: Fields on the Users Page

Field	Description
Username	Username of the user. Example: *abc@example.com*
First Name	First name of the user.
Last Name	Last name of the user.
Status	Indicates whether the user is enabled (can log in to CSO) or disabled (cannot log in to CSO).
Role	Depending on the scope selected, indicates the roles assigned to the tenant or OpCo tenant user. By default, this column lists only one role assigned to the user. When a user is assigned more than one role, a ***+<integer>*** icon (for example: +2) appears to the right of the role. The integer indicates the number of additional roles assigned to the user. Click on the integer to view the additional roles.

Table 12: Fields on the Users Page *(continued)*

Field	Description
Last Login	Date and time of the last login. The format is MM/DD/YYYY HH:MM. Example: 07/22/2017 20:07 A date and time is not displayed when a user has not logged in to the Customer Portal.

RELATED DOCUMENTATION

Switching the Tenant Scope | 10

Adding Tenant and OpCo Tenant Users

Use the Add Tenant User page and Add OpCo Tenant User page in the Customer Portal to add tenant users and OpCo tenant users respectively to Contrail Service Orchestration (CSO). After you add a user, the user receives an e-mail with the initial login credentials.

> **NOTE:** To add users, you should be assigned a role, such as Tenant Admin, that allows you to add users.

To add a tenant user or an OpCo tenant user:

1. Select **Administration > Users**.

   The Users page appears.

2. Click the add icon (+) or click **Add User button**. The Add User button appears when there are no users configured in the scope.

   In the Tenant scope, the Add Tenant User page appears. In the OpCo scope, the Add OpCo Tenant User page appears.

3. Complete the configuration as described in Table 13 on page 31.

4. Click **OK** to save the changes. If you want to discard the changes, click **Cancel**.

If you click OK, a confirmation message indicating that the user account is created appears and the user account is listed on the Users page.

To enhance the security related to login credentials, an automatically generated password is sent to the e-mail address that you have specified for the user. You are prompted to change the password when you login to the portal with the automatically generated password. For more information about changing the password on first login, see "Changing the Password on First Login" on page 17.

Table 13: Fields on the Add Tenant User and Add OpCo Tenant User Pages

Field	Description
First Name	Enter the first name as a string of alphanumeric characters, some special characters [underscore (_) and period(.)], and spaces. The maximum length allowed is 32 characters.
Last Name	Enter the last name as a string of alphanumeric characters, some special characters [underscore (_) and period(.)] and spaces. The maximum length allowed is 32 characters.
Username (E-mail)	Enter a valid e-mail address in the *user@domain* format.
Status	Click the toggle button to enable or disable the user. By default, the status is enabled. A user can log in to CSO only when the status is enabled.
Role	Select one or more roles (both predefined and custom) that you want to assign to the tenant or OpCo tenant user. The following predefined roles are available—Tenant Operator, Tenant Admin, and ConfigureSite. To know more about the predefined roles for tenant users and OpCo tenant users, see *Role-Based Access Control Overview*. Click the right-arrow icon to move the selected roles from the **Available** column to the **Selected** column. Note that you can use the search icon on the top right of each column to search for role names. Click the role name to preview the access privileges assigned to the tenant user.

Editing and Deleting Tenant and OpCo Tenant Users

IN THIS SECTION

- Editing Tenant and OpCo Tenant Users | 32
- Deleting Tenant and OpCo Tenant Users | 33

You can edit the information of tenant and OpCo tenant users and delete one or more users.

> **NOTE:** To edit and delete users, you should be assigned a role, such as Tenant Admin, that allows you to edit and delete users.

Editing Tenant and OpCo Tenant Users

To modify a tenant user or an OpCo tenant user:

1. Select **Administration > Users**.

   The Users page appears.

2. Select the user that you want to modify, and click the edit icon.

   In the Tenant scope, the Edit Tenant User page appears. In the OpCo scope, the Edit OpCo Tenant User page appears.

3. Modify the parameters following the guidelines provided in Table 13 on page 31.

   > **NOTE:** You cannot modify the **Username (E-mail)** field.

4. Click **OK** to save the changes or click **Cancel** to discard your changes.

   If you click OK, a confirmation message indicating that the user information is modified appears on top of the Users page.

Deleting Tenant and OpCo Tenant Users

To delete one or more tenant users and OpCo tenant users:

1. Select **Administration > Users**.

   The Users page appears.

2. Select the users that you want to delete and click the delete icon.

   An alert message appears, asking you to confirm the delete operation.

3. Click **Yes** to delete the users or click **No** to cancel the deletion.

   If you click **Yes**, a confirmation message indicating that the selected users account are deleted from CSO appears on top of the Users page.

Resetting the Password for Tenant Users

Users with the Tenant Administrator role can reset the password for tenant users. Also, users with the Update capability for Users objects can reset the password for tenant users.

To reset the password:

1. Select **Administration > Users** in Customer Portal.

   The Users page appears, displaying a list of tenant users.

2. Select the username for which you want to reset the password, and then select **More > Reset Password**.

   An alert message appears, asking you to confirm the reset password operation.

3. Click **Yes** to confirm the reset password operation.

   A confirmation message appears, indicating that the password has been successfully reset, and an e-mail with a new system-generated password is sent to the user.

   The user can use the new system-generated password to log in to CSO.

RELATED DOCUMENTATION

About the Users Page in Customer Portal | 28

Roles Overview

IN THIS SECTION

- Types of Roles | 34
- Role Scopes | 35
- Access Privileges | 35
- Relationship Between User, Roles, and Access Privileges | 35
- Benefits of role-based access control (RBAC) | 36

A role is a function assigned to a user that defines the tasks that the user can perform within CSO. Each user can be assigned one or more roles depending on the tasks that the user is expected to perform.

User roles enable you to classify users based on the privileges to perform tasks on CSO objects. Roles assigned to a user determine the tasks and actions that the user can perform.

This topic contains the following sections:

Types of Roles

There are two types of roles: predefined roles and custom roles.

- **Predefined roles**—System-defined roles with a set of predefined access privileges assigned to a user to perform tasks within the CSO application. Predefined roles are created in the system during CSO installation. For more information about predefined roles, see *Role-Based Access Control Overview*.
- **Custom roles**—Object-based user-defined roles with a set of access privileges assigned to a user to perform tasks within the CSO application. Objects include menu and submenu items (for example, Resources, Devices, Images, and POPs) in the CSO application, from which you can create, edit, clone, and delete objects.

  Custom roles can be created by:

  - An OpCo Administrator, or a Tenant Administrator.
  - A tenant user with the Create Role privilege. This user can create custom roles for tenant users.
  - An OpCo user with the Create Role privilege. This user can create custom roles for both OpCo and tenant users.

You can create custom roles and assign access privileges to each role by using the Roles page (**Administration > Roles**).

You can assign one or more roles to a user when you create or edit a user account. Each role can have one or more access privileges.

Role Scopes

A role scope defines the specific scope, which is assigned to the role, such as service provider, OpCo, or tenant. An OpCo Administrator can assign OpCo, and tenant roles to OpCo users and tenant roles to tenant users. A Tenant Administrator can assign tenant roles only to tenant users. A role can have the following scopes:

- **Tenant**—Represents a customer that can view, configure, and manage its sites through Customer Portal.
- **Operating Company (OpCo)**—Similar to a service provider that can manage its own tenants. Tenants managed by one OpCo are isolated from tenants of another OpCo. An OpCo can manage all activities related to its own tenants.

Access Privileges

The following access privileges and actions can be assigned to a user role:

- Read
- Create
- Update
- Delete
- Other actions (Example: For device templates object, the following other actions are supported: Clone and Edit Device Template) .

Relationship Between User, Roles, and Access Privileges

Figure 1 on page 36 shows the relationship between a user, user roles, and access privileges. A user can have one or more roles and each role can have one or more access privileges.

Figure 1: Relationship Between User, Roles, and Access Privileges

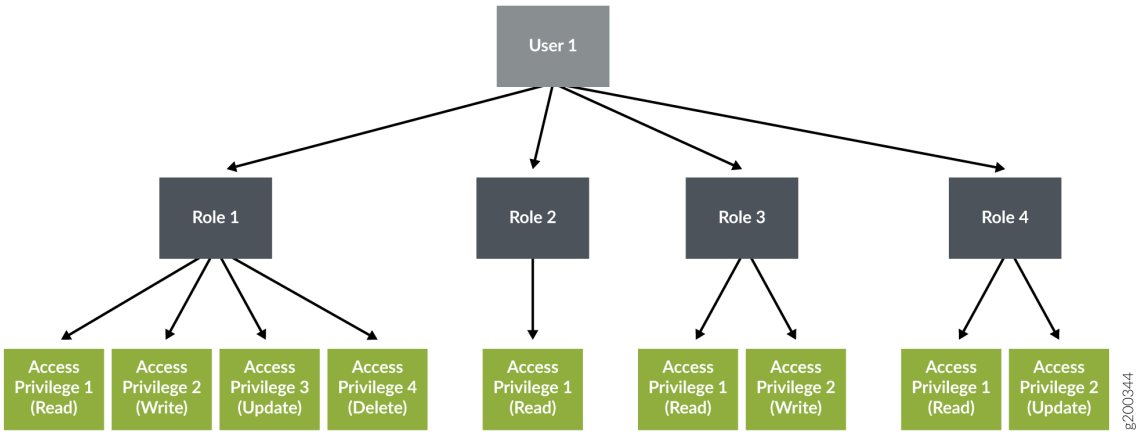

Benefits of role-based access control (RBAC)

- CSO provides pre-defined and user-defined set of roles for day-to-day system operations on the unified Administration and Customer portal.
- Controls which system users can view, read, write, and execute objects based on certain business and operation needs.
- Provides granular level access control on CSO objects within each navigation menu.
- Helps service providers in upselling advanced features to their tenants as a power user.
- CSO supports RBAC and authenticate users using local authentication and the external Single Sign On (SSO) server.

RELATED DOCUMENTATION

About the Tenant Roles Page | 37

Adding User-Defined Roles for Tenant Users | 38

Editing, Cloning, and Deleting User-Defined Roles for Tenant Users | 39

About the Tenant Roles Page

To access this page, select **Administration > Roles** in the Customer Portal.

You can use the Roles page to view a list of predefined (system-defined) and custom (user-defined) roles that can be assigned to tenant users. You can create, edit, or delete custom roles and clone both custom and predefined roles.

Tasks You Can Perform

You can perform the following tasks from this page:

- Create a custom role for the tenant users. See "Adding User-Defined Roles for Tenant Users" on page 38.
- Edit, clone, or delete a custom role. See "Editing, Cloning, and Deleting User-Defined Roles for Tenant Users" on page 39.

Field Descriptions

Table 14 on page 37 describes the fields on the Roles page.

Table 14: Fields on the Roles Page

Field	Description
Role Name	Displays the name of the role.
Role Scope	Displays the scope of the role.
Role Type	Displays whether the role is a predefined role or a custom role.
Created By	Displays the name of the user that created the role.

RELATED DOCUMENTATION

Roles Overview | 34

Role-Based Access Control Overview | 27

Adding User-Defined Roles for Tenant Users | 38

Editing, Cloning, and Deleting User-Defined Roles for Tenant Users | 39

Adding User-Defined Roles for Tenant Users

Use the Add Role page to create custom (user-defined) roles and assign access privileges (read, create, update, delete, and other actions) to the tenant user roles.

A Tenant Administrator or a user with the Create Role privilege can create custom roles for tenant users.

To create a custom role:

1. Select **Administration > Roles** in Customer Portal.

   The Roles page appears.

2. Click the add icon (+) to create a new role.

   The Add Role page appears.

3. Complete the configuration according to the guidelines provided in Table 15 on page 38.

4. Click **OK**.

   A new role is created and listed on the Roles page.

Table 15: Fields on the Add Role Page

Field	Description
Role Name	Enter a unique role name. The name can contain alphanumeric characters, underscore, period, and space.
Description	Enter a description for the role.
Role scope (Visibility)	Select the scope of the role. If you select the scope as Tenant, then the Privileges section of the page displays all the objects of Customer Portal.

Table 15: Fields on the Add Role Page *(continued)*

Field	Description
Privileges	**All Objects**—Displays the objects of the Customer Portal. You must select the check box against each object and then select the type of privileges (read, create, update, delete, and other actions (schedule, deploy, reboot, activate, retry, schedule update, schedule delete, and so on)) that you want to assign the user for the selected object. You can select one or more access privileges to assign to the tenant user role. **NOTE:** You must assign at least one access privilege to a role. If you select the first-level objects, the submenu items that belong to the main object and the corresponding access privileges are selected by default. The following access privileges can be assigned to a user role: • **Read**—Enables the user to read existing objects. • **Create**—Enables the user to create new objects. • **Update**—Enables the user to modify existing objects. • **Delete**—Enables the user to delete existing objects. You can also assign other actions to tenant roles. The other actions include retry, schedule update, schedule delete, activate, reboot, push license, RMA, deploy, schedule, start, disable, deploy, move, run, send, preview, renew, configure, and download.

RELATED DOCUMENTATION

Roles Overview | 34

Role-Based Access Control Overview | 27

About the Tenant Roles Page | 37

Editing, Cloning, and Deleting User-Defined Roles for Tenant Users | 39

Editing, Cloning, and Deleting User-Defined Roles for Tenant Users

IN THIS SECTION

- Editing Roles | 40
- Cloning Roles | 40
- Deleting Roles | 41

You can edit and delete custom (user-defined) roles for tenant users from the Roles page. This topic has the following sections:

> **NOTE:** You cannot modify or delete the predefined roles.

Editing Roles

To modify the parameters configured for a role.

1. Select **Administration > Roles**.

   The Roles page appears, displaying the existing role names.

2. Select the role that you want to edit and click the edit icon (pencil) to modify the parameters.

   The Edit Role page appears. The fields on the Edit Role page are available for editing.

   > **NOTE:** You cannot modify the role name and role scope.

3. Modify the role description and privileges as needed.

4. Click **OK** to save the changes.

   A confirmation message appears, indicating the status of the edit operation.

Cloning Roles

You can clone a role (both custom and predefined) when you want to quickly create a copy of an existing role and modify its access privileges.

> **NOTE:** You cannot modify the role name and role scope.

1. Select **Administration > Roles**.

   The Roles page appears, displaying the existing role names.

2. Select the role that you want to clone and then click the **Clone** button at the top-right corner of the page.

The Clone Role: *Role-Name* page appears.

3. Specify an appropriate name for the new clone role.

4. Click **OK** to save your changes.

   A confirmation message appears, indicating the status of the clone operation.

   The name of the clone role is displayed on the Roles page.

5. Select the new clone role and click the edit icon (pencil) to modify its parameters.

   The Edit Role page appears.

6. Select the objects, and modify the access privileges of the role, as needed.

7. Click **OK** to save your changes.

   A confirmation message appears, indicating the status of the edit operation.

Deleting Roles

To delete a role name:

1. Select **Administration > Roles**.

   The Roles page appears, displaying the existing role names.

2. Select the role name that you want to delete and then click the delete icon (X).

   An alert message appears, asking you to confirm the delete operation.

3. Click **Yes** to delete the selected role name.

   A confirmation message appears, indicating the status of the delete operation.

RELATED DOCUMENTATION

About the Tenant Roles Page | 37

Adding User-Defined Roles for Tenant Users | 38

Access Privileges for Role Scopes (Tenant and Operating Company)

This topic describes the access privileges for the tenant and Operating company (OpCo) role scopes. For more information about roles and role scopes, see “Roles Overview” on page 34.

Table 16 on page 43 shows the access privileges for operating company scope.

Table 17 on page 46 shows the access privileges for tenant scope.

Table 16: Access Privileges for Operating Company Scope

Role Scope	Menu Name	Actions	Other Actions
Operating company (OpCo)	**Monitor**		
	SP Geo Map	Read	-
	Tenants SLA Performance	Read	-
	Alerts	Read and Delete	-
	Alarms	Read	-
	Jobs	Read	Retry Schedule Update Schedule Delete
	Resources		
	POPs	Read, Edit, and Delete	-
	Provider Hub Devices	Read, Edit, and Delete	-
	Tenant Devices	Read	-
	Device Templates	Read, Create, Update, and Delete	Clone Edit Template
	Configuration Templates	Read, Edit, and Delete	-
	Images	Read	-
	Configuration		
	Site management	Edit and Delete	Add provider hub
	Path based steering profiles	Read, Edit, and Delete	-
		Read, Edit, and Delete	-

Table 16: Access Privileges for Operating Company Scope *(continued)*

Role Scope	Menu Name	Actions	Other Actions
	SDWAN Breakout Profiles		
	Application Signatures	Read, Edit, and Delete	-
	Network Services	-	Allocate Server and Detach Server
	Application SLA Profiles	Read, Create, Update, and Delete	-
	Application Traffic Type Profiles	Read	-
	Tenants	Read, Create, Update, and Delete	-
	Administration		
	Users	Read, Create, Update, and Delete	-
	Roles	Read, Create, Update, and Delete	-
	Audit Logs	-	Explore and Purge
	Authentication	Read, Create, Update, and Delete	-
	Licenses	Read, Create, Update, and Delete	Push License
	Dynamic Mesh	Edit	-
	Signature Database	Read	Download Signature Database
	SMTP	Read, Create, Update, and Delete	-
	Email Templates	Read and Update	-
	Terms of Use	Update	-
	Display Differences	Modify	-
	Help Menu (?)		

Table 16: Access Privileges for Operating Company Scope *(continued)*

Role Scope	Menu Name	Actions	Other Actions
	Getting Started	Read	-
	What's New	Read	-
	Quick Help	Read	-
	Help Center	Read	-
	FAQ	Read	-
	Release Notes	Read	-
	About	Read	-

Table 17: Access Privileges for Tenant Scope

Role Scope	Menu Name	Actions	Other Actions
Tenant	**Monitor**		
	Tenant GeoMap	Read	-
	Alerts	Read and Delete	Jump to Event Viewer
	Alarms	Read and Delete	
	Link Switch Events	Read	-
	Traffic Logs	Read	View only threat Show exact match Show raw log Create Alert Create Report Export to CSV
	Security Events	Read	Manage Filter Create Alert Create Report
	Application SLA Performance	Read	-
	Application Visibility	Read	-
	User Visibility	Read and Edit	-
	Threats Map (Live)	Read	-
	Jobs	Read	Retry Schedule Update Schedule Delete
	Resources		

Table 17: Access Privileges for Tenant Scope *(continued)*

Role Scope	Menu Name	Actions	Other Actions
	Site Management	Read, Create, and Delete	Configure Upgrade
	Devices	Read and Delete	Activate Push License Reboot RMA Traceroute Ping Discover APs Configure Stage-2
	Site Groups	Read, Create, Update, and Delete	-
	Mesh Tags	Read, Create, and Delete	-
	Site Templates	Read, Create, Clone, and Delete	-
	Device Templates	Read, Update, and Delete	Import device templates
	Configuration Templates	Read, Update, and Delete	-
	Images	Read	Upgrade History Deploy Stage
	Configuration		
	Firewall Policy	Read, Create, Update, and Delete	Deploy
	Schedule	Read, Create, Update, and Delete	-
	Default Settings	Edit	-

Table 17: Access Privileges for Tenant Scope *(continued)*

Role Scope	Menu Name	Actions	Other Actions
	Unified Threat Management	Read, Create, Update, and Delete	-
	NAT Policies	Read, Create, Update, and Delete	Deploy
	NAT Pools	Read, Create, Update, and Delete	Deploy
	IPS Profiles	Read, Create, Update, and Delete	Clone
	IPS Signatures	Read, Create, Update, and Delete	Clone Clear All Sections
	SSL Proxy Policy	Read, Create, Update, and Delete	Deploy Clone
	SSL Proxy Profiles	Read, Create, Update, and Delete	Clone
	SD-WAN Policy	Read, Create, Update, and Delete	Deploy Clone
	SLA Based Steering Profiles	Read, Create, Update, and Delete	Clone
	Path Based Steering Profiles	Read, Create, Update, and Delete	-
	Cloud Breakout Profiles	Read, Create, Update, and Delete	Assign Sites Detach Sites
	Port Profiles	Read, Create, Update, and Delete	Clone
	Authentication Profiles	Read, Create, Update, and Delete	Clone
	Firewall Filters	Read, Create, Update, and Delete	-
	Access Profiles	Read, Create, Update, and Delete	Clone
	RADIUS Server Profiles	Read, Create, Update, and Delete	Clone

Table 17: Access Privileges for Tenant Scope *(continued)*

Role Scope	Menu Name	Actions	Other Actions
	Address	Read, Create, Update, and Delete	-
	Department	Read, Create, and Delete	-
	Service	Read, Create, Update, and Delete	-
	Application Signature	Read, Create, Update, and Delete	Clone
	MAC	Read, Create, Update, and Delete	-
	Ports	Read, Create, Update, and Delete	-
	Protocols	Read, Create, Update, and Delete	-
	Network Services	Read, Update, and Delete	Start Disable
	Reports		
	Report Definitions - Security	Read, Create, Update, and Delete	Run/Preview Send Clone
	Report Definitions - SD-WAN	Read, Create, Update, and Delete	Run/Preview Send Clone
	Generated Reports -Security	Read and Delete	-
	Generated Reports SD-WAN	Read and Delete	-
	Administration		
	Users	Read, Create, Update, and Delete	-
	Roles	Read, Create, Update, and Delete	-

Table 17: Access Privileges for Tenant Scope *(continued)*

Role Scope	Menu Name	Actions	Other Actions
	Audit Logs		
	Device Licenses	Read, Create, Update, and Delete	Push License
	CSO Licenses	Read	
	Tenant Setting	Read, create, and update	
	Tenant Signature Database	Read	Install
	Certificates	Read, Create, Update, and Delete	-
	Identity Management	Read and Update	-
	Wi-Fi Settings	Read and Update	-
	Help Menu (?)		
	Getting Started	Read	-
	What's New	Read	-
	Quick Help	Read	-
	Help Center	Read	-
	FAQ	Read	-
	Release Notes	Read	-
	About	Read	-

RELATED DOCUMENTATION

About the Tenant Roles Page | 37

Role-Based Access Control Overview | 27

2

PART

Managing Sites, Site groups, and Site Templates

Managing Sites | **52**

Managing Site Groups | **196**

Managing Site Templates | **198**

Managing Mesh Tags | **217**

Managing Dynamic Mesh | **220**

CHAPTER 3

Managing Sites

IN THIS CHAPTER

- About the Sites Page | **53**
- Multihoming Overview | **55**
- Enterprise Hubs Overview | **55**
- Adding and Provisioning Switches to Provide LAN Capability to a Site Overview | **57**
- Adding Enterprise Hubs with SD-WAN Capability or SD-WAN and LAN Capabilities | **61**
- Adding Provider Hub Sites for SD-WAN Deployment | **81**
- Adding Cloud Spoke Sites for SD-WAN Deployment | **83**
- Provisioning a Cloud Spoke Site in AWS VPC | **91**
- Manually Adding On-Premise Spoke Sites | **95**
- Adding an On-Premise Spoke Site with Hybrid WAN Capability | **95**
- Adding an On-Premise Spoke Site with SD-WAN Capability | **100**
- Adding an On-Premise Spoke Site with SD-WAN and LAN Capabilities | **117**
- Add an On-Premise Spoke Site with LAN Capability | **136**
- Adding an On-Premise Spoke Site with Next Generation Firewall and LAN Capabilities | **151**
- Adding and Provisioning a Next Generation Firewall Overview | **158**
- Add a Switch to an Existing SD-WAN Site Or Next-Generation Firewall Site | **160**
- Add Switches to an Existing SD-LAN Site | **166**
- Enabling Integration with Mist Access Points | **177**
- Adding a Standalone Next Generation Firewall Site | **177**
- Managing LAN Segments on a Tenant Site | **182**
- Managing a Single Site | **188**
- Viewing the Sites History | **190**
- Edit Site Properties | **193**
- Deleting a Site | **194**

About the Sites Page

To access this page, click **Resources > Site Management**.

You can use the **Sites** page to view existing sites and to add different types of sites (such as spoke sites and hub sites) manually and by using a site template. You can also use this page to view site configuration and device activation information.

Tasks You Can Perform

You can perform the following tasks from this page:

- Add a provider hub site. See "Adding Provider Hub Sites for SD-WAN Deployment" on page 81.
- Manually add on-premise spoke sites. See "Manually Adding On-Premise Spoke Sites" on page 95.
- Add multiple on-premise spoke sites by using a site template. See "Adding On-Premise Spoke Sites by Using a Site Template" on page 199.
- Add an enterprise hub site. See "Adding Enterprise Hubs with SD-WAN Capability or SD-WAN and LAN Capabilities" on page 61.
- Add multiple sites by uploading a JSON file. See "Adding and Configuring Sites by Importing a JSON File" on page 215.
- Add a switch to an existing site. See "Add a Switch to an Existing SD-WAN Site Or Next-Generation Firewall Site" on page 160.
- Click on the site name to view the site details and to manage the site. See "Managing a Single Site" on page 188.
- View the jobs executed to add and delete sites for a tenant. see "Viewing the Sites History" on page 190.
- Delete a site. See "Deleting a Site" on page 194.
- View device activation logs. See "Viewing the History of Tenant Device Activation Logs" on page 257.

Field Descriptions

Table 18 on page 54 describes the fields on the **Sites** page.

Table 18: Fields on the Sites Page

Field	Description
Alert Icon	Alert associated with the site. The alert can be critical (indicated by a red icon), major (indicated by an orange icon), or minor (indicated by a yellow icon). **NOTE:** The alert icon is displayed only if there is an alert associated with the site. If there is no alert, no icon is displayed.
Site Name	Name of the tenant site. Click the name of the site to go to the *Site-Name* page where you can view the site details and configure parameters related to the site. See "Managing a Single Site" on page 188.
Location	Location of the tenant site.
Template	Displays whether an on-premise spoke site is associated with a site template (displays site template name) or not (–).
Sites Connected To	Number of sites to which the site is connected or N/A (not applicable) if no sites are connected. Mouse over the number to view the list of sites to which the site is connected.
Operational Status	Operational status (Up or Down) of the site.
Type	Indicates whether a site is an on-premise site or a cloud site.
Site Status	The current status of the site: • Created—Indicates that the site was added but not configured. • Configured—Indicates that the site was configured but not activated. • Provisioned—Indicates that the site is provisioned (activated). • Upgrade-Required—Indicates that the site needs to be upgraded. • Maintenance—Indicates that the site upgrade is in progress; any deployments that might occur because of other jobs are skipped when the site status is Maintenance.
Role	Indicates whether the site is a hub site or a spoke site.
Local Breakout	Indicates whether local breakout is enabled for at least one WAN link of the site. If it is enabled for at least one WAN link, the number of links on which local breakout is enabled is also displayed. Mouse over the number to view the list of links on which local breakout is enabled.
Version	Contrail Service Orchestration (CSO) version in which the site was added.

RELATED DOCUMENTATION

Enterprise Hubs Overview | 55

Adding and Provisioning Switches to Provide LAN Capability to a Site Overview | 57

Multihoming Overview

Multihoming is the ability of a spoke site to connect to two different hub devices in a hub and spoke topology, thereby providing redundancy. The hub devices function as primary and the secondary hub devices. If there are multiple spokes in the system, the same hub device may act as primary hub device for one spoke and secondary hub device for another spoke. That is, the selection of the primary and the secondary hub devices is only in the context of a spoke site. The spoke is connected to both the hub devices through an underlay network.

The hub devices can be SRX1500 or SRX4000 series routers. To enable multihoming for a site, you must select the hub and spoke topology when you create the tenant. If you enable multihoming for a site, you must specify a primary and back up site when you configure the site.

Traffic is switched from the primary hub to the secondary hub in the following scenarios:

- The primary hub is down
- The primary hub is up, but all the overlay tunnels between the spoke and the primary hub are down
- The tunnels are up, but the iBGP session between the primary hub and vRR is down. In this case, the failover occurs only after the BGP hold-time expires and the default route is withdrawn.

> **NOTE:** In addition to hub-level redundancy, you can provide VRR-level redundancy by creating two VRRs—primary and secondary—in two different redundancy groups.

Enterprise Hubs Overview

IN THIS SECTION

- Benefits of Enterprise Hubs | 57

An *enterprise hub* is an SD-WAN site that is used to carry site-to-site traffic between on-premise spoke sites and to break out backhaul (central breakout) traffic from on-premise spoke sites. An enterprise hub typically has a data center department behind it; however, this is not enforced in Contrail Service Orchestration (CSO). You add a enterprise hub from the **Sites** page.

You can add one or more enterprise hubs to act as central breakout (backhaul) nodes and then associate enterprise hubs with on-premise spoke sites. The enterprise hub that is associated with a spoke site functions like a data hub and performs the following functions:

- Before the creation of site-to-site tunnels, site-to-site traffic to or from a spoke site is sent through the enterprise hub. This traffic triggers the creation of the site-to-site tunnel based on dynamic mesh thresholds and matching mesh tags that you configure for the spoke site.
- If Internet-bound traffic from the spoke site (and all departments associated with the spoke site) is destined for central breakout (backhaul), the traffic first reaches the assigned enterprise hub and then breaks out from the enterprise hub.
- If a provider hub is associated with the spoke site, the provider hub works as a fallback option in case traffic cannot be sent through the enterprise hub.

> **NOTE:** You must attach an on-premise spoke site (with SDWAN capability) to a provider hub site or an enterprise hub site, or to both hub sites.

If a tenant has more than one enterprise hub configured, CSO statically meshes these sites with overlay tunnels so that the enterprise hubs can exchange routing information for the on-premise spoke sites with which they are associated. This enables the site-to-site communication between the spoke sites that are associated with different enterprise hubs.

The creation of static tunnels between one enterprise hub and another and between a enterprise hub and a spoke site depends on matching mesh tags. These static tunnels are created during the Zero Touch Provisioning (ZTP) workflow. For more information about mesh tags, see "Mesh Tags Overview" on page 217.

Enterprise hubs can have their own departments similar to other sites. If an enterprise hub does not have directly connected LAN segments in the departments used by the associated spoke sites, then CSO automatically pushes the appropriate department virtual routing and forwarding (VRF) instances to the enterprise hub for connectivity.

Benefits of Enterprise Hubs

- Because enterprise hubs can be used to carry backhaul (central breakout) traffic and are used as an anchor for site-to-site traffic, the volume of traffic sent to the provider hub (controlled by the service provider) is reduced.

RELATED DOCUMENTATION

Adding Enterprise Hubs with SD-WAN Capability or SD-WAN and LAN Capabilities | 61

Adding and Provisioning Switches to Provide LAN Capability to a Site Overview

You can use Contrail Service Orchestration (CSO) to provision, deploy, and monitor EX Series switches in branch deployments of enterprise networks. You can deploy an EX Series switch by connecting to a Customer Premise Equipment (CPE) (SRX Series devices only) functioning as a secure SD-WAN router or next-generation firewall. You can also connect the EX Series switch to a third-party Internet gateway device.

CSO Release 5.1.1 supports only EX2300, EX3400, EX4300, EX4600, and EX4650 switches as both, physical switches and Virtual Chassis.

You can provision a switch on a branch network by using CSO in one of the following ways:

- By adding a site with the switch and connecting it to an Internet gateway device.
- By adding a site with an SD-WAN CPE and the switch.
- By adding a site with a next-generation firewall site and the switch.
- By adding a site with an enterprise hub and the switch.
- By adding the switch to an SD-WAN CPE that is already provisioned and managed by CSO.
- By adding the switch to a next-generation firewall site that is already provisioned and managed by CSO.
- By adding the switch to an enterprise hub site that is already provisioned and managed by CSO.
- By adding one or more switches to an SD-LAN site that is already provisioned and managed by CSO.

Standalone Switch Overview

Figure 2 on page 58 shows a site with LAN capability managed by CSO.

Figure 2: Site With LAN Capability (Standalone Switch)

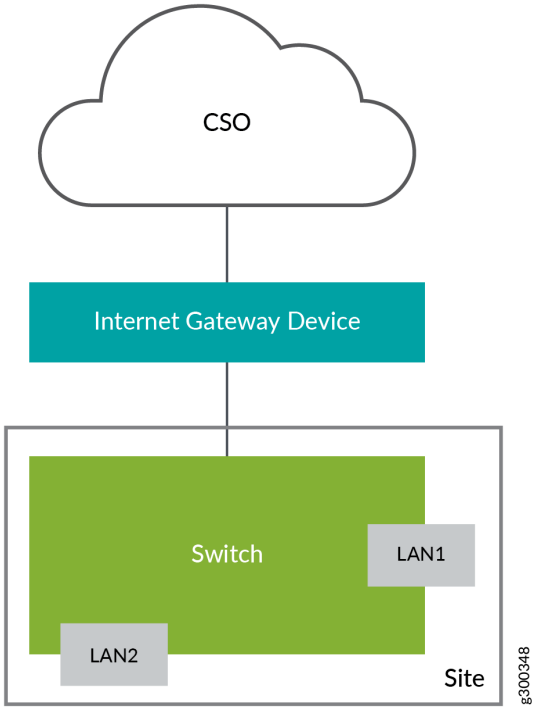

In Figure 2 on page 58, the EX Series switch is connected to CSO through an internet gateway. The gateway can be a device from a manufacturer other than Juniper Networks.

When provisioning a standalone switch (physical or Virtual Chassis), you can use either ZTP (if the EX Series switch supports Phone-Home client) or manually configure the stage-1 configuration on the switch. See “Add an On-Premise Spoke Site with LAN Capability” on page 136 for details.

NOTE:

- Only EX Series switches running 18.4R2.7 or 18.4R3.3 firmware support ZTP.
- EX4600 and EX4650 switches do not support Phone-Home client. You must disable ZTP and manually configure the stage-1 configuration on the switches.

Switch Behind a CPE or Next Generation Firewall Overview

Figure 3 on page 59 shows a site with SD-WAN and LAN capabilities managed by CSO.

Figure 3: Site with LAN and SD-WAN Capabilities

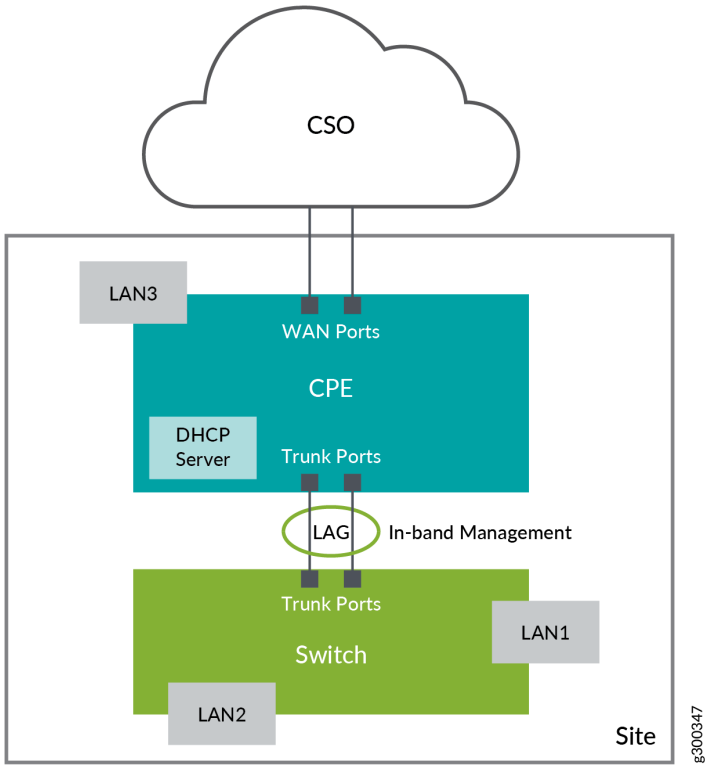

Figure 3 on page 59 shows an example of a switch configured behind a CPE where the switch is connected to two LAN segments (LAN1 and LAN2) and the CPE. The CPE is connected to a LAN segment (LAN3) and to the EX Series switch. The switch can also be connected to a next-generation firewall as shown in Figure 4 on page 60.

Figure 4: Site with LAN and Next-Generation Firewall Capabilities

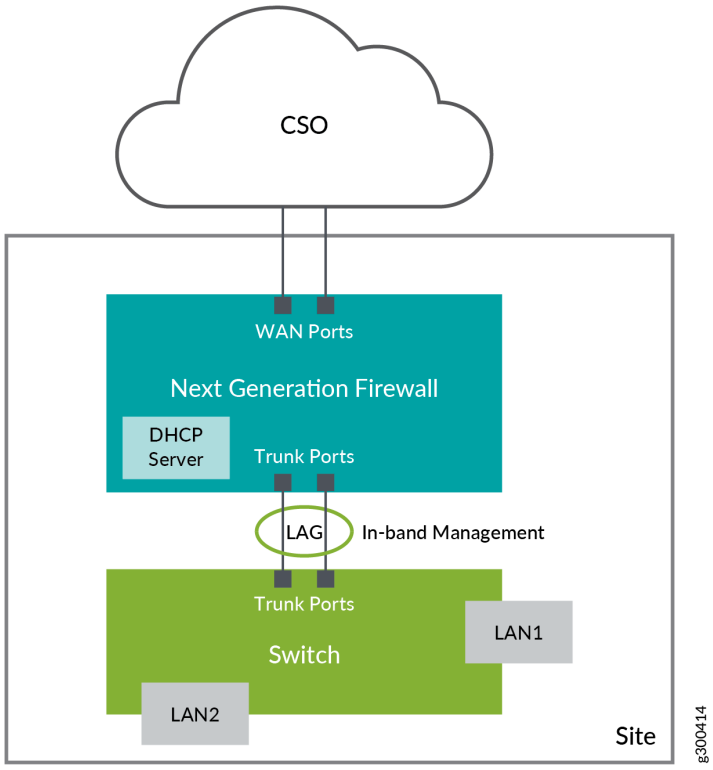

NOTE: You cannot add a LAN segment to the next-generation firewall by using CSO.

The switch and the CPE or firewall can be connected through a trunk port. However, you can use two trunk ports to connect the CPE and the switch and combine them to form a Link Aggregation Group (LAG) for higher throughput and redundancy. Traffic from LAN segments connected to the switch are routed to the CPE or firewall through the trunk ports for further routing into WAN.

You can manage the switch by in-band management, where in, the trunk ports carry the management traffic in addition to data.

NOTE: The ae0 port of the SRX Series device is configured as the trunk port for communication with the switch.

The DHCP server, configured on the CPE or firewall, runs on the trunk ports to:

- Allocate unique IP addresses to the access devices connected to the switch.
- Provide management connectivity to the switch.

During ZTP of a site with both WAN and LAN capabilities, the switch is provisioned after the CPE or firewall is provisioned.

When you add a switch to an already provisioned site, CSO redeploys the stage-2 configuration on the CPE or firewall to configure DHCP and LAG. The DHCP configuration enables management connectivity to the switch and allows CSO to discover and provision the switch.

Monitoring Switches Overview

You can monitor the following for an EX Series switch on the *Device-Name* page (**Resources > Devices**):

- Resource utilization (memory and CPU) on the switch for the past one hour, past eight hours, past one day, past one week, and past one month.
- Status of ports.
- Alerts and alarms generated on the switch for the past one hour, past eight hours, past one day, past one week, and past one month.
- Top Ports consuming the maximum bandwidth.
- Top Ports with the maximum number of errors.
- Top Ports with the maximum packet loss.

RELATED DOCUMENTATION

Adding an On-Premise Spoke Site with SD-WAN and LAN Capabilities | **117**

Adding an On-Premise Spoke Site with Next Generation Firewall and LAN Capabilities | **151**

Add an On-Premise Spoke Site with LAN Capability | **136**

Add a Switch to an Existing SD-WAN Site Or Next-Generation Firewall Site | **160**

Adding Enterprise Hubs with SD-WAN Capability or SD-WAN and LAN Capabilities | **61**

Adding Enterprise Hubs with SD-WAN Capability or SD-WAN and LAN Capabilities

An enterprise hub site is an SD-WAN site that is used to carry site-to-site traffic between on-premise spoke sites and to break out backhaul (central breakout) traffic from on-premise spoke sites. An enterprise hub *typically* has a data center department behind it; however, this is not enforced in Contrail Service Orchestration (CSO). The following device templates are supported for enterprise hubs:

- SRX as SD-WAN CPE (vSRX only)
- Dual SRX as SD-WAN CPEs (vSRX only)
- SRX-1500 as SD-WAN CPE
- Dual SRX1500 as SD-WAN CPEs
- SRX4x00 as SD-WAN CPE
- Dual SRX4x00 as SD-WAN CPEs

To add an enterprise hub site:

> **NOTE:** You can add enterprise hub sites only for tenants with real-time optimized SD-WAN mode.

1. Click **Resources > Site Management**.

   The Sites page appears.

2. Click **Add** and select **Add Enterprise Hub**.

   The **Add Enterprise Hub for *Tenant-Name*** page appears.

3. Do one of the following:

   - To add an enterprise hub with only SD-WAN capability, complete configuration settings according to guidelines provided in Table 19 on page 63.
   - To add an enterprise hub with both SD-WAN and LAN capabilities, complete configuration settings according to guidelines provided in Table 19 on page 63 for the SD-WAN capability and Table 21 on page 76 for LAN capability.

     > **NOTE:** Fields marked with an asterisk (*) are mandatory.

4. (Optional) You can review the configuration in the **Summary** tab and modify the settings, if required.

5. Click **OK**.

   The site activation job is initiated and the Site Activation: *Site-Name* page appears displaying the progress of the steps executed for activating the enterprise hub and the switch (when LAN capability is selected). The enterprise hub is activated first and then the process to activate the switch is initiated.

If you selected LAN capability for the enterprise hub site, go to step 6.

6. • If the Zero Touch Provisioning (ZTP) toggle button is enabled (default), CSO pushes the stage-1 and stage-2 configurations and provisions the switch.

   This process occurs immediately after the activation process, for which you entered the activation code or selected auto-activation.

   > **NOTE:** Stage-1 configuration is the initial configuration that allows basic connectivity to a device, which is pushed to the device.
   >
   > The configuration that is pushed to the device after it has connected to CSO is called stage-2 configuration.

   • If you disabled the Zero Touch Provisioning (ZTP) toggle button, you must manually configure the stage-1 configuration (as provided by CSO) on the switch.

   To manually configure the stage-1 configuration:

   a. On the **Site Activation: *Site-Name*** page, the **Click to copy stage-1 configuration** link appears after the Prestage Device step completes successfully.

   b. Click the **Click to copy stage-1 configuration** link.

      The stage-1 configuration page appears displaying the stage-1 configuration to be copied to the EX Series device.

   c. Copy the stage-1 configuration and log in to the console of the EX Series switch.

   d. Enter the configuration mode, paste, and commit the configuration.

      After the stage-1 configuration is committed, the switch has the outbound SSH configuration to connect with CSO.

      CSO then provisions the switch.

Table 19: Add Enterprise Hub for <Tenant-Name> Settings (WAN Capability)

Field	Description
General	
Site Information	

Table 19: Add Enterprise Hub for <Tenant-Name> Settings (WAN Capability) *(continued)*

Field	Description
Site Name	Enter a unique name for the site. You can use alphanumeric characters and hyphen (-); the maximum length is 10 characters.
Site Group	Select a site group to which you want to assign the site.
Site Capabilities	
WAN Capabilities	SD-WAN capability is selected by default. You cannot clear the selection.
LAN Capabilities	Select **LAN** if you want to include LAN capability in the enterprise hub site.
Configuration	
Primary Provider Hub	Select the provider hub site (or primary provider hub site in case of multihoming) to which you want to connect the enterprise hub site. If you do not specify a provider hub site, then the enterprise hub site can connect only to the on-premise spoke sites that are associated with the enterprise hub site. If you specify a provider hub site, then the enterprise hub site can also connect to the on-premise spoke sites to which that provider hub site is associated.
Secondary Provider Hub	Select the secondary provider hub site (in case of multihoming) to which you want to connect the enterprise hub site. When the primary provider hub is down, the enterprise hub connects to the secondary provider hub and the on-premise spoke sites to which that provider hub site is associated.
On-Demand Mesh Threshold	
Threshold for Tunnel Creation	Specify the threshold for the number of sessions (flows) closed (in a two-minute duration) between the enterprise hub and a destination site. When the number of sessions closed exceeds the specified threshold, a tunnel is created between the enterprise hub and the destination site. The default value is 5. For example, if you specify the Create Threshold as 5, dynamic mesh tunnels are created if the number of sessions closed between the enterprise hub and destination site exceeds 5 in 2 minutes.

Table 19: Add Enterprise Hub for <Tenant-Name> Settings (WAN Capability) *(continued)*

Field	Description
Threshold for Tunnel Deletion	Specify the threshold for the number of sessions closed (in a 15-minute duration) between the enterprise hub and a destination site. When the number of sessions closed is lower than the specified threshold, the tunnel between the enterprise hub and destination site is deleted. The default value is 2. For example, if you specify the number of sessions closed as 2, dynamic mesh tunnels between the enterprise hub and destination site are deleted if the number of sessions closed is lesser than or equal to 2.
Address and Contact Information	
Street Address	Enter the street address of the site.
City	Enter the name of the city where the site is located.
State/Province	Select the state or province where the site is located.
ZIP/Postal Code	Enter the postal code for the site.
Country	Select the country where the site is located. You can click the **Validate** button to verify the address that you specified: • The **site address verification successful** message is displayed if the address can be verified. You can click the **View location on a map** link to see the address location. • If the address cannot be verified, the **Site address could not be validated** message is displayed .
Contact Name	Enter the name of the contact person for the site.
Email	Enter the e-mail address of the contact person for the site.
Phone	Enter the phone number of the contact person for the site. Click **Next** to continue.
Advanced Configuration	

Table 19: Add Enterprise Hub for <Tenant-Name> Settings (WAN Capability) *(continued)*

Field	Description
Name Server IP List	Specify one or more IPv4 addresses of the DNS server. To enter more than one DNS server address, type the address, press Enter, and then type the next address, and so on. DNS servers are used to resolve hostnames into IP addresses.
NTP Server	Specify the fully qualified domain names (FQDNs) or IP addresses of one or more NTP servers. Example: ntp.example.net The site must have DNS reachability to resolve the FQDN during site configuration.
Select Timezone	Select the time zone of the site.
WAN	
Device Template	
Device Template	Select a device template, which contains information for configuring a device.
Device Information **NOTE:** Some fields in this section are displayed only if you select a dual CPE device template.	
Serial Number	For a single CPE device, enter the serial number of the CPE device. Serial numbers are case-sensitive.
Device Redundancy	For dual CPE device templates, displays Enabled indicating that redundancy is enabled. You cannot modify this field.
Primary Serial Number	For a dual CPE device, enter the serial number of the primary CPE device. The serial number is case sensitive.
Secondary Serial Number	For a dual CPE device, enter the serial number of the secondary CPE device. The serial number is case sensitive.

Table 19: Add Enterprise Hub for <Tenant-Name> Settings (WAN Capability) *(continued)*

Field	Description
Auto Activate	Click the toggle button to enable (default) or disable automatic activation of the CPE device. When you enable this field, zero-touch provisioning (ZTP) of the CPE device is automatically triggered after the site is added to CSO. The device template that you select determines whether this option is enabled or disabled by default.
Activation Code	For a single CPE device, if the automatic activation of the device is disabled, enter the activation code to manually activate the device. For information about manually activating the CPE device, see "Activate a Device" on page 229.
Primary Activation Code	For a dual CPE device, if the automatic activation of the device is disabled, enter the activation code to manually activate the primary CPE device.
Secondary Activation Code	For a dual CPE device, if the automatic activation of the device is disabled, enter the activation code to manually activate the secondary CPE device.
Boot image	Select the boot image from the drop-down list if you want to upgrade the image for the CPE device. The boot image is the latest build image uploaded to the image management system. The boot image is used to upgrade the device when the CSO starts the ZTP process. If the boot image is not provided, then the device skips the procedure to upgrade the device image. The boot image is populated based on the device template that you have selected while creating a site. See *Uploading a Device Image*.
WAN Links	
WAN_0 (*WAN-Interface-Name*)	This field is enabled by default. Enter parameters related to the WAN_0 (*WAN-Interface-Name*) link. Fields marked with an asterisk (*) must be configured to proceed.
Link Type	Select whether the link would be an MPLS link or Internet link.

Table 19: Add Enterprise Hub for <Tenant-Name> Settings (WAN Capability) *(continued)*

Field	Description
Egress Bandwidth	Enter the maximum bandwidth (in Mbps) that the CPE allows towards the WAN link. Range: 1 through 10,000.
Address Assignment	Displays the method of assigning an IP address to the WAN link (STATIC). You must provide the IP address prefix and the gateway address for the WAN link.
Static IP Prefix	Enter the IP address prefix of the WAN link.
Gateway IP Address	Enter the IP address of the gateway of the WAN service provider.
Public IP Address	Enter the public IPv4 address for the link. **NOTE:** This IP address should be provided only if the static IP prefix is a private IP address and 1:1 NAT is configured.
WAN Link (Primary or Secondary)	For dual CPE device templates, displays whether the WAN link is a primary link or a secondary link. You cannot modify this field.
Advanced Settings	
Provider	Enter the name of the service provider providing the WAN service.
Cost/Month	Enter the cost for using the WAN link per month and select the currency in which the cost is indicated from the adjacent drop-down list. Range: 1 through 10,000. In bandwidth-optimized SD-WAN, CSO uses this information to identify the least-expensive link to route traffic when multiple WAN links meet SLA profile parameters.

Table 19: Add Enterprise Hub for <Tenant-Name> Settings (WAN Capability) *(continued)*

Field	Description
Enable Local Breakout	Click the toggle button to enable local breakout on the WAN link. By default, local breakout is disabled. **NOTE:** • If you enable this option, the WAN link can be used for local breakout. The decision of whether traffic breaks out locally from the site depends on the breakout profile that is referenced in the SD-WAN policy intent. • If you do not enable local breakout on at least one WAN link for a single CPE connection plan and at least two WAN links for a dual CPE connection plan, then local breakout is disabled for the site.
Breakout Options	When the **Enable Local Breakout** field is enabled, select whether you want to use the WAN link for both breakout and WAN traffic (default) or only for breakout traffic.

Autocreate Source NAT Rule	Click the toggle button to enable or disable the automatic creation of source NAT rules. By default, this field is enabled when local breakout is enabled on the WAN link.
	
	Table 20 on page 75 explains how source NAT rules are automatically created on the WAN link. The automatically-created source NAT rules are implicitly defined and applied to the site and is not visible on the NAT Policies page.
	
	NOTE: You can manually override automatically created NAT rules, by creating a NAT rule within a particular rule-set. For example, to use a source NAT pool instead of an interface for translation, create a NAT rule within this particular rule-set, that includes the relevant department zone and WAN interface as the source and destination. For example:
	
	`Dept-Zone1 --> W1 : Translation=Pool-2`
	
	The manually created NAT rule is placed at a higher priority than the corresponding automatically created NAT rule.
	
	You can also add other fields (such as addresses, ports, protocols, and so on) as part of the source or destination endpoints. For example:
	
	`Dept-Zone1, Port 56578 --> W1: Translation=Pool-2`

Table 19: Add Enterprise Hub for <Tenant-Name> Settings (WAN Capability) *(continued)*

Field	Description
Translation	Select the type of NAT to use for the traffic on the WAN link: • **Interface**—Use interface-based NAT, which is the default. • **Pool**—Use pool-based NAT. If you select this option, you must specify the IP addresses that are to be used for the NAT pool. **NOTE:** No NAT is performed for tenant-owned public IP addresses that were added during the tenant addition workflow.
IP Addresses	For pool-based NAT, enter one or more IP addresses, subnets, or an IP address range. Separate multiple IP addresses by using commas and use a hyphen to denote a range; for example, 192.0.2.1-192.0.2.50.
Preferred Breakout Link	Click the toggle button to enable the WAN link as the most preferred breakout link. If you disable this option, then the breakout link is chosen using ECMP from the available breakout links.
BGP Underlay Options	**NOTE:** This setting can be configured only if the address assignment is static and local breakout is enabled. Click the toggle button to enable BGP underlay routing. When you enable BGP underlay routing, route advertisements to the primary PE node and, if configured, the secondary PE node occur as follows: • CSO advertises the WAN interface subnet. • If you configured pool-based translation, CSO advertises the NAT address pool. **NOTE:** If underlay BGP is enabled for a WAN link, then the routes learnt from BGP are installed for local breakout; CSO does not generate the static default route.
Primary Neighbor	Displays the IP address that you entered for the gateway for the WAN link.
Secondary Neighbor	If you want to provide PE resiliency, you can configure a secondary PE node. Enter the IP address of the secondary PE node. **NOTE:** If the primary PE node goes down, then the secondary PE is used as the next hop. When the primary PE comes back up, the route next hops are changed to the primary PE.

Table 19: Add Enterprise Hub for <Tenant-Name> Settings (WAN Capability) *(continued)*

Field	Description
eBGP Peer-AS-Number	Enter the autonomous system (AS) number for the external (EBGP) peer. **NOTE:** If the peer AS number is not configured or the peer AS number that is configured is the same as that of the CPE site, then the BGP type is assumed to be internal BGP (IBGP).
Local AS Number	Enter the local AS number for the WAN link. When you configure this parameter, the local AS number is used for eBGP peering instead of the global AS number configured for the device. **NOTE:** The local AS number must be different from the global AS and eBGP peer AS numbers.
Authentication	Select the BGP route authentication method to be used: • **None**—Indicates that no authentication should be used. This is the default. • **Use MD5**—Indicates that MD5 is to be used for authentication. If you choose this option, you must specify an authentication key.
Auth Key	If you specified that MD5 should be used for authentication, specify an MD5 authentication key (password), which is used to verify the authenticity of BGP packets.
Advertise Public LAN Prefixes	Click the toggle button to enable the advertisement of public LAN prefixes. This field is disabled by default. If the tenant has a public IP address pool configured and you enable the advertisement of public LAN prefixes, then for LAN segments that are created with a subnet that falls under the tenant public IP address pool, CSO advertises the LAN subnet to the BGP underlay. **NOTE:** When public LAN advertisement is enabled for the WAN link, public LAN prefixes are advertised through the BGP underlay towards MPLS or the Internet. If a site has two versions of the route installed for the same LAN prefix in the overlay and underlay, the overlay routes are always preferred over underlay.
Use For Fullmesh	Click the toggle button to specify whether the WAN link can be a part of a full mesh topology. A site can have all WAN links enabled for meshing. **NOTE:** You must enable at least one WAN link for full mesh.

Table 19: Add Enterprise Hub for <Tenant-Name> Settings (WAN Capability) *(continued)*

Field	Description

Mesh Overlay Link Type	When **Use for Fullmesh** field is enabled, select the type of mesh overlay link—GRE and GRE_IPSEC.
	• If the link type is Internet, the value for mesh overlay link type is GRE_IPSEC.
	• If the link type is MPLS, select one of the following options:
	  ◦ GRE-IPSEC
	  ◦ GRE

Mesh Tag	When the **Use for Fullmesh** field is enabled, select one or more mesh tags to be associated with the WAN link for creating tunnels. Matching mesh tags is one of the criteria used to form tunnels between sites that support meshing. For more information about mesh tags, see "Mesh Tags Overview" on page 217".
Connects to Hubs	Click the toggle button to specify that the WAN link of the site connects to a hub. **NOTE:** • For sites with a single CPE, you must enable at least one WAN link to connect to the hub so that OAM traffic can be transmitted. • For sites with a dual CPE, you must enable at least one WAN link per device to connect to the hub so that OAM traffic can be transmitted.
Use for OAM Traffic	If you have specified that the WAN link is connected to a hub, click the toggle button to enable sending the OAM traffic over the WAN link. This WAN link is then used to establish the OAM tunnel.
Overlay Peer Device	This field is displayed when the **Connects to Hubs** field is enabled and only a one provider hub (primary) is specified. Displays the peer hub device to which the site is connected.
Overlay Peer Interface	This field is displayed when the **Connects to Hubs** field is enabled and only a one provider hub (primary) is specified. Select the interface name of the hub device to which the WAN link of the site is connected.

Table 19: Add Enterprise Hub for <Tenant-Name> Settings (WAN Capability) *(continued)*

Field	Description
Overlay Tunnel Type 1	This field is displayed when the **Connects to Hubs** field is enabled and both primary and secondary hubs are specified. Select the mesh overlay tunnel type (GRE and GRE_IPSEC) for the tunnel to the primary hub. MPLS links can have both GRE and GRE_IPSEC as the overlay link type where as Internet links can have only GRE_IPSEC as the overlay link type.
Overlay Peer Device 1	This field is displayed when the **Connects to Hubs** field is enabled and both primary and secondary hubs are specified. Displays the primary peer hub device to which the site is connected.
Overlay Peer Interface 1	This field is displayed when the **Connects to Hubs** field is enabled and both primary and secondary hubs are specified. Select the interface name of the primary hub device to which the WAN link of the site is connected.
Overlay Tunnel Type 2	This field is displayed when the **Connects to Hubs** field is enabled and both primary and secondary hubs are specified. Select the mesh overlay tunnel type (GRE and GRE_IPSEC) for the tunnel to the secondary hub. MPLS links can have both GRE and GRE_IPSEC as the overlay link type where as Internet links can have only GRE_IPSEC as the overlay link type.
Overlay Peer Device 2	This field is displayed when the **Connects to Hubs** field is enabled and both primary and secondary hubs are specified. Displays the secondary peer hub device to which the site is connected.
Overlay Peer Interface 2	This field is displayed when the **Connects to Hubs** field is enabled and both primary and secondary hubs are specified. Select the interface name of the secondary hub device to which the WAN link of the site is connected.

Table 19: Add Enterprise Hub for <Tenant-Name> Settings (WAN Capability) *(continued)*

Field	Description
Backup Link	Select a backup link through which traffic can be routed when the primary (other) links are unavailable. You can select any link other than the default links or links that are configured exclusively for local breakout traffic. When a primary link comes back online, CSO monitors the performance on the primary link and when the primary link meets the SLA requirements, the traffic is switched back to the primary link. However, SLA data is not monitored for the backup link.
Default Link	Select one or more links that will be used for routing traffic in the absence of matching SD-WAN policy intents. A site can have multiple default links to the hub site. Default links are used primarily for overlay traffic but can also be used for local breakout traffic. However, a default link cannot be used exclusively for local breakout traffic. If you do not specify a default link, then equal-cost multipath (ECMP) is used to choose the link on which to route traffic.
Data VLAN ID	Enter a VLAN ID for the WAN link. Range: 2 through 4093. **NOTE:** • If you are configuring more than one WAN link on the same physical interface, only one WAN link can be untagged; for the remaining WAN links, you must configure a VLAN ID. • A combination of tagged and untagged on the same physical interface is supported only for single CPE devices. To enable the configuration of WAN links as logical interfaces in on-premise SD-WAN spoke sites, the SP Administrator user must modify the device template and configure the WAN ports as logical interfaces.
WAN_1 (*WAN-Interface-Name*)	Click the toggle button to enable or disable (default) the WAN link. When you enable the WAN link, fields related to the WAN link appear. Fields marked with an asterisk (*) must be configured to proceed. Refer to the fields described for WAN_0 (*WAN-Interface-Name*) for an explanation of the fields

Table 19: Add Enterprise Hub for <Tenant-Name> Settings (WAN Capability) *(continued)*

Field	Description
WAN_2 (*WAN-Interface-Name*)	Click the toggle button to enable or disable (default) the WAN link. When you enable the WAN link, fields related to the WAN link appear. Fields marked with an asterisk (*) must be configured to proceed. Refer to the fields described for WAN_0 (*WAN-Interface-Name*) for an explanation of the fields
WAN_3 (*WAN-Interface-Name*)	Click the toggle button to enable or disable (default) the WAN link. When you enable the WAN link, fields related to the WAN link appear. Fields marked with an asterisk (*) must be configured to proceed. Refer to the fields described for WAN_0 (*WAN-Interface-Name*) for an explanation of the fields
Management Connectivity	
IP Prefix	Enter an IPv4 address prefix for the loopback interface on the CPE device. The IP address prefix must be a /32 IP address prefix and must be unique across the entire management network. If you do not specify an IPv4 address prefix, CSO automatically assigns the IP prefix from the reserved pool 100.124.0.0/14.
Refer to Table 22 on page 78 for configuring LAN segments.	

Table 20: Automatic Creation of Source NAT Rules

Autocreate Source NAT Rule	Translation	NAT Rules Creation
Disabled	Not applicable (No NAT)	None.

Enabled	Interface-Based (Default)—CSO creates interface-based NAT rules.	Source NAT rules are automatically created, with each rule from a department zone to the WAN interface, with a translation of type interface. Each pair of [zone - interface] represents a rule-set.
		
		For example, the following department zone to (WAN link) W1 interface rule-set might be created:
		
		`Dept-Zone1 --> W1: Translation=Interface`
		`Dept-Zone2 --> W1: Translation=Interface`
		`Dept-Zone3 --> W1: Translation=Interface`

Table 20: Automatic Creation of Source NAT Rules *(continued)*

Autocreate Source NAT Rule	Translation	NAT Rules Creation

Enabled	Pool-Based—CSO automatically creates pool-based NAT rules	NAT source rules are automatically created, with each rule from a department zone to the WAN NAT pool with a translation of type pool.
		
		For example, a source NAT rule from department zone to NAT pool might be created:
		
		`Dept-Zone1 --> W1 : Translation=Pool-1`
		`Dept-Zone2 --> W1 : Translation=Pool-1`

Table 21: Add Enterprise Hub for <Tenant-Name> Settings (LAN Capability)

Field	Description
LAN **NOTE:** This tab is enabled only if you select **LAN** under LAN Capabilities in General Settings.	
Device Profile	
Device Name	Enter a name for the switch. You can use alphanumeric characters and hyphen (-). The maximum length allowed is 15 characters.
Device Type	Select the type of switch—EX2300, EX3400, and EX4300
Device Model	Select the model for the switch you specified in the Device Type. The models vary in the number and type of ports the switch contains. For example, If you selected EX3400, select a model such as EX3400-24P, EX3400-48P, EX3400-24T among others.
CPE Settings	
Trunk Ports	Select at least two trunk ports on the CPE device to connect with the switch. The trunk ports are used for carrying the following: • LAN traffic between the switch and the CPE. • Management traffic for in-band management of the switch.

Table 21: Add Enterprise Hub for <Tenant-Name> Settings (LAN Capability) *(continued)*

Field	Description
Switch Management Subnet	Specify the subnet that the DHCP can use to assign IP addresses. The DHCP server runs on the following ports: • Trunk ports to provide DHCP information to all devices connected to the switch and to the in-band management port, switch management port, and LAN ports on the CPE. • Out-of-band management port on the CPE to provide DHCP information to the management port on the switch. • LAN ports on the CPE to provide information to the devices connected to the CPE LAN ports.
Switch Details	
Serial Number	Specify the serial number of the switch.
Auto Activate	Click the toggle button to enable or disable automatic activation of the switch when the switch is detected by CSO (that is, management status of the device is Device_Detected). When you enable this field, zero-touch provisioning (ZTP) of the switch is automatically triggered when the device communicates with CSO. By default, auto activation for the switch is enabled or disabled if it is enabled or disabled for the CPE. **NOTE:** You must physically connect the switch to the CPE and power it on for the switch to be automatically activated when you enable this option.
Activation code	When the **Auto activate** field is disabled, enter the activation code to be used for manually activating the switch. . For information about manually activating a switch, see "Manually Activating a Switch" on page 232.
Zero Touch Provisioning	ZTP must be disabled for all EX Series switches for the CSO 5.0.0 release. The Stage-1 configuration must be copied and pasted onto the CLI of the switch during site activation. See "Step-by-Step Procedure" on page 62 for details.
LAN Segments	Displays the LAN segment that you configure on the switch. To add a LAN segment, click the + icon on the top, right corner of the LAN table. The Add LAN Segment page appears. See Table 22 on page 78.

Table 22: Add LAN Segment Settings

Field	Description
Name	Enter a name for the LAN segment. The name for a LAN segment should be a unique string of alphanumeric characters and some special characters (. -). No spaces are allowed and the maximum length is 15 characters.
Type **NOTE:** This field is displayed only for LAN segments associated with enterprise hub sites.	Select the type of LAN segment: • Directly Connected (default)—Indicates that the LAN segment is directly connected to the site. • Dynamic Routed—Indicates that the LAN segment is not directly connected to the site and is reachable by using a dynamic route. If you select this option, you must specify the dynamic routing information.
VLAN ID	Enter the VLAN ID for the LAN segment. Range: 2 through 4093.
Department	Select a department to which the LAN segment is assigned. Alternatively, click the **Create Department** link to create a new department and assign the LAN segment to it. See "Adding a Department" on page 767 for details. You can group LAN segments as departments for ease of management and for applying policies at the department-level. For LAN segments that are dynamically routed, you can assign only a data center department.
Protocol	For dynamically routed LAN segments, select the routing protocol (BGP or OSPF) to be used by the data center department to learn routes from the data center.
Advertise LAN Prefix	For dynamically routed LAN segments, click the toggle button to advertise the LAN prefix of the SD-WAN spoke site to the data center through the data center department associated with the enterprise hub. By default, the advertise LAN prefix field is disabled. **NOTE:** You must avoid overlapping IP addresses between the SD-WAN LAN network and the datacenter network.
Gateway Address/Mask	Enter a valid gateway IP address and mask for the LAN segment. This address will be the default gateway for endpoints in this LAN segment. For example: 192.0.2.8/24.

Table 22: Add LAN Segment Settings *(continued)*

Field	Description
DHCP	For directly connected LAN segments, click the toggle button to enable DHCP (default). You can enable DHCP if you want to assign IP addresses by using a DHCP server or disable DHCP if you want to assign a static IP address to the LAN segment. **NOTE:** If you enable DHCP, additional fields appear on the page.
Additional fields related to DHCP	
Address Range Low	Enter the starting IP address in the range of IP addresses that can be allocated by the DHCP server to the LAN segment.
Address Range High	Enter the ending IP address in the range of IP addresses that can be allocated by the DHCP server to the LAN segment.
Maximum Lease Time	Specify the maximum duration (in seconds) for which a client can request for and hold a lease on the DHCP server. Default: 1440 Range: 0 through 4,294,967,295 seconds.
Name Server	Specify one or more IPv4 addresses of the DNS server. To enter more than one DNS server address, type the address, press Enter, and then type the next address. **NOTE:** DNS servers are used to resolve hostnames into IP addresses.

Table 22: Add LAN Segment Settings *(continued)*

Field	Description

CPE Ports	• For sites with LAN capability, click the toggle button to include or exclude the CPE in the LAN segment.
	  ◦ When you include the CPE in the LAN segment:
	  ▪ CPE ports that you can include in the LAN segment are listed. Select the ports from the **Available** column and click the right-arrow to move the ports to the **Selected** column.
	  ▪ The **Switch Ports** field is disabled. CSO automatically assigns LAN ports on the Switch device and creates the same LAN segment on the Switch.
	  ◦ If you click to exclude the CPE from the LAN segment, you must specify the switch ports that connect with the LAN in the **Switch Ports** field. CSO automatically assigns LAN ports on the CPE device and creates the same LAN segment on the CPE device.
	  NOTE: You can select only one port if the CPE is a physical SRX Series device.
	• For sites without LAN capability, the CPE Ports field is disabled and the CPE ports that you can include in the LAN segment are listed.
	  Select the ports from the **Available** column and click the right-arrow to move the ports to the **Selected** column.

Switch Ports **NOTE:** This field is displayed only when LAN capability is selected for the enterprise hub.	If you disable the CPE ports field, select ports on the switch to be part of the LAN segment. The Switch ports and CPE ports are mutually exclusive. Select the ports from the **Available** column and click the right-arrow to move the ports to the **Selected** column.
BGP Configuration **NOTE:** This section is displayed only for dynamic routed LAN segments with BGP specified as the protocol.	
Authentication	Select the BGP route authentication method to be used: • None—Indicates that no authentication should be used. This is the default. • Use MD5—Indicates that MD5 is to be used for authentication. If you choose this option, you must specify an authentication key.
Peer IP Address	Enter the IP address of the BGP neighbor.
Peer AS Number	Enter the autonomous system (AS) number BGP neighbor.
Auth Key	If you specified that MD5 should be used for authentication, specify an MD5 authentication key (password), which is used to verify the authenticity of BGP packets.

Table 22: Add LAN Segment Settings *(continued)*

Field	Description
OSPF Configuration **NOTE:** This section is displayed only for dynamic routed LAN segments with OSPFspecified as the protocol.	
OSPF Area ID	Specify the OSPF area identifier to be used for the dynamic route.
Authentication	Select the OSPF route authentication method to be used: • Password—Indicates that password-based authentication should be used. If you choose this option, you must specify the password. (This is the default). • Use MD5—Indicates that MD5 is to be used for authentication. If you choose this option, you must specify an authentication key. • None—Indicates that no authentication should be used.
Password	Enter the password to be used to verify the authenticity of OSPF packets.
Confirm Password	Retype the password for confirmation purposes.
MD5 Auth Key ID	If you specified that MD5 should be used for authentication, enter the OSPF MD5 authentication key ID. Range: 1 through 255.
Auth Key	If you specified that MD5 should be used for authentication, enter an MD5 authentication key, which is used to verify the authenticity of OSPF packets.

RELATED DOCUMENTATION

About the Sites Page | 53

Enterprise Hubs Overview | 55

Add a Switch to an Existing SD-WAN Site Or Next-Generation Firewall Site | 160

Adding Provider Hub Sites for SD-WAN Deployment

A provider hub site represents an automation endpoint that is part of a data center or POP that is owned by the service provider. The provider hub site is connected to multiple spoke sites using the overlay connections. Provider hubs sites are logical entities in a multi-tenant device (provider hub device). You

add a provider hub site from the **Sites** page. This page describes how to add a provider hub site for a tenant.

To add a provider hub site:

1. Select **Resources > Site Management**.

   The Sites page appears.

2. Click **Add** and select **Add Provider Hub**.

   The **Add Provider Hub for *Tenant-Name*** page appears.

3. Complete the configuration settings according to the guidelines provided in Table 23 on page 82.

   NOTE: Fields marked with an asterisk (*) are mandatory.

4. Click **OK**.

   The newly added provider hub site is displayed on the **Sites** page.

Table 23: Fields on the Provider Hub for Tenant-Name Page

Field	Description
Configuration	
Service POP	Select the name of the point of presence (POP) for the site. A network POP is a location at which a service provider instantiates a network function, such as a virtualized network function (VNF).
Hub Device Name	Select the provider hub device for the hub site. The provider hub devices with **DATA_ONLY** capability and **DATA_AND_OAM** capability are listed. **NOTE:** You need not onboard a provider hub with OAM_ONLY capability. By default, the provider hub with OAM_ONLY capability will be imported.

RELATED DOCUMENTATION

About the Sites Page | 53

About the Site Groups Page | 196

Adding Cloud Spoke Sites for SD-WAN Deployment

A cloud spoke represents an automation endpoint (virtual machine (VM) or an EC2 Instance) running a Juniper Networks vSRX image in the Amazon Web Services(AWS) virtual private cloud (VPC). The cloud spoke sites are connected to the hub sites using the overlay connections. You create a cloud spoke site from the **Sites** page. This topic describes how to add a cloud spoke site for a tenant.

NOTE:

- You can add a cloud spoke site only in hub-and-spoke topology.
- To ensure that only hub-and-spoke topology is created, we recommend you to disable the DVPN configuration while adding the tenant.
- You cannot add a cloud spoke site in full mesh topology.

To add a cloud spoke site:

1. Select **Resources > Site Management**.

   The Sites page appears.

2. Click **Add** and select **Add Cloud Spoke**.

   The **Add On-Premise Spoke Site for *Tenant Name*** page appears.

3. Complete the configuration settings according to the guidelines provided in Table 24 on page 83.

   NOTE: Fields marked with an asterisk (*) are mandatory.

4. Review the configuration and modify the settings, if needed, from the **Summary** tab.

5. Click **OK**.

   The newly added cloud spoke site is displayed on the **Sites** page.

Table 24: Fields on the Add Cloud Spoke Site Page

Field	Description
General	
Site Information	

Table 24: Fields on the Add Cloud Spoke Site Page *(continued)*

Field	Description
Site Name	Enter a unique name for the site. Enter a unique string of alphanumeric characters and special character (-). The maximum length is 15 characters. Example: aws-cloud-spoke
Site Group	(Optional) Select a site group to which you want to assign the site. Example: cloud-spoke
Site Capabilities	
WAN Capabilities	Select **SD-WAN** to include SD-WAN capabilities in the cloud spoke site
Configuration	
Primary Provider Hub	Select the hub site to which the spoke site must connect.
Advanced Configuration	
Name Server IP List	Enter one or more IPv4 addresses of the DNS server. To enter more than one DNS server address, type address, press Enter, and then type the next address, and so on. DNS servers are used to resolve hostnames into IP addresses.
NTP Server	Enter the fully qualified domain names (FQDNs) or IP addresses of one or more NTP servers. Example: ntp.example.net The site must have DNS reachability to resolve the FQDN during site configuration.
Select Timezone	Select the time zone for the site.
WAN	
Device Template	Click a device template to select the plan for WAN connectivity. A device template contains information such as device family, a list of SD-WAN features supported, and the number of links supported. **NOTE:** vSRX as SD-WAN spoke in AWS template supports cloud spoke site for AWS VPC.
Cloud Information	

Table 24: Fields on the Add Cloud Spoke Site Page *(continued)*

Field	Description
Region	Select the region to which the site belongs. The regions in CSO are mapped to the regions in the AWS account. Example: Ohio
VPC ID	Enter the VPC ID from the AWS account. To obtain VPC ID: 1. Log in to your AWS account. 2. Search for the VPC service. 3. Click the VPC dashboard. 4. Select a VPC ID. Ensure that the VPC is connected to an Internet gateway. To check whether VPC is attached: 1. Log in to your AWS account. 2. Search for the VPC service. 3. Click the Internet Gateway dashboard. 4. Check whether the VPC state is **attached**. Example: vpc-6d810314
Management Subnet	Specify whether CSO must create a new subnet or use an existing subnet from the AWS account. The management subnet of vSRX is used to push the initial stage-1 configuration. The following options are available: • Use an existing subnet in AWS account • Create new
IP Prefix	Enter the management IP prefix. The first four IP addresses in the subnet are reserved by AWS. For example, IP addresses x.x.x.0/x through x.x.x.3/x are always reserved by AWS. Hence, provide an IP address prefix other than the reserved IP address prefix. Example: 105.0.1.5/24

Table 24: Fields on the Add Cloud Spoke Site Page *(continued)*

Field	Description
Device Information	
Activation Code	Enter the activation code of the primary device that your service provider supplied for the device. If you do not want to specify an activation code, on the Template Settings page, disable the ACTIVATION_CODE_ENABLED field and save the changes.
WAN Links	
WAN_0 (ge-0/0/0) WAN_1 (ge-0/0/1)	Select the check boxes to configure the WAN links. You can configure up to two WAN links per site that support SD-WAN.
Link Type	Displays the connection type for WAN underlays. Only Internet link is supported.
Egress Bandwidth	Enter the maximum bandwidth (in Mbps) to be allowed for a specific WAN link.
Address Assignment	Select the method of assigning an IP address to the WAN link—DHCP or STATIC. • If you select DHCP, the IP address is provided by using the DHCP server of the service provider of the WAN link. • If you select STATIC, you must provide the IP address prefix and the gateway address for the WAN link.
Static IP Prefix	If you configure the address assignment method as STATIC, enter the private IPv4 address of the WAN link from the subnet. For example, if the IPv4 CIDR address is 105.0.2.0/24 for a WAN interface in the AWS account, then enter any IP address within the subnet. The first four IP addresses in the subnet are reserved by AWS. Hence, provide an IP prefix other than the reserved IP prefix. Example: 105.0.2.12/24
Gateway IP	If you configured the address assignment method as STATIC, enter the IPv4 address for the gateway of the WAN service provider. Typically, the first IP address in the subnet is selected for gateway IP address. Example: 105.0.2.1
Elastic IP	Elastic IP address is a public, static IPv4 address designed for dynamic cloud computing. The public IP address is mapped to the private subnet IP using one-to-one NAT. You must allocate the IP addresses based on the number of WAN links that are enabled. For example, If two WAN links are enabled, then you must allocate two elastic IP addresses. Example: 34.213.255.184

Table 24: Fields on the Add Cloud Spoke Site Page *(continued)*

Field	Description
Advanced Settings	Based on the connectivity requirement, the following fields are populated:
Provider	Enter the name of the service provider (SP).
Cost/Month	Enter the cost per month of the subscribed bandwidth in the specified currency. In bandwidth-optimized SD-WAN, this information is used to identify the least-expensive link to route traffic when multiple WAN links meet SLA profile parameters.
Enable Local Breakout	Click the toggle button to enable or disable (default) local breakout on the WAN link. • If you enable this option, the WAN link can be used for local breakout. The decision of whether traffic breaks out locally from the site depends on the breakout profile that is referenced in the SD-WAN policy intent. • If you do not enable local breakout on at least one WAN link for a single CPE connection plan and at least two WAN links for a dual CPE connection plan, then local breakout is disabled for the site.
Breakout Options	Select whether you want to use the WAN link for both breakout and WAN traffic (default) or only for breakout traffic.

Table 24: Fields on the Add Cloud Spoke Site Page *(continued)*

Field	Description

Autocreate Source NAT Rule	If the WAN link is enabled for local breakout, you can click the toggle button to automatically create an interface-based source NAT rule on the WAN link. The automatically-created source NAT rule is implicitly defined and applied to the site and is not visible on the NAT Policies page.
	
	By default, this field is disabled.
	
	NOTE: If this option is enabled for a WAN interface W1 during the site addition workflow, a series of NAT source rules are automatically created. Each automatically created NAT rule is from a zone to the WAN interface, with a translation of type interface. Each pair of [zone - interface] represents a rule-set.
	
	For example, the following zone to W1 interface rule-set might be created:
	
	`Zone1 --> W1: Translation=Interface`
	`Zone2 --> W1: Translation=Interface`
	`Zone3 --> W1: Translation=Interface`
	
	To manually override any of these rules, you can create a NAT rule within a particular rule-set. For example, to use a source NAT pool instead of an interface for translation, create a NAT rule within this particular rule-set, that includes the relevant zone and WAN interface as the source and destination. For example:
	
	`Zone1 --> W1 : Translation=Pool-2`
	
	The manually created NAT rule is placed at a higher priority than the corresponding automatically created NAT rule.
	
	You can also add other fields (such as addresses, ports, protocols, and so on) as part of the source or destination endpoints. For example:
	
	`Zone1, Port 56578 --> W1: Translation=Pool-2`

Preferred Breakout Link	Click the toggle button to enable the WAN link as the preferred breakout link. If you disable this option, then the breakout link is chosen using ECMP from the available breakout links.
Use for OAM Traffic	If you have specified that the WAN link is connected to a hub, click the toggle button to enable sending the OAM traffic over the WAN link. This WAN link is then used to establish the OAM tunnel.
Overlay Tunnel Type	Select the mesh overlay tunnel type—GRE and GRE_IPSEC. MPLS links can have both GRE and GRE_IPSEC as the overlay link type where as Internet links can have only GRE_IPSEC as the overlay link type.

Table 24: Fields on the Add Cloud Spoke Site Page *(continued)*

Field	Description
Overlay Peer Device	Displays the peer hub device to which the site is connected.
Overlay Peer Interface	Select the interface name of the hub device to which the WAN link of the site is connected.
Backup Link	Select a backup link through which traffic can be routed when the primary links are unavailable. You cannot select the default link as the backup link. Note that you cannot assign the backup link for exclusive breakout traffic (the **Use only for breakout traffic** option). If local breakout is enabled for the site, the breakout traffic is also routed through the backup link when the breakout link is not available. When a primary link comes back online, CSO monitors the performance on the primary link and when the primary link meets the SLA requirements, the traffic is switched back to the primary link. However, note that the SLA data is not monitored for the backup link.
Default Links	Select the default links that must be used for routing traffic. The site can have multiple default links to the hub site as well as to the Internet. Default links are used primarily for overlay traffic but can be used for local breakout traffic as well. A default link cannot be used exclusively for local breakout traffic. The default link is optional and in case it is not chosen, all links are used through equal-cost multipath (ECMP).
Management Connectivity	
IP Prefix	Enter an IPv4 address prefix for the loopback interface on the CPE device. The IP address prefix must be a /32 IP address prefix and must be unique across the entire management network. If you do not specify an IPv4 address prefix, CSO automatically assigns the IP prefix from the reserved pool 100.124.0.0/14 Example: 192.0.2.10/32
LAN	Add at least one LAN segment.
LAN Segment	Displays the LAN segment that you configure on the switch. To add a LAN segment, click the **+** icon on the top, right corner of the LAN table. The Add LAN Segment page appears. See Table 25 on page 89.

Table 25: Fields on the Add LAN Segment Page

Field	Description
Add LAN Segment	

Table 25: Fields on the Add LAN Segment Page *(continued)*

Field	Description
Name	Enter a name for the LAN segment. The name for a LAN segment should be a unique string of alphanumeric characters. No spaces are allowed and the maximum length is 15 characters.
Department	Select a department to which the LAN segment is to be assigned. Alternatively, click the **Create Department** link to create a new department and assign the LAN segment to it. See "Adding a Department" on page 767 for details. You group LAN segments as departments for ease of management and for applying policies at the department-level.
Gateway Address/Mask	Enter a valid gateway IP address and mask for the LAN segment; for example, 192.0.2.8/24.
CPE Ports	Click the toggle button to include or exclude the CPE in the LAN segment. When you include the CPE in the LAN segment: • CPE ports that you can include in the LAN segment are listed. Select the ports from the **Available** column and click the right-arrow to move the ports to the **Selected** column. • The **Switch Ports** field is disabled. CSO automatically assigns LAN ports on the switch device and creates the same LAN segment on the switch. If you exclude the CPE from the LAN segment, you must specify the switch ports that connect with the LAN in the **Switch Ports** field. CSO automatically assigns LAN ports on the CPE device and creates the same LAN segment on the CPE device. **NOTE:** You can select only one port if the CPE is an SRX Series device.

RELATED DOCUMENTATION

Provisioning a Cloud Spoke Site in AWS VPC | **91**

About the Sites Page | **53**

About the Site Groups Page | **196**

Provisioning a Cloud Spoke Site in AWS VPC

IN THIS SECTION

- Add a Cloud Spoke Site | 91
- Download the Cloud Formation Template | 92
- Provision the Device on AWS Server | 93
- Activate the Device | 94

Use the following high-level steps to provision a vSRX cloud spoke site in Amazon Web Services (AWS) virtual private cloud (VPC).

Before you begin:

- Set up your Amazon Web Services (AWS) account.
- Identify the virtual private cloud (VPC) in which the AWS spoke site must be provisioned.
- Install licenses to use vSRX features. Choose any of the following AWS vSRX Image Licenses.
  - Bring Your Own License (BYOL)— If you plan to use a BYOL, then you must install the license on the device before deploying CSO SD-WAN functionality. See https://aws.amazon.com/marketplace/pp/B01LYWCGDX.
  - License included. See https://aws.amazon.com/marketplace/pp/B01NAUWN0G.
- Ensure that you have the supported software version for the AWS spoke.
- Reserve two elastic IP (public IP) addresses on AWS.

To set up and monitor your network:

Add a Cloud Spoke Site

To add a cloud spoke site:

1. Select **Resources > Site Management**.

   The Sites page appears.

2. Click **Add > Add Cloud Spoke**.

   The **Add On-Premise Spoke Site for *Tenant-Name*** page appears.

3. Specify the site information such as, site name, AWS region, VPC ID, management subnet, IP prefix and click **Next**.

4. Specify vSRX as SD-WAN spoke in AWS as the device template.

> **NOTE:**
> - Only hub-and-spoke topology is supported for AWS cloud spoke site.
> - Only Internet link is supported for WAN underlay connections.

5. Provide the WAN details and click **Next**.

   The WAN traffic page appears, displaying a set of values for the WAN link configuration.

6. Specify additional requirements and click **Next**.

7. Specify LAN segment information and click **Next**.

8. In the **Summary** tab, check the configuration and click **Edit** to modify the settings.

9. Click **OK** to save the changes.

   The new cloud spoke site that you created appears in the Sites page.

Download the Cloud Formation Template

To download the cloud formation template:

1. Click **Resources > Devices**.

   The Devices List page appears.

2. Select the device and click **Cloud Info Template**.

   The Cloud Info Template page appears.

3. Click **Download** to download the cloud formation template.

   The template is downloaded to your local computer in JSON format.

Provision the Device on AWS Server

CSO creates cloud formation template with stage-1 configuration bundled in JSON format. You must download this template and then upload to AWS to provision the vSRX. The cloud formation template creates the required resources such as subnet, interface, vSRX and so on and applies the stage-1 configuration.

To provision the device on AWS server:

1. Log in to your AWS account.
   - If you have already logged in to your AWS account, the Create Stack page appears.
   - If you are not logged into your AWS account, a new Web page opens in your browser, displaying the AWS login information. Log in to your AWS account.

     TIP: If you do not see the Create Stack page when you log in to or access your AWS account, then search for CloudFormation service.

   The Create Stack page appears.
2. Select **CloudFormation** > **Stacks** > **Create Stack** > **Upload a template to Amazon S3**.
3. Click **Choose File** and select the cloud formation template that you downloaded in JSON format .
4. Click **Next**.
5. Specify the Stack name. For example, Oregonstack.
6. In the Parameters section, specify the KeyName for your EC2 instance.
7. Click **Next**.
8. Select **I acknowledge that AWS CloudFormation might create IAM Resources**.
9. Click **Create**.

   The Create Stack pages displays a list of existing stacks and indicates that it is creating the stack that you requested. The create stack process takes up to 30 minutes. if the process does not complete in 30 minutes, a timeout occurs and you need to retry the process.

Activate the Device

To activate the device:

1. After the create stack process is complete, return to the Customer Portal and click **Next**.

   The Activate Device page displays a status indicating that CSO is detecting the provisioning agent. This process takes up to 30 minutes. if the process does not complete in 30 minutes, a timeout occurs and you need to retry the process.

   > **NOTE:** You need not download the cloud formation template again. You can log in to the Customer Portal, access the Activate Device page, enter the activation code and click **Next**. After the CREATE_COMPLETE message is displayed on the AWS server, click Next on the Activate Device page to proceed with device activation.

   If the spoke on AWS has been spawned successfully on AWS, it will contact CSO through outbound SSH connection. The device is detected and normal ZTP, process is triggered. The rest of the workflow is consistent with the normal on-premise workflow.

   On Device Activation page, the device is activated through the following steps:

   - Detecting the device
   - Applying stage-one configuration to the device
   - Bootstrapping of device
   - Activating the device

   After each successful step, you can see a green check mark. If any of these steps fails, a red exclamation mark appears.

2. After the activation process is complete, click **OK**.

   The Sites page appears. To see the device activation status, hover over the device icon on the Sites page.

RELATED DOCUMENTATION

Adding Cloud Spoke Sites for SD-WAN Deployment | 83

vSRX Deployment Guide for AWS

Manually Adding On-Premise Spoke Sites

An on-premise site can be added with one WAN capability (SD-WAN, or Hybrid WAN, or Next Gen Firewall), LAN capability, or both WAN and LAN capabilities.

> **NOTE:** The WAN and LAN capabilities that are displayed in the Add On-premise Spoke Site page are filtered based on the service types that are assigned to the tenant.

For more information on adding an on-premise spoke site with the following capabilities:

- WAN capability as SD-WAN, see "Adding an On-Premise Spoke Site with SD-WAN Capability" on page 100.
- WAN capability as Hybrid WAN, see "Adding an On-Premise Spoke Site with Hybrid WAN Capability" on page 95.
- WAN capability as Next Gen Firewall, see "Adding a Standalone Next Generation Firewall Site" on page 177.
- WAN capability as SD-WAN and LAN capability, see "Adding an On-Premise Spoke Site with SD-WAN and LAN Capabilities" on page 117.
- WAN capability as Next Gen Firewall and LAN capability, see "Adding an On-Premise Spoke Site with Next Generation Firewall and LAN Capabilities" on page 151.
- Only LAN capability, see "Add an On-Premise Spoke Site with LAN Capability" on page 136.

Adding an On-Premise Spoke Site with Hybrid WAN Capability

You add an on-premise spoke site with Hybrid WAN capability from the **Site** page. The Hybrid WAN sites can have a maximum of two WAN links. You cannot apply intent policies for Hybrid WAN sites.

To add an on-premise spoke site with Hybrid WAN capability:

1. Select **Resources > Site Management** .

   The Sites page appears.

2. Click **Add** and select **Add On-Premise Spoke (Manual)**.

   The Add Site for *Tenant-Name* page appears.

3. Complete the configuration settings in the General and WAN sections according to the guidelines provided in Table 26 on page 96.

4. Review the configuration and modify the settings, if needed, from the **Summary** tab.

5. Click **OK**.

   You are returned to the Sites page and a message indicating that the site creation job was triggered is displayed. You can click the job ID link to view the progress of the job. After the job is completed successfully, a confirmation message is displayed and the site that you added is displayed on the Sites page.

Table 26: Fields on the Add Spoke Site Page

Field	Description
General	
Site Information	
Site Name	Enter a site name. You can use any number of alphanumeric characters, including special characters. The maximum length is 10 characters.
Site Group	Select a site group to which you want to assign the site.
Site Capabilities	
WAN Capabilities	Select **Hybrid WAN** to include Hybrid WAN capability in the on-premise spoke site.
LAN Capabilities	This option is disabled for Hybrid WAN.
Address and Contact Information	
Street Address	Enter the street address of the site.
City	Enter the city where the site is located.
State/Province	Enter the state or province where the site is located.
ZIP/Postal Code	Enter the postal code for the locality of the site.

Table 26: Fields on the Add Spoke Site Page *(continued)*

Field	Description
Country	Select the country where the site is located. Click the **Validate** button to verify the address. • The **site address verification successful** message is displayed if the address is verified. You can click the **View location on a map** link to see the address location. • If the address cannot be verified, the **Site address could not be validated** message is displayed.
Contact Name	Enter the name of a contact person for the site.
Email	Enter the e-mail address of the contact person for the site.
Phone	Enter the phone number of the contact person for the site.
Advanced Configuration	
Name Server IP List	Specify one or more IPv4 addresses of the DNS server. To enter more than one DNS server address, type the address, press Enter, and then type the next address, and so on. DNS servers are used to resolve hostnames into IP addresses.
NTP Server	Specify the fully qualified domain names (FQDNs) or IP addresses of one or more NTP servers. Example: ntp.example.net The site must have DNS reachability to resolve the FQDN during site configuration.
Select Timezone	Select the time zone of the site.
Device Profile	
Device Series	Select the device series to which the CPE belongs—SRX, NFX150, or NFX250.

Table 26: Fields on the Add Spoke Site Page *(continued)*

Field	Description
Device Template	Select a device template for the selected device series. The device template contains information for configuring a device.
Device Information	
Device Model	Select the device model for NFX150 device.
Serial Number	Enter the serial number of the CPE device. The serial number is a 12-digit number present on the rear panel of the device. Serial numbers are case-sensitive.
Auto Activate	Click the toggle button to enable or disable automatic activation of the CPE device. When you enable this field, zero-touch provisioning (ZTP) of the CPE device is automatically triggered after the site is added to CSO. The device template that you select determines whether this option is enabled or disabled by default.
Activation Code	Enter the activation code of the CPE device that your service provider supplied. If you do not want to specify an activation code, then disable the ACTIVATION_CODE_ENABLED field in the device template and save the changes.
Boot image	Select the boot image from the drop-down list if you want to upgrade the image for the CPE device. The boot image is the latest build image uploaded to the image management system. The boot image is used to upgrade the device when the CSO starts the ZTP process. If the boot image is not provided, then the device skips the procedure to upgrade the device image. The boot image (NFX or SRX) is populated based on the device template that you have selected while creating a site. See *Uploading a Device Image*.
CPE Info	

Table 26: Fields on the Add Spoke Site Page *(continued)*

Field	Description
CPE AS Number	Specify the autonomous system(AS) number.
Access Info	
Router Name	Specify the router name.
Router AS Number	Specify the AS number for the router.
Management Connectivity	
OAM Traffic Information	Select this option if the management connectivity is initiated by Contrail Service Orchestration (CSO).
VLAN Id	Enter an OAM VLAN ID for in-band management of the site. If you specify an OAM VLAN ID, then an in-band OAM traffic reaches the site through the selected OAM interface. The range is 0 through 65535.
IP Prefix	Enter an IPv4 address prefix for the loopback interface on the CPE device. The IP address prefix must be a /32 IP address prefix and must be unique across the entire management network. If you do not specify an IPv4 address prefix, CSO automatically assigns the IP prefix from the reserved pool 100.124.0.0/14. • For NFX150 and NFX250 devices, if the USE_SINGLE_SSH_TO_NFX parameter is disabled in the device template, then enter the IP address prefix as /29 or lower based on the number of VNFs. • For all other devices, enter the IP address prefix as /32.
Gateway IP	If you configured the address assignment method as STATIC, enter the IP address of the gateway of the WAN service provider.
WAN links	
WAN_0	
Link Type	Select whether the link would be an MPLS link or Internet link.

Table 26: Fields on the Add Spoke Site Page *(continued)*

Field	Description
VLAN ID	Specify the identifier for the Layer 2 VLAN for the CPE device.
Local IP Prefix	Enter the local IP address prefix of the WAN link.
Remote IP Prefix	Enter the remote IP address prefix of the WAN link.
WAN_1	Refer to the fields described for WAN_0 for an explanation of the fields.
VRF Name	Specify the name of the virtual routing and forwarding (VRF) instance.
Ipsec Concentrator Name	Specify the name of the IPsec concentrator device.
Internet Gateway IP	If you specified that the device is an IPsec concentrator, then specify the IPv4 address of the Internet gateway.

RELATED DOCUMENTATION

About the Sites Page | **53**

About the Site Groups Page | **196**

Adding an On-Premise Spoke Site with SD-WAN Capability

An on-premise spoke represents an endpoint that is part of customer premise equipment (CPE) at some physical location such as branch office or point of sale location. Typically, these points are connected using overlay connections to hub sites. You add an on-premise spoke site from the **Sites** page. The following device templates are supported for on-premise spoke sites:

- NFX150 as SD-WAN CPE
- NFX250 as SD-WAN CPE
- Dual NFX250 as SD-WAN CPEs
- SRX as SD-WAN CPE
- Dual SRX as SD-WAN CPEs

- SRX4x00 as SD-WAN CPE
- Dual SRX4x00 as SD-WAN CPEs

To add an on-premise spoke site with only SD-WAN capability:

1. Select **Resources > Site Management**.

   The Sites page appears.

2. Click **Add** and select **Add On-Premise Spoke (Manual)**.

   The **Add On-Premise Spoke Site for *Tenant-Name*** page appears.

3. Complete the configuration settings according to the guidelines provided in Table 27 on page 101.

   NOTE: Fields marked with an asterisk (*) are mandatory.

4. (Optional) You can review the configuration in the **Summary** tab and modify the settings, if required.

5. Click **OK**.

   You are returned to the Sites page and a message indicating that the site creation job was triggered is displayed. You can click the job ID link to view the progress of the job. After the job is completed successfully, a confirmation message is displayed and the site that you added is displayed on the Sites page.

Table 27: Fields on the Add Site for Tenant-Name Page With only SD-WAN Capability

Field	Description
General	
Site Information	
Site Name	Enter a unique name for the site. You can use alphanumeric characters and hyphen (-); the maximum length is 10 characters.
Site Group	Select a site group to which you want to assign the site.
Site Capabilities	
WAN Capabilities	Select **SD-WAN** to include SD-WAN capability in the on-premise spoke site.

Table 27: Fields on the Add Site for Tenant-Name Page With only SD-WAN Capability *(continued)*

Field	Description
LAN Capabilities	You need not select this option because you are creating an on-premise spoke site with only SD-WAN capability.
Configuration	
Primary Provider Hub	Select the primary hub site to which this spoke site must connect.
Secondary Provider Hub	Select the secondary hub site to which this site must connect. This site connects to the secondary data hub site when the primary data hub is not reachable.
Primary Enterprise Hub	Select the enterprise hub with which you want to connect the spoke site. If you specify an enterprise hub, then the initial site-to-site traffic as well as the central breakout (backhaul) traffic (if applicable) is sent through the enterprise hub instead of the hub site.
Secondary Enterprise Hub	Select the secondary enterprise hub for this spoke site. The spoke site connects with secondary enterprise hub when the primary enterprise hub is not reachable.
On-Demand Mesh Threshold	
Threshold for Tunnel Creation	Enter the maximum number of sessions closed between the connected sites in a duration of two minutes at which full mesh is created between the two sites. The default value is 5. For example, if you specify the number of sessions as 5, dynamic mesh tunnels are created if the number of sessions closed between two spoke sites in 2 minutes exceeds 5.
Threshold for Tunnel Deletion	Enter the number of sessions closed between the connected sites in a duration of 15 minutes below which full mesh is deleted between the two sites. The default value is 8. For example, if you specify the number of sessions closed as 8, dynamic mesh tunnels are deleted if the number of sessions closed is lesser than or equal to 8.
Address and Contact Information	
Street Address	Enter the street address of the site.
City	Enter the city where the site is located.

Table 27: Fields on the Add Site for Tenant-Name Page With only SD-WAN Capability *(continued)*

Field	Description
State/Province	Select the state or province where the site is located.
ZIP/Postal Code	Enter the postal code for the site.
Country	Select the country where the site is located. Click the **Validate** button to verify the address. • The **site address verification successful** message is displayed if the address is verified. You can click the **View location on a map** link to see the address location. • If the address cannot be verified, the **Site address could not be validated** message is displayed .
Contact Name	Enter the name of the contact person for the site.
Email	Enter the e-mail address of the contact person for the site.
Phone	Enter the phone number of the contact person for the site.
Advanced Configuration	
Name Server IP List	Specify one or more IPv4 addresses of the DNS server. To enter more than one DNS server address, type the address, press Enter, and then type the next address, and so on. DNS servers are used to resolve hostnames into IP addresses.
NTP Server	Specify the fully qualified domain names (FQDNs) or IP addresses of one or more NTP servers. Example: ntp.example.net The site must have DNS reachability to resolve the FQDN during site configuration.
Select Timezone	Select the time zone of the site.
WAN	
Device Template	

Table 27: Fields on the Add Site for Tenant-Name Page With only SD-WAN Capability *(continued)*

Field	Description
Device Series	Select the device series to which the CPE belongs—SRX, NFX150, or NFX250. Based on the device series that you select, the supported device templates (containing information for configuring devices) are listed. Select a device template for the selected device series.
Device Information **NOTE:** Some fields in this section are displayed only if you select a dual CPE device template.	
Device Model	For NFX150 devices, select the device model number.
Serial Number	For a single CPE device, enter the serial number of the CPE device. Serial numbers are case-sensitive.
Device Redundancy	For dual CPE device templates, displays Enabled indicating that redundancy is enabled. You cannot modify this field.
Primary Serial Number	For a dual CPE device, enter the serial number of the primary CPE device. The serial number is case sensitive.
Secondary Serial Number	For a dual CPE device, enter the serial number of the secondary CPE device. The serial number is case sensitive.
Auto Activate	Click the toggle button to enable or disable automatic activation of the CPE device. When you enable this field, zero-touch provisioning (ZTP) of the CPE device is automatically triggered after the site is added to CSO. The device template that you select determines whether this option is enabled or disabled by default.
Activation Code	If the automatic activation of the device is disabled, enter the activation code to manually activate the device. The activation code is provided by the administrator who adds the site. For more information about manually activating a device, see "Activate a Device" on page 229
Primary Activation Code	For a dual CPE device, if the automatic activation of the device is disabled, enter the activation code to manually activate the primary CPE device.
Secondary Activation Code	For a dual CPE device, if the automatic activation of the device is disabled, enter the activation code to manually activate the secondary CPE device.

Table 27: Fields on the Add Site for Tenant-Name Page With only SD-WAN Capability *(continued)*

Field	Description
Boot Image	Select the boot image from the drop-down list if you want to upgrade the image for the CPE device. The boot image is the latest build image uploaded to the image management system. The boot image is used to upgrade the device when the CSO starts the ZTP process. If the boot image is not provided, then the device skips the procedure to upgrade the device image. The boot image (NFX or SRX) is populated based on the device template that you have selected while creating a site. See *Uploading a Device Image*.
WAN Links	
WAN_0 *WAN-Interface-Name*	This field is enabled by default. Enter parameters related to WAN_0. Fields marked with an asterisk (*) must be configured to proceed.
Link Type	Select whether the link would be an MPLS link or Internet link.
Access Type (NFX150 and NFX250)	If you select Internet as the link type, select the access type for the underlay link—Ethernet, LTE, ADSL, or VDSL. You can select the LTE, ADSL, or VDSL access type only for one WAN link. **NOTE:** • You cannot configure an access type (LTE, ADSL, and VDSL) if you are using the Dual SRX and Dual NFX device templates. By default, Ethernet is configured as the access type for the underlay link. • Currently, CSO supports only LTE and VDSL access types on SRX 300 line devices.
Egress Bandwidth	Enter the maximum bandwidth, in Mbps, allowed on the WAN link. Range: 1 through 10,000.
Address Assignment	Select the method of assigning an IP address to the WAN link—DHCP or STATIC. If you select STATIC, you must provide the IP address prefix and the gateway address for the WAN link.
Static IP Prefix	If you configured the address assignment method as STATIC, enter the IP address prefix of the WAN link.
Gateway IP	If you configured the address assignment method as STATIC, enter the IP address of the gateway of the WAN service provider.

Table 27: Fields on the Add Site for Tenant-Name Page With only SD-WAN Capability *(continued)*

Field	Description
WAN Link (Primary or Secondary)	For dual CPE device templates, displays whether the WAN link is a primary link or a secondary link. You cannot modify this field.
Advanced Settings	
Provider	Enter the name of the service provider (SP) providing the WAN service. Only alphanumeric characters and '_', '@', '.', '/', '#', '&', '+' and '-' are allowed. The maximum number of characters allowed in 15.
Cost/Month	Enter the cost for using the WAN link per month and select the currency in which the cost is indicated from the adjacent drop-down list. Range: 1 through 10,000. In bandwidth-optimized SD-WAN, CSO uses this information to identify the least-expensive link to route traffic when multiple WAN links meet SLA profile parameters.
Enable Local Breakout	Click the toggle button to enable local breakout on the WAN link. By default, local breakout is disabled. **NOTE:** • If you enable this option, the WAN link can be used for local breakout. The decision of whether traffic breaks out locally from the site depends on the breakout profile that is referenced in the SD-WAN policy intent. • If you do not enable local breakout on at least one WAN link for a single CPE connection plan and at least two WAN links for a dual CPE connection plan, then local breakout is disabled for the site.
Breakout Options	Select whether you want to use the WAN link for both breakout and WAN traffic (default) or only for breakout traffic.

Table 27: Fields on the Add Site for Tenant-Name Page With only SD-WAN Capability *(continued)*

Field	Description

Autocreate Source NAT Rule	Click the toggle button to enable or disable the automatic creation of source NAT rules. By default, this field is enabled when local breakout is enabled on the WAN link.
	
	Table 28 on page 113 explains how source NAT rules are automatically created on the WAN link. The automatically-created source NAT rules are implicitly defined and applied to the site and is not visible on the NAT Policies page.
	
	NOTE: You can manually override automatically created NAT rules, by creating a NAT rule within a particular rule-set. For example, to use a source NAT pool instead of an interface for translation, create a NAT rule within this particular rule-set, that includes the relevant department zone and WAN interface as the source and destination. For example:
	
	`Dept-Zone1 --> W1 : Translation=Pool-2`
	
	The manually created NAT rule is placed at a higher priority than the corresponding automatically created NAT rule.
	
	You can also add other fields (such as addresses, ports, protocols, and so on) as part of the source or destination endpoints. For example:
	
	`Dept-Zone1, Port 56578 --> W1: Translation=Pool-2`

Translation	Select the type of NAT to use for the traffic on the WAN link: • **Interface**—Use interface-based NAT, which is the default. • **Pool**—Use pool-based NAT. If you select this option, you must specify the IP addresses that are to be used for the NAT pool. **NOTE:** No NAT is performed for tenant-owned public IP addresses.
IP Addresses	For pool-based NAT, enter one or more IP addresses, subnets, or an IP address range. Separate multiple IP addresses by using commas and use a hyphen to denote a range; for example, 192.0.2.1-192.0.2.50.
Preferred Breakout Link	Click the toggle button to enable the WAN link as the most preferred breakout link. If you disable this option, then the breakout link is chosen using ECMP from the available breakout links.

Table 27: Fields on the Add Site for Tenant-Name Page With only SD-WAN Capability *(continued)*

Field	Description
BGP Underlay Options	**NOTE:** This setting can be configured only if the address assignment is static and local breakout is enabled. Click the toggle button to enable BGP underlay routing. When you enable BGP underlay routing, route advertisements to the primary PE node and, if configured, the secondary PE node occur as follows: • CSO advertises the WAN interface subnet. • If you configured pool-based translation, CSO advertises the NAT address pool. **NOTE:** If underlay BGP is enabled for a WAN link, then the routes learnt from BGP are installed for local breakout; CSO does not generate the static default route.
Primary Neighbor	Displays the IP address that you entered for the gateway for the WAN link.
Secondary Neighbor	If you want to provide PE resiliency, you can configure a secondary PE node. Enter the IP address of the secondary PE node. **NOTE:** If the primary PE node goes down, then the secondary PE is used as the next hop. When the primary PE comes back up, the route next hops are changed to the primary PE.
eBGP Peer-AS-Number	Enter the autonomous system (AS) number for the external (EBGP) peer. **NOTE:** If the peer AS number is not configured or the peer AS number that is configured is the same as that of the CPE site, then the BGP type is assumed to be internal BGP (IBGP).
Authentication	Select the BGP route authentication method to be used: • **None**—Indicates that no authentication should be used. This is the default. • **Use MD5**—Indicates that MD5 is to be used for authentication. If you choose this option, you must specify an authentication key.
Auth Key	If you specified that MD5 should be used for authentication, specify an MD5 authentication key (password), which is used to verify the authenticity of BGP packets.

Table 27: Fields on the Add Site for Tenant-Name Page With only SD-WAN Capability *(continued)*

Field	Description
Advertise Public LAN Prefixes	Click the toggle button to enable the advertisement of public LAN prefixes. This field is disabled by default. If the tenant has a public IP address pool configured and you enable the advertisement of public LAN prefixes, then for LAN segments that are created with a subnet that falls under the tenant public IP address pool, CSO advertises the LAN subnet to the BGP underlay. **NOTE:** When public LAN advertisement is enabled for the WAN link, public LAN prefixes are advertised through the BGP underlay towards MPLS or the Internet. If a site has two versions of the route installed for the same LAN prefix in the overlay and underlay, the overlay routes are always preferred over underlay.
Use For Fullmesh	Click the toggle button to specify whether the WAN link can be a part of a fullmesh topology. A site can have a maximum of three links enabled for meshing.
Mesh Overlay Link Type	When **Use for Fullmesh** field is enabled, select the type of mesh overlay link—GRE and GRE_IPSEC. If the link type is Internet, by default, the value for mesh overlay link type is GRE_IPSEC. If the link type is MPLS, select one of the following options: • GRE-IPSEC • GRE
Mesh Tag	When the **Use for Fullmesh** field is enabled, enter the tag to be associated with the WAN link for creating tunnels. You can assign only one tag to the link. Matching mesh tags is one of the criteria used to form tunnels between sites that support meshing. • For an on-premise spoke site, you can select one mesh tag. • For a enterprise hub you can select one or more mesh tags. For more information about mesh tags, see "Mesh Tags Overview" on page 217.
Connects to Hubs	Click the toggle button to specify that the WAN link of the site connects to a hub. **NOTE:** • For sites with a single CPE, you must enable at least one WAN link to connect to the hub so that OAM traffic can be transmitted. • For sites with a dual CPE, you must enable at least one WAN link per device to connect to the hub so that OAM traffic can be transmitted.

Table 27: Fields on the Add Site for Tenant-Name Page With only SD-WAN Capability *(continued)*

Field	Description
Use for OAM Traffic	If you have specified that the WAN link is connected to a hub, click the toggle button to enable sending the OAM traffic over the WAN link. This WAN link is then used to establish the OAM tunnel.
Overlay Tunnel Type	This field is displayed when the **Connects to Hubs** field is enabled and only a one provider hub (primary) is specified. Select the mesh overlay tunnel type (GRE and GRE_IPSEC) of the tunnel to the hub. MPLS links can have both GRE and GRE_IPSEC as the overlay link type where as Internet links can have only GRE_IPSEC as the overlay link type.
Overlay Peer Device	This field is displayed when the **Connects to Hubs** field is enabled and only a one provider hub (primary) is specified. Displays the peer hub device to which the site is connected.
Overlay Peer Interface	This field is displayed when the **Connects to Hubs** field is enabled and only a one provider hub (primary) is specified. Select the interface name of the hub device to which the WAN link of the site is connected.
Overlay Tunnel Type 1	This field is displayed when the **Connects to Hubs** field is enabled and both primary and secondary hubs are specified. Select the mesh overlay tunnel type (GRE and GRE_IPSEC) for the tunnel to the primary hub. MPLS links can have both GRE and GRE_IPSEC as the overlay link type where as Internet links can have only GRE_IPSEC as the overlay link type.
Overlay Peer Device 1	This field is displayed when the **Connects to Hubs** field is enabled and both primary and secondary hubs are specified. Displays the primary peer hub device to which the site is connected.
Overlay Peer Interface 1	This field is displayed when the **Connects to Hubs** field is enabled and both primary and secondary hubs are specified. Select the interface name of the primary hub device to which the WAN link of the site is connected.

Table 27: Fields on the Add Site for Tenant-Name Page With only SD-WAN Capability *(continued)*

Field	Description
Overlay Tunnel Type 2	This field is displayed when the **Connects to Hubs** field is enabled and both primary and secondary hubs are specified. Select the mesh overlay tunnel type (GRE and GRE_IPSEC) for the tunnel to the secondary hub. MPLS links can have both GRE and GRE_IPSEC as the overlay link type where as Internet links can have only GRE_IPSEC as the overlay link type.
Overlay Peer Device 2	This field is displayed when the **Connects to Hubs** field is enabled and both primary and secondary hubs are specified. Displays the secondary peer hub device to which the site is connected.
Overlay Peer Interface 2	This field is displayed when the **Connects to Hubs** field is enabled and both primary and secondary hubs are specified. Select the interface name of the secondary hub device to which the WAN link of the site is connected.
Backup Link	Select a backup link through which traffic can be routed when the primary (other) links are unavailable. You can select any link other than the default links or links that are configured exclusively for local breakout traffic. When a primary link comes back online, CSO monitors the performance on the primary link and when the primary link meets the SLA requirements, the traffic is switched back to the primary link. However, SLA data is not monitored for the backup link.
Default Link	Select one or more links that will be used for routing traffic in the absence of matching SD-WAN policy intents. A site can have multiple default links to the hub site. Default links are used primarily for overlay traffic but can also be used for local breakout traffic. However, a default link cannot be used exclusively for local breakout traffic. If you do not specify a default link, then equal-cost multipath (ECMP) is used to choose the link on which to route traffic.

Table 27: Fields on the Add Site for Tenant-Name Page With only SD-WAN Capability *(continued)*

Field	Description
Data VLAN ID	Enter a VLAN ID for the WAN link. Range: 2 through 4093. **NOTE:** • If you are configuring more than one WAN link on the same physical interface, only one WAN link can be untagged; for the remaining WAN links, you must configure a VLAN ID. • A combination of tagged and untagged on the same physical interface is supported only for single CPE devices.
WAN_1 *WAN-Interface-Name*	Click the toggle button to enable or disable the WAN link. When you enable the WAN link, fields related to the WAN link appear. Fields marked with an asterisk (*) must be configured to proceed. Refer to the fields described for *WAN_0 WAN-Interface-Name* for an explanation of the fields
WAN_2 *WAN-Interface-Name*	Click the toggle button to enable or disable the WAN link. When you enable the WAN link, fields related to the WAN link appear. Fields marked with an asterisk (*) must be configured to proceed. Refer to the fields described for *WAN_0 WAN-Interface-Name* for an explanation of the fields
WAN_3 *WAN-Interface-Name*	Click the toggle button to enable or disable the WAN link. When you enable the WAN link, fields related to the WAN link appear. Fields marked with an asterisk (*) must be configured to proceed. Refer to the fields described for *WAN_0 WAN-Interface-Name* for an explanation of the fields
Management Connectivity	
IP Prefix	Enter an IPv4 address prefix for the loopback interface on the CPE device. The IP address prefix must be a /32 IP address prefix and must be unique across the entire management network. If you do not specify an IPv4 address prefix, CSO automatically assigns the IP prefix from the reserved pool 100.124.0.0/14. • For NFX150 and NFX250 devices, if the USE_SINGLE_SSH_TO_NFX parameter is disabled in the device template, then enter the IP address prefix as /29 or lower based on the number of VNFs. • For all other devices, enter the IP address prefix as /32.
LAN	

Table 27: Fields on the Add Site for Tenant-Name Page With only SD-WAN Capability *(continued)*

Field	Description
Add LAN Segment	You must add at least one LAN segment for the on-premise site. To add a LAN segment: 1. Click the **+** icon. The Add LAN Segment page appears. 2. Complete the configuration settings according to the guidelines provided inTable 29 on page 114. 3. Click **Save**. The LAN segment is added and you are returned to the Add Site for *Tenant-Name* page.

Table 28: Automatic Creation of Source NAT Rules

Autocreate Source NAT Rule	Translation	NAT Rules Creation
Disabled	Not applicable (No NAT)	None.

Enabled	Interface-Based (Default)—CSO creates interface-based NAT rules.	Source NAT rules are automatically created, with each rule from a department zone to the WAN interface, with a translation of type interface. Each pair of [zone - interface] represents a rule-set.
		For example, the following department zone to (WAN link) W1 interface rule-set might be created:
		`Dept-Zone1 --> W1: Translation=Interface`
		`Dept-Zone2 --> W1: Translation=Interface`
		`Dept-Zone3 --> W1: Translation=Interface`

Enabled	Pool-Based—CSO automatically creates pool-based NAT rules	NAT source rules are automatically created, with each rule from a department zone to the WAN NAT pool with a translation of type pool.
		For example, a source NAT rule from department zone to NAT pool might be created:
		`Dept-Zone1 --> W1 : Translation=Pool-1`
		`Dept-Zone2 --> W1 : Translation=Pool-1`

Table 29: Fields on the Add LAN Segment page

Field	Description
Name	Enter a name for the LAN segment. The name for a LAN segment should be a unique string of alphanumeric characters and some special characters (. -). No spaces are allowed and the maximum length is 15 characters.
Type **NOTE:** This field is displayed only for LAN segments associated with enterprise hub sites.	Select the type of LAN segment: • Directly Connected (default)—Indicates that the LAN segment is directly connected to the site. • Dynamic Routed—Indicates that the LAN segment is not directly connected to the site and is reachable by using a dynamic route. If you select this option, you must specify the dynamic routing information.
VLAN ID	Enter the VLAN ID for the LAN segment. Range: 2 through 4093.
Department	Select a department to which the LAN segment is assigned. Alternatively, click the **Create Department** link to create a new department and assign the LAN segment to it. See "Adding a Department" on page 767 for details. You can group LAN segments as departments for ease of management and for applying policies at the department-level. For LAN segments that are dynamically routed, you can assign only a data center department.
Protocol	For dynamically routed LAN segments, select the routing protocol (BGP or OSPF) to be used by the data center department to learn routes from the data center.
Advertise LAN Prefix	For dynamically routed LAN segments, click the toggle button to advertise the LAN prefix of the SD-WAN spoke site to the data center through the data center department associated with the enterprise hub. By default, the advertise LAN prefix field is disabled. **NOTE:** You must avoid overlapping IP addresses between the SD-WAN LAN network and the datacenter network.
Gateway Address/Mask	Enter a valid gateway IP address and mask for the LAN segment. This address will be the default gateway for endpoints in this LAN segment. For example: 192.0.2.8/24.

Table 29: Fields on the Add LAN Segment page *(continued)*

Field	Description
DHCP	For directly connected LAN segments, click the toggle button to enable DHCP (default). You can enable DHCP if you want to assign IP addresses by using a DHCP server or disable DHCP if you want to assign a static IP address to the LAN segment. **NOTE:** If you enable DHCP, additional fields appear on the page.
Additional fields related to DHCP	
Address Range Low	Enter the starting IP address in the range of IP addresses that can be allocated by the DHCP server to the LAN segment.
Address Range High	Enter the ending IP address in the range of IP addresses that can be allocated by the DHCP server to the LAN segment.
Maximum Lease Time	Specify the maximum duration (in seconds) for which a client can request for and hold a lease on the DHCP server. Default: 1440 Range: 0 through 4,294,967,295 seconds.
Name Server	Specify one or more IPv4 addresses of the DNS server. To enter more than one DNS server address, type the address, press Enter, and then type the next address. **NOTE:** DNS servers are used to resolve hostnames into IP addresses.

Table 29: Fields on the Add LAN Segment page *(continued)*

Field	Description

CPE Ports	• For sites with LAN capability, click the toggle button to include or exclude the CPE in the LAN segment.
	  ◦ When you include the CPE in the LAN segment:
	  ▪ CPE ports that you can include in the LAN segment are listed. Select the ports from the **Available** column and click the right-arrow to move the ports to the **Selected** column.
	  ▪ The **Switch Ports** field is disabled. CSO automatically assigns LAN ports on the Switch device and creates the same LAN segment on the Switch.
	  ◦ If you click to exclude the CPE from the LAN segment, you must specify the switch ports that connect with the LAN in the **Switch Ports** field. CSO automatically assigns LAN ports on the CPE device and creates the same LAN segment on the CPE device.
	  NOTE: You can select only one port if the CPE is a physical SRX Series device.
	• For sites without LAN capability, the CPE Ports field is disabled and the CPE ports that you can include in the LAN segment are listed.
	  Select the ports from the **Available** column and click the right-arrow to move the ports to the **Selected** column.

Switch Ports **NOTE:** This field is displayed only when LAN capability is selected for the enterprise hub.	If you disable the CPE ports field, select ports on the switch to be part of the LAN segment. The Switch ports and CPE ports are mutually exclusive. Select the ports from the **Available** column and click the right-arrow to move the ports to the **Selected** column.
BGP Configuration **NOTE:** This section is displayed only for dynamic routed LAN segments with BGP specified as the protocol.	
Authentication	Select the BGP route authentication method to be used: • None—Indicates that no authentication should be used. This is the default. • Use MD5—Indicates that MD5 is to be used for authentication. If you choose this option, you must specify an authentication key.
Peer IP Address	Enter the IP address of the BGP neighbor.
Peer AS Number	Enter the autonomous system (AS) number BGP neighbor.
Auth Key	If you specified that MD5 should be used for authentication, specify an MD5 authentication key (password), which is used to verify the authenticity of BGP packets.

Table 29: Fields on the Add LAN Segment page *(continued)*

Field	Description
OSPF Configuration **NOTE:** This section is displayed only for dynamic routed LAN segments with OSPFspecified as the protocol.	
OSPF Area ID	Specify the OSPF area identifier to be used for the dynamic route.
Authentication	Select the OSPF route authentication method to be used: • Password—Indicates that password-based authentication should be used. If you choose this option, you must specify the password. (This is the default). • Use MD5—Indicates that MD5 is to be used for authentication. If you choose this option, you must specify an authentication key. • None—Indicates that no authentication should be used.
Password	Enter the password to be used to verify the authenticity of OSPF packets.
Confirm Password	Retype the password for confirmation purposes.
MD5 Auth Key ID	If you specified that MD5 should be used for authentication, enter the OSPF MD5 authentication key ID. Range: 1 through 255.
Auth Key	If you specified that MD5 should be used for authentication, enter an MD5 authentication key, which is used to verify the authenticity of OSPF packets.

RELATED DOCUMENTATION

About the Sites Page | 53

Adding an On-Premise Spoke Site with SD-WAN and LAN Capabilities

You can add an on-premise spoke site in CSO by provisioning a CPE and a switch behind the CPE to provide SD-WAN and LAN capabilities to the site. See "Switch Behind a CPE or Next Generation Firewall Overview" on page 58 for details.

To add a site with SD-WAN and LAN capabilities:

1. Select **Resources > Site Management**.

   The Sites page appears.

2. Click **Add** and select **Add On-Premise Spoke (Manual)**.

   The Add Site for *Tenant-Name* page appears.

3. Complete the configuration according to the guidelines provided in Table 30 on page 119.

   NOTE: Fields marked with an asterisk (*) are mandatory.

4. Review the configuration from the **Summary** tab.

   (Optional) click the Edit links within the summary to go directly to a specific page of the wizard and modify the configured settings.

5. Click **OK** to add the site.

   The site activation job is initiated and the Site Activation: *Site-Name* page appears displaying the progress of the steps executed for activating the CPE and the switch. The CPE is activated first and then the process to activate the switch is initiated.

6. - If the Zero Touch Provisioning (ZTP) toggle button is enabled (default), CSO pushes the stage-1 and stage-2 configurations and provisions the switch.

     This process occurs immediately after the activation process, for which you entered the activation code or selected auto-activation.

     NOTE: Stage-1 configuration is the initial configuration that allows basic connectivity to a device, which is pushed to the device.

     The configuration that is pushed to the device after it has connected to CSO is called stage-2 configuration.

   - If you disabled the Zero Touch Provisioning (ZTP) toggle button, you must manually configure the stage-1 configuration (as provided by CSO) on the switch.

To manually configure the stage-1 configuration:

a. On the **Site Activation: *Site-Name*** page, the **Click to copy stage-1 configuration** link appears after the Prestage Device step completes successfully.

b. Click the **Click to copy stage-1 configuration** link.

   The stage-1 configuration page appears displaying the stage-1 configuration to be copied to the EX Series device.

c. Copy the stage-1 configuration and log in to the console of the EX Series switch.

d. Enter the configuration mode, paste, and commit the configuration.

   After the stage-1 configuration is committed, the switch has the outbound SSH configuration to connect with CSO.

   CSO then provisions the switch.

Table 30: Fields on the Add Site for Tenant-Name Page(SD-WAN and LAN Capabilities)

Field	Description
General	
Site Information	
Site Name	Enter a unique name for the site. You can use alphanumeric characters and hyphen (-); the maximum length is 10 characters.
Site Group	Select a site group to which you want to assign the site.
Site Capabilities	
WAN Capabilities	Select **SD-WAN** to include SD-WAN capabilities in the spoke site.
LAN Capabilities	Select **LAN** to include SD-LAN capability in the spoke site.
Configuration	
Primary Provider Hub	Select the hub site (or primary hub site in case of multihoming) to which the spoke site must connect.

Table 30: Fields on the Add Site for Tenant-Name Page(SD-WAN and LAN Capabilities) *(continued)*

Field	Description
Secondary Provider Hub	Select the secondary hub site to which this site must connect. This site connects to the secondary data hub site when the primary data hub is down.
Primary Enterprise Hub	Select the primary enterprise hub with which you want to connect the spoke site. If you specify a enterprise hub, then the initial site-to-site traffic as well as the central breakout (backhaul) traffic (if applicable) is sent through the enterprise hub instead of the hub site.
Secondary Enterprise Hub	Select the secondary enterprise hub for this spoke site. The spoke site connects with secondary enterprise hub when the primary enterprise hub is down.
On-Demand Mesh Threshold	
Threshold for Tunnel Creation	Specify the threshold for the number of sessions (flows) closed (in a two-minute duration) between the on-premise spoke site and a destination site. When the number of sessions closed exceeds the specified threshold, a tunnel is created between the on-premise spoke site and the destination site. The default value is 5. For example, if you specify the number of sessions as 5, dynamic mesh tunnels are created if the number of sessions closed between two spoke sites in 2 minutes exceeds 5.

Table 30: Fields on the Add Site for Tenant-Name Page(SD-WAN and LAN Capabilities) *(continued)*

Field	Description
Threshold for Tunnel Deletion	Specify the threshold for the number of sessions closed (in a 15-minute duration) between the on-premise spoke site and a destination site. When the number of sessions closed is lower than the specified threshold, the tunnel between the on-premise spoke site and destination site is deleted. The default value is 2. For example, if you specify the number of sessions closed as 2, dynamic mesh tunnels are deleted if the number of sessions closed is lesser than or equal to 2.
Address and Contact Information	
Street Address	Enter the street address of the site.
City	Enter the city where the site is located.
State/Province	Select the state or province where the site is located.
ZIP/Postal Code	Enter the postal code for the site.
Country	Select the country where the site is located. Click the **Validate** button to verify the address. • The **site address verification successful** message is displayed if the address is verified. You can click the **View location on a map** link to see the address location. • If the address cannot be verified, the **Site address could not be validated** message is displayed .
Contact Name	Enter the name of the contact person for the site.
Email	Enter the e-mail address of the contact person for the site.
Phone	Enter the phone number of the contact person for the site.
Advanced Configuration	

Table 30: Fields on the Add Site for Tenant-Name Page(SD-WAN and LAN Capabilities) *(continued)*

Field	Description
Name Server IP List	Specify one or more IPv4 addresses of the DNS server. To enter more than one DNS server address, type the address, press Enter, and then type the next address, and so on. DNS servers are used to resolve hostnames into IP addresses.
NTP Server	Specify the fully qualified domain names (FQDNs) or IP addresses of one or more NTP servers. Example: ntp.example.net The site must have DNS reachability to resolve the FQDN during site configuration.
Select Timezone	Select the time zone of the site.
WAN	
Device Template	
Device Series	Select the device series to which the CPE belongs—SRX. NFX250. Based on the device series that you select, the supported device templates (containing information for configuring devices) are listed. Select a device template for the selected device series.
Device Information	
Serial Number	Enter the serial number of the CPE device.
Auto Activate	Click the toggle button to enable or disable automatic activation of the CPE device. When you enable this field, zero-touch provisioning (ZTP) of the CPE device is automatically triggered after the site is added to CSO. The device template that you select determines whether this option is enabled or disabled by default.

Table 30: Fields on the Add Site for Tenant-Name Page(SD-WAN and LAN Capabilities) *(continued)*

Field	Description
Activation Code	If you disable the Auto Activate field, enter the activation code for the CPE or firewall device. For information about activating a CPE or firewall device, see "Activate a Device" on page 229.
Boot Image	Select the boot image from the drop-down list if you want to upgrade the image for the CPE device. The boot image is the latest build image uploaded to the image management system. The boot image is used to upgrade the device when the CSO starts the ZTP process. If the boot image is not provided, then the device skips the procedure to upgrade the device image. The boot image is populated based on the device template that you have selected while creating a site. See *Uploading a Device Image*.
WAN Links	
WAN_0 WAN-Interface-Name	This field is enabled by default. Enter parameters related to WAN_0. Fields marked with an asterisk (*) must be configured to proceed.
Link Type	Select whether the link would be an MPLS link or Internet link.
Access Type (NFX150 and NFX250)	If you select Internet as the link type, select the access type for the underlay link—Ethernet, LTE, ADSL, or VDSL. You can select the LTE, ADSL, or VDSL access type only for one WAN link. **NOTE:** • You cannot configure an access type (LTE, ADSL, and VDSL) if you are using the Dual SRX and Dual NFX device templates. By default, Ethernet is configured as the access type for the underlay link. • Currently, CSO supports only LTE and VDSL access types on SRX 300 line devices.

Table 30: Fields on the Add Site for Tenant-Name Page(SD-WAN and LAN Capabilities) *(continued)*

Field	Description
Egress Bandwidth	Enter the maximum bandwidth (in Mbps) that the CPE or firewall allows over the WAN link. Range: 1 through 10,000.
Address Assignment	Select the method of assigning an IP address to the WAN link—DHCP or STATIC. • If you select DHCP, the IP address is provided by using the DHCP server of the service provider of the WAN link. • If you select STATIC, you must provide the IP address prefix and the gateway address for the WAN link.
Static IP Prefix	If you configured the address assignment method as STATIC, enter the IP address prefix of the WAN link.
Gateway IP Address	If you configured the address assignment method as STATIC, enter the IP address of the gateway of the WAN service provider.
Advanced Settings	
Provider	Enter the name of the service provider who is responsible for providing the WAN link.
Cost/Month	Enter the cost per month (in a specified currency) for the WAN link. Specify the currency from the adjacent drop-down list. Range: 1 through 10,000. In bandwidth-optimized SD-WAN, CSO uses this information to identify the least-expensive link to route traffic when multiple WAN links meet SLA profile parameters.

Table 30: Fields on the Add Site for Tenant-Name Page(SD-WAN and LAN Capabilities) *(continued)*

Field	Description
Enable Local Breakout	Click the toggle button to enable local breakout on the WAN link. By default, local breakout is disabled. **NOTE:** • If you enable this option, the WAN link can be used for local breakout. The decision of whether traffic breaks out locally from the site depends on the breakout profile that is referenced in the SD-WAN policy intent. • If you do not enable local breakout on at least one WAN link for a single CPE connection plan and at least two WAN links for a dual CPE connection plan, then local breakout is disabled for the site.
Breakout Options	Select whether you want to use the WAN link for both breakout and WAN traffic (default) or only for breakout traffic.

Table 30: Fields on the Add Site for Tenant-Name Page(SD-WAN and LAN Capabilities) *(continued)*

Field	Description

Autocreate Source NAT Rule	Click the toggle button to enable or disable the automatic creation of source NAT rules. By default, this field is enabled when local breakout is enabled on the WAN link.
	
	Table 31 on page 134 explains how source NAT rules are automatically created on the WAN link. The automatically-created source NAT rules are implicitly defined and applied to the site and is not visible on the NAT Policies page.
	
	NOTE: You can manually override automatically created NAT rules, by creating a NAT rule within a particular rule-set. For example, to use a source NAT pool instead of an interface for translation, create a NAT rule within this particular rule-set, that includes the relevant department zone and WAN interface as the source and destination. For example:
	
	`Dept-Zone1 --> W1 : Translation=Pool-2`
	
	The manually created NAT rule is placed at a higher priority than the corresponding automatically created NAT rule.
	
	You can also add other fields (such as addresses, ports, protocols, and so on) as part of the source or destination endpoints. For example:
	
	`Dept-Zone1, Port 56578 --> W1: Translation=Pool-2`

Translation	Select the type of NAT to use for the traffic on the WAN link: • **Interface**—Use interface-based NAT, which is the default. • **Pool**—Use pool-based NAT. If you select this option, you must specify the IP addresses that are to be used for the NAT pool. **NOTE:** No NAT is performed for tenant-owned public IP addresses.

Table 30: Fields on the Add Site for Tenant-Name Page(SD-WAN and LAN Capabilities) *(continued)*

Field	Description
IP Addresses	For pool-based NAT, enter one or more IP addresses, subnets, or an IP address range. Separate multiple IP addresses by using commas and use a hyphen to denote a range; for example, 192.0.2.1-192.0.2.50.
Preferred Breakout Link	Click the toggle button to enable the WAN link as the preferred breakout link. If you disable this option, then the breakout link is chosen using ECMP from the available breakout links.
BGP Underlay Options	**NOTE:** This setting can be configured only if the address assignment is static and local breakout is enabled. Click the toggle button to enable BGP underlay routing. When you enable BGP underlay routing, route advertisements to the primary PE node and, if configured, the secondary PE node occur as follows: • CSO advertises the WAN interface subnet. • If you configured pool-based translation, CSO advertises the NAT address pool. **NOTE:** If underlay BGP is enabled for a WAN link, then the routes learnt from BGP are installed for local breakout; CSO does not generate the static default route.
Primary Neighbor	Displays the IP address that you entered for the gateway for the WAN link.
Secondary Neighbor	If you want to provide PE resiliency, you can configure a secondary PE node. Enter the IP address of the secondary PE node. **NOTE:** If the primary PE node goes down, then the secondary PE is used as the next hop. When the primary PE comes back up, the route next hops are changed to the primary PE.

Table 30: Fields on the Add Site for Tenant-Name Page(SD-WAN and LAN Capabilities) *(continued)*

Field	Description
eBGP Peer-AS-Number	Enter the autonomous system (AS) number for the external (EBGP) peer. **NOTE:** If the peer AS number is not configured or the peer AS number that is configured is the same as that of the CPE site, then the BGP type is assumed to be internal BGP (IBGP).
Authentication	Select the BGP route authentication method to be used: • **None**—Indicates that no authentication should be used. This is the default. • **Use MD5**—Indicates that MD5 is to be used for authentication. If you choose this option, you must specify an authentication key.
Auth Key	If you specified that MD5 should be used for authentication, specify an MD5 authentication key (password), which is used to verify the authenticity of BGP packets.
Advertise Public LAN Prefixes	Click the toggle button to enable the advertisement of public LAN prefixes. This field is disabled by default. If the tenant has a public IP address pool configured and you enable the advertisement of public LAN prefixes, then for LAN segments that are created with a subnet that falls under the tenant public IP address pool, CSO advertises the LAN subnet to the BGP underlay. **NOTE:** When public LAN advertisement is enabled for the WAN link, public LAN prefixes are advertised through the BGP underlay towards MPLS or the Internet. If a site has two versions of the route installed for the same LAN prefix in the overlay and underlay, the overlay routes are always preferred over underlay.
Use For Fullmesh	Click the toggle button to specify whether the WAN link can be a part of a full mesh topology. A site can have a maximum of three links enabled for meshing.

Table 30: Fields on the Add Site for Tenant-Name Page(SD-WAN and LAN Capabilities) *(continued)*

Field	Description

Mesh Overlay Link Type	When **Use for Fullmesh** field is enabled, select the type of mesh overlay link—GRE and GRE_IPSEC:
	• If the link type is Internet, the value for mesh overlay link type is GRE_IPSEC.
	• If the link type is MPLS, select one of the following options:
	  ◦ GRE-IPSEC
	  ◦ GRE

Mesh Tag	When the **Use for Fullmesh** field is enabled, enter the tag to be associated with the WAN link for creating tunnels. You can assign only one tag to the link. Matching mesh tags is one of the criteria used to form tunnels between sites that support meshing. For more information about mesh tags, see "Mesh Tags Overview" on page 217.
Connects to Hubs	Click the toggle button to specify that the WAN link of the site connects to a hub. **NOTE:** • For sites with a single CPE, you must enable at least one WAN link to connect to the hub so that OAM traffic can be transmitted. • For sites with a dual CPE, you must enable at least one WAN link per device to connect to the hub so that OAM traffic can be transmitted.
Use for OAM Traffic	If you have specified that the WAN link is connected to a hub, click the toggle button to enable sending the OAM traffic over the WAN link. This WAN link is then used to establish the OAM tunnel.
Overlay Tunnel Type	Select the mesh overlay tunnel type—GRE and GRE_IPSEC. MPLS links can have both GRE and GRE_IPSEC as the overlay link type where as Internet links can have only GRE_IPSEC as the overlay link type.

Table 30: Fields on the Add Site for Tenant-Name Page(SD-WAN and LAN Capabilities) *(continued)*

Field	Description
Overlay Peer Device	Displays the peer hub device to which the site is connected.
Overlay Peer Interface	Select the interface name of the hub device to which the WAN link of the site is connected.
Backup Link	Select a backup link through which traffic can be routed when the primary (other) links are unavailable. You can select any link other than the default links or links that are configured exclusively for local breakout traffic. When a primary link comes back online, CSO monitors the performance on the primary link and when the primary link meets the SLA requirements, the traffic is switched back to the primary link. However, SLA data is not monitored for the backup link.
Default Link	Select one or more links that will be used for routing traffic in the absence of matching SD-WAN policy intents. A site can have multiple default links to the hub site. Default links are used primarily for overlay traffic but can also be used for local breakout traffic. However, a default link cannot be used exclusively for local breakout traffic. If you do not specify a default link, then equal-cost multipath (ECMP) is used to choose the link on which to route traffic.
Data VLAN ID	Enter a VLAN ID for the WAN link. Range: 2 through 4093. **NOTE:** • If you are configuring more than one WAN link on the same physical interface, only one WAN link can be untagged; for the remaining WAN links, you must configure a VLAN ID. • A combination of tagged and untagged on the same physical interface is supported only for single CPE devices.

Table 30: Fields on the Add Site for Tenant-Name Page(SD-WAN and LAN Capabilities) *(continued)*

Field	Description
WAN_1 WAN-Interface-Name	Click the toggle button to enable or disable the WAN link. When you enable the WAN link, fields related to the WAN link appear. Fields marked with an asterisk (*) must be configured to proceed. Refer to the fields described for *WAN_0 WAN-Interface-Name* for an explanation of the fields
WAN_2 WAN-Interface-Name	Click the toggle button to enable or disable the WAN link. When you enable the WAN link, fields related to the WAN link appear. Fields marked with an asterisk (*) must be configured to proceed. Refer to the fields described for *WAN_0 WAN-Interface-Name* for an explanation of the fields
WAN_3 WAN-Interface-Name	Click the toggle button to enable or disable the WAN link. When you enable the WAN link, fields related to the WAN link appear. Fields marked with an asterisk (*) must be configured to proceed. Refer to the fields described for *WAN_0 WAN-Interface-Name* for an explanation of the fields
Management Connectivity	
IP Prefix	Enter an IPv4 address prefix for the loopback interface on the CPE device. The IP address prefix must be a /32 IP address prefix and must be unique across the entire management network. If you do not specify an IPv4 address prefix, CSO automatically assigns the IP prefix from the reserved pool 100.124.0.0/14.
LAN	
Device Profile	
Device Name	Enter a name for the switch. You can use alphanumeric characters and hyphen (-). The maximum length allowed is 15 characters.
Device Type	Select the type of switch—EX2300, EX3400, EX4300, EX4600, and EX4650.

Table 30: Fields on the Add Site for Tenant-Name Page(SD-WAN and LAN Capabilities) *(continued)*

Field	Description
Device Model	Select the model for the switch you specified in the Device Type. The models vary in the number and type of ports the switch contains. For example, If you selected EX3400, select a model such as EX3400-24P, EX3400-48P, EX3400-24T among others.
CPE Settings	
Trunk Ports	Select at least two trunk ports on the CPE device to connect with the switch, which are used for the following: • LAN traffic between the switch and the CPE • Management traffic for in-band management of the switch. **NOTE:** The ae0 LAG interface of the SRX Series devices is used as the trunk port for communication with the switch.
Switch Management Subnet	Specify the subnet that the DHCP server can use to assign IP addresses. The DHCP server runs on the following ports: • Trunk ports to provide DHCP information to all devices connected to the switch and to the in-band management port, switch management port, and LAN ports on the CPE. • Out-of-band management port on the CPE to provide DHCP information to the management port on the switch. • LAN ports on the CPE to provide information to the devices connected to the CPE LAN ports.
Switch Details	
Serial Number	Specify the serial number of the switch.

Table 30: Fields on the Add Site for Tenant-Name Page(SD-WAN and LAN Capabilities) *(continued)*

Field	Description
Auto Activate	Click the toggle button to enable or disable automatic activation of the switch when the switch is detected by CSO (that is, management status of the device is Device_Detected). When you enable this field, zero-touch provisioning (ZTP) of the switch is automatically triggered when the device communicates with CSO. By default, auto activation for the switch is enabled, if it is enabled for the CPE and vice-versa. **NOTE:** You must physically connect the switch to the CPE and power it on for the switch to be automatically activated when you enable this option.
Activation code	When the **Auto activate** field is disabled, enter the activation code to be used for manually activating the switch. For information, see "Manually Activating a Switch" on page 232.
Zero Touch Provisioning	Click the toggle button to enable or disable zero-touch provisioning (ZTP) of the switch through ZTP. If you disable ZTP, you must manually copy and paste the Stage-1 configuration on the switch during site activation. See Step "Step-by-Step Procedure" on page 118 for details. **NOTE:** • Only EX Series devices running 18.4R2.7 firmware support ZTP. • EX4600 and EX4650 switches do not support Phone-Home client. You must disable ZTP and manually configure the stage-1 configuration on the switches.
LAN Segment	Displays the LAN segment that you configure on the switch. To add a LAN segment, click the **+** icon on the top, right corner of the LAN table. The Add LAN Segment page appears. See Table 32 on page 134.

Table 31: Automatic Creation of Source NAT Rules

Autocreate Source NAT Rule	Translation	NAT Rules Creation
Disabled	Not applicable (No NAT)	None.

Enabled	Interface-Based (Default)—CSO creates interface-based NAT rules.	Source NAT rules are automatically created, with each rule from a department zone to the WAN interface, with a translation of type interface. Each pair of [zone - interface] represents a rule-set.
		
		For example, the following department zone to (WAN link) W1 interface rule-set might be created:
		
		`Dept-Zone1 --> W1: Translation=Interface`
		`Dept-Zone2 --> W1: Translation=Interface`
		`Dept-Zone3 --> W1: Translation=Interface`

Enabled	Pool-Based—CSO automatically creates pool-based NAT rules	NAT source rules are automatically created, with each rule from a department zone to the WAN NAT pool with a translation of type pool.
		
		For example, a source NAT rule from department zone to NAT pool might be created:
		
		`Dept-Zone1 --> W1 : Translation=Pool-1`
		`Dept-Zone2 --> W1 : Translation=Pool-1`

Table 32: Fields on the Add LAN Segment Page when Adding a Switch along with CPE

Field	Description
Add LAN Segment	
Name	Enter a name for the LAN segment. The name for a LAN segment should be a unique string of alphanumeric characters and some special characters (. -). No spaces are allowed and the maximum length is 15 characters.
VLAN ID	Enter the VLAN ID for the LAN segment. Range: 2 through 4093.

Table 32: Fields on the Add LAN Segment Page when Adding a Switch along with CPE *(continued)*

Field	Description
Department	Select a department to which the LAN segment is to be assigned. Alternatively, click the **Create Department** link to create a new department and assign the LAN segment to it. See "Adding a Department" on page 767 for details. You group LAN segments as departments for ease of management and for applying policies at the department-level.
Gateway Address/Mask	Enter a valid gateway IP address and mask for the LAN segment; for example, 192.0.2.8/24.
DHCP	For directly connected LAN segments, click the toggle button to enable DHCP. DHCP is disabled by default. You enable DHCP if you want to assign IP addresses by using a DHCP server. You disable DHCP if you want to assign a static IP address to the LAN segment. **NOTE:** If you enable DHCP, fields related to DHCP-related parameters appear and must be configured.
[DHCP-Related Fields]	
Address Range Low	Enter the starting IP address in the range of IP addresses that can be allocated by the DHCP server to the LAN segment.
Address Range High	Enter the ending IP address in the range of IP addresses that can be allocated by the DHCP server to the LAN segment.
Maximum Lease Time	Specify the maximum duration (in seconds) for which a client can request for and hold a lease on a DHCP server. Range: 0 through 4,294,967,295.
Name Server	Specify or select one or more IPv4 addresses of the DNS server. To enter more than one DNS server address, type the address, press Enter, and then type the next address, and so on. DNS servers are used to resolve hostnames into IP addresses.

Table 32: Fields on the Add LAN Segment Page when Adding a Switch along with CPE *(continued)*

Field	Description
CPE Ports	Click the toggle button to include or exclude the CPE in the LAN segment. When you include the CPE in the LAN segment: • CPE ports that you can include in the LAN segment are listed. Select the ports from the **Available** column and click the right-arrow to move the ports to the **Selected** column. • The **Switch Ports** field is disabled. CSO automatically assigns LAN ports on the switch device and creates the same LAN segment on the switch. If you exclude the CPE from the LAN segment, you must specify the switch ports that connect with the LAN in the **Switch Ports** field. CSO automatically assigns LAN ports on the CPE device and creates the same LAN segment on the CPE device. **NOTE:** You can select only one port if the CPE is an SRX Series device.
Switch Ports	If you disable the CPE ports field, select ports on the switch that will be part of the LAN segment. Select the ports from the **Available** column and click the right-arrow to move the ports to the **Selected** column.

RELATED DOCUMENTATION

- Adding and Provisioning Switches to Provide LAN Capability to a Site Overview | 57
- Add an On-Premise Spoke Site with LAN Capability | 136
- Add a Switch to an Existing SD-WAN Site Or Next-Generation Firewall Site | 160
- Deleting a Site | 194

Add an On-Premise Spoke Site with LAN Capability

You can provision and monitor a switch (physical or Virtual Chassis) by adding an on-premise spoke site to CSO. See “Standalone Switch Overview” on page 57 for details.

To add an on-premise spoke site with a switch:

1. Select **Resources > Site Management**.

   The Sites page appears.

2. Click **Add** and select **Add On-Premise Spoke (Manual)**.

   The Add Site for *Tenant-Name* page appears.

3. Complete the configuration according to the guidelines provided in Table 33 on page 138.

   NOTE: Fields marked with an asterisk (*) are mandatory.

4. Review the configuration from the **Summary** tab.

   (Optional) click the Edit links within the summary to go directly to a specific page of the wizard and modify the configured settings.

5. Click **OK** to add the site.

   After you click OK, site activation is initiated and the **Site Activation: *Site-Name*** page appears.

6. - If the Zero Touch Provisioning (ZTP) toggle button is enabled (default), CSO pushes the stage-1 and stage-2 configurations and provisions the switch.

     This process occurs immediately after the activation process, for which you entered the activation code or selected auto-activation.

     NOTE: Stage-1 configuration is the initial configuration that allows basic connectivity to a device, which is pushed to the device.

     The configuration that is pushed to the device after it has connected to CSO is called stage-2 configuration.

   - If you disabled the Zero Touch Provisioning (ZTP) toggle button, you must manually configure the stage-1 configuration (as provided by CSO) on the switch.

     To manually configure the stage-1 configuration:

     a. On the **Site Activation: *Site-Name*** page, the **Click to copy stage-1 configuration** link appears after the Prestage Device step completes successfully.

     b. Click the **Click to copy stage-1 configuration** link.

The stage-1 configuration page appears displaying the stage-1 configuration to be copied to the EX Series device.

c. Copy the stage-1 configuration and log in to the console of the EX Series switch.

d. Enter the configuration mode, paste, and commit the configuration.

After the stage-1 configuration is committed, the switch has the outbound SSH configuration to connect with CSO.

CSO then provisions the switch.

Table 33: Fields on the Add Site for Tenant-Name Page (LAN Capability)

Field	Description
Site Information	
Site Name	Enter a unique name for the site. You can use alphanumeric characters and hyphen (-). The maximum length allowed is 10 characters.
Site Group	Select a site group to which you want to assign the site.
Site Capabilities	
LAN Capabilities	Select **LAN** to include LAN capabilities in the site.
Address and Contact Information	
Street Address	Enter the street address of the site.
City	Enter the city where the site is located.
State/Province	Enter the state or province where the site is located.
ZIP/Postal Code	Enter the postal code for the site.

Table 33: Fields on the Add Site for Tenant-Name Page (LAN Capability) *(continued)*

Field	Description
Country	From the list, select the country where the site is located. Click the **Validate** button to verify the address. • The **site address verification successful** message is displayed if the address is verified. You can click the **View location on a map** link to see the address location. • If the address cannot be verified, the **Site address could not be validated** message is displayed. If you enter the wrong address and click the **Validate** button to verify the address, the **Site address could not be validated message** is displayed.
Contact Name	Enter the name of the contact person for the site.
Email	Enter the e-mail address of the contact person for the site.
Phone	Enter the phone number of the contact person for the site.
Advanced Configuration	
Name Server IP List	Specify one or more IPv4 addresses of the DNS server. To enter more than one DNS server address, type the address, press Enter, and then type the next address, and so on. DNS servers are used to resolve hostnames into IP addresses.
NTP Server	Specify the IP addresses of one or more NTP servers.
Select Timezone	Select the time zone of the site from the list.
Device Profile	
Device Name	Enter a name for the switch. You can use alphanumeric characters and hyphen (-). The maximum length allowed is 15 characters.

Table 33: Fields on the Add Site for Tenant-Name Page (LAN Capability) *(continued)*

Field	Description
Device Type	Select the type of switch—EX2300, EX3400, EX4300, EX4600, and EX4650.
Device Model	Select the model for the switch you specified in the Device Type field. The models vary in the number and type of ports the switch contains. For example, If you selected EX3400, select a model such as EX3400-24P, EX3400-48P, EX3400-24T among others.
Switch Details	
Virtual Chassis	Click the toggle button to enable or disable (default) adding the switch as a Virtual Chassis. If you enable this toggle button, you must select the method of provisioning the Virtual Chassis. • Before you add a Virtual Chassis in CSO, ensure that the Virtual Chassis is setup. See "Step-by-Step Procedure" on page 148 for information about setting up a Virtual Chassis. • In Release 5.1.1, you cannot add a new member or change the roles assigned to the members after you onboard a Virtual Chassis. To change the roles, you must delete the Virtual Chassis, form a new Virtual Chassis, and then, onboard the new Virtual Chassis.

Table 33: Fields on the Add Site for Tenant-Name Page (LAN Capability) *(continued)*

Field	Description
Method	

Table 33: Fields on the Add Site for Tenant-Name Page (LAN Capability) *(continued)*

Field	Description
	Select the method of provisioning the Virtual Chassis: • Auto Provisioning: The Virtual Chassis automatically determines the roles (primary, backup, and line card) of the member devices. If you select this option, you must enter only the serial number of the primary device in the Master Serial Number field that appears. • Pre Provisioning: You can determine the roles (primary, backup, and line card) of the member devices in the Virtual Chassis. If you select this option, you must provide the serial number, device model, device type, and role of all the member devices of the Virtual Chassis in the fields that appear. **NOTE:** In the case of preprovisioning, the primary device must always be designated as Member 0. For both these methods, ensure that: • The devices in the Virtual Chassis are fully installed and ready to be configured in the site. In addition, all members must be powered on. This means that the output of the show virtual-chassis status command must display all the member devices of the Virtual Chassis and the devices must be in Present (**Prsnt**) state. **NOTE:** If you do not have access to the serial console port for preprovisioning, only the primary device must be powered on first. • The primary and backup member devices have internet access to the Juniper redirect server and CSO. • All members in the Virtual Chassis are running the same firmware (either JUNOS 18.4R2.7 or 18.4R3.3). • For EX3400, EX4300, EX4600, and EX4650 devices to act as a Virtual Chassis, all the corresponding member devices are interconnected through Virtual Chassis ports (VCPs). For EX2300 devices to act as a Virtual Chassis, the 10-Gbps Ethernet ports are configured as VCPs

Table 33: Fields on the Add Site for Tenant-Name Page (LAN Capability) *(continued)*

Field	Description
	manually and the member devices are interconnected.
Master Serial Number	If you selected Auto Provisioning, enter the serial number of the primary device (from the fully-formed Virtual Chassis). To obtain the serial number, log in to the CLI of any device that is part of the fully-formed Virtual Chassis, in operational mode, and enter **show virtual-chassis**. The list of the member devices in the Virtual Chassis, along with the serial number and role appear. The primary device is indicated as **Master** under **Role**. Alternatively, you can view the serial number on the barcode sticker, which is on the rear-panel of the switch.
Member *<member-number>*	If you selected Pre Provisioning, enter the serial numbers of all the devices (from the fully-formed Virtual Chassis or based on what roles you decide to assign each Virtual Chassis member), and also select the device type and model from the list. **NOTE:** • If you enable ZTP, you must enter the serial number of the primary device only in the Member 0 field. • If you do not have access to the serial console port of the virtual chassis, the first member that is powered on is considered the primary. Enter the serial number of this device in the Member 0 field. Click the Add (+) icon to add a member or the Remove (-) icon to remove a member. For information on the number of devices that can be added, see Table 36 on page 148. **NOTE:** The Routing Engine check box corresponding to Member 0 is always selected, indicating that Member 0 always acts as the primary. To select a member as backup, click the Routing Engine check box corresponding to that member; the remaining members act as line cards.

Table 33: Fields on the Add Site for Tenant-Name Page (LAN Capability) *(continued)*

Field	Description
Serial Number	If you disabled the Virtual Chassis toggle button, specify the serial number of the physical switch. To obtain the serial number, log in to the CLI of the switch in operational mode and enter **show chassis hardware**. Alternatively, you can view the serial number on the barcode sticker, which is on the rear-panel of the switch. The serial number is a case-sensitive, alphanumeric string.
Auto activate	Click the toggle button to enable (default) or disable automatic activation of the switch when the switch is detected by CSO (that is, management status of the device is Device_Detected). When you enable this field, zero-touch provisioning (ZTP) of the switch is automatically triggered when the device communicates with CSO. **NOTE:** The switch must be powered on for automatic activation when you enable this option.
Activation code	If you disabled the **Auto activate** field, enter the activation code to be used for manually activating the switch For information about manually activating a switch, see "Manually Activating a Switch" on page 232.
Zero Touch Provisioning	Click the toggle button to enable or disable zero-touch provisioning (ZTP) of the switch through ZTP. If you disable ZTP, you must manually copy and paste the Stage-1 configuration on the switch during site activation. See Step 5 for details. **NOTE:** • Only EX Series switches running 18.4R2.7 or 18.4R3.3 firmware support ZTP. • EX4600 and EX4650 switches do not support Phone-Home client. You must disable ZTP and manually configure the stage-1 configuration on the switches.

Table 33: Fields on the Add Site for Tenant-Name Page (LAN Capability) *(continued)*

Field	Description
Boot Image	Select the boot image from the list if you want to upgrade the image for the switch. The boot image is the latest device image that is uploaded to the image management system. The boot image is used to upgrade the device when the CSO starts the ZTP process. If the boot image is not provided, then the device skips the automatic upgrade procedure. The boot image is populated based on the device template that you have selected while creating a site. **NOTE:** This option is not available for a Virtual Chassis. To provision a Virtual Chassis in CSO, you must manually upgrade the image to either JUNOS 18.4R2.7 or 18.4R3.3.
LAN Segment	Displays the VLANs and their IDs that are configured on the switch. To add a VLAN, click the **+** icon on the top, right corner of the LAN table. The Add LAN Segment page appears. See Table 34 on page 147. Adding a VLAN while creating the site is optional.
Port Profile You are directed to this page only if you have added a physical switch or a preprovisioned Virtual Chassis. If you have added an autoprovisioned Virtual Chassis; you are automatically directed to the Summary page. This is because, in the case of autoprovisioning, port profiles can be configured only after provisioning the Virtual Chassis.	

Table 33: Fields on the Add Site for Tenant-Name Page (LAN Capability) *(continued)*

Field	Description
Interface List	Displays the list of interfaces present on the device. Optional: You can assign the ports to a port profile and VLAN from here. For more information on the fields displayed in the Interface List table, see Table 35 on page 147. To assign the ports to a port profile and VLAN: 1. Click **Edit Configuration** on the top-right corner, above the Interface List table. The Configuration page appears. 2. From the **Port Profile** list, select a port profile to be assigned to the port. **NOTE:** The port profile must already be created from the Port Profiles page (**Configuration > SD-LAN > Port Profiles**) for it to be listed here. 3. In the **VLAN** field, if the port is configured as a trunk port in the port profile, assign multiple VLANs by selecting the VLANs in the **Available** column and clicking the right-arrow to move them to the **Selected** column. If the port is configured as an access port in the port profile, you can assign only one VLAN. 4. From the **Native VLAN** list, select a VLAN that you want to configure as native. This option appears only if you select Trunk port profile from the Port Profile list. Optional: Click the **Search** icon to search for a specific port in the list.

Table 33: Fields on the Add Site for Tenant-Name Page (LAN Capability) *(continued)*

Field	Description
Access Profiles List	Displays the access profile configured on the device from the Access Profiles page (Configuration > SD-LAN > Access Profiles). For details of the fields displayed on the Access Profiles List table, see "About the Access Profiles Page" on page 692.

Table 34: Fields on the Add LAN Segment Page when Adding a Site With LAN Capability

Field	Description
Add LAN Segment	
Name	Enter a name for the VLAN. The name for a VLAN should be a unique string of alphanumeric characters and some special characters (. -). No spaces are allowed and the maximum length allowed is 15 characters.
VLAN ID	Enter the VLAN ID for the VLAN. Range: 2 through 4093.

Table 35 on page 147 describes the fields on the Interface List table.

Table 35: Fields on the Interface List Table on the Port Profile Page

Field	Description
Port	Name of the port.
VLAN ID	ID of the VLAN configured on the port. If the port is a trunk port, all the VLAN IDs configured on the port are displayed.
Native VLAN	ID of the VLAN configured to accept untagged frames.
Port Profile	Port profile used for configuring the port. if the port is configured manually, the column displays **Manually Configured**.

Table 36 on page 148 lists the supported device types, combinations in the non-mixed mode, and the total number of members, supported by each device type, in a Virtual Chassis.

Table 36: Supported Device Types, Modes, and Number of Members Allowed in a Virtual Chassis

Device Type	Non-mixed Virtual Chassis Support	Number of Members Allowed in the Virtual Chassis
EX2300	Combination of the same or different models of EX2300 switches.	Up to 4 members.
EX3400	Combination of the same or different models of EX3400 switches.	Up to 10 members.
EX4300	Combination of the same or different models of EX4300 switches.	Up to 10 members.
EX4600	Combination of the same or different models of EX4600 switches.	Up to 10 members.
EX4650	Combination of the same or different models of EX4650 switches.	Up to 2 members.

Before you autoprovision or preprovision a Virtual Chassis in CSO, ensure that the Virtual Chassis is setup.

- To setup a Virtual Chassis for autoprovisioning:

  1. Decide the number of member devices in the Virtual Chassis.

  2. If you've added EX3400, EX4300, EX4600, or EX4650 devices as Virtual Chassis, interconnect all the corresponding member devices through Virtual Chassis ports (VCPs).

     If you've added EX2300 devices as Virtual Chassis, configure the 10-Gbps Ethernet ports as VCPs manually (through CLI) and interconnect the member devices.

     NOTE: At this point, do not power on any member devices in the Virtual Chassis.

  3. Decide which member device acts as the primary device and power on only this device first.

NOTE:

- Remember the serial number of the primary device in the Virtual Chassis. This serial number is required during the site activation workflow to add this Virtual Chassis in CSO.
- For ZTP to be successful, the primary device should always be designated as Member 0. You must specify the same serial number in the Member 0 field in CSO.

4. Wait until the primary device completes booting.

   After booting is complete, the LCD panel on this device displays a menu that includes the JUNOS OS version loaded on the device, status of VCPs, status of power supplies, and so on.

5. Power on the remaining member devices one after the other.

6. Wait until all the member devices complete booting.

   After booting is complete, you can confirm that the Virtual Chassis is fully formed when all the LEDs on the VCPs are ON.

7. Connect the primary and backup device to the Internet through the management port or uplink port.

8. Verify the connectivity from the primary device to CSO or to any host on the Internet by using **ping** or **telnet** to Juniper redirect server on port 443.

- To setup a Virtual Chassis for preprovisioning:

  1. Decide the number of member devices in the Virtual Chassis.

  2. If you've added EX3400, EX4300, EX4600, or EX4650 devices as Virtual Chassis, interconnect all the corresponding member devices through Virtual Chassis ports (VCPs).

     If you've added EX2300 devices as Virtual Chassis, configure the 10-Gbps Ethernet ports as VCPs manually and interconnect the member devices.

     NOTE: At this point, do not power on any member devices in the Virtual Chassis.

  3. Decide which member device acts as the primary and which member device acts as the backup.

  4. Of the two devices, power on the device that you want to select as the primary (Member 0), and wait until it completes booting.

After booting is complete, the LCD panel on this device displays a menu that includes the JUNOS OS version loaded on the device, status of VCPs, status of power supplies, and so on.

NOTE:

- Remember the serial numbers of all the devices in the Virtual Chassis. These serial numbers will be needed in the site activation workflow to add this Virtual Chassis in CSO.
- For ZTP to be successful, the primary should always be designated as Member 0. You must specify the same serial number in the Member 0 field in CSO.

5. Power on the device that you want to select as the backup and wait until it completes booting.

   After booting is complete, the LCD panel on this device displays a menu that includes the JUNOS OS version loaded on the device, status of VCPs, status of power supplies, and so on.

6. Power on the remaining member devices one after the other.

7. Wait until all the member devices complete booting.

   After booting is complete, you can confirm that the Virtual Chassis is fully formed when all the LEDs on the VCPs are ON.

8. Connect the primary and backup device to the Internet through the management port or uplink port.

9. Verify the connectivity from the primary device to CSO or to any host on the Internet by using **ping** or **telnet** to Juniper redirect server on port 443

Now that the Virtual Chassis is setup, proceed to add the Virtual Chassis in CSO. See the **Switch Details** section in Table 33 on page 138 for details.

RELATED DOCUMENTATION

Adding and Provisioning Switches to Provide LAN Capability to a Site Overview | **57**

Add a Switch to an Existing SD-WAN Site Or Next-Generation Firewall Site | **160**

Deleting a Site | **194**

Adding an On-Premise Spoke Site with Next Generation Firewall and LAN Capabilities

You can add an on-premise spoke site with both firewall and LAN capabilities.

To add a site with both next generation firewall and LAN capabilities at the same time:

1. Select **Resources > Site Management**.

   The Sites page appears.

2. Click **Add** and select **Add On-Premise Spoke (Manual)**.

   The Add On-Premise Spoke Site for *Tenant-Name* page appears.

3. Complete the configuration settings according to the guidelines provided in Table 37 on page 152.

   NOTE: Fields marked with an asterisk (*) are mandatory.

4. Click **Next**.

   A summary page is displayed.

5. Review the configuration and modify the settings, if needed, from the Summary tab.

6. Click **OK** to add the site.

   The site activation job is initiated and the Site Activation: *Site-Name* page appears displaying the progress of the steps executed for activating the firewall device and the switch (when LAN capability is selected). The firewall device is activated first and then the process to activate the switch is initiated.

7. • If the Zero Touch Provisioning (ZTP) toggle button is enabled (default), CSO pushes the stage-1 and stage-2 configurations and provisions the switch.

   This process occurs immediately after the activation process, for which you entered the activation code or selected auto-activation.

> **NOTE:** Stage-1 configuration is the initial configuration that allows basic connectivity to a device, which is pushed to the device.
>
> The configuration that is pushed to the device after it has connected to CSO is called stage-2 configuration.

- If you disabled the Zero Touch Provisioning (ZTP) toggle button, you must manually configure the stage-1 configuration (as provided by CSO) on the switch.

  To manually configure the stage-1 configuration:

  a. On the **Site Activation: *Site-Name*** page, the **Click to copy stage-1 configuration** link appears after the Prestage Device step completes successfully.

  b. Click the **Click to copy stage-1 configuration** link.

     The stage-1 configuration page appears displaying the stage-1 configuration to be copied to the EX Series device.

  c. Copy the stage-1 configuration and log in to the console of the EX Series switch.

  d. Enter the configuration mode, paste, and commit the configuration.

     After the stage-1 configuration is committed, the switch has the outbound SSH configuration to connect with CSO.

     CSO then provisions the switch.

> **NOTE:** You can also add a site with LAN and next generation firewall capabilities using the site templates. For more information, see "Adding On-Premise Spoke Sites by Using a Site Template" on page 199.

Table 37: Fields on the Add On-Premise Spoke Site for Tenant-Name Page (Firewall and LAN)

Field	Description
General	
Site Information	

Table 37: Fields on the Add On-Premise Spoke Site for Tenant-Name Page (Firewall and LAN) *(continued)*

Field	Description
Site Name	Enter a unique name for the firewall site. You can use alphanumeric characters and hyphen (-); the maximum length is 10 characters.
Site Group	Select a site group to which you want to assign the site.
Site Capabilities	
WAN Capabilities	Select the WAN capabilities as **Next Gen Firewall** for the site.
LAN Capabilities	Select the LAN capability as **LAN** for the site.
Address and Contact Information	
Street Address	Enter the street address of the site.
City	Enter the name of the city where the site is located.
State/Province	Select the state or province where the site is located.
ZIP/Postal Code	Enter the postal code for the site.
Country	Select the country where the site is located. You can click the **Validate** button to verify the address that you specified: • The **site address verification successful** message is displayed if the address can be verified. You can click the **View location on a map** link to see the address location. • If the address cannot be verified, the **Site address could not be validated** message is displayed .
Contact Name	Enter the name of the contact person for the site.
Email	Enter the e-mail address of the contact person for the site.
Phone	Enter the phone number of the contact person for the site. Click **Next** to continue.
Advanced Configuration	

Table 37: Fields on the Add On-Premise Spoke Site for Tenant-Name Page (Firewall and LAN) *(continued)*

Field	Description
Name Server IP List	Enter one or more IPv4 addresses of the DNS server. To enter more than one DNS server address, type address, press Enter, and then type the next address, and so on. DNS servers are used to resolve hostnames into IP addresses.
NTP Server	Enter the fully qualified domain names (FQDNs) or IP addresses of one or more NTP servers. Example: ntp.example.net The site must have DNS reachability to resolve the FQDN during site configuration.
Select Timezone	Select the time zone for the site.
WAN	
Device Information	
Serial Number	Enter the serial number of the firewall device. Note that the serial numbers are case-sensitive.
Auto Activate	Click the toggle button to enable or disable automatic activation of the device. This option is enabled by default.
Activation Code	If the **Auto Activate** feature is disabled, enter the activation code to manually activate the device. The activation code is provided by the administrator who adds the site.
Zero Touch Provisioning	Click the toggle button to enable or disable Zero Touch Provisioning (ZTP). This option is enabled by default. If ZTP is enabled, the Boot Image field is displayed and you must select an image that supports the Phone-Home client. During ZTP, the image on the firewall device is upgraded to the image that you select for the Boot Image. If ZTP is disabled, you must manually copy (by using CLI), the Stage-1 configuration on to the firewall device.

Table 37: Fields on the Add On-Premise Spoke Site for Tenant-Name Page (Firewall and LAN) *(continued)*

Field	Description
Boot Image	When the Zero Touch Provisioning field is enabled, select the boot image from the drop-down list to upgrade the image on the firewall device to a version that supports Phone-Home client. The boot image is the device image that was previously uploaded to the image management system. The boot image is used to upgrade the device when the CSO starts the ZTP process. If the boot image is not provided, then the device skips the automatic upgrade procedure. The boot image is populated based on the device template that you have selected while creating a site. By default, the **Use Image on Device** option is selected.
In-band Management Port	Select the port that you want to configure as management interface and connect it to the management device. You can configure any of the ge-0/0/x ports, where x ranges from 0 to 14, as in-band management interfaces. This field is applicable only when a switch is behind a CPE (SD-WAN or a next generation firewall device).
Firewall Policies	Select the firewall policy that you want to deploy to the standalone firewall site. The firewall policy list is populated from the **Configuration > Firewall > Firewall Policy** page. Default: Factory_Default_Fw_Policy
NAT Policies	Select the NAT policy that you want to deploy to the standalone firewall site. The NAT policy list is populated from the **Configuration > NAT > NAT Policies** page. Default: Factory_Default_NAT_Policy
Import Configuration	Click the toggle button to automatically import firewall policies and NAT policies from a next generation firewall device to CSO. By default, this field is disabled. **NOTE:** This field is available only when Zero Touch Provisioning is disabled.
LAN	
Device Profile	
Device Name	Enter a name for the switch. You can use alphanumeric characters and hyphen (-). The maximum length allowed is 15 characters.
Device Type	Select the type of switch—EX2300, EX3400, EX4300, EX4600, and EX4650.

Table 37: Fields on the Add On-Premise Spoke Site for Tenant-Name Page (Firewall and LAN) *(continued)*

Field	Description
Device Model	Select the device model for the switch that you specified in the Device Type field. The models vary in the number and type of ports the switch contains. For example, If you selected EX3400, select a model such as EX3400-24P, EX3400-48P, EX3400-24T among others.
CPE Settings	
Trunk Ports	Select at least two trunk ports on the CPE device to connect with the switch. The trunk ports are used for the following: • LAN traffic between the switch and the CPE • Management traffic for in-band management of the switch.
Switch Management Subnet	Specify the subnet that the DHCP can use to assign IP addresses. The DHCP server runs on the following ports: • Trunk ports to provide DHCP information to all devices connected to the switch and to the in-band management port, switch management port, and LAN ports on the CPE. • Out-of-band management port on the CPE to provide DHCP information to the management port on the switch. • LAN ports on the CPE to provide information to the devices connected to the CPE LAN ports. Example: 192.0.2.0/24
Switch Details	
Serial Number	Specify the serial number of the switch The serial number is a 12-digit number present on the rear panel of the switch.
Auto Activate	Click the toggle button to enable or disable automatic activation of the switch. When you enable this field, zero-touch provisioning of the switch is automatically triggered when the device communicates with CSO. **NOTE:** You must physically connect the switch to the CPE device (firewall) and power it on for the switch to be automatically activated when you enable this option.

Table 37: Fields on the Add On-Premise Spoke Site for Tenant-Name Page (Firewall and LAN) *(continued)*

Field	Description
Activation code	When the **Auto activate** field is disabled, enter the activation code to be used for manually activating the switch. For information, see "Manually Activating a Switch" on page 232.
Zero Touch Provisioning	ZTP must be disabled for all EX Series switches for the CSO 5.0.0 release. The Stage-1 configuration must be copied and pasted onto the CLI of the switch during site activation. See "Step-by-Step Procedure" on page 151 for details.
LAN Segment	Displays the LAN segment configured on the switch. To add a LAN segment, click the + icon on the top, right corner of the LAN table. The Add LAN Segment page appears. Specify values for the LAN segment based on guidelines provided in Table 38 on page 157. Fields marked * are mandatory. **NOTE:** The same LAN segment is created on the CPE device (firewall)if the switch is connected to the CPE device (firewall) that is managed by CSO.

Table 38: Fields on the Add LAN Segment Page when Adding a Switch along with Next-Generation Firewall

Field	Description
Add LAN Segment	
Name	Enter a name for the LAN segment. The name for a LAN segment should be a unique string of alphanumeric characters and some special characters (. -). No spaces are allowed and the maximum length is 15 characters.
VLAN ID	Enter the VLAN ID for the LAN segment. Range: 2 through 4093
Gateway Address/Mask	Enter a valid gateway IP address and mask for the LAN segment; for example, 192.0.2.8/24.

Table 38: Fields on the Add LAN Segment Page when Adding a Switch along with Next-Generation Firewall *(continued)*

Field	Description
DHCP	For directly connected LAN segments, click the toggle button to enable DHCP. DHCP is disabled by default. You enable DHCP if you want to assign IP addresses by using a DHCP server. You disable DHCP if you want to assign a static IP address to the LAN segment. **NOTE:** If you enable DHCP, fields related to DHCP-related parameters appear and must be configured.
[DHCP-Related Fields]	
Address Range Low	Enter the starting IP address in the range of IP addresses that can be allocated by the DHCP server to the LAN segment.
Address Range High	Enter the ending IP address in the range of IP addresses that can be allocated by the DHCP server to the LAN segment.
Maximum Lease Time	Specify the maximum duration (in seconds) for which a client can request for and hold a lease on a DHCP server. Range: 0 through 4,294,967,295.
Name Server	Specify or select one or more IPv4 addresses of the DNS server. To enter more than one DNS server address, type the address, press Enter, and then type the next address, and so on. DNS servers are used to resolve hostnames into IP addresses.

RELATED DOCUMENTATION

Adding and Provisioning a Next Generation Firewall Overview | 158

Adding and Provisioning a Next Generation Firewall Overview

Overview

You can use Contrail Service Orchestration (CSO) to

- Add a firewall site for the next generation firewall device.
- Configure a CPE device (SRX Series services gateway) as a next generation firewall device.
- Add firewall policies for the standalone firewall site.
- Deploy the firewall policies for the standalone firewall site.

Topology

The topology to add an on-premise spoke site with next generation firewall capabilities is shown in Figure 5 on page 159.

Figure 5: On-premise spoke site with next generation firewall

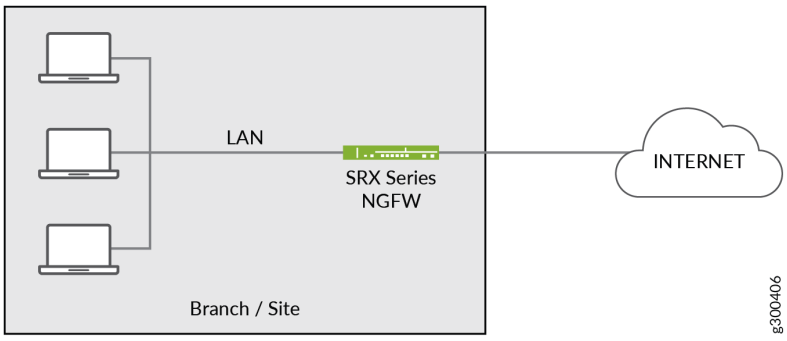

The topology to add an on-premise spoke site with next generation firewall and LAN capabilities is shown in Figure 6 on page 159.

Figure 6: On-premise spoke site with next generation firewall and LAN

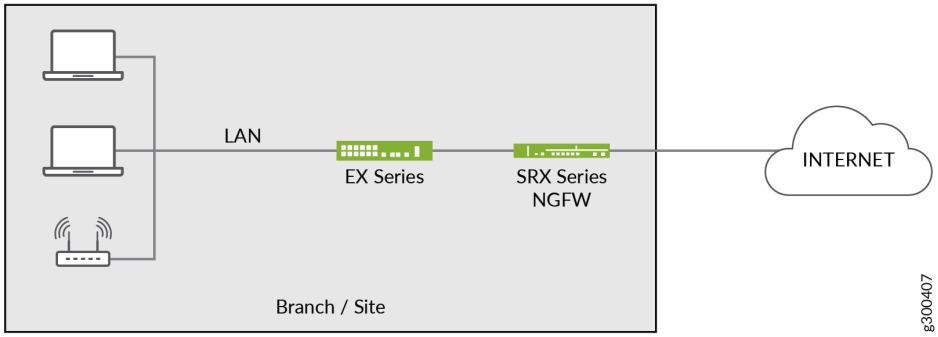

Workflow

The following workflow describes the steps that are required to set up a firewall site and provision the firewall device associated with the site.

To set up a next generation firewall site and provision the firewall device:

1. Add a standalone next generation firewall site. See "Adding a Standalone Next Generation Firewall Site" on page 177.

   To add a site with next generation firewall and switch, see "Adding an On-Premise Spoke Site with Next Generation Firewall and LAN Capabilities" on page 151.

   > **NOTE:** Before proceeding to the next step ensure that the ZTP process is complete and the firewall device status is set to **Provisioned** state.

2. Configure the firewall device. See "Configuring the Firewall Device" on page 299.

3. Add firewall policies for the site. See "Adding a Firewall Policy" on page 378.

4. Add firewall policy intents for the firewall policies that you added. See "Adding Firewall Policy Intents" on page 381.

5. Deploy firewall policies to the site. See "Deploying Firewall Policies" on page 441.

Add a Switch to an Existing SD-WAN Site Or Next-Generation Firewall Site

You can add one switch to an existing on-premise spoke site that already has a CPE or firewall device provisioned or an enterprise hub site, to provide LAN capability to the site. See "Switch Behind a CPE or Next Generation Firewall Overview" on page 58 for details.

To configure more than one switch with an SD-WAN CPE or Next-Generation firewall, add a separate SD-LAN site and add multiple switches to the site.

You can manage the connectivity and configuration between the switches and the CPE or Next-Generation firewall by creating uplink ports to the CPE device or Next-Generation Firewall, either manually or by using configuration templates. See "Add Switches to an Existing SD-LAN Site" on page 166 to add an SD-LAN site with multiple switches.

To add a switch to an existing site:

1. Select **Resources > Site Management**.

   The Sites page appears.

2. Do one of the following:

   - a. Select the site to which you want to add the switch.

     b. Click **Add > Add Switch**.

     The Add Switch page appears.

   - a. Click the *Site-Name* link of the site (to which you want to add the switch) in the Sites column.

     The *Site-Name* page appears.

     b. On the Devices tab, click **Add Switch**.

     The Add Switch page appears.

3. Complete the configuration according to the guidelines provided in Table 39 on page 162.

   NOTE: Fields marked with asterisk (*) are mandatory.

4. Review the configuration from the **Summary** tab.

5. (Optional) Click the Edit links within the summary to go directly to a specific page of the wizard and modify the configured settings.

6. Click **OK** to add the switch to the site.

   The site activation process is initiated and the Site Activation: *Site-Name* page appears displaying the progress of the steps executed for activating the CPE and the switch.

7. - If the Zero Touch Provisioning (ZTP) toggle button is enabled (default), CSO pushes the stage-1 and stage-2 configurations and provisions the switch.

     This process occurs immediately after the activation process, for which you entered the activation code or selected auto-activation.

> **NOTE:** Stage-1 configuration is the initial configuration that allows basic connectivity to a device, which is pushed to the device.
>
> The configuration that is pushed to the device after it has connected to CSO is called stage-2 configuration.

- If you disabled the Zero Touch Provisioning (ZTP) toggle button, you must manually configure the stage-1 configuration (as provided by CSO) on the switch.

  To manually configure the stage-1 configuration:

  a. On the **Site Activation: *Site-Name*** page, the **Click to copy stage-1 configuration** link appears after the Prestage Device step completes successfully.

  b. Click the **Click to copy stage-1 configuration** link.

     The stage-1 configuration page appears displaying the stage-1 configuration to be copied to the EX Series device.

  c. Copy the stage-1 configuration and log in to the console of the EX Series switch.

  d. Enter the configuration mode, paste, and commit the configuration.

     After the stage-1 configuration is committed, the switch has the outbound SSH configuration to connect with CSO.

     CSO then provisions the switch.

Table 39 on page 162 provides guidelines on using the fields on the Add Switch page.

Table 39: Fields on the Add Switch Page

Field	Description
Device Profile	
Device Name	Enter a name for the switch. You can use alphanumeric characters and hyphen (-). The maximum length allowed is 15 characters.
Device Type	Select the type of switch—EX2300, EX3400, EX4300, EX4600, and EX4650.

Table 39: Fields on the Add Switch Page *(continued)*

Field	Description
Device Model	Select the model for the switch you specified in the Device Type field. The models vary in the number and type of ports the switch contains. For example, If you selected EX3400, select a model such as EX3400-24P, EX3400-48P, EX3400-24T among others.
CPE Settings	
Trunk Ports	Select at least two trunk ports on the CPE device to connect with the switch, which are used for the following: • LAN traffic between the switch and the CPE or firewall. • Management traffic for in-band management of the switch. **NOTE:** The ae0 port of the SRX Series devices is used as the trunk port for communication with the switch.
Switch Management Subnet	Specify the subnet that the DHCP can use to assign IP addresses. The DHCP server runs on the following ports: • Trunk ports to provide DHCP information to all devices connected to the switch and to the in-band management port, switch management port, and LAN ports on the CPE or firewall. • Out-of-band management port on the CPE or firewall to provide DHCP information to the management port on the switch. • LAN ports on the CPE or firewall to provide information to the devices connected to the CPE or firewall LAN ports.
Switch Details	
Serial Number	Specify the serial number of the physical switch. You can either view the serial number on the label that is present on the rear panel of the switch or log in to the CLI of the switch in operational mode and enter **show chassis hardware**. The serial number is a case-sensitive, alphanumeric string.

Table 39: Fields on the Add Switch Page *(continued)*

Field	Description
Zero Touch Provisioning	Click to enable or disable ZTP on the switch. If you disable ZTP, you must manually copy and paste the Stage-1 configuration on the switch during site activation. See "Step-by-Step Procedure" on page 161 for details **NOTE:** • Only EX Series switches running 18.4R2.7 or 18.4R3.3 firmware support ZTP. • EX4600 and EX4650 switches do not support Phone-Home client. You must disable ZTP and manually configure the stage-1 configuration on the switches.
Boot Image	Select the boot image from the list if you want to upgrade the image for the switch. The boot image is the latest device image that is uploaded to the image management system. The boot image is used to upgrade the device when the CSO starts the ZTP process. If the boot image is not provided, then the device skips the automatic upgrade procedure. The boot image is populated based on the device template that you have selected while creating a site.
Auto activate	Click the toggle button to enable or disable automatic activation of the switch. When you enable this field, zero-touch provisioning of the switch is automatically triggered when the device communicates with CSO. **NOTE:** You must physically connect the switch to the CPE and power it on for the switch to be automatically activated when you enable this option.
Activation code	When the **Auto activate** field is disabled, enter the activation code to be used for manually activating the switch. For information, see "Manually Activating a Switch" on page 232.

Table 40 on page 164 describes the fields on the Create LAN Segment page.

The values that you configure here are populated on the LAN Segments section of the Add Switch page.

Table 40: Fields on the Create LAN Segment Page when Adding a Switch to an Existing Site

Field	Description
Add LAN Segment	
Name	Enter a name for the LAN segment. The name for a LAN segment should be a unique string of alphanumeric characters and some special characters (. -). No spaces are allowed and the maximum length allowed is 15 characters.

Table 40: Fields on the Create LAN Segment Page when Adding a Switch to an Existing Site *(continued)*

Field	Description
Type	Select the type of LAN segment: **NOTE:** This field is available only for SD-WAN sites. • Directly Connected (default)—Indicates that the LAN segment is directly connected to the site. • Dynamic Routed—Indicates that the LAN segment is not directly connected to the site and is reachable by using a dynamic route. If you select this option, you must specify the dynamic routing information.
VLAN ID	Enter the VLAN ID for the LAN segment. Range: 2 through 4093.
Department	Select a department to which the LAN segment is to be assigned. **NOTE:** This field is available only for SD-WAN sites. Alternatively, click the **Create Department** link to add a new department and assign the LAN segment to it. See "Adding a Department" on page 767 for details. You group LAN segments as departments for ease of management and for applying policies at the department-level. . **NOTE:** This field is not displayed when you add the switch to a site with next-generation firewall capability.
Gateway Address/Mask	Specify a valid gateway IP address and mask for the LAN segment; for example, 192.0.2.8/24.
DHCP	For directly connected LAN segments, click the toggle button to enable or disable DHCP on the LAN segment. DHCP is disabled by default. You enable DHCP if you want to assign IP addresses by using a DHCP server. You disable DHCP if you want to assign a static IP address to the LAN segment. **NOTE:** If you enable DHCP, fields related to DHCP-related parameters appear and must be configured.
[DHCP-Related Fields]	
Address Range Low	Enter the starting IP address in the range of IP addresses that can be allocated by the DHCP server to the LAN segment.

Table 40: Fields on the Create LAN Segment Page when Adding a Switch to an Existing Site *(continued)*

Field	Description
Address Range High	Enter the ending IP address in the range of IP addresses that can be allocated by the DHCP server to the LAN segment.
Maximum Lease Time	Specify the maximum duration (in seconds) for which a client can request for and hold a lease on a DHCP server. Range: 0 through 4,294,967,295 seconds.
Name Server	Specify or select one or more IPv4 addresses of the DNS server. To enter more than one DNS server address, type the address, press Enter, and then type the next address, and so on. DNS servers are used to resolve hostnames into IP addresses.
CPE Ports	Select ports on the switch to be part of the LAN segment. Select the ports from the **Available** column and click the right-arrow to move the ports to the **Selected** column. If you create a LAN segment on a switch when the switch is connected to the CPE device, CSO automatically assigns LAN ports on the CPE device and creates the same LAN segment on the CPE device.

RELATED DOCUMENTATION

- Add Switches to an Existing SD-LAN Site | **166**
- Adding and Provisioning Switches to Provide LAN Capability to a Site Overview | **57**
- Adding an On-Premise Spoke Site with SD-WAN and LAN Capabilities | **117**
- Adding an On-Premise Spoke Site with Next Generation Firewall and LAN Capabilities | **151**
- Adding Enterprise Hubs with SD-WAN Capability or SD-WAN and LAN Capabilities | **61**
- Add an On-Premise Spoke Site with LAN Capability | **136**

Add Switches to an Existing SD-LAN Site

From CSO Release 5.1.1 onward, you can add one or more EX Series switches to an existing SD-LAN site. Physical, standalone EX Series switches and EX Series Virtual Chassis are supported.

Figure 7 on page 167 shows an SD-LAN site, with multiple EX Series switches, managed by CSO.

Figure 7: SD-LAN Site with Multiple Switches

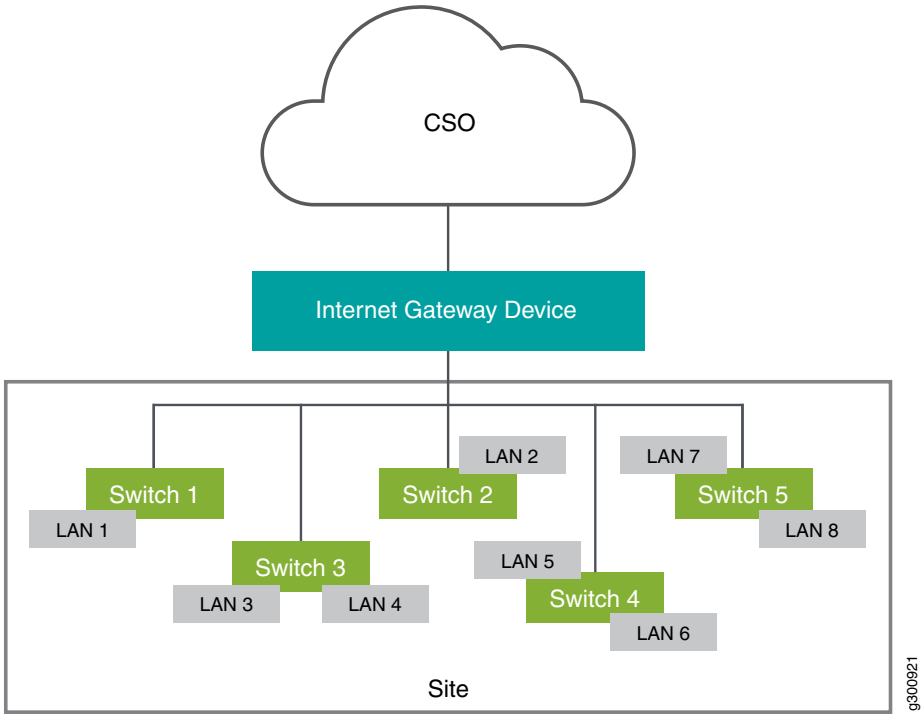

To add a switch to an existing site:

1. Select **Resources > Site Management**.

   The Sites page appears.

2. Do one of the following:

   - a. Select the site to which you want to add the switch.

     b. Click **Add > Add Switch**.

     The Add Switch page appears.

   - a. Click the *Site-Name* link of the site (to which you want to add the switch) in the Sites column.

     The *Site-Name* page appears.

     b. On the Devices tab, click **Add Switch**.

     The Add Switch page appears.

1. Complete the configuration according to the guidelines provided in Table 39 on page 162.

   NOTE: Fields marked with an asterisk (*) are mandatory.

2. Review the configuration from the **Summary** tab.

3. (Optional) Click the Edit links within the summary to go directly to a specific page of the wizard and modify the configured settings.

4. Click **OK** to add the switch to the site.

   The site activation process is initiated and the Site Activation: *Site-Name* page appears displaying the progress of the steps executed for activating the switch. If you add one or more EX Series switches, the progress of the steps executed for activating each switch is displayed.

5. - If the Zero Touch Provisioning (ZTP) toggle button is enabled (default), CSO pushes the stage-1 and stage-2 configurations and provisions the switch.

     This process occurs immediately after the activation process, for which you entered the activation code or selected auto-activation.

     NOTE: Stage-1 configuration is the initial configuration that allows basic connectivity to a device, which is pushed to the device.

     The configuration that is pushed to the device after it has connected to CSO is called stage-2 configuration.

   - If you disabled the Zero Touch Provisioning (ZTP) toggle button, you must manually configure the stage-1 configuration (as provided by CSO) on the switch.

     To manually configure the stage-1 configuration:

     a. On the **Site Activation: *Site-Name*** page, the **Click to copy stage-1 configuration** link appears after the Prestage Device step completes successfully.

     b. Click the **Click to copy stage-1 configuration** link.

        The stage-1 configuration page appears displaying the stage-1 configuration to be copied to the EX Series device.

c. Copy the stage-1 configuration and log in to the console of the EX Series switch.

d. Enter the configuration mode, paste, and commit the configuration.

   After the stage-1 configuration is committed, the switch has the outbound SSH configuration to connect with CSO.

   CSO then provisions the switch.

Table 39 on page 162 provides guidelines on using the fields on the Add Switch page.

Table 41: Fields on the Add Switch Page

Field	Description
Device Profile	
Device Name	Enter a name for the switch. You can use alphanumeric characters and hyphen (-). The maximum length allowed is 15 characters.
Device Type	Select the type of switch—EX2300, EX3400, EX4300, EX4600, and EX4650.
Device Model	Select the model for the switch you specified in the Device Type field. The models vary in the number and type of ports the switch contains. For example, If you selected EX3400, select a model such as EX3400-24P, EX3400-48P, EX3400-24T among others.
Switch Details	
Virtual Chassis	Click the toggle button to enable or disable (default) adding the switch as a Virtual Chassis. If you enable this toggle button, you must select the method of provisioning the Virtual Chassis. **NOTE:** • Before you add a Virtual Chassis in CSO, ensure that the Virtual Chassis is setup. See "Step-by-Step Procedure" on page 148 for information about setting up a Virtual Chassis. • In Release 5.1.1, you cannot add a new member or change the roles assigned to the members after you onboard a Virtual Chassis. To change the roles, you must delete the Virtual Chassis, form a new Virtual Chassis, and then, onboard the new Virtual Chassis.

Table 41: Fields on the Add Switch Page *(continued)*

Field	Description
Serial Number	If you disabled the Virtual Chassis toggle button, specify the serial number of the physical switch. You can either view the serial number on the label that is present on the rear panel of the switch or log in to the CLI of the switch in operational mode and enter **show chassis hardware**. The serial number is a case-sensitive, alphanumeric string.
Method	Select the method of provisioning the Virtual Chassis: • Auto Provisioning: The Virtual Chassis automatically determines the roles (primary, backup, and line card) of the member devices. If you select this option, you must enter only the serial number of the primary device in the Master Serial Number field that appears. • Pre Provisioning: You can determine the roles (primary, backup, and line card) of the member devices in the Virtual Chassis. If you select this option, you must provide the serial number, device model, device type, and role of all the member devices of the Virtual Chassis in the fields that appear. **NOTE:** In the case of preprovisioning, the primary device must always be designated as Member 0. For both these methods, ensure that: • The devices in the Virtual Chassis are fully installed and ready to be configured in the site. In addition, all member devices must be powered on. This means that the output of the show virtual-chassis status command must display all the member devices of the Virtual Chassis and the devices must be in Present (**Prsnt**) state. **NOTE:** If you do not have access to the serial console port for preprovisioning, only the primary device must be powered on first. • The primary and backup member devices have internet access to the Juniper redirect server and CSO. • All member devices in the Virtual Chassis are running the same firmware (either JUNOS 18.4R2.7 or 18.4R3.3). • For EX3400, EX4300, EX4600, and EX4650 devices to act as a Virtual Chassis, all the corresponding member devices are interconnected through Virtual Chassis ports (VCPs). For EX2300 devices to act as a Virtual Chassis, the 10-Gbps Ethernet ports are configured as VCPs manually and the member devices are interconnected.

Table 41: Fields on the Add Switch Page *(continued)*

Field	Description
Master Serial Number	If you selected Auto Provisioning, enter the serial number of the primary device (from the fully-formed Virtual Chassis). To obtain the serial number, log in to the CLI of any device that is part of the fully-formed Virtual Chassis, in operational mode, and enter **show virtual-chassis**. The list of the member devices in the Virtual Chassis, along with the serial number and role appear. The primary device is indicated as **Master** under **Role**. Alternatively, you can view the serial number on the barcode sticker, which is on the rear-panel of the switch.
Member *<member-number>*	If you selected Pre Provisioning, enter the serial numbers of all the devices (from the fully-formed Virtual Chassis or based on what roles you decide to assign each Virtual Chassis member), and also select the device type and model from the list. **NOTE:** • You must enter the serial number of the primary device only in the Member 0 field. • If you do not have access to the serial console port of the virtual chassis, the first member device that is powered on is considered the primary. Enter the serial number of this device in the Member 0 field. Click the Add (+) icon to add a member device or the Remove (-) icon to remove a member device. For information on the number of member devices that can be added, see Table 36 on page 148. **NOTE:** The Routing Engine check box corresponding to Member 0 is always selected, indicating that Member 0 always acts as the primary. To select a member device as backup, click the Routing Engine check box corresponding to that member device; the remaining member devices act as line cards.
Zero Touch Provisioning	Click to enable or disable ZTP on the switch. If you disable ZTP, you must manually copy and paste the Stage-1 configuration on the switch during site activation. See “Step-by-Step Procedure” on page 161 for details **NOTE:** • Only EX Series switches running 18.4R2.7 or 18.4R3.3 firmware support ZTP. • EX4600 and EX4650 switches do not support Phone-Home client. You must disable ZTP and manually configure the stage-1 configuration on the switches.

Table 41: Fields on the Add Switch Page *(continued)*

Field	Description
Boot Image	Select the boot image from the list if you want to upgrade the image for the switch. The boot image is the latest device image that is uploaded to the image management system. The boot image is used to upgrade the device when the CSO starts the ZTP process. If the boot image is not provided, then the device skips the automatic upgrade procedure. The boot image is populated based on the device template that you have selected while creating a site. **NOTE:** This option is not available for a Virtual Chassis. To provision a Virtual Chassis in CSO, you must manually upgrade the image to either JUNOS 18.4R2.7 or 18.4R3.3.
Auto activate	Click the toggle button to enable or disable automatic activation of the switch. When you enable this field, zero-touch provisioning of the switch is automatically triggered when the device communicates with CSO. **NOTE:** You must physically connect the switch to the CPE and power it on for the switch to be automatically activated when you enable this option.
Activation code	When the **Auto activate** field is disabled, enter the activation code to be used for manually activating the switch. For information, see "Manually Activating a Switch" on page 232.
Port Profile You are directed to this page only if you have added a physical switch or a preprovisioned Virtual Chassis. If you have added an autoprovisioned Virtual Chassis; you are automatically directed to the Summary page. This is because, in the case of autoprovisioning, port profiles can be configured only after provisioning the Virtual Chassis.	
Access Profiles List	Displays the access profile configured on the device from the Access Profiles page (Configuration > SD-LAN > Access Profiles). For details of the fields displayed on the Access Profiles List table, see "About the Access Profiles Page" on page 692.

Table 41: Fields on the Add Switch Page *(continued)*

Field	Description
Interface List	Displays the list of interfaces present on the device. You can assign the ports to a port profile and VLAN from here. For more information on the fields displayed in the Interface List table, see Table 35 on page 147. Optional: To assign the ports to a port profile and VLAN: 1. Click **Edit Configuration** on the top-right corner of the Interface List table. The Configuration page appears. 2. From the **Port Profile** list, select a port profile to be assigned to the port. **NOTE:** The port profile must already be created from the Port Profiles page (**Configuration > SD-LAN > Port Profiles**) for it to be listed here. 3. In the **VLAN** field, if the port is configured as a trunk port in the port profile, assign multiple VLANs by selecting the VLANs in the **Available** column and clicking the right-arrow to move them to the **Selected** column. If the port is configured as an access port in the port profile, you can assign only one VLAN. 4. From the **Native VLAN** list, select a VLAN that you want to configure as native. This option appears only if you select Trunk port profile from the Port Profile list. Optional: Click the **Search** icon to search for a specific port in the list.

Table 35 on page 147 describes the fields on the Interface List table.

Table 42: Fields on the Interface List Table on the Port Profile Page

Field	Description
Port	Name of the port.
VLAN ID	ID of the VLAN configured on the port. If the port is a trunk port, all the VLAN IDs configured on the port are displayed.
Native VLAN	ID of the VLAN configured to accept untagged frames.
Port Profile	Port profile used for configuring the port. If the port is configured manually, the column displays **Manually Configured**.

Table 36 on page 148 lists the supported device types, combinations in the non-mixed mode, and the total number of members, supported by each device type, in a Virtual Chassis.

Table 43: Supported Device Types, Modes, and Number of Members Allowed in a Virtual Chassis

Device Type	Non-mixed Virtual Chassis Support	Number of Members Allowed in the Virtual Chassis
EX2300	Combination of the same or different models of EX2300 switches.	Up to 4 members.
EX3400	Combination of the same or different models of EX3400 switches.	Up to 10 members.
EX4300	Combination of the same or different models of EX4300 switches.	Up to 10 members.
EX4600	Combination of the same or different models of EX4600 switches.	Up to 10 members.
EX4650	Combination of the same or different models of EX4650 switches.	Up to 2 members.

Before you autoprovision or preprovision a Virtual Chassis in CSO, ensure that the Virtual Chassis is setup.

- To setup a Virtual Chassis for autoprovisioning:

  1. Decide the number of member devices in the Virtual Chassis.

  2. If you've added EX3400, EX4300, EX4600, or EX4650 devices as Virtual Chassis, interconnect all the corresponding member devices through Virtual Chassis ports (VCPs).

     If you've added EX2300 devices as Virtual Chassis, configure the 10-Gbps Ethernet ports as VCPs manually (through CLI) and interconnect the member devices.

     NOTE: At this point, do not power on any member devices in the Virtual Chassis.

  3. Decide which member device acts as the primary device and power on only this device first.

NOTE:

- Remember the serial number of the primary device in the Virtual Chassis. This serial number is required during the site activation workflow to add this Virtual Chassis in CSO.
- For ZTP to be successful, the primary device should always be designated as Member 0. You must specify the same serial number in the Member 0 field in CSO.

4. Wait until the primary device completes booting.

   After booting is complete, the LCD panel on this device displays a menu that includes the JUNOS OS version loaded on the device, status of VCPs, status of power supplies, and so on.

5. Power on the remaining member devices one after the other.

6. Wait until all the member devices complete booting.

   After booting is complete, you can confirm that the Virtual Chassis is fully formed when all the LEDs on the VCPs are ON.

7. Connect the primary and backup device to the Internet through the management port or uplink port.

8. Verify the connectivity from the primary device to CSO or to any host on the Internet by using **ping** or **telnet** to Juniper redirect server on port 443.

- To setup a Virtual Chassis for preprovisioning:

1. Decide the number of member devices in the Virtual Chassis.

2. If you've added EX3400, EX4300, EX4600, or EX4650 devices as Virtual Chassis, interconnect all the corresponding member devices through Virtual Chassis ports (VCPs).

   If you've added EX2300 devices as Virtual Chassis, configure the 10-Gbps Ethernet ports as VCPs manually and interconnect the member devices.

   NOTE: At this point, do not power on any member devices in the Virtual Chassis.

3. Decide which member device acts as the primary and which member device acts as the backup.

4. Of the two devices, power on the device that you want to select as the primary (Member 0), and wait until it completes booting.

After booting is complete, the LCD panel on this device displays a menu that includes the JUNOS OS version loaded on the device, status of VCPs, status of power supplies, and so on.

NOTE:

- Remember the serial numbers of all the devices in the Virtual Chassis. These serial numbers will be needed in the site activation workflow to add this Virtual Chassis in CSO.
- For ZTP to be successful, the primary should always be designated as Member 0. You must specify the same serial number in the Member 0 field in CSO.

5. Power on the device that you want to select as the backup and wait until it completes booting.

   After booting is complete, the LCD panel on this device displays a menu that includes the JUNOS OS version loaded on the device, status of VCPs, status of power supplies, and so on.

6. Power on the remaining member devices one after the other.

7. Wait until all the member devices complete booting.

   After booting is complete, you can confirm that the Virtual Chassis is fully formed when all the LEDs on the VCPs are ON.

8. Connect the primary and backup device to the Internet through the management port or uplink port.

9. Verify the connectivity from the primary device to CSO or to any host on the Internet by using **ping** or **telnet** to Juniper redirect server on port 443

Now that the Virtual Chassis is setup, proceed to add the Virtual Chassis in CSO. See the **Switch Details** section in Table 39 on page 162 for details.

Release History Table

Release	Description
5.1.1	From CSO Release 5.1.1 onward, you can add one or more EX Series switches to an existing SD-LAN site. Physical, standalone EX Series switches and EX Series Virtual Chassis are supported.

RELATED DOCUMENTATION

Add a Switch to an Existing SD-WAN Site Or Next-Generation Firewall Site | 160

Enabling Integration with Mist Access Points

You can enable integration with the Mist access points to easily access and view Mist access points connected to the branch network. When integration with Mist access point is enabled, the connected access points are listed in the **Devices** tab of the **Resources** > **Site Management** > ***Site Name*** page. You can click the access point name to view the Mist access point details from the Mist portal that is integrated with CSO.

To enable integration with the Mist access point:

1. Select **Administration** > **WiFi Settings**.

   The WiFi Settings page appears.

2. Click the Enable toggle button to enable integration with Mist access points.

   The Login E-mail and Login Password fields appear.

3. In the Login E-mail page, enter the e-mail address that is the username for your Mist account.

4. In the Login Password page, enter the password for your Mist account.

5. Click **Save**.

After you enable integration and enter the login credentials, CSO adds the access point to the list of devices associated with a site. To view details about the access point, **Devices** tab of the **Resources** > **Site Management** > ***Site Name*** page and click the access point name. The Mist portal page for the selected device appears.

Adding a Standalone Next Generation Firewall Site

You add the standalone firewall site from the **Sites** page.

To add a standalone firewall site:

1. Select **Resources > Site Management**.

   The Sites page appears.

2. Click **Add** and select **Add On-Premise Spoke (Manual)**.

   The Add On-Premise Spoke Site for *Tenant-Name* page appears.

3. Complete the configuration settings according to the guidelines provided in Table 44 on page 179.

   NOTE: Fields marked with an asterisk (*) are mandatory.

4. Click **Next**.

   A summary page is displayed.

5. Review the configuration and modify the settings, if needed, from the Summary tab.

6. Click **OK** to add the site.

   The site activation job is initiated and the Site Activation: *Site-Name* page appears displaying the progress of the steps executed for activating the firewall site.

   You can click the view job details link in the pop-up to view the progress of the job.

7. - If you have disabled the Zero Touch Provisioning (ZTP) field for the firewall device, you must manually configure the stage-1 configuration on the firewall device.

     To manually configure the stage-1 configuration:

     a. On the Site Activation: *Site-Name* page, after the Prestage Device step completes successfully, the click to copy stage-1 config link appears .

     b. Click the **click to copy stage-1 config** link.

        The Stage-1 Configuration page appears displaying the stage-1 configuration.

     c. Copy the stage-1 configuration and log in to the CLI of the firewall device.

     d. Enter the configuration mode, paste, and commit the configuration.

        After the stage-1 configuration is committed, the firewall device establishes the outbound SSH connection to connect with CSO. After the firewall device is detected, CSO executes the bootstrap and provisioning processes and completes provisioning the firewall device. The standalone firewall site status is set to **Provisioned** in the Sites page.

   - If you have enabled the Zero Touch Provisioning field, CSO pushes the stage-1 and stage-2 configuration and provisions the firewall device. The standalone firewall site status is set to **Provisioned** in the Sites page.

NOTE: The firewall device is activated automatically, as you have already provided the activation code while creating the firewall site.

NOTE: You can also add a standalone firewall site using the site templates. For more information, see "Adding On-Premise Spoke Sites by Using a Site Template" on page 199.

Table 44: Fields on the Add On-Premise Spoke Site for Tenant-Name Page (Standalone Firewall)

Field	Description
General	
Site Information	
Site Name	Enter a unique name for the firewall site. You can use alphanumeric characters and hyphen (-); the maximum length is 10 characters.
Site Group	Select a site group to which you want to assign the site.
Site Capabilities	
WAN Capabilities	Select the WAN capability as **Next Gen Firewall** as you are adding a next generation firewall site.
Address and Contact Information	
Street Address	Enter the street address of the site.
City	Enter the name of the city where the site is located.
State/Province	Select the state or province where the site is located.
ZIP/Postal Code	Enter the postal code for the site.

Table 44: Fields on the Add On-Premise Spoke Site for Tenant-Name Page (Standalone Firewall) *(continued)*

Field	Description
Country	Select the country where the site is located. You can click the **Validate** button to verify the address that you specified: • The **site address verification successful** message is displayed if the address can be verified. You can click the **View location on a map** link to see the address location. • If the address cannot be verified, the **Site address could not be validated** message is displayed .
Contact Name	Enter the name of the contact person for the site.
Email	Enter the e-mail address of the contact person for the site.
Phone	Enter the phone number of the contact person for the site. Click **Next** to continue.
Advanced Configuration	
Name Server IP List	Enter one or more IPv4 addresses of the DNS server. To enter more than one DNS server address, type address, press Enter, and then type the next address, and so on. DNS servers are used to resolve hostnames into IP addresses.
NTP Server	Enter the fully qualified domain names (FQDNs) or IP addresses of one or more NTP servers. Example: ntp.example.net The site must have DNS reachability to resolve the FQDN during site configuration.
Select Timezone	Select the time zone for the site.
WAN	
Device Information	
Serial Number	Enter the serial number of the firewall device. Note that the serial numbers are case-sensitive.
Auto Activate	Click the toggle button to enable or disable automatic activation of the device. This option is enabled by default.
Activation Code	If the automatic activation of the device is disabled, enter the activation code to manually activate the device. The activation code is provided by the administrator who adds the site.

Table 44: Fields on the Add On-Premise Spoke Site for Tenant-Name Page (Standalone Firewall) *(continued)*

Field	Description
Zero Touch Provisioning	Click the toggle button to enable or disable Zero Touch Provisioning (ZTP). This option is enabled by default. If ZTP is enabled, the Boot Image field is displayed and you must select an image that supports the Phone-Home client. During ZTP, the image on the firewall device is upgraded to the image that you select for the Boot Image. If ZTP is disabled, you must manually copy (by using CLI), the Stage-1 configuration on to the firewall device.
Boot Image	When the Zero Touch Provisioning field is enabled, select the boot image from the drop-down list to upgrade the image on the firewall device to a version that supports Phone-Home client. The boot image is the device image that was previously uploaded to the image management system. The boot image is used to upgrade the device when the CSO starts the ZTP process. If the boot image is not provided, then the device skips the automatic upgrade procedure. The boot image is populated based on the device template that you have selected while creating a site. By default, the **Use Image on Device** option is selected.
In-band Management Port	Select the port that you want to configure as management interface and connect it to the management device. You can configure any of the ge-0/0/x ports, where x ranges from 0 to 14, as in-band management interfaces. This field is applicable only when a switch is behind a CPE (SD-WAN or next generation firewall device).
Firewall Policies	Select the firewall policy that you want to deploy to the standalone firewall site. The firewall policy list is populated from the **Configuration > Firewall > Firewall Policy** page. Default: Factory_Default_Fw_Policy
NAT Policies	Select the NAT policy that you want to deploy to the standalone firewall site. The NAT policy list is populated from the **Configuration > NAT > NAT Policies** page. Default: Factory_Default_NAT_Policy
Import Configuration	Click the toggle button to automatically import firewall policies and NAT policies from a next generation firewall device to CSO. By default, this field is disabled. **NOTE:** This field is available only when Zero Touch Provisioning is disabled.

RELATED DOCUMENTATION

Adding and Provisioning a Next Generation Firewall Overview | 158

Managing LAN Segments on a Tenant Site

IN THIS SECTION

- Adding LAN Segments | 182
- Deploying LAN Segments | 186
- Reassigning a Department to a LAN Segment | 187
- Deleting LAN Segments | 188

A network on a tenant site is divided into multiple LAN segments to improve traffic management and security. A LAN segment is a small portion of a LAN that is used by a work group. A grouping of multiple LAN segments form a department. LAN segments are separated by a bridge, router, or a switch.

You can view and manage LAN segments from the LAN tab of the *Site Name* page.

These topics describe how to manage LAN segments on a site.

Adding LAN Segments

You add LAN segments from the *Site Name* page.

To add a LAN segment:

1. Click **Resources > Site Management**.

   The Sites page appears.

2. Click the site for which you want to add the LAN segment.

   The *Site-Name* page appears.

3. Click the add icon (+) on the **LAN** tab.

   The Add LAN Segment page appears.

4. Complete the configuration settings according to the guidelines provided in Table 45 on page 183.

NOTE: Fields marked with an asterisk (*) are mandatory.

5. Click **OK**.

   You are returned to the *Site-Name* page, where the LAN segment that you added is displayed.

Table 45: Add LAN Segment Settings

Field	Description
Name	Enter a name for the LAN segment. The name for a LAN segment should be a unique string of alphanumeric characters and some special characters (. -). No spaces are allowed and the maximum length is 15 characters.
Type **NOTE:** This field is displayed only for LAN segments associated with enterprise hub sites.	Select the type of LAN segment: • Directly Connected (default)—Indicates that the LAN segment is directly connected to the site. • Dynamic Routed—Indicates that the LAN segment is not directly connected to the site and is reachable by using a dynamic route. If you select this option, you must specify the dynamic routing information.
VLAN ID	Enter the VLAN ID for the LAN segment. Range: 2 through 4093.
Department	Select a department to which the LAN segment is assigned. Alternatively, click the **Create Department** link to create a new department and assign the LAN segment to it. See "Adding a Department" on page 767 for details. You can group LAN segments as departments for ease of management and for applying policies at the department-level. For LAN segments that are dynamically routed, you can assign only a data center department.
Protocol	For dynamically routed LAN segments, select the routing protocol (BGP or OSPF) to be used by the data center department to learn routes from the data center.

Table 45: Add LAN Segment Settings *(continued)*

Field	Description
Advertise LAN Prefix	For dynamically routed LAN segments, click the toggle button to advertise the LAN prefix of the SD-WAN spoke site to the data center through the data center department associated with the enterprise hub. By default, the advertise LAN prefix field is disabled. **NOTE:** You must avoid overlapping IP addresses between the SD-WAN LAN network and the datacenter network.
Gateway Address/Mask	Enter a valid gateway IP address and mask for the LAN segment. This address will be the default gateway for endpoints in this LAN segment. For example: 192.0.2.8/24.
DHCP	For directly connected LAN segments, click the toggle button to enable DHCP (default). You can enable DHCP if you want to assign IP addresses by using a DHCP server or disable DHCP if you want to assign a static IP address to the LAN segment. **NOTE:** If you enable DHCP, additional fields appear on the page.
Additional fields related to DHCP	
Address Range Low	Enter the starting IP address in the range of IP addresses that can be allocated by the DHCP server to the LAN segment.
Address Range High	Enter the ending IP address in the range of IP addresses that can be allocated by the DHCP server to the LAN segment.
Maximum Lease Time	Specify the maximum duration (in seconds) for which a client can request for and hold a lease on the DHCP server. Default: 1440 Range: 0 through 4,294,967,295 seconds.
Name Server	Specify one or more IPv4 addresses of the DNS server. To enter more than one DNS server address, type the address, press Enter, and then type the next address. **NOTE:** DNS servers are used to resolve hostnames into IP addresses.

Table 45: Add LAN Segment Settings *(continued)*

Field	Description

CPE Ports	• For sites with LAN capability, click the toggle button to include or exclude the CPE in the LAN segment.
	  ◦ When you include the CPE in the LAN segment:
	  ▪ CPE ports that you can include in the LAN segment are listed. Select the ports from the **Available** column and click the right-arrow to move the ports to the **Selected** column.
	  ▪ The **Switch Ports** field is disabled. CSO automatically assigns LAN ports on the Switch device and creates the same LAN segment on the Switch.
	  ◦ If you click to exclude the CPE from the LAN segment, you must specify the switch ports that connect with the LAN in the **Switch Ports** field. CSO automatically assigns LAN ports on the CPE device and creates the same LAN segment on the CPE device.
	  NOTE: You can select only one port if the CPE is a physical SRX Series device.
	• For sites without LAN capability, the CPE Ports field is disabled and the CPE ports that you can include in the LAN segment are listed.
	  Select the ports from the **Available** column and click the right-arrow to move the ports to the **Selected** column.

Switch Ports **NOTE:** This field is displayed only when LAN capability is selected for the enterprise hub.	If you disable the CPE ports field, select ports on the switch to be part of the LAN segment. The Switch ports and CPE ports are mutually exclusive. Select the ports from the **Available** column and click the right-arrow to move the ports to the **Selected** column.
BGP Configuration **NOTE:** This section is displayed only for dynamic routed LAN segments with BGP specified as the protocol.	
Authentication	Select the BGP route authentication method to be used: • None—Indicates that no authentication should be used. This is the default. • Use MD5—Indicates that MD5 is to be used for authentication. If you choose this option, you must specify an authentication key.
Peer IP Address	Enter the IP address of the BGP neighbor.
Peer AS Number	Enter the autonomous system (AS) number BGP neighbor.
Auth Key	If you specified that MD5 should be used for authentication, specify an MD5 authentication key (password), which is used to verify the authenticity of BGP packets.

Table 45: Add LAN Segment Settings *(continued)*

Field	Description
OSPF Configuration **NOTE:** This section is displayed only for dynamic routed LAN segments with OSPFspecified as the protocol.	
OSPF Area ID	Specify the OSPF area identifier to be used for the dynamic route.
Authentication	Select the OSPF route authentication method to be used: • Password—Indicates that password-based authentication should be used. If you choose this option, you must specify the password. (This is the default). • Use MD5—Indicates that MD5 is to be used for authentication. If you choose this option, you must specify an authentication key. • None—Indicates that no authentication should be used.
Password	Enter the password to be used to verify the authenticity of OSPF packets.
Confirm Password	Retype the password for confirmation purposes.
MD5 Auth Key ID	If you specified that MD5 should be used for authentication, enter the OSPF MD5 authentication key ID. Range: 1 through 255.
Auth Key	If you specified that MD5 should be used for authentication, enter an MD5 authentication key, which is used to verify the authenticity of OSPF packets.

Deploying LAN Segments

After you create a LAN segment and assign it to a department, you must deploy the LAN segment. You can deploy LAN segments from the *Site Name* page.

To deploy one or more LAN segments:

1. Click the **LAN** tab.

2. Select one or more LAN segments that you want to deploy and click **Deploy**.

   A Deploy LAN Segment job is created.

> **NOTE:** If a Deploy LAN Segment job is in progress for a site, wait for the job to finish before triggering another Deploy LAN Segment job.
>
> If you attempt to trigger a Deploy LAN segment job when another one is running, the job fails with a message indicating that the previous LAN segment deployment job is in progress.

3. Click **More** > **Deploy History** to view job status and deployment history of the LAN segment.

   The **Deploy LAN Segment History** page displayed.

   Alternatively, you can verify the status of the job from the **Monitor** > **Jobs** page.

Reassigning a Department to a LAN Segment

You can reassign the department assigned to a LAN segment from the *Site Name* page. .

> **NOTE:** Departments are not applicable for SD-LAN sites.

To reassign a department:

1. Click the **LAN** tab.

2. Select a LAN segment and click **Re-assign Department**.

   The Re-assign Department page appears.

   > **NOTE:** You cannot reassign a LAN segment that is already assigned to a department and is deployed.

3. Select the department to which the LAN segment is to be assigned.

4. Click **Deploy**.

   The success message **Re-assign department succeeded.** is displayed.

5. Click **OK**.

   The LAN segment with the newly assigned department is displayed on the tenant site page.

Deleting LAN Segments

You can delete a LAN segments from the *Site Name* page.

To delete a LAN segment:

1. Select a LAN segment and click the delete icon (**X**) icon on the **LAN** tab.

   The Delete LAN Segment page appears.

2. Click **OK** to confirm deletion.

   The LAN segment is deleted.

Managing a Single Site

You can use the **Sites** page to view the site details and to manage the site configurations for a single site. To access the page, click **Resources > Site Management** and then click the site that you want to manage.

You can perform the following tasks from this page:

- On the **Overview** tab, view detailed information about the tenant site, such as geographical location, connection details, device details, alarms, and alerts.

- On the **WAN** tab, view detailed information about the WAN links, such as topology of the hub-site WAN links, total number of hub and spoke links, total number of applications, link utilization details, link metrics based on throughput, and the maximum bandwidth capacity of a WAN link in a site. Hover over the WAN link to view bandwidth capacity.

  You can add or delete a mesh tunnel between a source site and destination site. For more information about creating and deleting mesh tunnels between a source site and destination site, see "Adding On-Demand Mesh Tunnels" on page 222 and "Deleting On-Demand Mesh Tunnels" on page 223.

  For sites owned by a tenant in a full mesh topology, you can view all the WAN link connections between WAN interfaces in all the sites. Click a site to see all connections between its WAN interfaces.

  NOTE: This tab is available only for SD-WAN sites.

- On the **Services** tab, view services, deploy network services, start a service, and disable services for a tenant site. You can also view the topology of the site.

To deploy a network service to a site, select the service, and then select an attachment point in the topology graphic. Alternatively, drag and drop the network service to an attachment point in the topology graphic.

> **NOTE:** This tab is available only for SD-WAN sites.

- On the **Policies** tab, view the following details:

  > **NOTE:** This tab is available only for SD-WAN sites.

  - List of all policies applicable to a tenant site. Click the policy name to view the rules that are applicable for the tenant site. Click the edit icon at the end of the row to edit a policy. You are taken to the **Configuration > Policy** page, where you can edit the policies.
  - Details about the tenant user who last updated the policy.
  - Time when the policy was last updated.
  - Deployment status of the policy—deployed or not deployed.
  - Number of rules applicable to the site compared to the total number of rules applicable to the tenant.
- On the **LAN** tab, view, create, deploy, and delete a LAN segment. In addition, you can use this tab to reassign a LAN segment to a different department. See "Managing LAN Segments on a Tenant Site" on page 182. You can also view all the devices in a LAN segment for an SD-LAN site and deploy one or more of these devices.

  > **NOTE:** Departments are not applicable for SD-LAN sites.

- On the **Devices** tab, view and manage the devices in a site. See "About the Devices Page" on page 268.

  > **NOTE:** If you have added only one Virtual Chassis to an SD-LAN site and want to delete the Virtual Chassis, delete the site from the Site Management page.
  >
  > If you have more than one Virtual Chassis in your site, you can delete the Virtual Chassis from the Devices tab.

To delete one or more EX Series switches from an SD-LAN site or to delete an EX Series switch from an SD-WAN site or Next-Generation Firewall site that has LAN capability:

a. Select the devices that you want to delete and click the Delete icon.

   The Delete Switch Options page appears.

b. (Default) Select the Load Recovery Configuration option

   The base configuration of the switches is restored.

c. Click **OK**.

   A job is initiated to delete the switches.

   If the job is successful, the switches are deleted from the site.

RELATED DOCUMENTATION

About the Sites Page | 53

Dynamic Mesh Tunnels Overview | 221

Viewing the Sites History

IN THIS SECTION

Viewing Jobs Initiated to Add and Configure Sites | 190

Viewing Jobs Initiated to Delete Sites | 191

Viewing Jobs Initiated to Add and Configure Sites

You can view the jobs initiated to add and configure all the sites in a tenant from the Sites page.

To view the jobs executed to add and configure sites:

1. Click **Resources > Site Management**.

   The Sites page appears.

2. On the Sites page, click **More > View Add Site History**.

   The Add/Configure Site History page appears. The Add/Configure Site History pane lists the jobs executed for adding and configuring all the sites in a tenant.

Table 46 on page 191 describes the fields displayed on the Add/Configure Site History page.

Table 46: Fields on the Add/Configure Site History Page

Field	Description
In Progress	Number of jobs in progress.
Success	Number of jobs that were successfully executed.
Failure	Number of jobs that failed during execution.
Name	Name of the job executed to add a site. Click the hyperlinked name to open the Site History Tasks page, where the tasks executed to complete the job are listed. Click the task link to open the Job Status page, where you can view the date and time when various tasks were executed to complete the job.
Start Date	Date and time that the job was initiated.
End Date	Date and time that the job finished executing. If the job is in-progress or failed, no date is displayed.
Status	Indicates whether the job completed successfully (Success) or not (Failed).
Log	Link to the log generated for the job. Click the link to view the logs on the Job Status page. The Job Status page displays the date and time when the job and various tasks associated with the job were executed in chronological order.

Viewing Jobs Initiated to Delete Sites

To view jobs initiated to delete sites from a tenant:

1. Click **Resources > Site Management**.

   The Sites page appears.

2. On the Sites page, click **More > View Delete History**.

The Delete History page appears. The Delete History page lists the all the jobs executed for deleting sites from a tenant.

Table 47 on page 192 details the fields displayed on the Delete Site History page.

Table 47: Fields on the Site Delete History Page

Field	Description
In Progress	Number of jobs in progress.
Success	Number of jobs that were successfully executed.
Failure	Number of jobs that failed during execution.
Name	Name of the job executed to delete a site. Click the hyperlinked name to open the Site History Tasks page, where the tasks executed to complete the job are listed. Click the task link to open the Job Status page, where you can view the date and time when various tasks were executed to complete the job.
Start Date	Date and time that the job was initiated.
End Date	Date and time that the job finished executing. If the job is in-progress or failed, no date is displayed.
Status	Indicates whether the job completed successfully (Success) or not (Failed).
Log	Link to the log generated for the job. Click the link to view the logs on the Job Status page. The Job Status page displays the date and time when the job and various tasks associated with the job were executed in chronological order.

SEE ALSO

About the Sites Page | 53

Deleting a Site | 194

Edit Site Properties

You, as a tenant administrator, can modify the following properties configured for a site from the Sites page (**Resources > Site Management**):

- Address and Contact Information—Street Address, City, State/Province, ZIP/Postal Code, Country, Contact Name, Email, and Phone Number.
- Advanced Configuration—Name Server IP List, NTP Server, and Timezone.
- In-band Management Port (available only for sites with Next-generation Firewall capability).

To modify the properties configured for a site:

1. Select **Resources > Site Management**.

   The Sites page appears.

2. Select the site whose properties you want to modify and click the **Edit** icon (pencil).

   The Edit *Site-Name* page appears, displaying the same fields that are presented when you add a site.

3. Modify the site properties as needed. For more information on these properties, see "Manually Adding On-Premise Spoke Sites" on page 95.

4. Click **OK** to save the changes or click **Cancel** to discard the changes.

   If you click OK, the changes that you made for the site are saved. A job to edit the site is triggered and the job link appears on the Sites page.

5. (Optional) You can click the job link to view details of the job (including job status, start date and time, and end date and time) on the **Site Edit Details** page. Alternatively, you can view the status of the job on the Jobs (**Monitor > Jobs**) page.

   If the job is completed successfully, a confirmation message appears on top of the Sites page.

Site Type	Editable Fields
On-premise spoke site	• Address and Contact Information—Street Address, City, State/Province, ZIP/Postal Code, Country, Contact Name, Email, and Phone Number. • Advanced Configuration—Name Server IP List, NTP Server, and Timezone. • In-band Management Port (available only for sites with Next-generation Firewall capability).
Cloud spoke site	Advanced Configuration—Name Server IP List, NTP Server, and Timezone.

Site Type	Editable Fields
Enterprise hub site	• Address and Contact Information—Street Address, City, State/Province, ZIP/Postal Code, Country, Contact Name, Email, and Phone Number. • Advanced Configuration—Name Server IP List, NTP Server, and Timezone.

RELATED DOCUMENTATION

About the Sites Page | 53

Adding an On-Premise Spoke Site with SD-WAN Capability | 100

Deleting a Site

You can delete a site from the Sites page (**Resources > Site Management**) in the Customer Portal.

The following are applicable when you want to delete a site:

- You can delete only one site at a time.
- You cannot delete a gateway or hub site when spoke sites are connected to it. You must first delete the spoke sites associated with a gateway site or hub site and then delete the gateway or hub site.
- You cannot delete a site that is associated with site-specific policies. You must first delete the intents in the policy that are associated with the site, redeploy the policy, and then delete the site.
- You must delete service instances before attempting to delete the site.

To delete a site:

1. Click **Resources > Site Management**.

   The Sites page appears.

2. Select the site that you want to delete and click the delete (trash can) icon.

   The Remove Site Options dialog box appears.

3. Select one of the following options:

   - **Load Recovery Configuration**—Use this option to back up any custom configuration on the device. CSO restores the custom configuration when you reinstall the device.
   - **Zeorize the device**—Use this option to reset the device to the factory default configuration.

4. Click **OK**.

A job to remove the site is created. After the job completes successfully, a message appears on top of the Sites page indicating that the site is deleted.

If the delete site operation is successful, configurations associated with the site are deleted from the hub, a peer site, and virtual Route Reflectors (VRRs).

NOTE:

- The time taken to delete a site is dependant on the device template attached to the site.
- If the delete site operation is unsuccessful, you cannot add a new site using the same site name.

RELATED DOCUMENTATION

About the Sites Page | 53

Contrail Service Orchestration Monitoring and Troubleshooting Guide

CHAPTER 4

Managing Site Groups

IN THIS CHAPTER

- About the Site Groups Page | 196
- Creating Site Groups | 197

About the Site Groups Page

To access this page, click **Resources > Site Groups**.

You can use the **Site Groups** page to view, create, and delete site groups for a tenant. Site groups enable you to group sites logically, thereby easing site management. You can use site groups to apply policies at the site group level.

You must be a Tenant Administrator user to access the **Site Groups** page.

Tasks You Can Perform

You can perform the following tasks from this page:

- View details of existing site groups. Click the details icon that appears when you hover over the name of a site group or select **More > Detailed View**.
- Create site groups. See "Creating Site Groups" on page 197.
- Edit site groups. Select a site group and click the edit icon.
- Delete site groups. To delete a site group, select it on the Site Groups page and click the delete (**X**) icon.

Field Descriptions

Table 48 on page 197 shows the descriptions of the fields on the **Site Groups** page.

Table 48: Fields on the Site Groups Page

Field	Description
Name	Displays the name of the site group.
Sites	Displays the names of the sites that are members of a site group.

RELATED DOCUMENTATION

Creating Site Groups | 197

Creating Site Groups

You can use the **Create Site Group** page to create a new site group for a tenant and add sites to it.

To create a site group:

1. Click **Resources > Site Groups**.

   The Site Groups page appears.

2. Click the add icon (**+**).

   The **Create Site Group** page appears.

3. Enter a unique name for the site group.

4. From the list of sites in the **Available** column, select the sites that you want to include in the new group and click the greater-than icon (**>**).

   The selected sites are moved to the **Selected** column.

5. Click **OK**. If you want to discard your changes, click **Cancel** instead.

   The new site group is displayed on the **Site Groups** page.

RELATED DOCUMENTATION

About the Site Groups Page | 196

CHAPTER 5

Managing Site Templates

IN THIS CHAPTER

- About the Site Templates Page | 198
- Adding On-Premise Spoke Sites by Using a Site Template | 199
- Cloning, Editing, and Deleting Site Templates | 201
- Adding a Site Template | 203
- Adding and Configuring Sites by Importing a JSON File | 215

About the Site Templates Page

To access this page, click **Resources > Templates > Site Templates**.

You can use the Site Templates page to add site templates, edit, clone, delete, and view existing site templates.

A site template enables you to specify values for many of the attributes used to add a site. You can then use the site template to add multiple sites that share the same set of values (for attributes already specified in the site template) and only specify values for site-specific attributes.

Site templates are available only for on-premise spoke sites.

Tasks You Can Perform

You can perform the following tasks from this page:

- Add a site template. See "Adding a Site Template" on page 203
- Edit, clone, or delete site templates. See "Cloning, Editing, and Deleting Site Templates" on page 201.
- View site template details. The site templates are displayed in card format. For each site template you can view the following details:
  - Template name
  - Published by
  - Capability type

- Number of devices attached
- Number of WAN links, if applicable
- Number of sites that the device template is attached to
- Date and time that the template was last updated

Adding On-Premise Spoke Sites by Using a Site Template

Using a site template, you can add on-premise spoke sites in bulk either manually or by importing the JavaScript Object Notation (JSON) file that contains on-premise site attributes.

To add on-premise spoke sites by using a site template:

1. Select **Resources > Site Management**.

   The Sites page appears.

2. Click **Add** and select **Add On-Premise Spoke Site (Using Site Template)**.

   The Add Spoke Site page appears listing the existing site templates.

3. Select the site template and click **Continue**.

   The Add Spoke Site page appears.

4. The **Site Template** field displays the name of the site template that you have selected. If you want to change the site template, click the **Change** link and select another site template of your preference from the Add Spoke Site page.

5. Do one of the following to add on-premise spoke sites:
   - To add on-premise spoke sites in bulk by importing the JSON file:
     a. Select **Import from file** in the Site Data field.
     b. (Optional) Click **Download sample JSON file** to download a sample JSON template and use it to specify site data that you can later import.
     c. Click **Browse** to upload the JSON file.
     d. Navigate to the folder and select the JSON file.
     e. Click **Open**.
   - To manually add on-premise spoke sites in bulk, select **Add Manually** in the Site Data field.

The Site 0 tab appears listing the fields based on the capabilities that were selected for the site template. For example, if the site template that you selected has only LAN capability, only general information fields and fields related to LAN capability are displayed.

6. Complete the configuration for Site0.

   For more information on the fields for adding an on-premise spoke site with the following capabilities:

   - WAN capability as SD-WAN, see "Adding an On-Premise Spoke Site with SD-WAN Capability" on page 100.
   - WAN capability as Hybrid WAN, see "Adding an On-Premise Spoke Site with Hybrid WAN Capability" on page 95.
   - WAN capability as Next Gen Firewall, see "Adding a Standalone Next Generation Firewall Site" on page 177.
   - WAN capability as SD-WAN and LAN capability, see "Adding an On-Premise Spoke Site with SD-WAN and LAN Capabilities" on page 117.
   - WAN capability as Next Gen Firewall and LAN capability, see "Adding an On-Premise Spoke Site with Next Generation Firewall and LAN Capabilities" on page 151.
   - Only LAN capability, see "Add an On-Premise Spoke Site with LAN Capability" on page 136.

7. Click the plus icon (+) to add more sites and complete the configuration for each site.

8. Review the sites.

   If there are validation errors, an error icon appears in the left pane (next to the site name). You must ensure that all errors are resolved before proceeding.

9. (Optional) You can remove a site by clicking the **X** icon when you hover over the site name in the left pane.

10. Click **Save**.

    A confirmation message is displayed indicating that the job is created for adding sites in bulk.

RELATED DOCUMENTATION

Adding a Site Template | **203**

About the Site Templates Page | **198**

Cloning, Editing, and Deleting Site Templates | **201**

Cloning, Editing, and Deleting Site Templates

IN THIS SECTION

- Cloning Site Templates | 201
- Editing Site Templates | 202
- Deleting Site Templates | 202

You can clone, edit, or delete a site template.

Cloning Site Templates

You can clone a site template when you want to quickly create a copy of an existing site template.

To clone a site template:

1. Select **Resources > Templates > Site Templates**.

   The Site Template page appears.

2. Select the site template that you want to clone.

3. Click **Clone**.

   The Clone Site Template page appears.

4. Specify a unique name for the site template that can contain alphanumeric characters and hyphens (-); the maximum length is 32 characters.

5. Click **OK** to save your changes.

   A confirmation message appears, indicating the status of the clone operation.

   The site template that you have cloned is displayed on the Site Templates page. You can modify the cloned site template as needed by clicking the edit icon.

Editing Site Templates

> **WARNING:** You cannot edit a site template if you have associated the site template with a site.

To edit a site template:

1. Select **Resources > Templates > Site Templates**.

   The Site Template page appears.

2. Select the site template that you want to edit.

3. Click the edit icon (pencil) to modify the parameters.

   The Edit Site Template page appears.

4. Edit the fields, as needed. Modify the fields according to the guidelines provided in "Adding a Site Template" on page 203.

   > **NOTE:** You cannot edit the site template name.

5. Click **OK** to save your changes.

   A confirmation message appears, indicating the status of the edit operation. You can use the modified site template for creating multiple on-premise spoke sites.

Deleting Site Templates

> **WARNING:** You cannot delete a site template if you have associated the site template with a site.

To delete a site template:

1. Select **Resources > Templates > Site Templates**.

   The Site Template page appears.

2. Click the site template that you want to delete and click the delete icon.

An alert message appears, asking you to confirm the delete operation.

3. Click **Yes** to delete the site template.

   A confirmation message appears, indicating the status of the delete operation.

RELATED DOCUMENTATION

- About the Site Templates Page | **198**
- Adding a Site Template | **203**
- Adding On-Premise Spoke Sites by Using a Site Template | **199**

Adding a Site Template

You can add a site template for an on-premise spoke site. A site template can be added with one WAN capability (SD-WAN, or Hybrid WAN, or Next Gen Firewall), LAN capability, or both WAN and LAN capabilities.

> **NOTE:** If you select the WAN capability as Hybrid WAN you cannot select the LAN capability.

To add a site template:

1. Select **Resources > Templates > Site Templates**.

   The Site Templates page appears.

2. Click the plus icon (+).

   The Add Site Template page appears.

3. Complete the configuration according to the guidelines in Table 49 on page 204.

   The fields that are displayed in the Add Site Template page are based on the LAN and WAN capabilities that you choose. The last column of Table 49 on page 204 indicates the capabilities for which a field is applicable.

NOTE: Fields marked with * are mandatory.

4. Click **OK**.

The site template is added and listed in the Site Templates page. You can use the site template to add multiple on-premise spoke sites.

Table 49: Fields on the Add Site Template Page

Field	Description	Applicable To
General Tab		
Template Name	Specify a unique name for the site template that can contain alphanumeric characters and hyphens (-); the maximum length is 32 characters.	• SD-WAN • Hybrid-WAN • Next Gen Firewall • LAN
Description	Enter a description for the site template; the maximum length is 512 characters.	• SD-WAN • Hybrid-WAN • Next Gen Firewall • LAN
Site Information		
Site Group	Select a site group to which you want to assign the template. Example: sdwan-spoke	• SD-WAN • Hybrid-WAN • Next Gen Firewall • LAN
Site Capabilities		

Table 49: Fields on the Add Site Template Page *(continued)*

Field	Description	Applicable To
WAN Capabilities	Select one of the following WAN capabilities to include LAN capabilities for the site template: • SD-WAN • Hybrid-WAN • Next Gen Firewall **NOTE:** • You must select at least one WAN capability or a LAN capability. • The WAN capabilities that are displayed here are filtered based on the service type that are assigned to the tenant.	• SD-WAN • Hybrid-WAN • Next Gen Firewall
LAN Capabilities	Select **LAN** to include LAN capabilities for the site template. This field is disabled if the WAN capability is Hybrid WAN.	LAN
Configuration		
Primary Provider Hub	Select the provide hub site (or primary provider hub site in case of multihoming) to which the spoke site must connect.	SD-WAN
Secondary Provider Hub	Select the secondary provider hub site to which this site must connect. This site connects to the secondary provider hub site when the primary provider hub is down.	SD-WAN
Primary Enterprise Hub	Select the primary enterprise hub with which you want to connect the spoke site. If you specify a enterprise hub, then the initial site-to-site traffic as well as the central breakout (backhaul) traffic (if applicable) is sent through the enterprise hub instead of the hub site.	SD-WAN

Table 49: Fields on the Add Site Template Page *(continued)*

Field	Description	Applicable To
Secondary Enterprise Hub	Select the secondary enterprise hub for this spoke site. The spoke site connects with secondary enterprise hub when the primary enterprise hub is down.	SD-WAN
Create Threshold	Enter the maximum number of sessions closed between the connected sites in a duration of two minutes at which full mesh is created between the two sites. The default value is 5. For example, if you specify the number of sessions as 5, dynamic mesh tunnels are created if the number of sessions closed between two spoke sites in 2 minutes exceeds 5.	SD-WAN
Delete Threshold	Enter the number of sessions closed between the connected sites in a duration of 15 minutes below which full mesh is deleted between the two sites. The default value is 2. For example, if you specify the number of sessions closed as 2, dynamic mesh tunnels are deleted if the number of sessions closed is lesser than or equal to 2.	SD-WAN
Address and Contact Information		
Street Address	Enter the street address of the site.	• SD-WAN • Hybrid-WAN • Next Gen Firewall • LAN
City	Enter the city where the site is located.	• SD-WAN • Hybrid-WAN • Next Gen Firewall • LAN

Table 49: Fields on the Add Site Template Page *(continued)*

Field	Description	Applicable To
State/Province	Select the state or province where the site is located.	• SD-WAN • Hybrid-WAN • Next Gen Firewall • LAN
ZIP/Postal Code	Enter the postal code for the site.	• SD-WAN • Hybrid-WAN • Next Gen Firewall • LAN
Country	Select the country where the site is located. Click the **Validate** button to verify the address. The **site address verification successful** message is displayed if the address is correct. You can click the **View location on a map** link to see the address location. If you enter the wrong address and click the **Validate** button to verify the address, the **Site address could not be validated message** is displayed .	• SD-WAN • Hybrid-WAN • Next Gen Firewall • LAN
Contact Name	Enter the name of the contact person at the site.	• SD-WAN • Hybrid-WAN • Next Gen Firewall • LAN
Email	Enter the e-mail address of the contact person at the site.	• SD-WAN • Hybrid-WAN • Next Gen Firewall • LAN
Phone	Enter the phone number for the site.	• SD-WAN • Hybrid-WAN • Next Gen Firewall • LAN
Advanced Configuration		

Table 49: Fields on the Add Site Template Page *(continued)*

Field	Description	Applicable To
Name Server IP List	Specify one or more IPv4 addresses of the DNS server. To enter more than one DNS server address, type the address, press Enter, and then type the next address, and so on. DNS servers are used to resolve hostnames into IP addresses.	• SD-WAN • Hybrid-WAN • Next Gen Firewall • LAN
NTP Server	Specify the fully qualified domain names (FQDNs) or IP addresses of one or more NTP servers. Example: ntp.example.net The site must have DNS reachability to resolve the FQDN during site configuration.	• SD-WAN • Hybrid-WAN • Next Gen Firewall • LAN
Select Timezone	Select the time zone in which the site is located from the drop-down list.	• SD-WAN • Hybrid-WAN • Next Gen Firewall • LAN
WAN Tab		
Device Template		
Device Series	Select the device series to which the CPE belongs (SRX, NFX150, or NFX250) and select a device template for the selected device series. The device template contains information for configuring a device.	• SD-WAN • Hybrid-WAN
Device Model	For NFX150 devices, select a device model from the list. Device models are listed based on the connection plan that you select.	• SD-WAN • Hybrid-WAN

Table 49: Fields on the Add Site Template Page *(continued)*

Field	Description	Applicable To
Auto Activate	Click the toggle button to enable or disable automatic activation of the CPE when the CPE is detected by CSO (management status of the device is Device_Detected). When you enable this field, zero-touch provisioning of the device is automatically triggered after the site with the CPE is added to CSO.	• SD-WAN • Hybrid-WAN • Next Gen Firewall
Pre Staged	Click the toggle button to use the preconfigured settings for the firewall device. The preconfigured settings are as follows: • Device Template—NGSRXZTP • In-band Management Port—ge-0/0/0 port • Firewall Policies—Factory_Default_Fw_Policy • NAT Policies—Factory_Default_NAT_Policy	Next Gen Firewall
Boot Image	Select the boot image from the drop-down list if you want to upgrade the image for the CPE device. The boot image is the latest build image uploaded to the image management system. The boot image is used to upgrade the device when the CSO starts the ZTP process. If the boot image is not provided, then the device skips the procedure to upgrade the device image. The boot image (NFX or SRX) is populated based on the device template that you have selected while adding a site. See *Uploading a Device Image*.	• SD-WAN • Hybrid-WAN • Next Gen Firewall
In-band Management Port	Select the port that you want to configure as management interface and connect it to the management device. You can configure any of the ge-0/0/*x* ports, where x ranges from 0 to 14, as in-band management interfaces.	Next Gen Firewall

Table 49: Fields on the Add Site Template Page *(continued)*

Field	Description	Applicable To
Firewall Policies	Select the firewall policy that you want to deploy. The firewall policy list is populated from the **Configuration > Firewall > Firewall Policy** page.	• SD-WAN • Hybrid-WAN • Next Gen Firewall
NAT Policies	Select the NAT policy that you want to deploy to the standalone firewall site. The NAT policy list is populated from the **Configuration > NAT > NAT Policies** page.	• SD-WAN • Hybrid-WAN • Next Gen Firewall
CPE AS Number	Specify the autonomous system (AS) number for the CPE device on the site.	Hybrid-WAN
Router Name	Specify the router name.	Hybrid-WAN
Router AS Number	Specify the AS number for the router in the point of presence (POP).	Hybrid-WAN
WAN 0	Click the toggle button to enable or disable this WAN link. By default, the WAN_0 link is enabled. When you enable a WAN link, fields related to the WAN link appear. Fields marked with an asterisk (*) must be configured to proceed.	• SD-WAN • Hybrid-WAN
Link Type	Select the underlay network type (MPLS or Internet) of the WAN link that is connected to the on-premise spoke site.	• SD-WAN • Hybrid-WAN
Egress Bandwidth	Enter the maximum bandwidth (in mega bits per second [Mbps]) to be allowed for the WAN link. Range: 1 through 10,000	SD-WAN
Address Assignment	Select the method for IP address assignment. The options available are: • DHCP—Select DHCP to assign IP address by using a DHCP server. • STATIC—Select STATIC to assign a static IP address.	SD-WAN

Table 49: Fields on the Add Site Template Page *(continued)*

Field	Description	Applicable To
Advanced Settings		
Provider	Enter the name of the service provider who is responsible for providing the WAN link.	SD-WAN
Cost/Month	Enter the cost per month (in the specified currency) of the subscribed bandwidth. Range: 1 through 10,000	SD-WAN
Enable Local Breakout	Click the toggle button to enable local breakout on the WAN link. By default, local breakout is disabled.	SD-WAN
Use For Fullmesh	Click the toggle button to specify that the WAN link is part of a fullmesh topology.	SD-WAN
Connects To Hubs	Click the toggle button to specify that the WAN link of the site connects to a hub. **NOTE:** • For sites with a single CPE, you must enable at least one WAN link to connect to the hub so that OAM traffic can be transmitted. • For sites with a dual CPE, you must enable at least one WAN link per device to connect to the hub so that OAM traffic can be transmitted.	SD-WAN
Backup Link	Select a backup link through which traffic can be routed when the primary (other) links are unavailable.	SD-WAN
Default Link	Select one or more links to be used for routing traffic in the absence of matching SD-WAN policy intents.	SD-WAN
Data VLAN Id	Enter the VLAN ID that is associated with the data link. A data VLAN identifier is an integer. Range: 0 through 65,535	SD-WAN

Table 49: Fields on the Add Site Template Page *(continued)*

Field	Description	Applicable To
VLAN ID	Specify the identifier for the Layer 2 VLAN for the CPE device.	Hybrid-WAN
WAN 1	Click the toggle button to enable or disable this WAN link. By default, the WAN 1 link is disabled. Refer to the fields described for WAN 0 for an explanation of the fields.	• SD-WAN • Hybrid-WAN
WAN 2	Click the toggle button to enable or disable this WAN link. By default, the WAN 2 link is disabled. Refer to the fields described for WAN 0 for an explanation of the fields	SD-WAN
WAN 3	Click the toggle button to enable or disable this WAN link. By default, the WAN 3 link is disabled. Refer to the fields described for WAN 0 for an explanation of the fields	SD-WAN
LAN Tab		
Device Type	Select the type of switch—EX2300, EX3400, EX4300, EX4600, and EX4650.	LAN
Device Model	Select the model for the switch that you chose in the Device Type.	LAN
LAN Segments	Displays the LAN segment on the switch. To add a LAN segment, click the + icon on the top, right corner of the grid. The Add LAN Segment page appears. See Table 50 on page 213.	LAN

Table 49: Fields on the Add Site Template Page *(continued)*

Field	Description	Applicable To
Auto Activate Switch	Click the toggle button to enable or disable automatic activation of the switch when the switch is detected by CSO (that is, management status of the device is Device_Detected). When you enable this field, zero-touch provisioning of the switch is automatically triggered after the site with the switch is added to CSO. **NOTE:** The device template that you select determines whether this option is enabled or disabled by default.	LAN

Table 50: Fields on the Add LAN Segment Page

Field	Description
Add LAN Segment	
Name	Enter a name for the LAN segment. The name for a LAN segment should be a unique string of alphanumeric characters and some special characters (. -). No spaces are allowed and the maximum length is 15 characters.
VLAN ID	Enter the VLAN ID for the LAN segment. Range: 2 through 4093.
Department	Select a department to which the LAN segment is to be assigned. Alternatively, click the **Create Department** link to create a new department and assign the LAN segment to it. See “Adding a Department” on page 767 for details. You group LAN segments as departments for ease of management and for applying policies at the department-level.
Gateway Address/Mask	Enter a valid gateway IP address and mask for the LAN segment; for example, 192.0.2.8/24.

Table 50: Fields on the Add LAN Segment Page *(continued)*

Field	Description
DHCP	For directly connected LAN segments, click the toggle button to enable DHCP. DHCP is disabled by default. You enable DHCP if you want to assign IP addresses by using a DHCP server. You disable DHCP if you want to assign a static IP address to the LAN segment. **NOTE:** If you enable DHCP, fields related to DHCP-related parameters appear and must be configured.
[DHCP-Related Fields]	
Address Range Low	Enter the starting IP address in the range of IP addresses that can be allocated by the DHCP server to the LAN segment.
Address Range High	Enter the ending IP address in the range of IP addresses that can be allocated by the DHCP server to the LAN segment.
Maximum Lease Time	Specify the maximum duration (in seconds) for which a client can request for and hold a lease on a DHCP server. Range: 0 through 4,294,967,295.
Name Server	Specify or select one or more IPv4 addresses of the DNS server. To enter more than one DNS server address, type the address, press Enter, and then type the next address, and so on. DNS servers are used to resolve hostnames into IP addresses.
CPE Ports	Click the toggle button to include or exclude the CPE in the LAN segment. When you include the CPE in the LAN segment: • CPE ports that you can include in the LAN segment are listed. Select the ports from the **Available** column and click the right-arrow to move the ports to the **Selected** column. • The **Switch Ports** field is disabled. CSO automatically assigns LAN ports on the switch device and creates the same LAN segment on the switch. If you exclude the CPE from the LAN segment, you must specify the switch ports that connect with the LAN in the **Switch Ports** field. CSO automatically assigns LAN ports on the CPE device and creates the same LAN segment on the CPE device. **NOTE:** You can select only one port if the CPE is an SRX Series device.

Table 50: Fields on the Add LAN Segment Page *(continued)*

Field	Description
Switch Ports	If you disable the CPE ports field, select ports on the switch that will be part of the LAN segment. Select the ports from the **Available** column and click the right-arrow to move the ports to the **Selected** column.

RELATED DOCUMENTATION

- About the Site Templates Page | **198**
- Cloning, Editing, and Deleting Site Templates | **201**
- Adding On-Premise Spoke Sites by Using a Site Template | **199**

Adding and Configuring Sites by Importing a JSON File

You can add and configure one or more sites in CSO by uploading a JavaScript Object Notation (JSON) file that contains the parameters for adding and configuring the sites.

Before you begin, ensure that the JSON file contains all the parameters required for each site that you want to add or configure.

> **NOTE:** We recommend that you create the JSON file by referring to the sample JSON file provided by CSO. Refer to Step 3 of the procedure for for downloading and modifying the sample JSON file.

To add and configure multiple sites by uploading a JSON file:

1. Click **Resources > Site Managemnet**.

   The Sites page appears.

2. Click **Add** and select **Import Sites**.

   The **Import Sites** page appears.

3. (Optional) Download a sample JSON file by clicking the **Download Sample JSON** link. Edit the parameters based on your requirements and save the file.

4. Click **Browse** and navigate to the directory that contains the JSON file.

5. Select the file and click **Open**.

6. Click **Import**.

   A message is displayed indicating that the file is imported to CSO. After the file is imported to CSO successfully, you are returned to the Sites page, where you can view the newly added sites:

   - If you enabled automatic activation for the sites, CSO activates and provisions the sites and sets the status of the sites to PROVISIONED.
   - If you did not enable automatic activation, CSO sets the status of the sites to CONFIGURED. You must manually activate the sites after which CSO provisions the sites. See "Activate a Device" on page 229 or "Activating Dual CPE Devices (Device Redundancy)" on page 254.

After the sites are activated and provisioned, you can install certificates and create policies for the sites. See the *Managing Policies, Profiles, and Proxies* section in this guide for details.

RELATED DOCUMENTATION

About the Sites Page | **53**

Manually Adding On-Premise Spoke Sites | **95**

CHAPTER 6

Managing Mesh Tags

IN THIS CHAPTER

- Mesh Tags Overview | 217
- About the Mesh Tags Page | 218
- Creating User-defined Mesh Tags | 219

Mesh Tags Overview

A mesh tag is a label that you associate with a WAN link of a spoke site. Mesh tags provide you the flexibility to establish overlay tunnels between WAN links of two different spoke sites. If WAN links are associated with same mesh tags, CSO creates a VPN tunnel between WAN links of spoke sites (enterprise hub to enterprise hub, on-premise spoke site to enterprise hub, on-premise spoke site to on-premise spoke site).

> **NOTE:** Mesh tags are applicable only for SD-WAN sites in the Real-time optimized mode (Full mesh).

Mesh tags can be predefined (MPLS and Internet) or user-defined. You can create user-defined mesh tags on the **Administration > Mesh Tags** page.

> **NOTE:** With mesh tags, you can connect two WAN links even if the link types (MPLS and Internet) are different.

For example, consider that a tenant has two sites—Site A and Site B. Site A has four WAN links (WAN_A0 through WAN_A3) and Site B has three WAN links (WAN_B0 through WAN_B2). WAN_A0 and WAN_B0 are associated with MPLS (predefined mesh tag), and WAN_A1 and WAN_B1 are associated with Internet (predefined mesh tag).

A tunnel is established between WAN-A0 and WAN-B0 because they are associated with the same predefined mesh tags.

NOTE: You can associate mesh tags for up to three WAN links.

RELATED DOCUMENTATION

About the Mesh Tags Page | 218
Creating User-defined Mesh Tags | 219

About the Mesh Tags Page

To access this page, select **Resources > Mesh Tags** in Customer Portal.

Use this page to view predefined mesh tags and create user-defined mesh tags. Mesh tags connect WAN links of different sites. CSO creates a tunnel between WAN links if they are associated with the same mesh tag (user-defined or predefined).

Tasks You Can Perform

You can perform the following tasks from this page:

- Create a mesh tag—See "Creating User-defined Mesh Tags" on page 219.
- Delete a mesh tag. Select a mesh tag and click the delete icon.

Field Descriptions

Table 51 on page 218 describes the fields on the Mesh Tags page.

Table 51: Fields on the Mesh Tag Page

Field	Description
Name	Displays the name of the mesh tag.
Tenant	Displays the tenant name to which the site is associated.
Type	Displays whether the mesh tag is a predefined or a user-defined.

RELATED DOCUMENTATION

Mesh Tags Overview | 217

Creating User-defined Mesh Tags

CSO creates a tunnel between WAN links of two different sites if they are associated with the same mesh tag (user-defined or predefined).

To create a user-defined mesh tag:

1. Select **Resources > Mesh Tags**.

   The Mesh Tags page appears.

2. Click the add icon (+) to create a mesh tag.

   The Create New Mesh Tag page appears.

3. Complete the configuration according to the guidelines provided in Table 52 on page 219.

   NOTE: Fields marked with * are mandatory.

4. Click **OK** to save the changes.

   The user-defined mesh tag is created.

   If you want to discard your changes, click **Cancel** instead.

Table 52: Fields on the Create New Mesh Tag Page

Field	Description
Name	Enter a unique name for the mesh tag.
Description	Enter a description for the mesh tag.

RELATED DOCUMENTATION

About the Mesh Tags Page | 218

CHAPTER 7

Managing Dynamic Mesh

IN THIS CHAPTER

- Dynamic Mesh Tunnels Overview | **221**
- Adding On-Demand Mesh Tunnels | **222**
- Deleting On-Demand Mesh Tunnels | **223**

Dynamic Mesh Tunnels Overview

In releases earlier than CSO 4.1.0, all static tunnels are established between spoke sites during the Zero Touch Provisioning (ZTP) process.

However starting with Release 4.1.0, during ZTP, only the following static tunnels are established:

- Between an on-premise spoke site and the corresponding enterprise hub (primary enterprise hub or secondary enterprise hub)
- Between an on-premise spoke site and the provider hub (primary provider hub or secondary provider hub)
- Between two enterprise hubs

Therefore, the communication between two on-premise spoke sites is established only through the enterprise hub or the provider hub.

CSO dynamically creates or deletes a mesh tunnel (without passing through an enterprise hub or a provider hub) between two spoke sites, if:

- The number of sessions closed between two spoke sites crosses the configured threshold value, and
- The WAN links of spoke sites have matching mesh tags. For more information, see "Mesh Tags Overview" on page 217.

> **NOTE:** The dynamic mesh feature is applicable only for SD-WAN sites in Real Time-Optimized mode (Full mesh).

The tenant administrator can modify the default threshold value on the following pages:

- The **Administration > Tenant Settings** page (Dynamic Mesh section) of Customer Portal (global level)
- The Add On-Premise Spoke Site page (site-level)
- The Add Enterprise Hub page (site-level)

The threshold value that you specify at site-level takes precedence over the global-level threshold values.

That is, the threshold value that you specify on the Add Site page (on-premise or enterprise hub) overrides the threshold value that you specified on the Tenant Settings page of Customer Portal.

CSO allows you to manually create or delete dynamic mesh tunnels between a source site and a destination site by using the Add On-Demand Mesh Tunnel or Delete On-Demand Mesh Tunnel pages in Customer Portal.

RELATED DOCUMENTATION

View and Edit Tenant Settings | 20

Adding On-Demand Mesh Tunnels

An OpCo administrator or the Tenant administrator can create a mesh tunnel between a source site and destination site by using the CSO GUI in Customer Portal.

To create a mesh tunnel between the source site and a destination site:

1. Select **Resources > Site Management.**

   The Sites page appears.

2. Click the site to which you want to add mesh tunnels.

   The *Site Name* page appears.

3. On the WAN tab, click **+**.

   The Add On-Demand Mesh Tunnel page appears.

4. Complete the configuration according to the guidelines Table 53 on page 222.

5. Click **Ok** to save the changes. If you want to discard the changes, click **Cancel** instead.

   A tunnel is created between the source site and a destination site.

Table 53: Fields on On-demand VPN Tunnel page

Field	Description
Add Tunnel Threshold	
Source Site	Displays the name of the source site.
Destination Site	Select the destination site from the list.
Delete Tunnel	
Displays the threshold value (sessions closed) below which a tunnel is deleted between two sites.	

Table 53: Fields on On-demand VPN Tunnel page *(continued)*

Field	Description
Enable Threshold	By default, this toggle button is disabled. That is, even if the threshold value (sessions closed) is met, CSO does not delete mesh tunnels. You have to manually delete the mesh tunnel that you created.
Sessions Closed	Displays the number of sessions closed for 15 minutes.

RELATED DOCUMENTATION

Dynamic Mesh Tunnels Overview | 221

Deleting On-Demand Mesh Tunnels

A user with either OpCo administrator or Tenant administrator roles can delete a VPN tunnel between a source site and destination site by using the CSO GUI in Customer Portal.

To delete a mesh tunnel between a source site and a specific destination site:

1. Select **Resources > Site Management.**

   The Sites page appears.

2. Click the site from which you want to delete mesh tunnels.

   The *Site Name* page appears.

3. On the WAN tab, click **+**.

   The Delete On-Demand Mesh Tunnel page appears.

4. Complete the configuration according to the guidelines Table 54 on page 224.

5. Click **Ok** to save the changes. If you want to discard the changes, click **Cancel** instead.

   A VPN tunnel is deleted between the source site and a specific destination site.

Table 54: Fields on the Delete On-Demand Mesh Tunnel page

Field	Description
Delete Tunnel Threshold	
Source Site	Displays the name of the source site.
Destination Site	Select the destination site from the list.
Delete Tunnel Displays the threshold value (sessions closed) below which a tunnel is deleted between two spokes.	
Enable Threshold	By default, this toggle button is disabled. That is, even if the threshold value (sessions closed) is met, CSO does not create mesh tunnels. You have to manually create the mesh tunnel that you deleted.
Sessions Closed	Displays the number of sessions closed for 2 minutes. .

RELATED DOCUMENTATION

Dynamic Mesh Tunnels Overview | 221

3

PART

Managing Devices and Resources

Managing Devices | **226**

Managing Device Images | **263**

Managing Resources | **266**

Managing Configuration Templates | **316**

Managing Licenses | **342**

Managing Signature Database and Certificates | **346**

Managing Juniper Identity Management Service | **363**

CHAPTER 8

Managing Devices

IN THIS CHAPTER

- Device Redundancy Support Overview | **227**
- Activate a Device | **229**
- Manually Activating a Switch | **232**
- Manage an EX Series Switch | **233**
- Managing Ports on an EX Series Switch | **245**
- Activating Dual CPE Devices (Device Redundancy) | **254**
- Viewing the History of Tenant Device Activation Logs | **257**
- Zero Touch Provisioning Overview | **259**
- Workflow for Onboarding a Device Using ZTP | **260**

Device Redundancy Support Overview

Contrail Service Orchestration (CSO) provides support for spoke device redundancy for large enterprise SD-WAN on-premise spoke sites. You can configure an SD-WAN site with two CPE devices to act as primary and secondary devices and protect the site against device and link failures. If the primary device fails, the secondary device takes over the traffic processing.

> **NOTE:** You must use the same device model for both primary and secondary devices and the devices must have the same version of Junos OS installed.

The following SD-WAN features are not supported for device redundancy:

- LTE WAN backup link
- Service chain support

> **NOTE:** Device redundancy is supported only on SD-WAN deployments.

Prerequisites for SRX Series Devices

The prerequisites to configure an SD-WAN site with dual CPE SRX Series devices are as follows:

- For SRX Series, you need to form the cluster manually by connecting two SRX Series devices together using a pair of the same type of Ethernet connections. To create an SRX cluster, see Chassis Cluster Feature Guide for SRX Series Devices.
- Log in to any one of the SRX Series devices, copy the **Stage-1** configuration from the **Sites** page and paste it into the console screen and commit the configuration.

Supported Connection Plans

The following connection plans are supported for device redundancy:

- Dual NFX250 as SD-WAN CPEs—Supports dual CPE NFX Series devices on an SD-WAN site.
- Dual SRX as SD-WAN CPEs—Supports dual CPE SRX Series devices on an SD-WAN site.
- Dual SRX4x00 as SD-WAN CPEs—Supports SRX 4100 and SRX4200 devices as dual CPE devices in an SD-WAN site.

Create and Configure an SD-WAN Site

You can create and configure an SD-WAN site with dual CPE devices and the two devices back up each other, with one node acting as the primary device and the other as the secondary device. The workflow to add and configure a site with dual CPE devices is similar to the single CPE device. For more information about creating and configuring a site with dual CPE devices, see "Adding an On-Premise Spoke Site with SD-WAN Capability" on page 100.

Dual CPE Devices Logical Topology for NFX Network Services Platform

Figure 8 on page 228 shows the logical topology of the NFX Series dual CPE devices.

Figure 8: Dual CPE Device Topology - NFX Network Services Platform

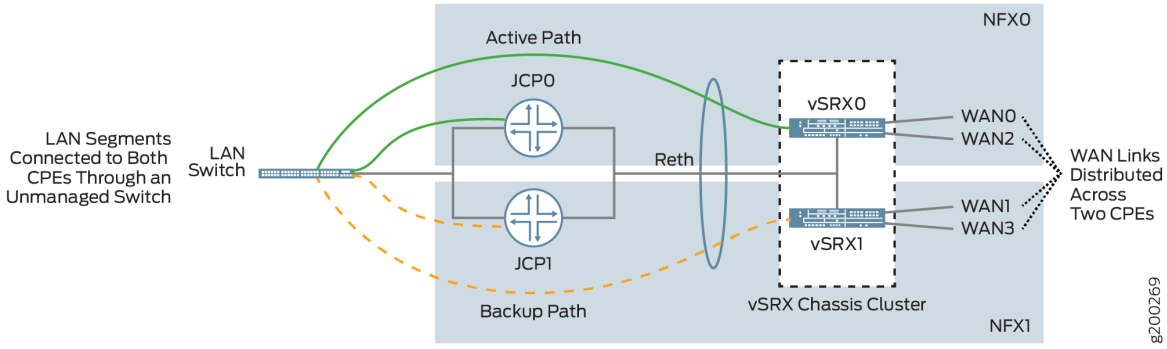

You can form a cluster using two NFX Series devices. The front panel ports of the NFX Series devices are used to interconnect two NFX Series devices and to carry the control and fabric interconnect traffic between the two NFX250 devices.

The Junos Control Plane (JCP) component acts as a switch, controls the front panel ports, and sends the traffic which arrives from the LAN or WAN to the NFX Series devices. On the LAN, the active/backup mechanism is used and if the primary device fails, the secondary device takes over processing of traffic. On the WAN, the active/active mechanism is used and all four WAN links are active and distributed across two NFX Series devices.

Dual CPE Devices Logical Topology for SRX Series Gateway Devices

Figure 9 on page 229 shows the logical topology of the SRX Series dual CPE devices.

Figure 9: Dual CPE Device Topology - SRX Series Devices

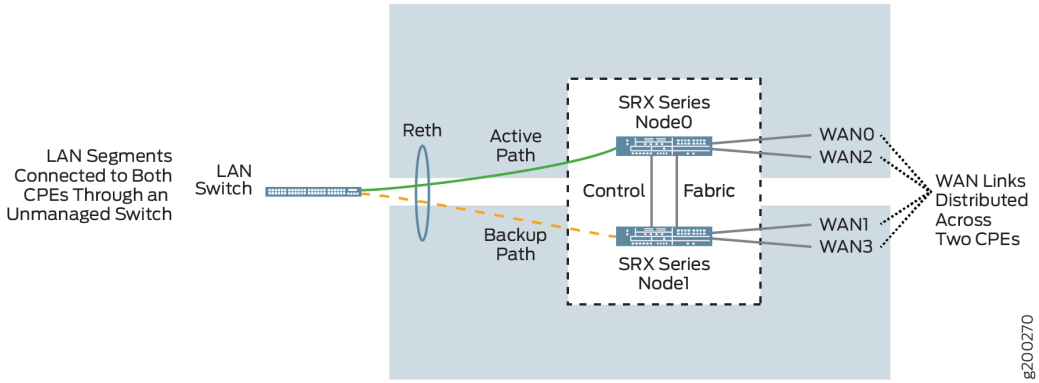

You can form a cluster using two SRX devices. A chassis cluster is formed between these nodes and performs as a single logical router. On the LAN, the active/backup mechanism is used and if the primary device fails, the secondary device takes over traffic processing. On the WAN, the active/active mechanism is used and all four WAN links are active and distributed across two NFX Series device.

> **NOTE:** On SRX 4100 and SRX4200 devices, out of the eight 1-Gigabit Ethernet/10-Gigabit Ethernet, a maximum of two ports are used for WAN links, and the remaining ports are used for LAN connectivity. The HA ports are used only for forming the cluster.

RELATED DOCUMENTATION

- Adding an On-Premise Spoke Site with SD-WAN Capability | 100
- Activating Dual CPE Devices (Device Redundancy) | 254

Activate a Device

IN THIS SECTION

- Manually Activate a Device That Supports Phone-Home Client | 230
- Manually Activate a Device That Does Not Support Phone-Home Client | 231

If the automatic activation of a site is disabled, you must manually activate the device that is associated with the site. Use the following procedures (as applicable) to manually activate CPE devices, enterprise hub devices, and cloud spoke devices.

NOTE:

You can manually activate a device only when:

- The status of the site (with which the device is associated) is Configured.
- The management status of the device is Device_Detected.

Manually Activate a Device That Supports Phone-Home Client

If a device supports the phone-home client, the ZTP toggle button is enabled by default in the corresponding device template.

To manually activate a device that supports the phone-home client:

1. Do one of the following:

   - a. Select **Resources > Devices**.

     The Devices page appears.

     b. Select the device that you want to activate and click **Activate Device**.

     The **Activate Device** page appears. The Activate Device page consists of the Device Information and Device Activation tabs. Proceed to step 2.

   - a. Select **Resources > Site Management**.

     The Sites page appears.

     b. Click the ***Site-Name*** link for the site with which the device is associated.

     The ***Site-Name*** page appears.

     c. On the **Devices** tab, select the device that you want to activate and click **Activate Device**.

     The **Activate Device** page appears. The Activate Device page consists of the Device Information and Device Activation tabs.

2. On the Device Information tab, enter the activation code for the device in the **Activation Code** field.

The activation code that you enter must match the activation code that is provided during the site addition workflow.

3. Click **Next**.

   The Device Activation tab appears displaying the progress of the device activation.

4. After the device is activated, click **OK**.

   The Sites page appears. If the device is successfully activated, the Management Status of the device and site change to **PROVISIONED**.

Manually Activate a Device That Does Not Support Phone-Home Client

If the device does not support the phone-home client, you must push the stage-1 configuration to the device manually from the Devices page and then manually activate the device.

SRX4100, SRX4200, vSRX, and SRX dual-CPE devices do not support the phone-home client.

To activate dual-CPE devices, see "Activating Dual CPE Devices (Device Redundancy)" on page 254.

To push the stage-1 configuration manually:

1. Select **Resources > Devices**.

   The Devices page appears.

2. Select the device and click **Stage1 Config**.

   The **Copy the following config and deploy it to the device** page appears.

3. Click **Copy**.

   A confirmation message appears indicating that the stage-1 configuration is copied to the clipboard.

4. Log in to the device console and enter the Junos OS configuration mode.

5. Paste and then, commit the configuration.

   The Management Status of the device on the Devices page changes to **Device_Detected**.

6. Proceed to activate the device. The rest of the workflow is the same as the one you encounter when you manually activate a device that supports the phone-home client. Refer to the "Manually Activate a Device That Supports Phone-Home Client" on page 230 for details.

RELATED DOCUMENTATION

Zero Touch Provisioning Overview | 259

Manually Activating a Switch

You can manually activate a switch if you have disabled automatic activation while onboarding the switch to CSO.

> **NOTE:** Before activating, ensure that the switch is powered on and connected to the CPE or firewall (if the switch is behind a CPE or firewall).

To manually activate the switch:

1. Select **Resources > Site Management**.

   The Sites page appears.

2. On the Sites page, click on the site that you want to activate.

   The detailed view of the site appears.

   > **NOTE:** You can activate a site that is in the CONFIGURED state.

3. Click the **Devices** tab.

4. Select the switch that you added to the site and click **Activate Device** to activate the switch.

   The Activate Device page appears.

5. On the Activate Device page, enter the activation code for the switch. The activation code must match the activation code that you provided during the site addition workflow.

6. Click **Next**.

   The progress of switch activation is displayed.

7. After the switch is activated, click **OK**.

The Sites page appears. The status of the switch is set to PROVISIONED if the switch is successfully activated. After the switch is activated, you can manage the switch by using CSO.

RELATED DOCUMENTATION

Activate a Device | 229

Zero Touch Provisioning Overview | 259

Manage an EX Series Switch

IN THIS SECTION

- View the Chassis Information of an EX Series Switch | 234
- View Information about an EX Series Switch | 237
- View Information about Ports on an EX Series Switch | 239
- Deploying an Access Profile on a Switch | 243
- Dissociating an Access Profile | 244

You can use the *Device-Name* page to view and manage an EX Series switch (physical and Virtual Chassis).

To view the chassis information of a member device in an EX Series Virtual Chassis, click the member device. The chassis view displays the ports and the status of the ports on the member device.

To access this page, do one of the following:

- To access the *Device-Name* page from the **Resources > Site Management** menu:

  1. Click **Resources > Site Management**.

     The Sites page appears with a list of sites displayed under the Site Name column.

  2. Click a *site* from this list.

     The Site Management page appears with a list of devices displayed under the Device Name column of the Devices tab.

  3. Select an *EX Series switch* from the list.

The *Device-Name* page appears.

- To access the *Device-Name* page from the **Resources > Devices** menu:

  1. Click **Resources > Devices**.

     The Devices page appears.

  2. Select an EX Series switch in the Device Name column of the Devices List.

     The *Device-Name* page appears.

You can perform the following actions from this page:

View the Chassis Information of an EX Series Switch

The chassis view displays the device model and all the ports on an EX Series switch.

You can perform the following actions from the chassis view dashlet that appears on this page:

> **NOTE:** The chassis view is refreshed after every 60 seconds.

- View information about ports—Hover over a port on the chassis view to view general information (such as administrative status, link status, and link mode) about the port.

  SeeTable 55 on page 235 for more details.

  > **NOTE:** The ports on the chassis view are color coded depending on the admin and link statuses:
  >
  > - Green—If the admin status and link status are up.
  > - Red—If the admin status is up and the link status is down.
  > - Dark Gray—If the admin status and link status are down.
  > - Light Gray—If the port is not configured as part of any LAN segment.

- View additional details of a port—Click a port to view additional details of the port.

  The Port Overview tab appears. Table 60 on page 242 describes the fields on the Port Overview tab.

- View details of system meters—Hover over a system meter to view more information from the trays that appear. SeeTable 56 on page 236 for more details.

- Perform various actions on an EX Series switch—Click **Actions** on the chassis view dashlet.

A list of all actions that you can perform on the switch is displayed. See Table 57 on page 236 for more details.

Table 55 on page 235 describes the fields on the port view pane.

Table 55: Fields on the Port View Pane

Field	Description
Admin Status	Administrative status of the port: • Green—Indicates that the admin status is up (enabled). • Gray—Indicates that the admin status is down (disabled).
Link Status	Operational status of the link or connection to the port: • Green—Indicates that the connection to the port is up. • Red—Indicates that the connection to the port is down.
Port Mode	Indicates the mode in which the port operates: • Access (default)—Only one VLAN is configured on the port. • Trunk—One or more VLANs are configured on the port. (Optional) A native VLAN may also be configured.
Link Mode	Mode in which the link to the port operates—Half-duplex or Full-duplex.
Power Consumption	Power consumed by the port, in watts (W).
PoE Status	Indicates whether the port is configured to transmit electrical power through an Ethernet cable (ON) or not (OFF).
Negotiated Speed	Current negotiated speed (in Kbps, Mbps, and Gbps) of the port.
VLAN	ID of the VLAN configured on the port. Range: 1 through 4094.
Input Bandwidth Utilization	Bandwidth (in %) consumed by the incoming packets on the port.
Output Bandwidth Utilization	Bandwidth (in %) consumed by the outgoing packets on the port.
Input Drops	Number of incoming packets dropped by the port due to congestion.
Output Drops	Number of outgoing packets dropped by the port due to congestion.
Input Errors	Number of errors in the incoming packets.

Table 55: Fields on the Port View Pane *(continued)*

Field	Description
Output Errors	Number of errors in the outgoing packets.

Table 56 on page 236 describes the system meters available on the chassis view dashlet. The system meters display current data.

> **NOTE:** The UI polls the CSO database every 30 seconds and the database polls the devices every five minutes.

Table 56: System Meters on the Chassis View Dashlet

System Meter	Description
CPU	CPU utilization (in %) in the switch.
Memory	Memory (in %) utilized in the switch.
Storage	Storage space (in %) allocated to the logical partitions of the switch.
Fan	Details of the fan used on the switch.
Temperature	Temperature details of the components in the available FPC.
LEDs	Severity level of the Alarms, System, and Primary LEDs.
Power	Details of the power supplies for the switch.

Table 57 on page 236 describes the actions that you can perform on an EX Series switch.

Table 57: Options on the Actions List

Action	Description
Ping	Select this option to ping a remote host to verify the connectivity between the EX Series switch and the remote host. See "Identifying Connectivity Issues by Using Ping" on page 291 for more information.

Table 57: Options on the Actions List *(continued)*

Action	Description
Traceroute	Select this option to execute the traceroute command from the EX Series switch, to view the path a packet travels to reach the remote host. See "Identifying Connectivity Issues by Using Traceroute" on page 295 for more information.
Reboot Device	Select this option to reboot the switch. See "Rebooting a CPE Device" on page 287 for more information. **NOTE:** This option is available only for a physical EX Series switch.
Reboot Member	Select a member device in the Virtual Chassis that you want to reboot and click **Reboot Member**. A reboot job is triggered to reboot the selected member. (Optional) You can view the status of the job on the Jobs (**Monitor > Jobs**) page. **NOTE:** This option is available only for an EX Series Virtual Chassis.
Reboot All	Select this option to reboot all member devices in the Virtual Chassis. A reboot job is triggered to reboot all the members. (Optional) You can view the status of the job on the Jobs (**Monitor > Jobs**) page. This option is available only for an EX Series Virtual Chassis.
View ARP Table	The switch uses the ARP table to map MAC addresses to IP addresses of the ports on the switch. If you select this option, the View ARP Details page appears with details, such as MAC addresses, IP addresses, Interface names, and flags associated with the switch.
View MAC Table	The switch uses the MAC table to map MAC addresses to specific ports on the switch. If you select this option, the View MAC Details page appears with details, such as MAC addresses, Interface names, and flags associated with the switch.

View Information about an EX Series Switch

Click the **Overview** tab to view information about an EX Series switch.

You can select one of the following options as the time span to view details about the recent alarms, PoE, resource utilization, and physical box storage:

- Past 1 hour
- Past 8 hours
- Past 1 day
- Past 1 week
- Past 1 month

Table 58 on page 238 describes the dashlets on the Overview tab. The graphical representations on this tab display trends based on historical data.

Table 58: Dashlets on the Overview Tab

Dashlet	Description
Port Link Status	Graphical representation (Donut chart) of the link status. Hover over the chart to view the number and percentage of links that are up and down. You can click the chart to view all the ports on the switch, on the Port Details page. You can also search for a port or filter the list based on port name and link status (up or down).
Recent Alarms	Recent alarms (Critical, Major, and Minor) generated on the switch. Click the *View All Alarms* link to view information about all the alarms, on the Alarms page. See *About the Alarms Page* for more information.
Details	Details (such as serial number, management IP address, OS version, and device template) of the switch.
Resource Utilization	Graphical representation of memory and CPU utilized in the switch, for the selected time span.
Current System Users	Details (such as name, duration, and login time) of the system users who are currently logged in to the switch. Click the *More Details* link on this dashlet to view additional information (such as username and session type) about the current users, on the View Details page. You can search and sort the information on this page as per your requirement, by using the Search and Filter icons, respectively.

Table 58: Dashlets on the Overview Tab *(continued)*

Dashlet	Description
PoE	Graphical representation of the power consumed by each PoE interface, in Watts (W). **NOTE:** This graph is displayed only for P models of EX Series switches.
Top Ports by Input Bandwidth	Graphical representation of the top 10 ports on which the incoming packets consume the maximum bandwidth.
Top Ports by Output Bandwidth	Graphical representation of the top 10 ports on which the outgoing packets consume the maximum bandwidth.
Top Ports with Input Errors	Graphical representation of the top 10 ports with the highest number of errors in incoming packets.
Top Ports with Input Errors	Graphical representation of the top 10 ports with the highest number of errors in outgoing packets.
Top Ports with Input Packet Loss	Graphical representation of the top 10 ports that drop the highest number of incoming packets.
Top Ports with Output Packet Loss	Graphical representation of the top 10 ports that drop the highest number of outgoing packets.
Licenses	Details of licenses (such as license name and description) installed on the switch. Click the *More Details* link on this dashlet to view additional information about the licenses, on the Device License Files page. . See *About the Device License Files Page* for more information.
Physical Box Storage	Graphical representation of the storage space (in %) allocated to the logical partitions of the switch.

View Information about Ports on an EX Series Switch

Click the **Ports** tab to view information about each port on an EX Series switch.

Table 59 on page 240 describes the fields on the Ports tab.

You can perform the following tasks on this tab:

- Search for a specific port by using keywords—Click the Search icon to search for a port by entering partial or full text of the keyword in the text box.

  The search results are displayed on the same tab.

- Filter the data displayed on the tab—Click the Filter icon to apply a quick filter. The filtered results are displayed on the same tab.
- Show or hide columns that contain information about the ports—Click the Show or Hide Columns icon to select or clear columns that you want to display or hide on the tab.
- View additional details of a port—Click a port in the Port column to view additional details of the port, on the Port Overview tab that appears.

  Table 60 on page 242 describes the fields on the Port Overview tab.

Table 59: Fields on the Ports Tab

Field	Description
Interface List	
Port	Name of the port. Click each port to view additional information about the port, on the Port Overview page. SeeTable 60 on page 242 for details of dashlets that appear on the Port Overview page.
Admin Status	Indicates the administrative status of the port: • Up—if the port is enabled. • Down—If the port is disabled.
Link Status	Indicates the status of the link or connection to the port: • Up—If the connection to the port is up. • Down—If the connection to the port is down.
MTU	Maximum transmission unit (MTU) size (in bytes) on the ports. Default: 1500 bytes.
Negotiated Speed	Current negotiated speed (in Kbps, Mbps, and Gbps) of the port.

Table 59: Fields on the Ports Tab *(continued)*

Field	Description
Interface List	
Port Mode	Operating mode of the port: • Trunk • Access
Link Mode	Mode in which the port operates—Half-duplex or Full-duplex.
Media Type	Type of transmission medium—Copper or Fiber.
Power Consumption	Power consumed by the port, in Watts (W).
PoE (Power Over Ethernet)	Indicates whether the port is configured to transmit electrical power through an Ethernet cable (ON) or not (OFF).
Input Bandwidth Utilization	Bandwidth (in %) consumed by the incoming packets on the port.
Output Bandwidth Utilization	Bandwidth (in %) consumed by the outgoing packets on the port.
Input Drops	Number of incoming packets dropped by the port due to congestion.
Output Drops	Number of outgoing packets dropped by the port due to congestion.
Input Errors	Number of errors in the incoming packets.
Output Errors	Number of errors in the outgoing packets.
VLAN ID	ID of the VLAN configured on the port. Range: 1 through 4094.
Auto Negotiation	Indicates whether the interface speed is auto-negotiated (Enabled) or is fixed based on an explicitly configured value (Disabled).
Port Profile	Port profile used for configuring the port.
Deployment Status	Deployment Status of the port profile. • Success—The port profile is successfully deployed on the port. • Pending—The port profile is in the process of being deployed on the port. • Failed—The port profile failed to deploy on the port.

Table 60 on page 242 describes the dashlets available on the Port Overview tab.

You can select one of the following options as the time span for which you want to view the graph for these dashlets:

- Past 1 hour
- Past 8 hours
- Past 1 day
- Past 1 week
- Past 1 month

> **NOTE:** The dashlets on the Port Overview tab are refreshed after every 30 seconds. The date and time of the last refresh appear at the bottom-left corner on each dashlet.

Table 60: Dashlets on the Port Overview Tab

Dashlet	Description
Details	Details (such as port number, admin status, and link mode) of the port that you selected.
Utilization	Graphical representation of CPU utilized by the selected port (in terms of input and output) for the selected time span.
Errors	Graphical representation of the number of errors in the incoming (input) and outgoing (output) packets for the selected time span. You can select either the past 1 hour, 8 hours, 1 day, 1 week, or 1 month as the time span for which you want to view the graph.
Packet Loss	Graphical representation of packet loss in incoming (input) and outgoing (output) packets for the selected time span. You can select either the past 1 hour, 8 hours, 1 day, 1 week, or 1 month as the time span for which you want to view the graph.
Bytes	Graphical representation of the MTU (in bytes) for incoming (input) and outgoing (output) packets for the selected time span. You can select either the past 1 hour, 8 hours, 1 day, 1 week, or 1 month as the time span for which you want to view the graph.

Table 60: Dashlets on the Port Overview Tab *(continued)*

Dashlet	Description
Packets	Graphical representation of the number of incoming (input) and outgoing (output) packets for the selected time span. You can select either the past 1 hour, 8 hours, 1 day, 1 week, or 1 month as the time span for which you want to view the graph.

Deploying an Access Profile on a Switch

An access profile defines the list of RADIUS servers to be used for authentication and accounting. You can deploy only one access profile on a switch.

To deploy an access profile on a switch:

1. Select **Resources > Devices** and click the switch on which you want to deploy the access profile.

   The *Devices* page appears.

2. Click the **Device Settings** tab.

   The access profiles configured are listed in a tabular format.

3. Select the access profile that you want to deploy on the switch.

4. In the **Type** field:
   - Click **Run now** to deploy the access profile immediately.
   - Click **Schedule at a later time** to schedule a time for deploying the access profile.

     If you select the Schedule at a later time option, enter the date and time when you want to deploy, in the Date and Time fields that appear when you select the option.

5. Click **Deploy**.

   If you select the Run now option, a job is created to deploy the profile immediately; otherwise, the job to deploy is created on the date and at the time that you schedule.

Dissociating an Access Profile

To dissociate an access profile from a switch:

1. Select **Resources > Devices** and click the switch on which you want to deploy the access profile.

   The *Devices* page appears.

2. Click the **Device Settings** tab.

   The access profiles configured are listed in a tabular format and the access profile that is deployed on the switch is selected.

3. Do one of the following:

   - Select another profile, if you want to deploy another profile on the switch.
   - Click on the selected profile to clear the selection, if you don't want to deploy any access profile on the switch.

4. In the **Type** field:

   - Click **Run now** to either deploy another access profile or dissociate the access profile immediately.
   - Click **Schedule at a later time** to schedule a time for the deployment.

     If you select the Schedule at a later time option, enter the date and time when you want to deploy, in the Date and Time fields that appear when you select the option.

5. Click **Deploy**.

   If you select the Run now option, a job is created for the deployment immediately; otherwise, the deployment job is created on the date and at the time that you schedule.

   After the access profile is deployed successfully, a message indicating that the switch is deployed successfully appears on top of the page. A link, View Logs, to view the job logs is also displayed.

6. (Optional) Click the **View Logs** link to view the details of the job to deploy the access profile.

   The Deploy Access Profile Details page appears. You can view the date and time the access profile was deployed and the status of the deployment job on this page.

RELATED DOCUMENTATION

Managing a Single CPE Device | 284

Managing Ports on an EX Series Switch | 245

Managing Ports on an EX Series Switch

IN THIS SECTION

- View Port Details | 245
- Enable Ports | 246
- Disable Ports | 247
- Deploy or Redeploy a Port Profile | 247
- Configure EX Series Switch Ports Overview | 248
- Edit Configuration of Ports | 251
- Dissociate a Profile from a Port | 253

You can manage an EX Series switch port from the Ports tab of the *Devices* page.

View Port Details

You can view the following details of a port on the Ports page:

- General information: Admin status, link status, port mode, and VLAN IDs configured on the port.
- Link Settings: Auto negotiation, flow control, MTU, Speed, and link mode.

To view the details of a port:

1. Do one of the following:
   - To view the general information and the link configuration of the port, select the port and click **More > Details**. Alternatively, right-click the port and click **Details**.

     A panel, Detail for *port name*, displaying the port details, appears on the right side of the ports grid.
   - To view statistics of the port, click the port link.

     The Port Overview pane appears displaying the following statistics of the port in graphical format for the past 1 hr, 8 hrs, 1 day, 1 week, and 1 month duration:
     - Port utilization—CPU utilized by the selected port (in terms of input and output) for the selected time span.
     - Packet transmitted through the port—Number of incoming (input) and outgoing (output) packets for the selected time span.
     - Number of Bytes transmitted through the port—MTU (in bytes) for incoming (input) and outgoing (output) packets for the selected time span.
     - Packet loss at the port—Packet loss in incoming (input) and outgoing (output) packets for the selected time span.
     - Errors on the port—Number of errors in the incoming (input) and outgoing (output) packets for the selected time span.

   Click the **Back to Device** link to view to the Ports tab.

Enable Ports

You enable a port to allow traffic through the port. When you enable a port, the Admin status of the port is changed to Up.

> **NOTE:** You can enable one or more ports at the same time.

To enable one or more ports:

1. Select the ports and click **More > Enable Port(s)**. Alternatively, right-click the ports and click **Disable Port(s)**.

   A job to enable the ports is initiated.

After a port is enabled, the Admin Status is changed to Up.

NOTE: The device is monitored once every five minutes. Therefore, it takes upto five minutes for the change in the Admin Status to reflect on the CSO GUI.

Disable Ports

You disable a port to block traffic through the port. When you disable a port, the Admin status and Link status of the port are changed to Down.

NOTE: You can disable one or more ports at the same time.

To disable one or more ports:

1. Select the ports and click **More > Disable Port(s)**. Alternatively, right-click the port and click **Disable Port(s)**

   A job to disable the ports is initiated.

After a port is disabled, the Admin and Link statuses for the port are set to Down.

NOTE: The device is monitored once every five minutes. Therefore, it takes upto five minutes for the change in the Admin Status and Link Status to reflect on the GUI.

Deploy or Redeploy a Port Profile

A port profile defines the authentication settings for the port and other port parameters such as flow control, link mode, storm control, MAC limit, and so on. The behavior of the port is defined by the values for parameters defined in the port profile. You can deploy a port profile on one or more ports at the same time.

You must redeploy a port profile when:

- the port profile that is assigned to a port is modified.
- the authentication profile or firewall filter that is assigned to the port profile is modified.

When a port profile or a profile associated with the port profile is modified, the deployment status of the port is changed to Pending Deployment.

The changes made to the port profile or the associated authentication profile or firewall filter are applied on the ports only when you redeploy the modified port profile.

To deploy or redeploy a port profile on one or more switch ports:

1. Select the port and click **More > Deploy**.

   The Deploy page appears.

2. For the **Type** field, do one of the following:
   - Click **Run now** to deploy the port profile immediately.
   - Click **Schedule at a later time** to deploy the port profile later.

     If you select this option, enter the date and time when you want to deploy in the Date and Time fields that appear.
3. Click **OK**.

   If you select the Run now option, a job is created to deploy the profile immediately; otherwise, the job to deploy is created on the date and at the time that you scheduled.

   When you deploy a port profile, the deployment status of the ports is set to Pending Deployment indicating that the profile is associated with the port. When the profile is in the process of being committed on the ports, the deployment status changes to In progress. If the deployment job completes successfully, the deployment status of the ports is set to Success and if the job fails, the deployment status is set to Failed.

Configure EX Series Switch Ports Overview

Starting from Release 5.1.0, you can use CSO to configure and monitor the ports of the following EX Series switches: EX2300, EX3400, EX4300, EX4600, and EX4650

> **NOTE:** You can configure and monitor ports on both a physical device and members of a virtual chassis. You cannot configure and monitor EX4600 and EX4650 Series switches, if the switches are configured as a virtual chassis.

You can add the following profiles to CSO and deploy them on the switch to configure the switch and the switch ports:

- Port profiles to define the behavior of a port. Port profiles allow you to provision multiple ports on a switch with the same set of attributes at the same time. A port profile includes the following:
  - Authentication profile (optional)
  - Firewall filters (Optiional)

  - Link settings
  - Storm control settings
  - Power over Ethernet (PoE) settings
  - Port security settings
- Authentication profiles to implement network access control (NAC).

  An authentication profile defines the authentication method, fallback options, and other settings such as number of retries, maximum number of authentication requests that can be allowed for a supplicant, authentication server timeout, and so on related to the communication between the switch and the supplicant (a user or device such as printer).

  An authentication profile may or may not be referenced in a port profile. However, you can assign the authentication profile to a port when you configure the port manually.
- Firewall filters to deny or permit network access to supplicants based on the filter terms.

  Firewall filters may or may not be referenced in a port profile. However, you can assign the firewall filters to a port when you configure the port manually.
- Access profiles to define the list of RADIUS servers to be used for authentication.

  An access profile is deployed on a switch and is referenced by an authentication profile when dot1x authentication is configured on the switch port.
- RADIUS server profiles to define the RADIUS server for authentication and accounting. You define the RADIUS server IP address, password, authorization ports, accounting ports, retry counts, and server timeout in this profile.

  A RADIUS server profile is referenced by an access profile and deployed on the switch through the access profile.

Configure Switch Ports

You can configure the ports either by using a port profile or manually.

By using a port profile, you can configure multiple ports of the switch at the same time.

To configure the ports by using a port profile:

1. (Optional) Create an authentication profile. See “Add Authentication Profiles” on page 684 for details.
2. (Optional) Create a firewall filter. See “Add Firewall Filters” on page 707 for details.
3. Create a port profile. See “Add Port Profiles” on page 674 for details.

4. Assign and deploy the port profile on one or more switch ports. See “Edit Configuration of Ports” on page 251 for details.

5. Deploy the port profile on the switch ports. See “Deploy or Redeploy a Port Profile” on page 247 for details.

To configure a port manually:

> **NOTE:** If you want to configure 802.1x authentication or apply firewall filters on a port while configuring the port manually, ensure that the authentication profile and the firewall filters are already configured.

1. In the customer portal, select **Resources > Devices**.

   The Devices page appears.

2. Click the switch for which you want to configure ports.

   The *switch* page appears.

3. On the Ports tab, select the ports that you want to configure and click **More > Edit Configuration**.

   The Edit Port Configuration page appears.

4. Edit the port parameters and deploy the configuration on the port. Refer to the instructions in the “Edit Configuration of Ports” on page 251 topic for completing and deploying the port configuration.

When you deploy a port profile, the deployment status of the ports is set to Pending Deployment indicating that the profile is assigned to the ports. When the profile is in the process of being committed on the ports, the deployment status changes to In Progress. If the deployment job completes successfully, the deployment status of the ports is set to Success and if the job fails, the deployment status is set to Failed.

WHAT'S NEXT

> After you configure a switch port, you can allow traffic through the switch ports and start monitoring the port.

Life Cycle of a Port Profile

The life cycle of a port profile is as follows:

1. Add a port profile to CSO.
2. Assign the port profile to a port.

   When you assign the port profile, the deployment status of the port is set to Pending Deployment indicating that the profile is assigned to the port.
3. Deploy the port profile on a port.

   During the deployment, that is when the configuration is committed on the port, the deployment status is changed to In Progress. If the deployment job completes successfully, the deployment status of the port is set to Success; otherwise, the deployment status is set to Failed.
4. Edit the port profile.

   When you edit the port profile or any profile associated with the port profile, the deployment status of the port profile is set to Pending Deployment.
5. Redeploy the port profile for the changes to reflected in the port configuration.

   During the redeployment, the deployment status of the port is changed to In Progress. If the deployment job completes successfully, the deployment status of the port is set to Success; otherwise, the deployment status is set to Failed.
6. Dissociate the port profile.
7. Delete the port profile.

SEE ALSO

Configure a Firewall Filter for an EX Series Switch | **706**

Edit Configuration of Ports

You can edit the configuration of a port either by using a port profile or manually.

If a profile is already deployed, edit the profile, and then redeploy for the changes to take effect.

When using a port profile to edit ports, you can:

- assign a port profile or modify the profile assigned to the port
- modify the VLANs assigned to the port

To edit the configuration of one or more ports:

1. Select one or more ports and click **More > Edit Configuration**. Alternatively, right-click the ports and click **Edit Configuration**.

   The Edit Port(s) page appears.

2. Do one of the following:

   For **Options**, select one of the following:

   - To edit the port configuration by using a port profile:

     a. Click **Use Port Profile**.

        The Port Profile drop down list and an option to select VLAN appears.

     b. Select a port profile that you want to assign to the port from the **Port Profile** drop down list.

     c. In the VLAN field:

        - If the port is configured as a trunk port in the port profile, assign multiple VLANs by selecting the VLANs in the Available column and clicking the right-arrow to move them to the Selected column.
        - If the port is configured as an access port in the port profile, assign a single VLAN by selecting the VLAN in the Available column and clicking the right-arrow to move it to the Selected column.

   - To edit the port configuration manually:

     a. Click **Manually edit port configurations.**

        The port parameters such as port mode, link mode, flow control appear.

     b. Edit the parameters by referring to Table 227 on page 674.

3. Click **Next**.

   Deployment options appear.

4. (Optional) Select the **Do not deploy** option if you want to only save the edited configuration in CSO, but not push and commit the configuration to the switch.

5. For the **Type** field, do one of the following:

   - Click **Run now** to save the edited configuration in CSO and commit the edited configuration on the switch.
   - Click **Schedule at a later time** to schedule a time to commit the edited configuration on the switch.

If you select the Schedule at later time option, enter the date and time when you want to deploy, in the fields that appear when you select the option.

6. Click **Next**.

   A summary of all the port profile parameter appears.

7. (Optional) Click **Edit** to revisit the settings and make further changes.

8. Click **OK**.

   If you select the Run now option, a job is created to deploy the profile at once; otherwise, the job to deploy is created on the date and at the time that you scheduled.

   During the deployment (that is when the deploy is executing), the Deployment status of the port is set to Pending Deployment, on the Ports tab of the *Devices* page.

   If the deployment job completes successfully:

   - A message appears on the top of the Ports page indicating that the deployment is successful.
   - Deployment Status of the port is set to Deployed.
   - The Port Profile column displays:
     - the profile name when a port profile is used for editing.
     - Manually Configured when the configuration is manually edited.

Dissociate a Profile from a Port

You can dissociate a profile from a port when you want to delete the profile or assign another profile to the port. The dissociate action removes the association between the port profile and one or more ports in CSO. You must perform the Deploy action on the ports from which the port profile was dissociated to remove the configuration deployed on the port and move the ports back to the default state.

To dissociate a profile from a port:

1. Select the ports and click **More > Dissociate Profile**. Alternatively, right-click the port and click **Dissociate Profile**.

   The Dissociate Port Configuration dialog box appears.

2. Click **OK** to dissociate the port profile.

   After the port profile is dissociated successfully, the Deployment Status on the Ports tab of the Devices page changes to Pending Deployment.

3. Select the port and click **More > Deploy** to commit the configuration on the switch.

The port profile is dissociated from the port and the port is returned to the factory default settings.

RELATED DOCUMENTATION

- Adding and Provisioning Switches to Provide LAN Capability to a Site Overview | 57
- Configure EX Series Switch Ports Overview | 248
- Configure a Firewall Filter for an EX Series Switch | 706
- Manage an EX Series Switch | 233

Activating Dual CPE Devices (Device Redundancy)

You can activate a device after the device status is changed to **EXPECTED** in the Sites page. When you see the device status is **EXPECTED**, it indicates that the device is ready to be activated. If you see the device status as **Undefined**, contact your service provider for assistance.

> **NOTE:** You must activate both the primary and the secondary devices simultaneously.
>
> You must use the same device model for both primary and secondary devices and the devices must have the same version of Junos OS installed.

To activate dual CPE devices used as a cluster:

1. Log in to Customer Portal.
2. Select **Sites**.

   The Sites page appears.
3. Click on the *Site Name*.

   The *Site Name* page appears.
4. On **Devices** tab, select the cluster device and click **Activate Device**.

   The Activate Device page appears. The Activate Device page consists of Device Information and Device Activation tabs.

NOTE: You can also activate the device through **Resource > Devices** page.

5. On **Device Information** page, complete the configuration according to the guidelines provided in Table 61 on page 255.

Table 61: Fields on the Activate Device Page

Field	Description
Site Name & Type	View the name of the site on which the CPE device is activated.
Connected Region	View the name of the region to which the CPE device is connected.
Primary Device Serial Number	View the serial number of the primary CPE device.
Primary Device Activation Code	Enter the activation code of the primary device that your service provider supplied for the device. **NOTE:** If you do not want to specify an activation code, on the **Resources > Edit Template > Template Settings** page, disable the ACTIVATION_CODE_ENABLED field and save the changes.
Secondary Device Serial Number	View the serial number of the secondary CPE device.
Secondary Device Activation Code	Enter the activation code of the secondary device that your service provider supplied for the device. **NOTE:** If you do not want to specify an activation code, on the **Resources > Edit Template > Template Settings** page, disable the ACTIVATION_CODE_ENABLED field and save the changes.

6. Click **Next**.

   The Activate Device page appears.

7. On **Activate Device** page, the cluster device (both primary and secondary) is activated through the following steps:

   - Device is detected
   - Stage-one configuration apply on device is successful
   - Bootstrap of device success
   - Activation of device is successful

- Device is modelled and is expected to be activated
- Device is active
- Device gateway component is spawned
- Device gateway router is put into cluster mode.
- Device is successful provisioned

After each successful step, you can see a green check mark. If any of these steps fail, a red exclamation mark appears.

8. After the activation process is complete, click **OK**.

   The *Site Name* page appears. If the device activation is successful, the management status of the cluster device is changed to **PROVISIONED**. You can also see the following device states:

   - **EXPECTED**—Device is ready for activation.
   - **ACTIVE**—Device is authenticated but not yet operational.
   - **ACTIVATION_FAILED**—Device is not authenticated.
   - **GWR_SPAWNED**—Device gateway component spawning is successful.
   - **GWR_SPAWN_FAILED** —Device gateway component spawning fails.
   - **PROVISIONED**—Device is operational.
   - **PROVISION_FAILED**—Device failed to become operational. Contact your service provider for assistance.

   > **NOTE:** The **GWR_SPAWNED** and **GWR_SPAWN_FAILED** statuses are not applicable for dual CPE SRX Series Services Gateway devices.

RELATED DOCUMENTATION

Device Redundancy Support Overview | 227

Activate a Device | 229

About the Sites Page | 53

Adding an On-Premise Spoke Site with SD-WAN Capability | 100

About the Certificates Page | 350

Viewing the History of Tenant Device Activation Logs

You can use the Activation Logs page to view the history of device activation logs. You can also view the details of the activation logs and their status.

To view the tenant device activation logs:

1. Click **Resources > Tenant Devices**.

   The Tenant Devices page appears, which list all devices.

2. Select a device and click **More > Activation Logs**.

   The Activation Logs page is displayed. Table 62 on page 257 describes the fields on the Activation Logs page.

3. Click a task name.

   The ZTP Logs page appears. Table 63 on page 258 describes the fields on the ZTP Logs page.

4. Click the Task Name.

   The Job Status page appears. Table 64 on page 258 describes the fields on the Job Status page.

5. Click **OK** to return to the previous page.

Table 62: Fields on the ZTP History Page

Field	Description
In progress	View the number of activated tasks that are in progress.
Success	View the number of activated tasks that are successful.
Failure	View the number of activated tasks that have failed.
Name	View the name of the task. Example: csp.tssm_ztp-Juniper-site-17-NFX-250-8052cc9451914be28c7c98fb64fd0db3
Start Date	View the start date and time of the task.
End Date	View the end date and time of the task.

Table 62: Fields on the ZTP History Page *(continued)*

Field	Description
Status	View the status of the task to know whether the task succeeded or failed.
Log	View the import logs. Click a log to access more detailed information about the imported log.

Table 63: Fields on the ZTP Logs Page

Field	Description
Task Name	View the ID created for the task. Example: install-license-to-device
Status	View the status of the task to know whether the task succeeded or failed.

Table 64: Fields on the Job Status Page

Field	Description
Name	View the name of the task.
Actual Start Time	View the start date and time of the task.
User	View the name of the user who activated the task.
End Time	View the end date and time of the task.
State	View the status of the task to know whether the task succeeded or failed.

RELATED DOCUMENTATION

Workflow for Onboarding a Device Using ZTP | 260

Zero Touch Provisioning Overview

Zero Touch Provisioning (ZTP) enables you to configure and provision devices automatically, and thus reduces the manual intervention required for adding devices to a network.

In Contrail Service Orchestration (CSO), the ZTP of a device involves the following high-level steps:

1. Activate the device that is associated with the site.
2. The device contacts CSO through the Redirect Server or Phone Home Client (PHC) and the stage-1 (initial) configuration is automatically applied to the device, allowing CSO to establish a secure management connection with the device. The process of applying the stage-1 configuration to the device is called as bootstrapping.
3. CSO automatically applies the provisioning configuration on the device after the completion of the bootstrapping process.

   The provisioning configuration is generated by CSO and is applied on a device to make it functional and ready for the intended functionality. For example, provisioning configuration can include IPSec or GRE tunnel configurations, virtual route reflector (vRR) configuration, routing configuration, and so on.

   For additional functionality, you can create the stage-2 configuration and apply the configuration to the device by using the CSO GUI. For example, the stage-2 configuration can include LAN configuration, firewall policies, and so on.

After the ZTP process is complete, the device is provisioned.

Devices Supported

In CSO, you can provision the following devices (including dual CPE devices, if applicable), by using ZTP:

- NFX150 and NFX250 Series
- SRX300, SRX320, SRX340, SRX345, SRX550 High Memory (SRX550M), and SRX1500
- vSRX on an x86 server
- EX2300, EX3400, EX4300, EX4600, and EX4650 Series switches

Benefits

Using ZTP offers the following benefits:

- Simplified, faster, and automated deployment of configurations.
- Auto-generated configurations that are more accurate.
- Faster scaling of the network because you need not manually apply configuration on each device in the network.

RELATED DOCUMENTATION

Workflow for Onboarding a Device Using ZTP | 260

Workflow for Onboarding a Device Using ZTP

Zero Touch Provisioning (ZTP) enables you to configure and provision devices automatically, minimizing the manual intervention required for adding devices to a network.

This topic provides a sequential list of tasks that you need to perform for successfully onboarding a device to the network by using ZTP:

1. From Customer Portal, add an on-premise spoke site or an enterprise site, and associate a device.

   For more information on adding an on-premise spoke site with the following capabilities:

   - WAN capability as SD-WAN, see "Adding an On-Premise Spoke Site with SD-WAN Capability" on page 100.
   - WAN capability as Hybrid WAN, see "Adding an On-Premise Spoke Site with Hybrid WAN Capability" on page 95.
   - WAN capability as Next Gen Firewall, see "Adding a Standalone Next Generation Firewall Site" on page 177.
   - WAN capability as SD-WAN and LAN capability, see "Adding an On-Premise Spoke Site with SD-WAN and LAN Capabilities" on page 117.
   - WAN capability as Next Gen Firewall and LAN capability, see "Adding an On-Premise Spoke Site with Next Generation Firewall and LAN Capabilities" on page 151.
   - Only LAN capability, see "Add an On-Premise Spoke Site with LAN Capability" on page 136.

   For more information on adding an enterprise hub, see "Adding Enterprise Hubs with SD-WAN Capability or SD-WAN and LAN Capabilities" on page 61.

2. Activate the device:

   - If you have enabled the **Auto Activate** field while adding an on-premise site or an enterprise hub, ZTP of the device is automatically triggered after the site is added to CSO.
   - If you have disabled the **Auto Activate** field while adding an on-premise site or an enterprise hub, you must manually activate the device.

     To manually activate the device:

a. Select **Resources > Site Management**.

   The **Sites-Name** page appears.

b. On the Sites page, click the site that you want to activate.

   The detailed view of the site appears.

   > **NOTE:** You can activate a site that is in the CONFIGURED state.

c. Click the **Devices** tab.

d. Select the device that you added to the site and click **Activate Device** to activate the device.

   The Activate Device page appears.

e. On the Activate Device page, enter the activation code for the device. The activation code must match the activation code that you provided during the site addition workflow.

f. Click **Next**.

   The progress of device activation is displayed.

g. After the device is activated, click **OK**.

   The Sites page appears.

- If you have to activate a vSRX or SRX4X00 Services Gateway devices:

  1. Select **Resources > Site Management**.

     The Sites page appears.

  2. Click on the site that you want to activate.

     The *Site-Name* page appears.

  3. On the Devices tab, select the device that you want to activate and click **Stage1 Config**.

     A new page appears displaying the stage-1 configuration of the device.

  4. Click **Copy to Clipboard** to copy the stage-1 configuration of the device.

  5. Log in to the CLI of the device and enter the configuration mode.

  6. Paste the stage-1 configuration and commit.

The Phone-Home client or the Redirect Server authenticates the device and establishes a communication between the device and CSO.

After the device activation is complete, CSO applies the stage-1 configuration. The status of the device is changed from Expected to Active, which indicates the device is authenticated but not yet operational.

3. After authenticating the device, CSO automatically triggers a job to push the provisioning and stage-2 (optional) configurations.

   You can use the Activation Logs page (**Resources > Tenant Devices > More >Activation Logs**) to view bootstrap logs (stage-1 configuration and device activation) and ZTP logs (provisioning and stage-2 configurations) and their status.

   After the job is completed successfully:

   - The provisioning configuration and stage-2 configuration (optional) are applied.
   - The device state changes from Active to Provisioned, which indicates that the device is fully functional.

The newly-added device is provisioned and is onboarded to the network. You can apply SD-WAN and security policies, if applicable.

RELATED DOCUMENTATION

Zero Touch Provisioning Overview | 259

CHAPTER 9

Managing Device Images

IN THIS CHAPTER

- Device Images Overview | 263
- About the Device Images Page | 263
- Deleting Device Images | 264

Device Images Overview

An image management system provides full lifecycle management of images for all network devices, including CPE device and virtualized network function (VNF) images. A *device image* is a software installation package for the CPE device or an image for a virtual application that runs on the device. For example, for a NFX Series device platform, you require an NFX software image and a software image for the vSRX application that provides security functions and routing on the device.

RELATED DOCUMENTATION

About the Device Images Page | 263

About the Device Images Page

To access this page, click **Resources > Images**.

You can use the Images page to view the list of device images that are available in tenant's network.

Tasks You Can Perform

You can perform the following tasks from this page:

- View details about a device image. Click the details icon that appears when you hover over the name of an image or click **More > Details**.

- Show or hide columns about the device image. Click the **Show Hide columns** icon in the top right corner of the page and select columns that you want to view on the page.
- Search for a device image. Click the Search icon in the top right corner of the page to search for a device image.

Field Descriptions

Table 65 on page 264 shows the fields on the Images page.

Table 65: Fields on the Device Images Page

Field	Description
Image Name	View the name of the device image. Example: juniper_srx_v1.tgz
Type	View the type of the device image. Example: VNF Image
Version	View the version number of the device image. Example: 1.1
Vendor	View the vendor name of the device. Example: Juniper
Size	View the size of the device image. Example: 14 KB

RELATED DOCUMENTATION

Device Images Overview | 263

Deleting Device Images

You can delete one or more device images from the Device Images page.

To delete a device image:

1. Select **Resources > Images**.

   The Images page appears with a list of device images.

2. Select the device image that you want to delete and then click the delete icon (X).

   The Confirm Delete page appears.

3. Click **Yes** to confirm.

   The Delete Success messages is displayed.

   The device image is deleted.

RELATED DOCUMENTATION

About the Device Images Page | 263

CHAPTER 10

Managing Resources

IN THIS CHAPTER

- Multidepartment CPE Device Support | **267**
- About the Devices Page | **268**
- Perform Return Material Authorization (RMA) for a Single-CPE Device or an EX Series Device | **271**
- Performing Return Material Authorization (RMA) for Dual-CPE Devices | **274**
- Granting RMA for a Device | **278**
- Managing a Single CPE Device | **284**
- Rebooting a CPE Device | **287**
- Configuring APN Settings on CPE Devices | **288**
- Identifying Connectivity Issues by Using Ping | **291**
- Identifying Connectivity Issues by Using Traceroute | **295**
- Remotely Accessing a Device CLI | **297**
- Configuring the Firewall Device | **299**
- About the Physical Interfaces Page | **300**
- About the Logical Interfaces Page | **301**
- Adding a Logical Interface | **302**
- Editing, Deleting, and Deploying Logical Interfaces | **305**
- Adding a Security Zone | **306**
- Adding a Routing Instance | **309**
- About the Static Routes Page | **310**
- Adding a Static Route | **311**
- Editing, Deleting, and Deploying Static Routes | **314**

Multidepartment CPE Device Support

Multitenancy enables a single NFX Series device to be mapped to serve across multiple departments within a single tenant. Each department has its own Layer 3 VPN and all Layer 3 VPNs are carried over to the hub using a shared overlay. The traffic is segregated to each department. A single overlay of IPsec or generic routing encapsulation (GRE) tunnels is used to carry all department traffic from the site through MPLS-based traffic separation.

Multitenancy is a cost-effective approach where the cost of a device and its maintenance is shared among multiple departments across a tenant. With multitenant device support, a dedicated share of the device is allocated to each department, and the data is kept private from the other tenants that access the same device.

> **NOTE:** Only users with the Tenant Administrator role have access to the Customer Portal GUI.

The tenant administrator can perform the following tasks:

- Manage and monitor all policies and dashboards for all departments.
- Manage applications in the dashboard for each tenant.
- Create SD-WAN and security policies for each tenant and monitor the dashboard at the site level or at the department level.
- View or select SD-WAN or security services on the shared CPE device through the management portal.
- View the shared CPE device and its services and networks even though the WAN links might be shared by multiple departments.

The service provider administrator can see all departments within the CPE device and activate the device.

RELATED DOCUMENTATION

About the SLA Performance of a Single Tenant Page | 848

Viewing the SLA Performance of a Site | 851

About the Devices Page

To access this page, click **Resources > Devices**.

You can use the Devices page to view the list of available CPE devices at the customer premises. You can also view information about each CPE device in the network.

Tasks You Can Perform

You can perform the following tasks from this page:

- Quickly view activation data created for CPEs in the widgets that appear at the top of the page. See Table 66 on page 269.
- Manage a single CPE. See "Managing a Single CPE Device" on page 284.
- Reboot a CPE device. See "Rebooting a CPE Device" on page 287.
- Push licenses to devices. Select the devices and click **Push License**.

  The Push License page appears displaying the list of licenses uploaded in CSO. Select the license(s) which you want to push to the selected devices. Click **Push Licenses** to push the licenses to the selected devices. To cancel the action, click **Cancel**.

  For information on pushing licenses to devices, see *Pushing a License to Devices*.
- Perform Return Material Authorization (RMA) to replace a device that is faulty or not reachable. You can perform RMA for a single-CPE or a dual-CPE device.
  - For information on performing RMA on devices, see "Perform Return Material Authorization (RMA) for a Single-CPE Device or an EX Series Device" on page 271
  - For information on performing RMA on dual-CPE devices, see "Performing Return Material Authorization (RMA) for Dual-CPE Devices" on page 274
- View details about a CPE . Click the details icon that appears when you hover over the row for a device or click **More > Details**.
- Show or hide columns about the CPE. Click the **Show Hide columns** icon in the top right corner of the page and select columns that you want to view on the page.
- Filter and sort CPE devices. Click a column name to sort the CPE devices based on the column name.

  Click the filter icon and select whether you want to show or hide column filters or apply a quick filter.

  NOTE: Sorting and filtering is applicable only to some fields.

- Search for a CPE. Click the Search icon in the top right corner of the page to search for a specific CPE.

You can enter partial text or full text of the keyword in the text box and press Enter.

The search results are displayed on the same page.

Field Descriptions

- Table 66 on page 269 describes widgets on the Devices page.
- Table 67 on page 269 describes the fields on the Devices page.

Table 66: Widgets on the Devices Page

Widget	Description
CPE by Status	Displays the management status of the CPE devices deployed in the cloud. • Pending Activation—Number of CPE devices that are yet to connect to the regional server. • Activation Failed—Number of CPE devices that could not connect to the regional server. • Expected—Number of CPE devices that are yet to connect to the regional server. • Active—Number of CPE devices that have downloaded images, but are not yet configured. • Provisioned—Number of CPE devices on which IPsec tunnels are fully operational. • Provision Failed—Number of CPE devices failed as the vSRX was not instantiated properly.

Table 67: Fields on the Devices Page

Field	Description
Device Name	Displays the name of the device. Example: sunny-NFX-250
Tenant	Displays the name of the tenant. Example: tenant-blue
Site Name	Displays the name of the tenant site. Example: site-blue-white

Table 67: Fields on the Devices Page *(continued)*

Field	Description
Management Status	Displays the management status of the CPE devices deployed in the cloud. • EXPECTED—Regional server has the activation details for the CPE device, but CPE device has not yet established a connection with the server. • DEVICE DETECTED—Device is configured and is reachable by CSO. After the user enters the activation code for the device, the activation code is validated and device is authenticated. • RMA—CPE device has been tagged for RMA as a result of the user applying the **Initiate RMA** action on the device. • ACTIVE—ZTP is initiated, CPE device has downloaded images, but not yet configured, and stage-1 configuration is pushed to the device. • PROVISIONED—ZTP is complete, and IPsec tunnel is established and operational on the device. • PROVISION_FAILED—Multiple factors lead to failure in provisioning a device. If any of the steps in ZTP fails or if any process fails as a part of device activation, then provision fails. For example, CPE device provisioning fails when the vSRX is not instantiated properly.
Model	Displays the name of the device model. Example: NFX
Active Services	Displays the number of services that are activated for the device. Example: 3
Operational Status	Displays whether the device is up or down.
Location	Displays the name of the location. Example: San Jose, CA
Status Message	Displays the latest status message. Example: IPsec provision success
WAN Links	Displays the number of WAN links. Example: 2

Table 67: Fields on the Devices Page *(continued)*

Field	Description
POP Name	Displays the name of the POP. Example: pop_blue
Image Name	Displays the name of the device image file. Example: install_nfx_fmpm_agent_1_0.sh
OS Version	Displays the Junos OS Release version. Example: 15.1X49-D40
Serial Number	Displays the serial number of the device. Example: XXXXXXXXXXXX
UUID	Displays the universally unique identifier (UUID) of the device. Example: xxxxxxxx-xxxx-xxxx-xxx-xxxxxxxxxxxx

RELATED DOCUMENTATION

Managing a Single CPE Device | 284

Perform Return Material Authorization (RMA) for a Single-CPE Device or an EX Series Device

Sometimes, due to hardware failure, a device managed by Contrail Service Orchestration (CSO) needs to be returned to the vendor for repair or replacement. In such situations, you perform Return Material Authorization (RMA) to back up the configuration of the faulty device, recall the faulty device and replace it with a new or restored device, push the required configuration to the replacement device, and activate it in order for CSO to recognize and manage the replacement device.

NOTE:

- When you request RMA on a site that is associated with a single-CPE device or an EX Series device and that has a version earlier than the CSO version, the site version is upgraded to the CSO version. The site version is upgraded as part of the device activation and zero touch provisioning (ZTP) process of the replacement device that is performed after RMA.
- You can initiate the RMA workflow only when the EX Series device is behind an SRX Series device or an internet gateway.

RMA can be initiated by an administrator with the RMA privilege.

To return a faulty device and replace it with a new or restored device using RMA:

1. Select **Resources** > **Devices**.

   The **Devices** page appears displaying all the devices and clusters.

2. Select the faulty device and click **More** > **Initiate RMA**.

   A confirmation page appears requesting for confirmation to go ahead with the initiate RMA process for the device. Click **Yes** to confirm RMA for the device.

   Click **No** to cancel the process.

   NOTE:

   - The **Initiate RMA** option is enabled for a device only if the management status is **PROVISIONED**.
   - For an EX virtual chassis, expand the virtual chassis and select the member device for which you want to initiate RMA.
   - On the Sites page (**Resources** > **Site Management**), the **Site Status** for the device for which you performed **Initiate RMA**, will remain **PROVISIONED**, however, you will see a red colored **RMA** tag beside the current status to indicate that RMA has been initiated for this device.

   If you click **Yes**, the RMA process is initiated for the selected device. The management status of the device changes to **RMA**. Once you put a device in the RMA state, you have to start the process for getting the new replacement device (referred to as the new device). This process must be performed outside of CSO.

3. After you receive the new device, provide the details of the new device by clicking **More** > **Grant RMA**. See “Granting RMA for a Device” on page 278.

> **NOTE:** The **Grant RMA** option is enabled only if the management status is **RMA**.

4. Activate the new device by selecting the device and clicking **Activate Device**. Enter the **Activation Code** of the device to activate the device for usage.

   When the device is activated, its **Management Status** changes to **PROVISIONED**. For a CPE device, an **RMA** tag is also displayed beside the management status indicating that the RMA process is not yet complete. Hover over the RMA tag to see the additional steps that you need to perform to complete the RMA process.

5. Manually push the following configuration to the newly provisioned device:

   > **NOTE:**
   > - In SD-WAN deployments, once the new device is in the **PROVISIONED** state, you can proceed to configure the device by manually pushing application signatures, certificates, and policies.
   > - In hybrid WAN deployments, service chains are restored automatically.

   - Licenses—If the new device is a physical SRX device, you must generate a new license and upload it to CSO.
   - Application Signatures—Push the application signatures to the new device. See "About the Application Signatures Page" on page 754.
   - Certificates—Import and install the required certificates on the new device. See "Importing a Certificate" on page 352 and "Installing and Uninstalling Certificates" on page 354.
   - Policies—Push the defined firewall and NAT policies to the new device. See "About the Firewall Policy Name Page" on page 377 and "About the NAT Policies Page" on page 558.

For a CPE device, to complete the RMA process, you have to remove the RMA tag manually. To remove the RMA tag, hover over the **PROVISIONED (RMA)** tag, select the check box indicating that you have completed all the steps for RMA, and click **OK**.

RELATED DOCUMENTATION

About the Devices Page | **268**

Granting RMA for a Device | **278**

Performing Return Material Authorization (RMA) for Dual-CPE Devices

IN THIS SECTION

- Performing RMA for an NFX Cluster | 274
- Performing RMA for an SRX Cluster | 276

You can perform RMA of devices in an NFX or SRX cluster when the devices fail or need to be replaced with new devices. During the RMA process, the configuration of a faulty device is backed up and the faulty device is shipped back to Juniper Networks. Juniper Networks ships a new device to the customer for replacing the faulty device. At the customer site, the new device is installed, the configuration backed up from the faulty device is pushed to the new device, and the new device is activated. CSO then manages the new device.

> **NOTE:** On a site associated with an NFX dual-CPE device, if the site version is lower than that of the CSO version, you can perform RMA only at the cluster level. After RMA of the cluster, the version of the site is upgraded to that of CSO as part of the device activation and ZTP process of the replacement devices in the cluster.

The following sections discuss how you can perform RMA for an NFX or SRX cluster:

Performing RMA for an NFX Cluster

Starting from CSO Release 4.1.0 onward, when an NFX250 device in a dual CPE cluster fails, you can perform RMA for only the failed device. Previously, RMA had to be done for both the devices in the cluster.

To perform RMA for an NFX cluster:

1. Select **Resources** > **Devices**.

   The **Devices** page appears displaying all the devices and clusters.

   Alternatively,

   a. Select **Resources** > **Site Management**.

      The Sites page appears.

   b. Click a site that contains the NFX cluster for which you want to perform RMA.

The *site* page appears

c. On the *site* page, click the **Devices** tab to view the devices and clusters installed at the site.

2. Do one of the following:

   - To perform RMA at the cluster lever, select the cluster.
   - To perform RMA for a single device in the cluster, select the device.

   A confirmation dialog box appears. Click **Yes** to confirm RMA for the selected NFX cluster or device.

   Click **No** to cancel the process.

   If you click **Yes**, the RMA process is initiated for the selected NFX cluster or device. The **Management Status** of the device changes to **RMA**.

   NOTE:
   - The **Initiate RMA** option is enabled for an NFX cluster only if the management status is **PROVISIONED**.
   - On the Sites page (**Resources** > **Site Management**), the status of the site, where the NFX cluster or the device for which you initiated RMA is installed, remains as **PROVISIONED**. However, a red-colored **RMA** tag appears beside the **Site Status** to indicate that RMA is initiated for a cluster or device at the site.

   After the NFX cluster or device is in the **RMA** state, you can raise a device replacement request for one or more faulty devices in the NFX cluster. This action is performed outside of CSO.

3. After you receive the new devices or device, click **More** > **Grant RMA** to enter the details of the new devices or device. See "Granting RMA for a Device" on page 278.

   NOTE: The **Grant RMA** option is enabled only when the management status of the device is RMA.

4. To activate the devices in the NFX cluster, do one of the following:

   - To activate both the devices in the cluster, select the cluster (in case of whole cluster is RMAed).
   - To activate a single device in the cluster, select the device (in case of a single device is RMAed).

5. Click **Activate Device**.

   The Activate Device page appears.

6. Enter the **Activation Code** for the new device or devices.

   When the device is activated, its **Management Status** changes to **PROVISIONED**. An **RMA** tag is displayed beside the management status indicating that the RMA process is not yet complete. Mouse over the RMA tag to see the additional steps that you need to perform to complete the RMA process.

7. Manually push the following configuration to the newly provisioned devices:

   NOTE:
   - In SD-WAN deployments, once the new devices are in the **PROVISIONED** state, you can proceed to configure the devices by manually pushing application signatures, certificates, and policies.
   - In hybrid WAN deployments, service chains are restored automatically.
   - When you perform RMA on a single device, CSO restores certificates on the new device. However, application signatures and policies need to be pushed manually.

   - Application Signatures—Push the application signatures to the replaced device. See "About the Application Signatures Page" on page 754.
   - Certificates—Import and install the required certificates on the replaced devices. See "Importing a Certificate" on page 352 and "Installing and Uninstalling Certificates" on page 354.
   - Policies—Push the defined firewall and NAT policies to the replaced devices. See "About the Firewall Policy Name Page" on page 377 and "About the NAT Policies Page" on page 558.

To complete the RMA process, you have to remove the RMA tag manually. To remove the RMA tag, mouse over the **PROVISIONED (RMA)** tag, select the check box indicating that you have completed all the steps for RMA, and click **OK**.

Performing RMA for an SRX Cluster

For an SRX cluster, you can perform RMA on a member device of the cluster. That is, you can select the faulty device from the SRX cluster and perform RMA on it individually. You cannot perform RMA for an SRX cluster at the cluster level.

To return a faulty device in an SRX cluster and replace it with a new or restored device by using RMA:

1. Select **Resources** > **Devices**.

   The **Devices** page appears displaying all the devices and clusters.

   Alternatively,

   a. Select **Resources** > **Site Management**.

      The Sites page appears.

b. Click a site that contains the SRX cluster for which you want to perform RMA.

   The *Site-Name* page appears.

c. Click the **Devices** tab to view the devices and clusters installed at the site.

2. Select the faulty device in the SRX cluster and click **More** > **Initiate RMA**.

   A confirmation page appears requesting for confirmation to initiate the RMA process for the device.

3. Click **Yes** to confirm RMA for the device.

   You are returned to the Devices page where a confirmation message appears, indicating that the RMA process is initiated.

   NOTE:
   - The **Initiate RMA** option is enabled only for a device with the management status **PROVISIONED**.
   - On the Sites page (**Resources** > **Site Management**), the status of the site, where the device for which you initiated RMA is installed, remains **PROVISIONED**. However, a red-colored **RMA** tag appears beside the current site status to indicate that RMA is initiated for a cluster or device at the site.

4. After the management status of the device changes to RMA, you must raise a replacement request for the faulty device in the SRX cluster. This process is performed outside of CSO.

5. After you receive the new device, perform the following actions on the new device through CLI.

   NOTE:
   - At this point, the new device is not onboarded to CSO.
   - Before performing the following actions through CLI, ensure that the new device is connected to the Internet.

   - Download licenses on the new device.
   - Install application signatures on the new device.

     If you do not install application signatures on the new device, configuration between the devices will not be synchronized when you pair the devices to form the cluster. Therefore, the cluster will be in Disabled state.

   See Application Security User Guide for Security Devices for details.

6. Pair the existing device with the new device to form the cluster. See SRX Series Chassis Cluster Configuration Overview for details on the tasks you must perform on the new device.

   After the cluster is formed, configuration is automatically synchronized between the devices in the cluster.

7. Click **More** > **Grant RMA** to provide details of the new device. See "Granting RMA for a Device" on page 278 for details.

   > **NOTE:** The **Grant RMA** option is enabled only for a device with the management status **RMA**.

8. To complete the RMA process, you must remove the RMA tag manually. To remove the RMA tag, hover over the **PROVISIONED (RMA)** tag, select the check box indicating that you have completed all the steps for RMA, and click **OK**.

   The RMA process is now complete. You can start using the new device.

RELATED DOCUMENTATION

- About the Devices Page | **268**
- Perform Return Material Authorization (RMA) for a Single-CPE Device or an EX Series Device | **271**
- Granting RMA for a Device | **278**

Granting RMA for a Device

IN THIS SECTION

- Grant RMA for a Single-CPE Device or an EX Series Device | **279**
- Grant RMA for a Dual-CPE Device | **280**
- Grant RMA for an SRX Device within an SRX Cluster | **282**

To grant RMA for a device, you must specify the information such as serial number, about the new device that is received as a replacement for a faulty device. Grant RMA can be performed by an administrator with the RMA privilege.

Grant RMA for a Single-CPE Device or an EX Series Device

Before you grant RMA for a device, ensure that:

- You have received a new device for replacing the faulty device.
- You have the serial number and the activation code for the new replacement device (referred to as new device).

To grant RMA for a device:

1. Select **Resources** > **Devices**.

   The **Devices** page appears displaying all the devices and clusters.

   Alternatively,

   a. Select **Resources** > **Site Management**.

      The Sites page appears.

   b. Click a site that contains the single-CPE device or EX Series device for which you want to perform RMA.

      The *site* page appears

   c. On the *site* page, click the **Devices** tab to view the devices and clusters installed at the site.

2. Select the faulty device for which you initiated RMA and click **More** > **Grant RMA**.

   The **Grant RMA for Device** page appears.

   > **NOTE:** The **Grant RMA** option is enabled only when the **Management Status** of the device is **RMA**.

3. Complete the configuration according to the guidelines provided in Table 68 on page 280.

4. Click **OK** to perform the grant RMA process.

When you perform **Grant RMA** for a device, a job is created to perform the following tasks:

- The device-related configuration is backed up to the CSO database, and the existing device is recalled and the new device is added to the network.
- The management status of the device changes to **EXPECTED** on the **Devices** page. For a CPE device, an **RMA** tag is also displayed beside the management status to indicate that the RMA process is not yet complete. Mouse over the RMA tag to see additional steps that you need to perform for completing the RMA process.

  On the Sites page (**Resources** > **Site Management**), the status of the site, where the device for which you performed **Grant RMA** is installed, changes to **Expected**.

> **NOTE:** You can see the progress of this job on the **Monitor** > **Jobs** page. This job might take around 15 minutes to complete.

To complete the RMA process and start using the new device, you must activate the device using the **Activate Device** option. See step 5 in "Perform Return Material Authorization (RMA) for a Single-CPE Device or an EX Series Device" on page 271.

Table 68 on page 280 provides guidelines on using the fields on the **Grant RMA for Device** panel.

Table 68: Fields on the Grant RMA for Single-CPE Device Page

Field	Description
Customer Name	Displays the name of the tenant who is performing RMA.
Site Name	Displays the name of site in which the faulty device is present.
Device Name	Displays the name of the faulty device that will be replaced with a new one through the **Grant RMA** process.
Serial Number	Enter the serial number of the new device. The serial number is case sensitive. Example: DD2316AF0177
Activation Code	Enter the activation code for the new device. You receive the activation code (Example: 545454) from the service provider, outside of CSO.

Grant RMA for a Dual-CPE Device

Before you perform **Grant RMA for a Device**, ensure that:

- You have received the replacement for the faulty devices.

- You have the serial numbers and the activation codes of the new devices.

To perform **Grant RMA** for a dual-CPE device(cluster):

1. Select **Resources** > **Devices**.

   The **Devices** page appears displaying all the devices and clusters.

   Alternatively,

   a. Select **Resources** > **Site Management**.

      The Sites page appears.

   b. Click a site that contains the dual-CPE device for which you want to perform RMA.

      The *site* page appears

   c. On the *site* page, click the **Devices** tab to view the devices and clusters installed at the site.

2. Select the cluster or the device in the cluster for which you initiated RMA and click **More** > **Grant RMA**.

   The **Grant RMA for Device** page appears.

   > **NOTE:** The **Grant RMA** option is only enabled if the **Management Status** of the device is **RMA**.

3. Complete the configuration according to the guidelines provided in Table 69 on page 282.

4. Click **OK** to complete the grant RMA process.

   When you grant RMA, the following actions are performed:

   - The cluster-related configuration is backed-up to the CSO database, and the devices in the cluster are recalled and the new or restored devices are added to the network.
   - The management status of the cluster changes to **EXPECTED** on the **Devices** page. An **RMA** tag is also displayed beside the management status indicating that the RMA process is not yet complete. Mouse over the RMA tag to see the additional steps that you must perform to complete the RMA process.

     On the Sites page (**Resources** > **Site Management**), the status of the site, where the cluster or device for which you performed **Grant RMA** is installed, changes to **Expected**.

> **NOTE:** You can see the progress of this job on the **Monitor** > **Jobs** page. This job might take around 15 minutes to complete.

Table 69 on page 282 provides guidelines on using the fields on the **Grant RMA for Device** panel.

Table 69: Fields on the Grant RMA for Dual-CPE Device Page

Field	Description
Customer Name	Displays the name of the tenant performing RMA.
Site Name	Displays the name of site in which the faulty device is present.
Device Name	Displays the name of the faulty device that will be replaced with new or restored devices through the **Grant RMA** process.
Primary Serial Number	Enter the serial number of the new primary device. The serial number is case sensitive. Example: DD2316AF0177
Primary Activation Code	Enter the activation code for the new primary device. You will receive the activation code from the service provider, outside of CSO. Example: 545454
Secondary Serial Number	Enter the serial number of the new secondary device. The serial number is case sensitive. Example: DD2316AF0145
Secondary Activation Code	Enter the activation code for the new secondary device. You receive the activation code from your service provider, outside of CSO. Example: 545476

Grant RMA for an SRX Device within an SRX Cluster

Before you perform **Grant RMA for a Device**, ensure that:

- You have received the replacement for the faulty device.
- You have the serial number and the activation code of the new device.

To perform **Grant RMA** for a device:

1. Select **Resources** > **Devices**.

   The **Devices** page appears displaying all the devices and clusters.

   Alternatively,

   a. Select **Resources** > **Site Management**.

The Sites page appears.

b. Click a site that contains the SRX device for which you want to perform RMA.

The *site* page appears

c. On the *site* page, click the **Devices** tab to view the devices and clusters installed at the site.

2. Select the defective device for which you initiated RMA and click **More** > **Grant RMA**.

The **Grant RMA for Device** page appears.

> **NOTE:** The **Grant RMA** option is enabled only when the **Management Status** of the device is **RMA**.

3. Complete the configuration according to the guidelines provided in Table 70 on page 283.

4. Click **OK** to perform the grant RMA process.

When you perform **Grant RMA** for a device, a job is created to update the device object in CSO with the serial number and activation code of the new device.

On the Sites page (**Resources** > **Site Management**), the status of the site, where the device for which you performed **Grant RMA** is installed, changes to **PROVISIONED**.

> **NOTE:** You can see the progress of this job on the **Monitor** > **Jobs** page.

Table 70 on page 283 provides guidelines on using the fields on the **Grant RMA for Device** page.

Table 70: Fields on the Grant RMA for Device Page (for SRX Device in an SRX Cluster)

Field	Description
Customer Name	Displays the name of the tenant who is performing RMA.
Site Name	Displays the name of site in which the faulty device is present.
Device Name	Displays the name of the faulty device that will be replaced with a new one through the **Grant RMA** process.
Serial Number	Enter the serial number of the new device. The serial number is case sensitive. Example: DD2316AF0177

Table 70: Fields on the Grant RMA for Device Page (for SRX Device in an SRX Cluster) *(continued)*

Field	Description
Activation Code	Enter the activation code for the new device. You receive the activation code (Example: 545454) from the service provider, outside of CSO.

RELATED DOCUMENTATION

- About the Devices Page | **268**
- Perform Return Material Authorization (RMA) for a Single-CPE Device or an EX Series Device | **271**
- Performing Return Material Authorization (RMA) for Dual-CPE Devices | **274**

Managing a Single CPE Device

You can use the Devices page to view and manage a single customer premises equipment (CPE) device at the tenant site. To access this page, click **Resources > Tenant Devices > *Device-Name***.

To manage a single CPE device at the tenant site:

1. Click **Overview** tab.

   View the following information on the Overview tab:

   - Geographical location of the device at the tenant site.
   - Aggregate throughput of the device.
   - Recent alerts for the device.
   - Recent Alarms for the device.
   - License details of the device.
   - Details of the device, such as serial number, management IP address, OS version, device template, tenant name, site name, and site location.

2. Click **Policies** tab.

   View the following information on the Policies tab:

   - List of all policies applicable to a CPE device.
     - Click a policy name to view the rules that are applicable for the CPE device.

- Click the edit icon at the end of the row to edit a policy. You are taken to the **Configuration > Policy** page, where you can edit the policies.

- Details about the tenant user who last updated the policy.
- Time when the policy was last updated.
- Deployment status of the policy.
- Number of rules applicable to the device compared to the total number of rules applicable to the tenant.

3. Click **Configuration Template** tab.

> **NOTE:** The **Configuration** tab that was available in earlier releases for stage-2 template-based configuration is renamed as **Configuration Template**.

Complete the configuration according to the guidelines provided in Table 71 on page 285, Table 72 on page 286, and Table 73 on page 286.

You can also perform the following operations from this page.

- Click **Deploy** to deploy the configuration.
- Click **Rollback** to rollback to previous working configuration.
- Click **Deployment History** to view the deployment history.
- Click **View Changes** to view the changes between current and previous configuration.
- Click **Render Template** to render the configuration in CLI format.
- Click **Save** to save the changes.

4. Click **Configuration** tab (applicable for firewall devices).

You can perform the following operations from this page.

- Click **Physical Interfaces** tab to view and manage the physical interfaces for the device.
- Click **Security Zone** tab to view and manage the security zones for the device.
- Click **Routing Instance** tab to view and manage the routing instances for the device.

For information about managing an EX Series switch, see “Manage an EX Series Switch” on page 233.

Table 71: Fields on the SRX Root Password Tab

Field	Description
Edit Settings for SRX root password	

Table 71: Fields on the SRX Root Password Tab *(continued)*

Field	Description
Password	Enter your new password. The password that you set must be between 6 and 21 characters long, and it must include at least one lowercase letter, one uppercase letter, one special character, and one number.
Confirm Password	Reenter the password for confirmation.

Table 72: Fields on the NFX Root Password Tab

Field	Description
Edit Settings for NFX root password	
Password	Enter your new password. The password that you set must be between 6 and 21 characters long, and it must include at least one lowercase letter, one uppercase letter, one special character, and one number.
Confirm Password	Reenter the password for confirmation.

Table 73: Fields on the APN Configuration Tab

Field	Description
Edit Settings for APN Configuration	
Use default APN settings	Click the toggle button to change the default APN settings. • Enabled(Default)—Select this option to use the default APN setting that is shipped along with the CPE device. • Disabled—Select this option to configure the APN settings.
APN Settings	
APN Name	Enter the access point name (APN) of the gateway router. The name can contain alphanumeric characters and special characters.

Table 73: Fields on the APN Configuration Tab *(continued)*

Field	Description
SIM Change Required	Click the toggle button to change the SIM card. • Enabled(Default)—Select this option to change the APN settings and to use a new SIM card. • Disabled—Select this option to change the APN settings without changing the SIM card.
Authentication Method	Select the authentication method for the APN configuration. • PAP— Select to use Password Authentication Protocol (PAP) authentication. This is the default option. • CHAP— Select to use Challenge Handshake Authentication Protocol (CHAP) authentication. • None—Select to indicate that there is no authentication method.
Authentication Info	
SIP User ID	Enter the Session Initiation Protocol (SIP) user ID for authentication if you have selected the APN authentication method as either **PAP** or **CHAP**.
SIP Password	Enter the SIP password for authentication if you have selected the APN authentication method as either **PAP** or **CHAP**.

RELATED DOCUMENTATION

About the Devices Page | 268

Rebooting a CPE Device

You need to reboot a CPE device if the device is down, or if all troubleshooting options fail.

To reboot a CPE device:

1. Select **Resources** > **Devices**.
2. Select the CPE device that you want to reboot and select **More** > **Reboot**.

   A Device Reboot job link is created and the Status Message column displays the status as **Reboot in-progress**.

> **NOTE:** If you reboot a tenant device, deployments that are in progress are stopped.

3. (Optional) Click the **Device Reboot** link to view the device reboot logs.

4. (Optional) You can view the job status on the **Monitor** > **Jobs** page.

You can view the status of reboot in the Status Message column.

On successful reboot of the CPE device, the Status Message column displays the status as **Reboot Succeeded**.

If a CPE device is not reachable or if the reboot time exceeds the timeout value, the reboot fails and the Status Message column displays the status as **Reboot Failed**.

> **NOTE:** The timeout value for rebooting a CPE device is 14 minutes.

RELATED DOCUMENTATION

About the Devices Page | 268

Configuring APN Settings on CPE Devices

IN THIS SECTION

- Configuring APN Settings with SIM Change on CPE Devices | 289
- Configuring APN Settings without SIM Change on CPE Devices | 291

You can configure Access Point Name (APN) settings on the following devices, with or without SIM change. You can change the APN settings either to use a private APN with the current LTE service provider or to use a different LTE service provider.

NOTE: You can only insert a SIM card in the SIM1 slot of the LTE Mini-Physical Interface Module (Mini-PIM).

Following is the list of devices on which you can configure APN settings:

- NFX Series—NFX150 and NFX250 CPE devices
- SRX Series—SRX320, SRX340, and SRX345 CPE devices

Configuring APN Settings with SIM Change on CPE Devices

To configure APN settings with SIM change:

1. Select **Resources > Devices**.

   The Devices page appears.

2. Click the device that you want to configure.

   The *Device-Name* page appears.

3. Click the **Configuration Template** tab and change the APN settings according to the guidelines provided in Table 74 on page 290.

4. Click **Save**.

   The new settings are applied after one minute.

5. Remove the USB dongle from the CPE device, change the SIM card, and re-insert the USB dongle.

   The system checks for the new APN settings every minute.

   - If the applied APN settings are compatible with the new SIM card—The LTE WAN link and its tunnels go down after one minute and remain down till the new SIM card is inserted. The LTE dongle LED indicates that the connection is down during this period. Maximum one minute after the new SIM is inserted, the LTE dongle LED indicates that the connection is up. The LTE WAN link and its tunnels come up automatically.
   - If the applied APN settings are not compatible with the new SIM card—The LTE WAN link and its tunnels go down after one minute and remain down even after the new SIM card is inserted. The LTE dongle LED indicates that the connection is down even after the new SIM is inserted.

6. To revert to the old SIM, remove the USB dongle, replace the current SIM with the previous SIM, and re-insert the dongle.

The system checks for the new APN settings every minute. Maximum one minute after the old SIM is inserted, the LTE dongle LED indicates that the connection is up (using the old SIM and old APN). The LTE WAN link and its tunnels come up automatically

Table 74: Fields for the APN Configuration Settings on the Configuration Template Tab

Field	Description
Edit Settings for APN Configuration	
Use default APN settings	Click the toggle button to enable (default) or disable the default APN settings. • If you enable this option, the default APN settings that are shipped along with the CPE device are used for configuring the APN. • If you disable this option, you must configure the APN settings manually.
APN Settings	
APN Name	Enter the access point name (APN) of the gateway router.
SIM Change Required	Click the toggle button to enable or disable changing the SIM card: **NOTE:** You can change the SIM card either to use a different LTE service provider or to use a private APN with the current LTE service provider. • (Default) Enable this option to change the APN settings and to use a new SIM card. • Disable this option to change the APN settings without changing the SIM card.
Authentication Method	From the list, select one of the following authentication methods for the APN configuration as configured by the service provider: • (Default) PAP—Select this option to use Password Authentication Protocol (PAP) as the authentication method. • CHAP—Select this option to use Challenge Handshake Authentication Protocol (CHAP) authentication as the authentication method. • None—Select this option if you do not want to use any authentication method.
Authentication Information	
SIP User ID	Enter the Session Initiation Protocol (SIP) user ID for authentication.
SIP Password	Enter the SIP password for authentication.

Configuring APN Settings without SIM Change on CPE Devices

To configure APN settings without SIM change:

1. Select **Resources > Devices**.

   The Devices page appears.

2. Click the device that you want to configure.

   The *Device-Name* page appears.

3. Click the **Configuration Template** tab and change the APN settings according to the guidelines provided in Table 74 on page 290.

4. Click **Save**.

   The new settings will be applied after one minute.

   - If the applied APN settings are valid—In CSO, the LTE WAN link and its associated tunnels go down momentarily and then, get re-established automatically.
   - If the applied APN settings are invalid—After one minute, the LTE dongle LED indicates that the connection is down. In CSO, the LTE WAN link and its associated tunnels go down.

     After two minutes, the LTE dongle LED indicates that the connection is up (using old APN). In CSO, the LTE WAN link and its tunnels come up automatically.

RELATED DOCUMENTATION

About the Devices Page | 268

Identifying Connectivity Issues by Using Ping

You can use Contrail Service Orchestration (CSO) to perform a ping operation from a device (provider hub, tenant device, CPE device, EX switch, enterprise hubs, or next-generation firewall device) to a remote host for identifying issues in connectivity with the remote host.

When you ping a remote host from a device, an Internet Control Message Protocol (ICMP) packet is sent to the remote host. By analyzing the results of the ping operation, you can identify the possible device connectivity issues between the remote host and the device.

NOTE: In Contrail Service Orchestration (CSO) Release 5.0, the following devices support ping:

- EX Series: EX2300, EX3400, EX4300, EX4600, and EX4650
- NFX Series: NFX150, NFX250
- SRX Series: SRX300, SRX320, SRX340, SRX345, SRX1500, SRX4100, SRX4200, SRX4600
- vSRX

To perform the ping operation:

1. Select **Resources** > **Devices**.

   The Devices page appears.

2. Select a device from the list of devices displayed and click **More** > **Ping**.

   The Ping page appears.

   NOTE: You can ping from a device only when its operational status is Up in CSO.

3. Complete the configuration according to the guidelines provided in Table 75 on page 292.

   NOTE: Fields marked with an asterisk (*) are mandatory.

4. Click **Ping** to initiate the ping request.

   A job is created and the Ping Progress page appears. After the device sends the ping packets, the Ping Result page appears. . If the ping operation is successful, the Ping Result page displays the parameters specified in Table 76 on page 294.

   If the ping operation fails, the Ping Result page displays an appropriate error message (such as **No response** or **No route to host**), indicating that there is an issue in the connectivity to the remote host.

Table 75: Fields on the Ping page

Field	Description
Remote Host	Enter the IPv4 address or hostname of the remote host.

Table 75: Fields on the Ping page *(continued)*

Field	Description
Ping Request Packets	Enter the number of ping request packets to be sent to the remote host. Default: 5. Range: 1 through 300.
Advanced	
Source Interface	Select the source interface on the device through which you want to send the ping request to the remote host. If you do not select a source interface, ping requests are sent on all interfaces. To clear the selected interface, click **Clear All** and select another interface.
Hostname Resolution	Click the toggle button to enable or disable (default) the display of hostname of the hops along the path to the remote host.
Rapid Ping	Click the toggle button to enable or disable (default) sending ping requests rapidly. If you enable this option, the source device sends a minimum of 100 ping request packets per second or sends a packet as soon as a response to the previous packet is received, whichever is greater. • If the source device does not receive a response for 500 ms, timeout is considered. • If the source device receives a response within 500 ms, the next ping request packet is sent immediately. **NOTE:** The ping results are displayed in a single consolidated message instead of individual messages for each ping request packet sent.
Packet Fragmentation	Click the toggle button to enable or disable (default) the fragmenting of ping request packets. If packet fragmentation is disabled, ping packets with the maximum transmission unit (MTU) greater than 1500 bytes are dropped.
Packet Size (bytes)	Enter the size (in bytes) of the ping request packet. Default: 56 bytes. Range: • 1 through 1,472 bytes, if packet fragmentation is disabled. • 1 through 65,468 bytes, if packet fragmentation is enabled.

Table 75: Fields on the Ping page *(continued)*

Field	Description
Wait Time (seconds)	Enter the time (in seconds) for which the source device waits for a response to the ping request packet. The source device considers the remote host as not reachable after the wait time elapses. Default: 10 seconds. Range: 0 through 600 seconds.
Incoming Interface	Click the toggle button to include or exclude (default) information (on the Ping Result page) about the interface on the source device that receives the ping responses.
Routing Instance	Select a specific routing instance that the ping request packets can use to reach the remote host. The ping result displays the information about the connectivity between the source device and the remote host based on the selected routing instance. To clear the selected routing instance, click **Clear All** and select another routing instance.

Table 76: Fields on the Ping Result page

Field	Description
Packet Loss	Displays the percentage of ping packets sent for which the source device did not receive a response.
Round Trip Time Taken (in µs)	Displays the following information about the duration (in microseconds) between the time when the source device sends the ping request and the time when the source device receives a response from the remote host. Displays the following: • Minimum: The minimum time taken to receive a response for a ping request packet. • Maximum: The maximum time taken to receive a response for a ping request packet. • Average: The average time taken to receive a response for all the ping request packets sent in a ping operation. • Standard Deviation: The variation of the round trip time from the mean round trip time.
Details	
Sequence	Sequence number of all the ping request packets.

Table 76: Fields on the Ping Result page *(continued)*

Field	Description
Result	Result of the ping request packets—Success or Failure.
Incoming Interface	Interface on the source device on which the responses are received for the ping requests. This data appears if you have enabled the Incoming Interface option on the Ping page.
Time Taken	Time taken (in microseconds) to receive response to a ping request packet.

Identifying Connectivity Issues by Using Traceroute

You can use Contrail Service Orchestration (CSO) to perform a traceroute operation from a device (provider hub, tenant device, CPE device, EX switch, enterprise hubs, or next-generation firewall device) to the remote host. Traceroute helps you view the path that a packet travels to reach the remote host. The result is useful in identifying the point of network failure in the path between the source device and remote host.

> **NOTE:** In Contrail Service Orchestration (CSO) Release 5.0, the following devices support traceroute:
>
> - EX Series: EX2300, EX3400, EX4300, EX4600, and EX4650
> - NFX Series: NFX150, NFX250
> - SRX Series: SRX300, SRX320, SRX340, SRX345, SRX1500, SRX4100, SRX4200, SRX4600
> - vSRX

To perform traceroute operation:

1. Select **Resources** > **Devices**.

   The Devices page appears.

2. Select a device from the list of devices displayed and Click **More** > **Traceroute**.

   The Traceroute page appears.

3. Complete the configuration according to the guidelines provided in Table 77 on page 296.

Fields marked with an asterisk (*) are mandatory.

4. Click **Traceroute** to initiate the traceroute operation.

   .A job is created and a Traceroute Progress page appears. If the traceroute operation is successful, the Traceroute Result page displays the traceroute parameters specified in Table 78 on page 297.

   If the traceroute operation fails, the Traceroute Result page displays an appropriate error message (such as **No response** or **No route to host**).

Table 77: Fields on the Traceroute page

Field	Description
Remote Host	Enter the IPv4 address or hostname of the remote host.
Maximum Hops	Specify the maximum number of network devices that a packet can pass through to reach the remote host. Default: 30. Range: 1 through 255. If the number of hops to reach the remote host exceeds the set value, the traceroute packet is dropped.
Advanced	
Source Interface	Select a source interface on the source device from which you want to send the packets to the remote host. Click **Clear All** to remove the selected interface and select another interface.
Hostname Resolution	Click the toggle button to enable or disable (default) the display of hostname of the hops in the path to the remote host.
Wait Time (seconds)	Enter the time until which the device waits for a response from the remote host to a packet sent before considering timeout. Default: 10 seconds. Range: 0 through 86,399 seconds.

Table 77: Fields on the Traceroute page *(continued)*

Field	Description
Routing Instance	Select a routing instance that the traceroute reqest packets can use to reach the remote host. The trace result displays the route information based on the configured routing instance type. To clear the selected routing instance, click **Clear All** and select another routing instance.

Table 78 on page 297 lists the parameters on the Traceroute Result page when the traceroute operation is successful.

Table 78: Fields on the Traceroute Result page

Field	Description
Hop	Hostname or IPv4 address of the network devices that the packet passed through to reach the remote host.
Time Taken by Packet 1 Time Taken by Packet 2 Time Taken by Packet 3	Duration (in microseconds) between the time from when the source device sends a packet, and the time it received a response from the hops and the remote host.

Remotely Accessing a Device CLI

You can use the Devices page to remotely access the CLI of CPE devices and EX Series switches, and run **show** operational commands.

You can use the **Show** operational commands to monitor a device, verify system operation, view details of specific components on a device, and so on.

To access the Devices page:

1. Select **Resources > Devices**.

   The Devices page appears.

2. Select a device from the Devices List.

> **NOTE:** You can only select a device whose operational status is marked **Up**.

3. Click **More**.

   A list of actions that you can perform on the device appears.

> **NOTE:** For dual CPE devices, the **Remote Console** option is disabled for a parent cluster device. Only member devices can select this option to access the device CLI.

4. Select the **Remote Console** option to access the device CLI.

   The Remote Terminal browser window appears, displaying the **CONNECTING TO DEVICE. PLEASE WAIT FOR PROMPT** message.

> **NOTE:** You can automatically log in to the device through the Remote Terminal browser window, without entering a username and password. If you access the device CLI through the remote terminal, root user log in is disabled.

   - If the connection is successfully established, the CLI prompt appears on the browser window. Proceed to Step 1.
   - If the connection is not established, the **Remote console connection was closed. Please close this window and open the remote console again** message appears on the browser window.

5. Enter a **show** operational command to view information about current system configuration, log files, routing tables, and so on.

   The output for the show command that you entered, appears on the same browser window.

6. Close the Remote Terminal browser window to disconnect from the device.

   The Devices page appears.

> **NOTE:** The session times out if the session remains idle for more than two minutes (default) and you are automatically logged out of the device. The **Remote console connection was closed. Please close this window and open the remote console again** message appears on the browser window.

Configuring the Firewall Device

Zones, physical interfaces, and routing instances are the basic building blocks of firewall policy and NAT policy. You can configure them from the **Resources > Devices > *Device-Name* > Configuration** page.

> **NOTE:** The **Configuration** tab that was available in earlier releases for stage-2 template-based configuration is renamed as **Configuration Template**.

To configure the firewall device:

1. Select **Resources > Devices** .

   The Devices page appears.

2. Click the device name that you want to configure.

   The *Device-Name* page appears

3. Click the **Configuration** tab.

   The physical interfaces, routing instances, and zones tabs appear.

4. Complete the configuration settings according to the guidelines provided inTable 79 on page 299.

5. Click **OK** to save the changes.

   The newly created physical interfaces, routing instances, and zones are displayed in the relevant tabs in the Configurations page.

Table 79: Fields on the Device Configuration Page

Field	Description
Physical Interfaces	
Interface Name	Name of the physical interface on the device.
Logical Interfaces	Click **View/Configure** to view or configure the logical interfaces associated with the physical interface on the device. To view and add logical interfaces for a physical interface, see “About the Logical Interfaces Page” on page 301 and “Adding a Logical Interface” on page 302.
Zones	

Table 79: Fields on the Device Configuration Page *(continued)*

Field	Description
Name	Name of the zone that you use for firewall policies and NAT policies. To add a new security zone, see "Adding a Security Zone" on page 306.
Interfaces	Interfaces associated with the zone.
Screen	Screen name for the security zone.
Description	Description for the zone.
Routing Instances	
Name	Name of the routing instances for security configuration. A routing instance is a collection of routing tables, interfaces, and routing protocol parameters.
Static Route	Click **View/Configure** to view or configure the static routes associated with a routing instance on the device.
Interfaces	Name of the interface over which the traffic flows.
Instance Type	Type of routing instance.
Description	Description of the routing instance.

About the Physical Interfaces Page

To access this page, click **Resources > Devices > *Device-Name* > Configuration**.

Use this page to view or edit the physical interfaces on the device. You can also view and configure the logical interfaces associated with the physical interfaces.

Tasks You Can Perform

You can perform the following tasks from this page:

- View details about the physical interfaces. Click the **View/Configure** button to view or configure the logical interfaces for the physical interface. See "About the Logical Interfaces Page" on page 301.
- Show or hide columns about the physical interface. Click the **Show Hide columns** icon in the top right corner of the page and select columns that you want to view on the page.

- Search for a physical interface. Click the Search icon in the top right corner of the page to search for an interface.

  You can enter partial text or full text of the keyword in the text box and press Enter.

  The search results are displayed on the same page.

Field Descriptions

Table 80 on page 301 shows the fields on the Physical Interface page.

Table 80: Fields on the Physical Interface Page

Field	Description
Interface Name	Name of the physical interface on the device.
Logical Interfaces	Click View/Configure to view or configure the logical interfaces associated with the physical interface on the device. To view and add logical interface for a physical interface, see "About the Logical Interfaces Page" on page 301 and "Adding a Logical Interface" on page 302.
MTU	Maximum transmission unit (MTU) size (in bytes) on the physical interface.
Speed	Speed in gigabytes per second (Gbps), at which data is transferred for the physical interface.
Description	Description of the physical interface.

About the Logical Interfaces Page

To access this page, click **Resources > Devices > *Device-Name* > Configuration > Physical Interfaces > View/Configure**.

Use this page to view, create, edit, or delete logical interfaces associated with a physical interface on the device.

Tasks You Can Perform

You can perform the following tasks from this page:

- View details about the logical interfaces. See Table 81 on page 302 for descriptions of the fields on the logical interfaces page.
- Add a logical Interface. See "Adding a Logical Interface" on page 302.
- Edit, delete, or deploy logical interfaces. See "Editing, Deleting, and Deploying Logical Interfaces" on page 305.
- Clear all selected logical interfaces. Select the logical interface and then right-click or click **More > Clear All Selections**.
- Show or hide columns that contain information about the logical interface. Click the **Show Hide columns** icon in the top right corner of the page and select columns that you want to view on the page.
- Search for logical interfaces using keywords. Click the search icon. Enter partial text or full text of the keyword in the search bar and press Enter. The search results are displayed.

Field Descriptions

Table 81 on page 302 shows the fields on the Logical Interfaces page.

Table 81: Fields on the Logical Interfaces Page

Field	Description
Interface Name	Displays the name of the logical interface associated with a physical interface on the device.
IPv4 Address	Displays the IPv4 address for the logical interface.
IPv6 Address	Displays the IPv6 address for the logical interface.
VLAN ID	Displays the VLAN ID for the 802.1q VLAN tags.
Description	Displays the description of the logical interface.

Adding a Logical Interface

For a physical interface to function, you must configure at least one logical interface on that device. . You can also configure other logical interface properties.

The logical properties of an interface are the characteristics that do not apply to the physical interface. Logical properties include:

- IP address or addresses associated with the interface. A logical interface can be configured with an IPv6 address, IPv4 address, or both. The IP specification requires a unique address on every interface of each system attached to an IP network, so that traffic can be correctly routed. Individual hosts such as home computers must have a single IP address assigned. Devices must have a unique IP address for every interface.
- Virtual LAN (VLAN) tagging

To create logical interface for a physical interface:

1. Select **Resources > Devices** .

   The Devices page appears.

2. Click the device name that you want to configure.

   The *Device-Name* page appears

3. Click the **Configuration** tab.

   The Physical Interfaces, Routing Instances, and Zones tab appears.

4. Click **Physical Interfaces** tab.

   The Physical Interfaces page appears.

5. Select a physical interface and click **View/Configure**.

   The Logical Interfaces page appears.

6. Click the plus icon **(+)** .

   The Create Logical Interface page appears.

7. Complete the configuration settings according to the guidelines provided inTable 82 on page 303.

8. Click **OK** to save the changes.

Table 82: Fields on the Logical Interfaces Page

Field	Description
Basic Information	
Unit	Enter the number of the logical interface. Range: 0 through 2,147,483,647.

Table 82: Fields on the Logical Interfaces Page *(continued)*

Field	Description
Description	Enter a description for the logical interface. The maximum number of characters is 255.
VLAN ID	Enter the VLAN ID for the 802.1q VLAN tags.
IPv4 Address	Click the plus icon **(+)** . The Add - Address(IPv4) page appears.
IPv4 Address	Enter an IPv4 address for the logical interface.
Subnet	Enter the subnet for the IPv4 address. Range: 0 through 32
Primary	Select this check box to mark the IPv4 address as the primary address of the protocol on the logical interface. If the logical unit has more than one IP address, the primary IP address is used by default as the source address when packet transfer originates from the interface and the destination address does not indicate the subnet.
Preferred	Select this check box to specify that the IPv4 address is the preferred address for the logical interface. If you configure more than one IP address on the same subnet, the preferred source address is chosen by default as the source address when you initiate frame transfers to destinations on the subnet.
IPv6 Address	Click the plus icon **(+)** . The Add - Address(IPv6) page appears.
IPv6 Address	Enter an IPv6 address for the logical interface.
Subnet	Enter the subnet for the IPv6 address. Range: 0 through 32
Primary	Select this check box to specify that the IPv6 address is the primary address of the protocol on the logical interface. If the logical unit has more than one IP address, the primary IP address is used by default as the source address when packet transfer originates from the interface and the destination address does not indicate the subnet.
Preferred	Select this check box to specify that the IPv6 address is the preferred address for the logical interface. If you configure more than one IP address on the same subnet, the preferred source address is chosen by default as the source address when you initiate frame transfers to destinations on the subnet.

Editing, Deleting, and Deploying Logical Interfaces

IN THIS SECTION

- Editing Logical Interfaces | 305
- Deleting Logical Interfaces | 306
- Deploying Logical Interfaces | 306

You can edit, delete, and deploy logical interfaces from the **Logical Interfaces** page.

Editing Logical Interfaces

To modify the parameters configured for a logical interface:

1. Select **Resources > Devices >*Device-Name* > Configuration > Physical Interface > View/Configure**.

   The **Logical Interfaces** page appears.

2. Select the logical interface that you want to edit, and then click on the edit icon (pencil symbol) on the right top corner of the table, or right-click and select **Edit**.

   The **Edit Logical Interface** page appears.

3. Modify the parameters according to the guidelines provided in "Adding a Logical Interface" on page 302.

   NOTE: You can modify only some fields when you are editing a logical interface.

4. Click **OK** to save the changes. If you want to discard your changes, click **Cancel** instead.

If you click **OK**, the modified values appear on the **Logical Interfaces** page.

Deleting Logical Interfaces

To delete a logical interface:

> **NOTE:** You cannot delete a logical interfaces that is associated with a physical interface on the device.

1. Select **Resources > Devices >*Device-Name* > Configuration > Physical Interface > View/Configure**.

   The **Logical Interfaces** page appears.

2. Select one or more logical interfaces that you want to delete and then click the delete icon.

   A page requesting confirmation for the deletion appears.

3. Click **Yes** to delete the selection. If you do not want to delete, click **Cancel** instead.

If you click **Yes**, the selected logical interface is deleted from the **Logical Interfaces** page.

Deploying Logical Interfaces

To deploy a logical interface:

1. Select **Resources > Devices >*Device-Name* > Configuration > Physical Interface > View/Configure**.

   The **Logical Interfaces** page appears.

2. Select one or more logical interfaces that you want to deploy and then click **Deploy**.

   A job is created. Click the job link or go to the Jobs page and view the status of the deployment.

RELATED DOCUMENTATION

Adding a Logical Interface | 302

Adding a Security Zone

A security zone is a collection of one or more network segments requiring the regulation of inbound and outbound traffic through policies. Security zones are logical entities to which one or more interfaces are

bound. You can define multiple security zones, the exact number of which you determine based on your network needs.

An interface for a security zone can be thought of as a doorway through which TCP/IP traffic can pass between that zone and any other zone. Through the policies you define, you can permit traffic between zones to flow in one direction or in both. With the routes that you define, you specify the interfaces that traffic from one zone to another must use. Because you can bind multiple interfaces to a zone, the routes you chart are important for directing traffic to the interfaces of your choice. An interface can be configured with an IPv4 address, IPv6 address, or both.

Security zones have the following properties:

- Policies—Active security policies that enforce rules for the transit traffic, in terms of what traffic can pass through the firewall, and the actions that need to take place on the traffic as it passes through the firewall.
- Screens—A Juniper Networks stateful firewall secures a network by inspecting, and then allowing or denying, all connection attempts that require passage from one security zone to another. For every security zone, you can enable a set of predefined screen options that detect and block various kinds of traffic that the device determines as potentially harmful.
- TCP-RST—When this feature is enabled, the system sends a TCP segment with the RESET flag set when traffic arrives that does not match an existing session and does not have the SYNchronize flag set.
- Interfaces—List of interfaces in the zone.

Use this page to configure zones and assign interfaces to them.

To create a security zone:

1. Select **Resources > Devices** .

   The Devices page appears.

2. Click the device name that you want to configure.

   The *Device-Name* page appears

3. Click the **Configuration** tab.

   The Physical Interfaces, Routing Instances, and Zones tab appears.

4. Click **Zones** tab.

   The Zones page appears.

5. Click the plus icon **(+)** .

   The Add New Zone page appears.

6. Complete the configuration settings according to the guidelines provided inTable 83 on page 308.

7. Click **OK** to save the changes.

Table 83: Fields on the Add New Zone Page

Field	Description
General Information	
Name	Enter a unique string of alphanumeric characters, and some special characters, such as dashes, and underscores.The maximum length is 31 characters.
Description	Enter a description for the zone; the maximum length is 900 characters.
Application Tracking	Select the checkbox to maintain application usage statistics on a device.
Interfaces	From the list of interfaces in the Available column, select the interfaces that you want to include in the new zone and click the greater-than icon (>). The selected interfaces are moved to the Selected column.
System Services	From the list of system services in the Available column, select the system services that you want to include in the new zone and click the greater-than icon (>). The selected system services are moved to the Selected column.
Is Except	Select the checkbox to disable specific incoming system service traffic, only when **all** system services option is defined.
Protocols	From the list of protocols in the Available column, select the protocols that you want to include in the new zone and click the greater-than icon (>). The selected protocols are moved to the Selected column.
Is Except	Select this option to disable specific incoming protocol traffic, only when **all** protocols option is defined.
Traffic Control Options	
TCP RST	Select the checkbox to enable sending TCP packets with the RST (reset) flag set to 1 in response to TCP packets with any flag other than SYN set and that do not belong to an existing session.
Screen	Enter a predefined security screen for a security zone to detect and block various kinds of traffic that the device determines as potentially harmful.

Table 83: Fields on the Add New Zone Page *(continued)*

Field	Description
Interface Services and Protocols	View the summary of interface, services and protocols for your device.

Adding a Routing Instance

A routing instance is a collection of routing tables, interfaces, and routing protocol parameters. The set of interfaces belongs to the routing tables, and the routing protocol parameters control the information in the routing tables. There can be multiple routing tables for a single routing instance—for example, unicast IPv4, unicast IPv6, and multicast IPv4 routing tables can exist in a single routing instance. Routing protocol parameters and options control the information in the routing tables.

Each routing instance consists of sets of the following:

- Routing tables
- Interfaces that belong to these routing tables (optional, depending upon the routing instance type)
- Routing option configurations

You can configure the following routing instances:

- Forwarding—Use this routing instance type for filter-based forwarding applications. For this instance type, there is no one-to-one mapping between an interface and a routing instance. All interfaces belong to the default instance inet.0.
- Virtual router—Use this for non-VPN-related applications. There are no virtual routing and forwarding (VRF) import, VRF export, VRF target, or route distinguisher requirements for this instance type.

To create a routing instance:

1. Select **Resources > Devices** .

   The Devices page appears.

2. Click the device name that you want to configure.

   The *Device-Name* page appears

3. Click the **Configuration** tab.

   The Physical Interfaces, Routing Instances, and Zones tab appears.

4. Click **Routing Instances** tab.

   The Routing Instances page appears.

5. Click the plus icon **(+)** .

   The Create Routing Instance page appears.

6. Complete the configuration settings according to the guidelines provided inTable 84 on page 310.

7. Click **OK** to save the changes.

Table 84: Fields on the Create Routing Instance Page

Field	Description
General Settings	
Name	Enter a unique string of alphanumeric characters, dashes, and underscores. The maximum length is 31 characters.
Description	Enter a description for the zone; the maximum length is 900 characters.
Instance Type	Select the routing instance type from the list. Select virtual-router for non-VPN-related application. Select forwarding for filter-based forwarding applications where interfaces are not associated with instances.
Interfaces	From the list of interfaces in the Available column, select the interfaces that you want to include in the new routing instance and click the greater-than icon (>). The selected interfaces are moved to the Selected column.

About the Static Routes Page

To access this page, click **Resources > Devices > *Device-Name* > Configuration > Routing Instances> View/Configure**.

Use this page to view, create, edit, or delete static routes for the routing instance.

Tasks You Can Perform

You can perform the following tasks from this page:

- View details about the static route. See Table 85 on page 311 for descriptions of the fields on the static routes page.
- Add a static route. See "Adding a Static Route" on page 311.
- Edit, delete, or deploy static routes. See "Editing, Deleting, and Deploying Static Routes" on page 314.
- Clear all selected static routes. Select the static route and then right-click or click **More > Clear All Selections**.
- Show or hide columns that contain information about the static route. Click the **Show Hide columns** icon in the top right corner of the page and select columns that you want to view on the page.
- Search for static routes using keywords. Click the search icon. Enter partial text or full text of the keyword in the search bar and press Enter. The search results are displayed.

Field Descriptions

Table 85 on page 311 shows the fields on the Static Routes page.

Table 85: Fields on the Static Routes Page

Field	Description
IP Address	View the IP address of the static route.
Next Hop	View the Ipv4 or IPv6 address for the next hop.
Next Table	View the name of the next routing table to the destination.
Metric	View the metric value that signifies the cost for an access route, for the next hop
Status	View the status of the static route.

Adding a Static Route

Routes that are permanent fixtures in the routing and forwarding tables are often configured as static routes. These routes generally do not change, and often include only one or very few paths to the destination.

To create a static route in the routing table, you must, at minimum, define the route as static and associate a next-hop address with it. The static route in the routing table is inserted into the forwarding table when

the next-hop address is reachable. All traffic destined for the static route is transmitted to the next-hop address for transit.

You can specify options that define additional information about static routes that is included with the route when it is installed in the routing table. All static options are optional.

To create a static route for a routing instance:

1. Select **Resources > Devices** .

   The Devices page appears.

2. Click the device name that you want to configure.

   The *Device-Name* page appears

3. Click the **Configuration** tab.

   The Physical Interfaces, Routing Instances, and Zones tab appears.

4. Click **Routing Instances** tab.

   The Routing Instances page appears.

5. Select a routing instance and click **View/Configure** link in the **Static Route** column.

   The Static Routes page appears.

6. Click the plus icon **(+)** .

   The Create Static Route page appears.

7. Complete the configuration settings according to the guidelines provided in Table 86 on page 312.

8. Click **OK** to save the changes.

Table 86: Fields on the Create Static Route Page

Field	Description
Basic Information	
IP Address	Enter the IPv4 or IPv6 address depending on the type of IP address specified.
Subnet	Enter the subnet for the IPv4 address or the prefix for the IPv6 address.
Next Hop	

Table 86: Fields on the Create Static Route Page *(continued)*

Field	Description
IP Address	Enter an IPv4 or IPv6 address for the next hop depending on the type of IP address specified for the static route.
Interface	Select the interface name to be used as the next hop.
Qualified Next Hop	
IP Address	Enter an IPv4 or IPv6 address for the qualified next hop depending on the type of IP address specified for the static route.
Interface	Select the interface name to be used as the qualified next hop.
Preference	Enter the preference for the qualified next hop; the lower the number, the higher the route preference.
Metric	Enter a metric value to signify the cost for an access route for the qualified next hop.
Next Table	
Next Table	Select the name of the next routing table to the destination.
Advanced Options	
Preference	Enter the preference for the next hop; the lower the number, the higher the route preference. Range: 0 through 2147483647
Metric	Enter a metric value which signifies the cost for an access route, for the next hop Range: 0 through 2147483647
Discard	Specify whether to drop packets to destination; send no ICMP unreachable
Resolve Choices	Select whether indirectly connected next hops must be resolved (Resolve) or not (Do not resolve). Select None if no action is required.
Retain Choices	Select whether the route must be retained (Retain) or deleted from the forwarding table (Do not retain) when the routing protocol process shuts down normally. Select None if no action is required.

Table 86: Fields on the Create Static Route Page *(continued)*

Field	Description
Install Choices	Select whether the route must be installed in the forwarding table (Install) or not (Do not install). Select None if no action is required.
Re-advertise Choices	Select whether the route must be re-advertised by routing protocols (Re-advertise) or not (Do not re-advertise). Select None if no action is required.

Editing, Deleting, and Deploying Static Routes

IN THIS SECTION

- Editing Static Routes | **314**
- Deleting Static Routes | **315**
- Deploying Static Routes | **315**

You can edit, delete, and deploy static routes from the **Static Routes** page.

Editing Static Routes

To modify the parameters configured for a static route:

1. Select **Resources > Devices >*Device-Name* > Configuration > Routing Instances > View/Configure**.

   The **Static Routes** page appears.

2. Select the static route that you want to edit, and then click on the edit icon (pencil symbol) on the right top corner of the table, or right-click and select **Edit**.

   The **Edit static route** page appears.

3. Modify the parameters according to the guidelines provided in “Adding a Static Route” on page 311.

4. Click **OK** to save the changes. If you want to discard your changes, click **Cancel** instead.

If you click **OK**, the modified values appear on the **Static Routes** page.

Deleting Static Routes

To delete a static route:

1. Select **Resources > Devices >*Device-Name* > Configuration > Routing Instances > View/Configure**.

   The **Static Routes** page appears.

2. Select one or more static routes that you want to delete and then click the delete icon.

   A page requesting confirmation for the deletion appears.

3. Click **Yes** to delete the selection. If you do not want to delete, click **Cancel** instead.

If you click **Yes**, the selected static route is deleted from the **Static Routes** page.

Deploying Static Routes

To deploy a static route:

1. Select **Resources > Devices >*Device-Name* > Configuration > Routing Instances > View/Configure**.

   The **Static Routes** page appears.

2. Select one or more static routes that you want to deploy and then click **Deploy**.

   A job is created. Click the job link or go to the Jobs page and view the status of the deployment.

RELATED DOCUMENTATION

Adding a Static Route | 311

CHAPTER 11

Managing Configuration Templates

IN THIS CHAPTER

- About the Configuration Templates Page | 316
- Edit, Clone, and Delete Configuration Templates | 319
- Deploy Configuration Templates to Devices | 322
- Preview and Render Configuration Templates | 327
- Import Configuration Templates | 328
- Assign Configuration Templates to Device Templates | 330
- Add Configuration Templates | 332
- View the Configuration Deployed on Devices | 339

About the Configuration Templates Page

IN THIS SECTION

- Tasks You Can Perform | 317
- Field Descriptions | 317

To access this page, click **Resources > Templates > Configuration Templates** in Customer Portal.

You can use the Configuration Templates page to view and manage configuration templates.

NOTE: In CSO releases before Release 5.1.0, configuration templates are called stage-2 configuration templates.

Tasks You Can Perform

In Customer Portal, users with the Tenant Administrator role can perform the following tasks from this page, while users with operator roles only have read capabilities.

- Clone a configuration template—"Edit, Clone, and Delete Configuration Templates" on page 319.
- Deploy a configuration template on one or more devices—See "Deploy Configuration Templates to Devices" on page 322.
- Preview and render a configuration template—See "Preview and Render Configuration Templates" on page 327.
- View the details a configuration template—Select a configuration template and click **More** > **Template Details** or mouse over the configuration template click the Detailed View icon. The Detail for *Template-Name* pane appears on the right side of the page. See Table 88 on page 318 for an explanation of the fields.
- Import a configuration template—See "Import Configuration Templates" on page 328.
- Assign a configuration template to a device template—See "Assign Configuration Templates to Device Templates" on page 330.
- Add a configuration template—See "Add Configuration Templates" on page 332.
- Edit or delete configuration templates—See "Edit, Clone, and Delete Configuration Templates" on page 319.
- View the configuration deployed on one or more devices—See "View the Configuration Deployed on Devices" on page 339.
- Search for configuration templates by using keywords—Click the search icon and enter the search term in the text box and press Enter. The search results are displayed on the same page.
- Sort configuration templates—Click a column name to sort the configuration templates based on the column name.

  NOTE: Sorting and filtering is applicable only to some fields.

- Show or hide columns—Click the **Show Hide Columns** icon at the top right corner of the page and select the columns that you want displayed on the Configuration Templates page.

Field Descriptions

Table 87 on page 318 displays the description of the fields on the Configuration Templates page and Table 88 on page 318 displays the description of the fields on the Detail for *Template-Name* Pane.

Table 87: Fields on the Configuration Templates Page

Field	Description
Name	Name of the configuration template.
Family	Device family to which the configuration template belongs.
Description	Description of the configuration template.
Deployed Devices	Number of devices on which the configuration template was deployed. If the configuration template is not yet deployed on any devices then a blank cell is displayed. Click the ***number-of-devices*** link to view the configuration (for that configuration template) deployed on devices. See "View the Configuration Deployed on Devices" on page 339.
Last Updated	Date and time on which the template was last updated.
Owner	Depending on who added the configuration template, displays the following • **System**—If the template is predefined or added by the Service Provider administrator • ***Tenant-Name***—Name of the tenant if the template was added by an tenant administrator.

Table 88: Fields on the Detail for <Template-Name> Pane

Field	Description
General tab	
Name	See Table 87 on page 318.
Description	See Table 87 on page 318.
Family	See Table 87 on page 318.
Format	Format used by the configuration template: • CLI • XML (Extensible Markup Language)
Details tab	**NOTE:** If you want to add a new configuration template based on an existing one, you can copy the three files from the Details tab, modify the files as needed, and use the Import Configuration Template page to import a new template.
Jinja Template	Displays the configuration in Jinja Template language syntax.
Data Model	Displays the Yang data model (configuration schema).

Table 88: Fields on the Detail for <Template-Name> Pane *(continued)*

Field	Description
View Def	Displays the View Def (GUI configuration).

RELATED DOCUMENTATION

About the Devices Page | 268

Edit, Clone, and Delete Configuration Templates

IN THIS SECTION

- Edit a Configuration Template | 319
- Clone a Configuration Template | 320
- Delete a Configuration Template | 321

Users with the Tenant Administrator role can modify the parameters of existing configuration templates, clone existing configuration templates, and delete configuration templates that are no longer being used.

NOTE: In CSO releases before Release 5.1.0, configuration templates are called stage-2 configuration templates.

Edit a Configuration Template

Users with the Tenant Administrator role can edit only the templates that they added (created).

To modify a configuration template:

1. Select **Resources > Templates > Configuration Templates**.

   The Configuration Templates page appears.

2. Select the configuration template that you want to modify and click the edit (pencil) icon.

   The Edit Configuration Template page appears. The fields on this page are same as the fields that you configure in the Add Configuration Template workflow.

3. Modify the fields as needed. You can modify all the fields except the name of the configuration template.

   Refer to "Add Configuration Templates" on page 332 for an explanation of the fields.

   > **NOTE:** Fields marked with an asterisk (*) are mandatory.

4. Click **OK**.

   The modifications are saved and you are returned to the Configuration Templates page, where a confirmation message is displayed. If the configuration template was previously deployed on a device or assigned to a device template, then you must redeploy the configuration template for the changes to take effect.

Clone a Configuration Template

To clone a configuration template:

1. Select **Resources > Templates > Configuration Templates**.

   The Configuration Templates page appears.

2. Select the configuration template that you want to clone and click **Clone**.

   The Clone Configuration Template page appears.

3. In the **Template Name** field, enter a unique template name that can only contain alphanumeric characters and hyphens up to a maximum of 15 characters.

4. Click **OK**.

   You are returned to the Configuration Templates page and a confirmation message appears at the top of the page indicating the status of the clone operation.

After a template is cloned successfully, you can modify the template if needed. See the preceding section for details.

Delete a Configuration Template

To delete a configuration template:

> **NOTE:**
> - You cannot delete predefined configuration templates.
> - You can delete a configuration template only if the following conditions hold good:
>   - You added (created) the template.
>   - The template is not assigned to a device template.
>   - The template is not deployed on a device.

1. Select **Resources > Templates > Configuration Templates**.

   The Configuration Templates page appears.

2. Select the configuration template that you want to delete and click the **X** (delete) icon.

   You are asked to confirm the delete operation.

3. Click **Yes**.

   You are returned to the Configuration Templates page and a popup appears indicating whether the deletion was successful or not.

RELATED DOCUMENTATION

Preview and Render Configuration Templates | 327

Deploy Configuration Templates to Devices

IN THIS SECTION

- Deploy from the Configuration Templates Page | 322
- Deploy from the Devices Page | 326

In Customer Portal, users with the Tenant Administrator role can deploy a configuration template directly on one or more devices that were previously activated. This enables you to deploy configuration templates added after a device was activated or to deploy additional configuration to devices.

You can deploy configuration templates to devices from the Configuration Templates or the Devices pages.

> **NOTE:** In CSO releases before Release 5.1.0, configuration templates are called stage-2 configuration templates.

Deploy from the Configuration Templates Page

To deploy a configuration template to one or more devices:

1. Select **Resources > Templates > Configuration Templates**.

   The Configuration Templates page appears.

2. Select the configuration template that you want deploy and click **Deploy to Devices**.

   The Deploy Template *Template-Name* To Devices page appears.

3. Complete the configuration according to the guidelines provided in Table 89 on page 323

   > **NOTE:** Fields marked with an asterisk (*) are mandatory.

4. Click **OK**.

The settings that you entered are saved and you are returned to the Configuration Templates page. A confirmation message appears indicating that a job was created. For each device, a separate job is triggered to deploy the configuration.

You can view the status of the jobs from the Jobs page (**Monitor > Jobs**).

Table 89: Deploy Template <Template-Name> To Devices Settings

Setting	Guideline
Select Devices	
Configuration Template	Displays the name of the configuration template that you are deploying; you cannot modify this field.
Component Name	This field is displayed only for NFX250 devices. Select the component of the NFX250 device on which to deploy the template: • JCP—Junos Control Plane • JDM—Junos Device Manager • GWR-Gateway Router

Table 89: Deploy Template <Template-Name> To Devices Settings *(continued)*

Setting	Guideline
Devices	You can specify the devices on which you want to deploy the configuration template in the following ways: • By adding the devices manually: 1. From the list of devices displayed, select one or more devices by clicking the check box next to each device name. **NOTE:** You can search for devices or filter the list of devices displayed. • By uploading a comma-separated values (CSV) file containing the device information: **NOTE:** You must ensure that the CSV file is in the format that CSO can read and that the number of device records is 200 or lower. You can download a sample file by clicking the **Download Sample CSV File** button. 1. Click **Upload CSV File**. The Upload CSV File page appears. 2. Click **Browse** to open the file selection dialog, select a file, and click **Open**. The name of the file that you selected is displayed in the CSV File field. 3. Click **OK**. You are returned to the previous page where the devices that you imported are selected and displayed in the table. Click **Next**. You are taken to the Configure Global Parameters or the Configure Device Parameters tab.
Configure Global Parameters	**NOTE:** This tab is displayed only if the configuration template contains parameters that are global in scope. Specify the global parameters that are common to all the devices that you selected in the preceding step. After you are done, click **Next**. You are taken to the Configure Device Parameters tab.
Configure Device Parameters	

Table 89: Deploy Template <Template-Name> To Devices Settings *(continued)*

Setting	Guideline
Devices	The devices that you selected in the preceding step are displayed in the Devices table, and the first device is selected by default. For each device, the device name, device family, operational status, and the configuration status are displayed. When you first arrive on this tab, the configuration status for each device is *Not configured*. The *Device-Name* Parameters pane on the right displays the input parameters (from the configuration template) that you can specify for each device. After you specify the values for one device, you can select a different device and enter the configuration values. • If the configuration template contains validations for the parameters, CSO validates the values you entered for the device and changes the configuration status to Valid and displays a green check mark (√). • If the configuration template does not contain any validations, CSO changes the configuration status to Valid and displays a green check mark (√). • If the values that you entered do not match the validation, the configuration status displays Invalid. **NOTE:** You can optionally delete a device by selecting the device and clicking the delete (trash can) icon. After you specify the input parameter values for all the devices and ensure that the configuration status of all devices is Valid, click **Next**. You are taken to the Summary tab.
Summary	
Devices	The devices that you selected in the preceding step are displayed in the Devices table, and the first device is selected by default. For each device, the device name, device family, and operational status are displayed. For each device, the *Device-Name* Configuration pane on the right displays the actual configuration that will be deployed on the device. After you review the configuration for all the devices, click **Next**. You are taken to the Deploy tab.
Deploy	

Table 89: Deploy Template <Template-Name> To Devices Settings *(continued)*

Setting	Guideline
Deployment Schedule	Specify whether the configuration should be deployed on devices immediately(**Deploy now**) or deployed later (**Deploy later**). If you choose to deploy the configuration later, you must enter the date (in MM/DD/YYYY format) and time (in HH:MM:SS 24-hour or AM/PM format) that you want the deployment to occur.

Deploy from the Devices Page

To deploy a configuration template to one or more devices:

1. Select **Resources > Devices**.

   The Devices page appears.

2. Select the device on which you want deploy and click **More > Deploy Configuration**.

   NOTE: The devices that you select must belong to the same device family. If you select devices from different device families, CSO displays an error message.

   The Deploy Configuration to Devices page appears displaying the selected device in the Devices table.

3. From the **Configuration Template** field, choose the configuration template that you want to deploy.

   The configuration templates displayed are filtered based on the device family of the devices that you selected.

4. The rest of the deploy workflow is the same as you encounter if you initiate the deployment from the Configuration Templates page. Complete the configuration according to the guidelines provided in Table 89 on page 323

   NOTE: Fields marked with an asterisk (*) are mandatory.

5. Click **OK**.

The settings that you entered are saved and you are returned to the Devices page. A confirmation message appears indicating that a job was created. For each device, a separate job is triggered to deploy the configuration.

You can view the status of the jobs from the Jobs page (**Monitor > Jobs**).

RELATED DOCUMENTATION

Assign Configuration Templates to Device Templates | 330

Preview and Render Configuration Templates

In Customer Portal, users with the Tenant Administrator role can use the Preview workflow to validate a configuration template by entering values for the configuration template and then rendering the template to view the configuration.

Although this is not mandatory, we recommend that you use this workflow to validate a configuration template before attaching it to a device template or deploying it on a device.

> **NOTE:** In CSO releases before Release 5.1.0, configuration templates are called stage-2 configuration templates.

To preview and render a configuration template:

1. Select **Resources > Templates > Configuration Templates**.

   The Configuration Templates page appears.

2. Select the configuration template that you want to check and click **Render Configuration**.

   The Preview Configuration page appears displaying the parameters configured for the template.

3. Specify values for the parameters as needed.

   > **NOTE:** Fields marked with an asterisk (*) are mandatory.

4. After you have entered the necessary parameters, click **Render**.

The Rendered Config page appears displaying the configuration rendered based on the configuration template and the values that you specified.

5. Check if the configuration was rendered correctly.

   If the configuration was not rendered correctly, you can modify the configuration template as needed. See "Edit, Clone, and Delete Configuration Templates" on page 319.

6. Click **OK**.

   You are returned to the Preview Configuration Template page.

7. Click **Cancel** to exit the Preview Configuration Template page.

   You are returned to the Configuration Templates page. You can assign the configuration template to one or more device templates.

RELATED DOCUMENTATION

Assign Configuration Templates to Device Templates | 330

Import Configuration Templates

In Customer Portal, users with the Tenant Administrator role can import a configuration template by specifying the parameters using a configuration template file (Jinja template language), Yang model file (schema for the configuration), and the Viewdef file (configuration of the UI).

> **NOTE:** In CSO releases before Release 5.1.0, configuration templates are called stage-2 configuration templates.

To import a configuration template:

1. Select **Resources > Templates > Configuration Templates**.

   The Configuration Templates page appears.

2. Select **More** > **Import**.

   The Import Configuration Template page appears.

3. Complete the configuration according to the guidelines provided in Table 90 on page 329.

> **NOTE:** Fields marked with an asterisk (*) are mandatory.

4. Click **OK**.

   You are returned to the Configuration Templates page and a popup appears displaying the status of the import operation.

5. Click **OK** to close the popup.

   You are returned to the Configuration Templates page.

If the configuration template is imported successfully, you can validate the configuration template by using the Preview workflow and then assign the configuration template to a device template or deploy it on a device.

Table 90: Import Configuration Template Settings

Setting	Guideline
Template Name	Enter a unique name that can only contain alphanumeric characters and hyphens; 15-character maximum.
Description	Enter a description for the configuration template.
Output Config Format	Select the output configuration format for the template: • CLI (default) • XML
Device Family	Select the device family for which you are adding the template; for example, juniper-nfx.
Configuration Template File	Specify the file containing the configuration (in Jinja Template language syntax) by clicking the **Browse** button to navigate to the directory where the configuration template file is located and selecting the file.
Yang Model File	Specify the Yang data model (configuration schema) file by clicking the **Browse** button to navigate to the directory where the Yang model file is located and selecting the file.

Table 90: Import Configuration Template Settings *(continued)*

Setting	Guideline
Viewdef File	Specify the Viewdef file, which contains the configuration for the UI, by clicking the **Browse** button to navigate to the directory where the Viewdef file is located and selecting the file.

RELATED DOCUMENTATION

Preview and Render Configuration Templates | 327

Deploy Configuration Templates to Devices | 322

Assign Configuration Templates to Device Templates

In Customer Portal, users with the Tenant Administrator role can assign a configuration template to one or more device templates. Associating a configuration template with a device template enables you to deploy additional configuration on the device during ZTP and after the device is activated.

> **NOTE:** In CSO releases before Release 5.1.0, configuration templates are called stage-2 configuration templates.

To assign a configuration template to one or more device templates:

1. Select **Resources > Templates > Configuration Templates**.

   The Configuration Templates page appears.

2. Select the configuration template that you want to assign and select **More** > **Assign to Device Template**.

   The Assign Configuration Template to Device Templates page appears.

3. Complete the configuration according to the guidelines provided in Table 91 on page 331

   > **NOTE:** Fields marked with an asterisk (*) are mandatory.

4. Click **OK**.

You are returned to the Configuration Templates page and a popup appears indicating whether the assignment is successful or has failed. If the assignment failed, you can retry the assignment or contact Juniper Networks support.

If the assignment is successful, the configuration parameters are displayed in the *Device-Name* page and you can enter values for the configuration and deploy the changes on the device.

Table 91: Assign Configuration Template to Device Template Settings

Setting	Guideline
Template Settings	
Template	Displays the name of the configuration template that you are assigning; you cannot modify this field.
Display Name	Enter the name that you want displayed on the *Device-Name* page.
Component Name	For NFX250 devices, select the component name to which the configuration should be deployed. The components that are currently supported are: • Juniper Device Manager (JDM)—JDM is a Linux container that manages software components. • Juniper Control Plane (JCP)—JCP is the Junos VM running on the hypervisor. Administrators can use JCP to configure the network ports of the NFX250 device. JCP is used to configure the switching and routing function on the NFX250 device. • Gateway Router (GWR)—vSRX as a gateway provides the same capabilities as Juniper Networks SRX Series Services Gateways in a virtual form factor.
Auto Deploy	Specify whether the configuration should be deployed automatically on the device during the zero touch provisioning (ZTP) process. The options are: • **Yes**—Deploy the configuration automatically on the device during ZTP. • **No** (Default)—Don't deploy the configuration is not deployed automatically on the device during ZTP. • **Same as global settings**—Use the same settings as the one configured in the device template.
Device Templates	
Select Device Templates	The list of device templates to which you can assign the configuration template are displayed in a grid along with some information about the template. CSO displays only those device templates whose device family matches the device family of the configuration template. Select one or more device templates to which you want to assign the configuration template.

RELATED DOCUMENTATION

Deploy Configuration Templates to Devices | 322

Add Configuration Templates

NOTE: In CSO releases before Release 5.1.0, configuration templates are called stage-2 configuration templates.

In Customer Portal, users with the Tenant Administrator role can add a configuration template by providing the device configuration using the Jinja template language syntax.

NOTE:

- Before you add the configuration template, ensure that you have the device configuration ready.
- We recommend that you use a working configuration on the device to add the configuration template.

To add a configuration template:

1. Select **Resources > Templates > Configuration Templates**.

   The Configuration Templates page appears.

2. Click the **+** (add) icon.

   The Add Configuration Template page (wizard) appears.

   NOTE: Fields marked with an asterisk (*) are mandatory.

3. Configure the fields on the Basic Information tab according to the guidelines provided in Table 92 on page 333.

   Click **Next** to go to the Templatize Config tab.

4. Add the configuration on the Templatize Config tab. Refer to Table 93 on page 334 for an explanation of the actions on this tab

Click **Next** to go to the Customize Variables tab, where the template parameters that you configured are displayed in a grid.

5. You can modify the parameters for one or more rows as follows:

   a. Click inside a row to modify the values for the fields in that row according to the guidelines provided in Table 94 on page 334.

   b. Click **✓** (check mark) to save your changes or **X** to discard the changes.

   After you are finished, click **Next** to go to the Generated UI tab, where the UI for the variables that you entered is generated and displayed after a few seconds.

6. Perform one or more actions on this tab, as explained in Table 95 on page 336.

7. Click **OK**.

   The configuration template is added and you are returned to the Configuration Templates page, where a confirmation message is displayed. You can assign the configuration template to device templates or deploy the template on devices.

Table 92: Basic Information Settings (Add Configuration Template Page)

Setting	Guideline
Template Name	Enter a unique name that can only contain alphanumeric characters and hyphens; 15-character maximum.
Description	Enter a description for the configuration template.
Configuration Format	Select the output configuration format for the template: • CLI (default) • XML
Device Family	Select the device family for which you are adding the template; for example, juniper-nfx.
	Click **Next** to continue.

Table 93: Templatize Config Actions (Add Configuration Template Page)

Action	Description
View the sample configuration and detected template variables	The Templatize Config tab displays a sample Jinja configuration template in an inline editor on the left and the corresponding template variables in the Detected Template Variables pane on the right.
Download a sample configuration	You can download a sample configuration by clicking the **here** link near the top of the tab.
Add the device configuration	To add the device configuration: 1. In the inline editor, copy and paste the device configuration ensuring that the syntax follows the Jinja Template language. CSO detects the template variables corresponding to the configuration that you entered and displays them in the Detected Template Variables pane. 2. Check that the template variables detected match the configuration that you added to the template: • (Optional) If the template variables do not match, check the Jinja syntax that you used for the template configuration and make any changes needed in the inline editor as indicated in the first step.
	If the template variables detected match the configuration that you added to the template, click **Next** to continue. CSO validates the Jinja template and displays an error message if there are any errors.

Table 94: Customize Variables Settings (Add Configuration Templates Page)

Setting	Guideline
Detected Variables	Displays the name of the detected variable. You cannot edit the name.

Table 94: Customize Variables Settings (Add Configuration Templates Page) *(continued)*

Setting	Guideline
Data Type	Select the data type for the variable: **NOTE:** You cannot specify a data type for a top-level variable that has one or more leaf nodes under it. • **String** (default)—If the variable is a string of characters. • **Boolean**—If the variable is a boolean value (true or false). • **Number**—If the variable is a number. • **Enumeration**—If the variable is an enumerated values with all strings or all numbers. In the Add enumeration for *variable-name* that appears, enter the list of enumerated values for the variable, separating each one by entering a space or pressing Enter. Click **OK** to save the values that you entered and go back to the Customize Variables tab. The values you entered are displayed in the Enumerated Values field. • **IPv4**—If the variable is an IPv4 address. • **IP Prefix**—If the variable is an IPv4 prefix.
Key	Select this check box if the variable is to be used as a key. Keys are unique identifiers used in defining list entries in the Yang data hierarchy. They help distinguish one list entry from another. **NOTE:** For lists, you must define a key.
Required	Select this check box if you want the parameter to be mandatory.
Default Value	Enter a default value for the variable.
Pattern	For data types string or number, specify one of the following: • If the data type of the variable is string, specify the regular expression (regex pattern) ; for example, ^[a-z][A-Z]. • If the data type is number, specify the range in the format *Starting Number...Ending Number*. For example, 1...100.
Enumerated Values	If you selected entered enumerated values, the values are displayed here. Click inside the field to edit the enumerated values (in the Edit enumeration for *variable-name* page that appears).
Scope	**NOTE:** This field can be configured only for the root-level (top-level) node. Select the scope of the parameter: • **Device**, which means that the parameter is specific to each device. This is the default. • **Global**—The parameter is common across devices.

Table 94: Customize Variables Settings (Add Configuration Templates Page) *(continued)*

Setting	Guideline
Description	Enter a meaningful description for the variable; for example, Gateway IP address.
	After you finish customizing the parameters, click **Next**.

Table 95: Generated UI Actions (Add Configuration Template Page)

Action	Description
Reorder the UI	Drag and drop individual fields, grids, or sections to change the order in which the variables appear in the UI.
Modify the settings for a section	1. Hover over a section and click the settings icon (gear). The Section setting for *section-name* page appears. 2. Modify the following fields, as needed; fields marked with an asterisk (*) are mandatory. • **Label**—Enter the label that you want displayed in the UI. • **Collapsed**—Click the toggle button to expand (default) or collapse the section in the UI. 3. Click **OK** to save your changes. You are returned to the Generated UI tab and the modifications that you made are displayed on the UI.

Table 95: Generated UI Actions (Add Configuration Template Page) *(continued)*

Action	Description
Modify the settings for a field	1. Hover over a field and click the settings (gear) icon to modify the settings for a field. The Input setting for *field-name* page appears. 2. Modify the following fields, as needed; fields marked with an asterisk (*) are mandatory. • **Label**—Enter the label that you want displayed in the UI. • **Input type**—Select whether you want the field to be a text box (**Input text**), list (**input dropdown**), or text area (**input textarea**). If you select **input dropdown**, the Resources and Multiple Selection fields appear. • **Place holder**—Enter the text, which provides guidance to the user, that you want displayed (as ghost text) in the field on the UI. • **Hidden**—Click the toggle button to hide the field in the UI; by default, a field is displayed. • **Resource**—For a list (dropdown), if the administrator created resources for the configuration template, then you can select a resource or choose not to use resources (**No Resource**, which is the default). If there are no resources for the configuration template, then **No Resource** is displayed and can't be modified. • **Multiple Selection**—For a list (dropdown), click the toggle button to enable the selection of more than one items or not (default). • **Event listener**—If you want a field to be conditionally displayed based on an event, select **Data change**; if not, select **No listener** (default). If you select **Data change**, the following fields appear: ◦ **Event behavior**—Displays the event behavior (**function**) based on which the field is conditionally displayed. ◦ **Change path**—Click the **Select path** link to select the variable (that you want to conditionally display) from the viewdef tree in the popup page, and click **OK**. The selected path is displayed. ◦ **Event function**—Enter a JavaScript function that will be used to determine if the variable selected in the **Change path** field is displayed in the UI or not. 3. Click **OK** to save your changes. You are returned to the Generated UI tab and the modifications that you made are displayed on the UI.

Table 95: Generated UI Actions (Add Configuration Template Page) *(continued)*

Action	Description
Modify the settings for a grid	1. Hover over the grid area and click the settings (gear) icon to modify the settings for a field. The Grid setting for *Grid-name* page appears. 2. Modify the following fields, as needed; fields marked with an asterisk (*) are mandatory. • **Title**—Enter the title that you want displayed in the UI. • **Height**—Enter the height (in pixels) of the grid or click the up or down arrows to specify a height. • **Columns**—You can modify the following for each field: ◦ **Header**—Click inside the cell, modify as needed, and click **√** (check mark) to save your changes ◦ **Values**—Click inside the cell that you want to modify. The Settings for column *column-name* page appears. a. **Resource**—For a list (dropdown), if resources are present for the configuration template, then you can select a resource or choose not to use resources (**No Resource**, which is the default). If there are no resources for the configuration template, then **No Resource** is displayed and can't be modified. b. If you chose No Resource, enter one or more values (in JavaScript Object Notation [JSON] format) in the **Values** field. c. Click **OK**. You are returned to the Grid setting for *Grid-name* page d. Click **OK** to save your changes. You are returned to the Generated UI tab and the modifications that you made are displayed on the UI. 3. Click **OK** to save your changes. You are returned to the Generated UI tab and the modifications that you made are displayed on the UI.
Reset the generated UI	Click **Undo all edits** to discard the changes that you made and undo the changes made in the UI.

Table 95: Generated UI Actions (Add Configuration Template Page) *(continued)*

Action	Description
Validate and render the configuration	Validating the template and rendering the configuration enables you to check the configuration template that you added. To validate and render a configuration template: 1. Click **Validation**. The Validate Template page appears. 2. Enter values for the different parameters in the configuration template. 3. Click **Render**. The Rendered Config page appears, where CSO displays the configuration that was rendered based on the values that you entered. 4. Check if the configuration was rendered correctly. If the configuration was not rendered correctly, you can go back and make modifications as needed. 5. Click **OK**. You are returned to the Validate Template page. 6. Click **Cancel** to exit the template validation workflow. You are returned to the Generated UI page.
	After you are finished checking the configuration, click **Next** to continue.

RELATED DOCUMENTATION

Preview and Render Configuration Templates | 327

View the Configuration Deployed on Devices

In Customer Portal, for any configuration template, users with administrator or operator roles can view the configuration deployed on one more devices.

NOTE: In CSO releases before Release 5.1.0, configuration templates are called stage-2 configuration templates.

To view the configuration deployed on one or more devices:

1. Select **Resources > Templates > Configuration Templates**.

   The Configuration Templates page appears.

2. Navigate to the Devices column of the configuration template for which you want to view the deployed configuration, and click the **Number-of-devices** link.

   The Device Configuration page appears.

   Table 96 on page 340 explains the fields on this page.

3. After you have viewed the deployed configurations, click **OK**.

   You are returned to the Configuration Templates page.

Table 96: Device Configuration Page Fields

Setting	Guideline

Devices	The devices on which the configuration was deployed are displayed in a table. For each device, the following fields are displayed:
	• Device Name
	• Device Family
	• Operational Status, indicating whether the device is up or down.
	• Configuration State:
	  ◦ CREATED, indicating that the deployment hasn't started.
	  ◦ DEPLOYED, indicating that the configuration was successfully deployed.
	  ◦ DEPLOYING, indicating that the deployment of the configuration is in progress.
	  ◦ DEPLOY_FAILED, indicating that the deployment of the configuration failed.
	• Deployment Date, indicating the date and time on which the deployment was triggered.
	• Job—Click the **View logs** link for a device to view the deployment history for that device. The Deployment History page appears displaying the number of jobs in progress, number of successful jobs, and number of failed jobs in addition to a table listing some details of the job. You can drill down further by clicking the **Regional Log** and **Log** links.

Table 96: Device Configuration Page Fields *(continued)*

Setting	Guideline
Device-Name Configuration	Select a device by clicking the check box corresponding to the row: • For each device on which the configuration deployed successfully, table, this pane displays the configuration that is deployed on the device. • For each device on which the configuration deployment is in progress, DEPLOYING is displayed.

RELATED DOCUMENTATION

About the Configuration Templates Page | 316

CHAPTER 12

Managing Licenses

IN THIS CHAPTER

- About the Device Licenses Page | 342
- About the CSO Licenses Page | 343

About the Device Licenses Page

To access this page, click **Administration > Licenses > Device Licenses**.

You can use the Licenses page to view information about uploaded device licenses for virtual network services from your local file system. The license key is required to enable application-based routing, application monitoring, and other vSRX security features.

Tasks You Can Perform

You can perform the following tasks from this page:

- View details about a device license. Click the details icon that appears when you hover over the name of an image or click **More > Details**.
- Show or hide columns about the device license. Click the **Show Hide columns** icon in the top right corner of the page and select columns that you want to view on the page.
- Search for a device license. Click the Search icon in the top right corner of the page to search for a device license.

  You can enter partial text or full text of the keyword in the text box and press Enter.

  The search results are displayed on the same page.

Field Descriptions

Table 97 on page 343 describes the fields on the Device License Files page.

Table 97: Fields on the Device License Files Page

Field	Description
File Name	View the filename of the license. Example: license_Image_v1
Build	View the build name of the license. Example: 1
Version	View the version number of the license. Example: 1.1
Vendor	View the vendor name of the license. Example: Juniper Networks
Family	Select the device family of the license. Example: SRX
Model	View the model number of the license. Example: 1
Description	View the description of the license. Example: The license is applicable for SRX340 device.
Uploaded By	View the administrator who uploaded the license. Example: test_admin
Last Uploaded	View the date and time when the license was uploaded. Example: 11/18/2016 19:15

About the CSO Licenses Page

To access this page, click **Administration > Licenses > CSO Licenses** in Customer Portal.

You can use the CSO Licenses page to view information about the existing CSO licenses assigned to a tenant.

Tasks You Can Perform

You can perform the following tasks from this page:

- Group CSO licenses by sales order or SKUs:
  - Click **Group By** and select **Sales Order** to group CSO licenses by sales orders. By default, CSO licenses are grouped by sales order.
  - Click **Group By** and select **SKU** to group CSO licenses by SKUs.
- Search for CSO licenses by using keywords—Click the search icon and enter the search term in the text box and press Enter. The search results are displayed on the same page.

  You can search using license SKU, sales order, type, tier, or device class.
- Sort CSO licenses—Click a column name to sort based on the column name.

  NOTE: Sorting is applicable only to some fields.

- Show or hide columns—Click the **Show Hide Columns** icon at the top right corner of the page and select the columns that you want displayed on the CSO Licenses page.

Field Descriptions

Table 98 on page 344 describes the fields on the CSO Licenses page.

Table 98: Fields on the CSO Licenses page

Field	Description
License SKU	Displays the CSO license SKU name; for example, S-CSO-C-S1-A-3.
Sales Order	Sales order number; for example, 15563238.
Type	Displays whether the license is for an on-premise installation or for a cloud-hosted CSO installation.
Tier	Support tier associated with the license; for example, Standard.
Device Class	Class of the Juniper device associated with the license; for example, B-class.

Table 98: Fields on the CSO Licenses page *(continued)*

Field	Description
SSRN	Software support reference number (SSRN), which is necessary to identify your purchase order when you contact Juniper Networks for support
Start Date	Date (in MMM DD , YYYY format) from which the license is valid; for example, Aug 29, 2019.
End Date	Date (in MMM DD , YYYY format) up to which the license is valid. CSO calculates the end date based on the validity of the license SKU.
Device Quantity	Total number of on-premise devices that the tenant is authorized to use.
Available	Available number of devices that the tenant can add for the license.
Assigned	This field is not applicable to Customer Portal.

RELATED DOCUMENTATION

About the Device Licenses Page | 342

CHAPTER 13

Managing Signature Database and Certificates

IN THIS CHAPTER

- Signature Database Overview | 346
- About the Signature Database Page | 347
- Installing Signatures | 348
- Certificates Overview | 349
- About the Certificates Page | 350
- Importing a Certificate | 352
- Installing and Uninstalling Certificates | 354
- About the VPN Authentication Page | 355
- Modify PKI Settings for All Sites | 358
- Modify PKI Settings for Selected Sites | 361

Signature Database Overview

The signature database that Juniper provides contains application and intrusion prevention system (IPS) signatures:

- Application signatures are definitions of predefined attacks and applications, and can be used to identify applications for tracking firewall policies and quality-of-service (QoS) prioritization.
- IPS signatures are definitions of predefined attack object and attack object groups that you can use in IDP policies to match traffic against known attacks.

Users with the tenant administrator role can install the active signature database on one or more devices.

RELATED DOCUMENTATION

About the Signature Database Page | 347

Installing Signatures | 348

About the Signature Database Page

To access this page, select **Administration** > **Signature Database**.

Use the Signature Database page to install the active signature database, which contains intrusion prevention system (IPS) and application signatures, on one or more devices. The signature database contains definitions of attacks and application, which are used in defining IPS profile rules and application firewall rules. These attack objects and groups are designed to detect known attack patterns and protocol anomalies within the network traffic.

Tasks You Can Perform

You can perform the following task from this page:

- Install signatures on one or more devices. See "Installing Signatures" on page 348.
- Show or hide columns—Click the **Show Hide Columns** icon at the top right corner of the page and select the columns that you want displayed on the Signature Database page.

Field Descriptions

Table 99 on page 347 describes the fields on this page.

Table 99: Fields on the Signature Database Page

Field	Description
Active Database	
Database Version	Version of signature database.
Publish Date	Date and time (YYYY-MM-DD HH:MM:SS 24-hour format) when the signature database was published.
Installed Device Count	Number of devices on which the signature database was successfully installed.
Detectors	Version numbers of the detector engines associated with the signature database. Click the *detector-versions* link to view the detector details. The Detector Details for *Signature-Database-Version* page appears displaying (in a table) the platform, OS version, and version of the detectors for the signature database. Click **Close** to return to the Signature Database page.

Table 99: Fields on the Signature Database Page *(continued)*

Field	Description
Action	Click the **Install on device** link to install the signatures on one or more devices. The Install Signatures page appears. See "Installing Signatures" on page 348. **NOTE:** This field appears only for users with the Tenant Administrator role.

RELATED DOCUMENTATION

Signature Database Overview | 346

Installing Signatures

Users with the tenant administrator role can install the active signature database on one or more devices. Signatures must be present on the device for application firewall or intrusion prevention system (IPS) features to be used. If you do not install the signature database on a device, the deployment of IPS profiles or application firewall will fail.

NOTE:

- Before you install the signature database on the device, ensure that the IPS license is installed on the device. If the IPS license is not installed, only the application signatures will be installed when the signature database installation is triggered.
- You can install the signature database on the following devices: NFX150, NFX250, SRX Series, and vSRX.

To install the active signature database:

1. Select **Administration > Signature Database**.

   The Signature Database page appears.

2. Click **Install Signatures**.

   The **Install Signatures** page appears displaying the signature database version and the devices on which you can install the signature database.

3. Select the check boxes corresponding to the devices on which you want to install the signature database.

You can also search for, filter, or sort the devices displayed in the table.

4. From the **Type** field:
   - Select **Run now** to immediately trigger the installation of the signature database on the devices that you selected.
   - Select **Schedule at a later time** to install the signature database later and specify a date and time at which you want the installation to be triggered.
5. Click **OK**.
   - If you specified that the database must be installed immediately, a job is triggered and in the Job Tasks page that appears, the tasks associated with the signature database installation are displayed. Click **OK** to exit and return to the Signature Database page.
   - If you specified that the database must be installed later, a job is created and you are returned to the Signature Database page. A confirmation message (with the job ID) is displayed at the top of the page.

After the signature database is installed successfully, you can deploy the firewall policy (that references IPS profiles or application signatures) on the device.

RELATED DOCUMENTATION

Signature Database Overview | 346

About the Signature Database Page | 347

Certificates Overview

SSL uses public–private key technology that requires a private key paired with an authentication certificate for the SSL service. An SSL certificate includes identifying information such as a public key and a signature issued by a certificate authority (CA).

CAs are entities that validate identities and issue certificates. A CA can issue multiple certificates in the form of a tree structure. A root certificate is the topmost certificate of the tree, the private key of which is used to sign other certificates. All certificates immediately below the root certificate inherit the signature or trustworthiness of the root certificate. This is somewhat like the notarizing of an identity. You can configure a root CA certificate by first obtaining a root CA certificate (by either generating a self-signed one or importing one) and then applying it to an SSL proxy profile.

NOTE: SSL certificates are used for the SSL forward proxy feature in CSO.

RELATED DOCUMENTATION

SSL Forward Proxy Overview | 636

About the SSL Proxy Profiles Page | 653

About the Certificates Page

To access this page, select **Administration > Certificates** in Customer Portal.

Use this page to view and manage SSL certificates. You can import a root certificate or a trusted certificate (directly from a file or by pasting the content) and install a certificate on a site.

Tasks You Can Perform

You can perform the following tasks from this page:

- View information about the existing certificates; see Table 100 on page 351.
- Import a certificate—Select **More > Import Certificate**. See "Importing a Certificate" on page 352.
- View the sites on which a certificate is installed—Select a certificate and then select **More > View Installed Sites**.

  The View Installed Sites page appears, displaying the list of sites on which the selected certificate is installed. Click **OK** to close the page and return to the Certificates page.
- Install a certificate on a site—Select a certificate and then select **More > Install Certificate**. See "Installing and Uninstalling Certificates" on page 354.
- Uninstall a certificate from a site—Select a certificate and then select **More > Uninstall Certificate**. See "Installing and Uninstalling Certificates" on page 354.
- View details about a certificate—Select a certificate and then select **More > Detailed View**. The Detailed View page appears. See Table 101 on page 351 for an explanation of fields on this page.

Field Descriptions

Table 100 on page 351 displays the fields on the Certificates page.

Table 100: Fields on the Certificates Page

Field	Description
Certificate Name	Name of the certificate.
Type	Type of the certificate: • Root certificate • Trusted certificate
Description	Description of the certificate.

Table 101: Fields on the Detailed View Page

Field	Description
Certificate Name	See Table 100 on page 351.
Type	See Table 100 on page 351.
Valid From	Date and time (UTC) from which the certificate is valid.
Valid Upto	Date and time (UTC) until which the certificate is valid.
Serial Number	Serial number of the certificate.
Signature Algorithm	Algorithm used to sign the certificate.
Issuer Details	Details of the authority that issued the certificate, including details such as name, country, organization, and so on.
Version	X.509 version of the certificate.

RELATED DOCUMENTATION

About the SSL Proxy Profiles Page | 653

Importing a Certificate

You can import an SSL certificate (directly from a file or by pasting the content) from the Import Certificate page.

> **NOTE:** If you want to use the SSL proxy feature, you must import at least one root certificate for a tenant; the certificate can be used in one or more sites.

To import a certificate:

1. Select **Administration > Certificates** in Customer Portal.

   The Certificates page appears.

2. Select **More > Import Certificate**.

   The Import Certificate page appears.

3. Complete the configuration according to the guidelines provided in Table 102 on page 352.

   > **NOTE:** Fields marked with * are mandatory.

4. Click **OK** to import the certificate.

   You are taken to the Certificates page. If the certificate content that you imported is validated successfully, a confirmation message is displayed; if not, an error message is displayed.

After importing a certificate, you can use it when you create an SSL proxy profile.

Table 102: Import Certificate Settings

Setting	Guideline
Certificate Name	Enter the certificate name, which must be a unique string of alphanumeric characters and some special characters (_ -). No spaces are allowed and the maximum length is 32 characters.
Certificate Type	Select an option to specify whether the certificate that you are importing is a root certificate (**Root CA**) or a trusted certificate (**Trusted CA**).

Table 102: Import Certificate Settings *(continued)*

Setting	Guideline
Passphrase	Enter the passphrase to protect the private key or key pair of the Privacy-Enhanced Mail (PEM) certificate file.
Description	Enter a description for the certificate.
Certificate Content	Select an option to specify whether you want to import the certificate content from a file or whether you want to paste the certificate content.
Paste Certificate Content	Depending on the method you choose in the preceding field, do one of the following: • Import the certificate content directly from a file—Click **Browse**, and in the File Upload dialog, select a file and click **Open**. The filename of the file that you uploaded is displayed. • Paste the certificate content directly from a file—Copy the certificate content from the file and paste it in the text box. **NOTE:** • The following certificate file extensions are supported: **.cert**, **.pem**, and **.txt**. • The certificate content must be in the X.509 ASCII format. • If the certificate type is **Root CA**, then the both the certificate content and private key must be specified.

The following is an example of root certificate content.

```
-----BEGIN PRIVATE KEY-----
AbcXyz123AbcXyz123AbcXyz123AbcXyz123AbcXyz123AbcXyz123AbcXyz123A
AbcXyz123AbcXyz123AbcXyz123AbcXyz123AbcXyz123AbcXyz123AbcXyz123A
AbcXyz123AbcXyz123AbcXyz123AbcXyz123AbcXyz123AbcXyz123AbcXyz123A
AbcXyz123AbcXyz123AbcXyz123AbcXyz123AbcXyz123AbcXyz123AbcXyz123A
AbcXyz123AbcXyz123AbcXyz123AbcXyz123AbcXyz123AbcXyz123AbcXyz123A
AbcXyz123AbcXyz123AbcXyz123AbcXyz123AbcXyz123AbcXyz123AbcXyz123A
-----END PRIVATE KEY-----
-----BEGIN CERTIFICATE-----
DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456A
DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456A
DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456A
DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456A
DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456A
DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456A
DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456A
```

```
DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456A
DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456A
-----END CERTIFICATE-----
```

RELATED DOCUMENTATION

Installing and Uninstalling Certificates | 354
Creating SSL Forward Proxy Profiles | 655

Installing and Uninstalling Certificates

IN THIS SECTION

Installing a Certificate | 354
Uninstalling a Certificate | 355

You can install and uninstall certificates from the Certificates page. This topic has the following sections:

Installing a Certificate

Use the Install Certificate page to install certificates on one or more sites.

To install a certificate on one or more sites:

1. Select **Administration > Certificates** in Customer Portal.

   The Certificates page appears, displaying the existing certificates.

2. Select the certificate that you want to install, and then select **More > Install Certificate**. Alternatively, right-click a certificate and select **Install Certificate**.

   The Install Certificate page appears, displaying a list of sites.

3. Select the sites on which you want to install the certificate.

4. Click **Install** to install the certificate on the selected sites.

You are taken to the Certificates page. A job is created and a confirmation message appears with the ID of the job. Click the job ID to go to the Jobs page, where you can view the status of the job.

Uninstalling a Certificate

If a certificate's validity has expired or if you want to remove a certificate from a site, you can uninstall the certificate from that site.

To uninstall a certificate from one or more sites:

1. Select **Administration > Certificates** in Customer Portal.

   The Certificates page appears, displaying the existing certificates.

2. Select the certificate that you want to uninstall, and then select **More > Uninstall Certificate**. Alternatively, right-click a certificate and select **Uninstall Certificate**.

   The Uninstall Certificate page appears, displaying only those sites on which the certificate was previously installed.

3. Select the sites from which you want to uninstall the certificate.

4. Click **Uninstall** to uninstall the certificate from the site.

   You are taken to the Certificates page. A job is created and a confirmation message appears with the ID of the job. Click the job ID to go to the Jobs page, where you can view the status of the job.

RELATED DOCUMENTATION

Importing a Certificate | 352

About the VPN Authentication Page

IN THIS SECTION

- Tasks You Can Perform | 356
- Field Descriptions | 356

Contrail Service Orchestration (CSO) establishes secure IPsec Virtual Private Network (VPN) tunnels to connect sites after authenticating the tunnel endpoints. CSO authenticates tunnel endpoints by using either preshared keys or Public Key Infrastructure (PKI) certificates.

Service Provider (SP) and Operating Company (OpCo) Administrators can configure the authentication type when the tenant is onboarded.

If PKI certificate is configured as the authentication type, then tenant administrators can modify the PKI settings from the VPN Authentication page (**Administration > Certificate Management > VPN Authentication**) after the tenant is onboarded.

> **NOTE:** The VPN Authentication page is displayed only for tenants with SD-WAN service that are configured with PKI as the authentication type.

Tasks You Can Perform

- View information about the existing certificates for all provisioned sites in the tenant. See Table 103 on page 357.
- Change the Certificate Authority (CA) server settings (URL, password, and CRL Server URL) for the tenant. See “Modify PKI Settings for All Sites” on page 358.
- Change the Certificate Revocation List (CRL) URL of certificates for the tenant. See “Modify PKI Settings for All Sites” on page 358.
- Change the method of renewing PKI certificates for all provisioned sites in the tenant. See “Modify PKI Settings for All Sites” on page 358.
- Change the method of renewing PKI certificates for one or more provisioned sites in the tenant. See “Modify PKI Settings for Selected Sites” on page 361.
- Manually renew certificates for one or more provisioned sites in the tenant. See“Modify PKI Settings for Selected Sites” on page 361.
- Search for certificates by using keywords. Click the **Search** icon to enter the search term in the text box and press **Enter**. The search results are displayed on the same page.
- Show or hide columns. Click the **Show Hide Columns** icon at the top right corner of the grid and select the columns that you want displayed on the VPN Authentication page.

Field Descriptions

Table 103 on page 357 describes the fields on the VPN Authentication page.

Table 103: Fields on the VPN Authentication page

Field	Description
Tenant-Level Settings for PKI Certificates	
Certificate Renewal	
Current Tenant Setting	Renewal method currently configured for PKI certificates of the tenant.
Next Renew Check Time	• If the **Auto Renew Certificate** toggle button on the Edit Tenant Certificate page is enabled, displays the date and time at which the next renewal check is scheduled. • If the **Auto Renew Certificate** toggle button on the Edit Tenant Certificate page is disabled, displays N/A (not applicable).
Next CRL check time	Date and time at which the next CRL check is scheduled.
Last CRL update time	Date and time at which the CRL was last updated.
Details of Certificates	
Tenant Name	Name of the tenant.
Common Name	Name of the PKI certificate.
Certificate ID	ID of the PKI certificate.
Serial Number	Serial number of the PKI certificate.
Used In	Name of the site with which the PKI certificate is associated.
Device	Name of the device with which the PKI certificate is associated.

Table 103: Fields on the VPN Authentication page *(continued)*

Field	Description
Status	Expiration status of the PKI certificate: • If you set the certificate to be renewed automatically, the status displayed depends on the renewal period that you selected from the Edit Certificate Settings for Tenant page. For example, if you selected the renewal period as 1 month, the Status field displays **Less than 1 month before expiry**. • If you set the certificate to be manually renewed, the status displayed depends on the expiration notification time for the certificate (Status: **Less than 2 weeks before expiry**). • If the expiration date of the certificate does not meet the expiration notification time yet, the Status field displays –. • If the certificate has expired, the Status field displays **Expired**.
Expires on	Date and time at which the PKI certificate expires.
Renewal Method	Renewal method of the PKI certificate: • Auto • Manual

RELATED DOCUMENTATION

View and Edit Tenant Settings | 20

Modify PKI Settings for All Sites

The VPN authentication settings for a tenant are configured when the tenant is onboarded. If PKI Certificate is configured as the authentication type, tenant administrators can modify the PKI settings even after adding sites for the tenant. The changed settings are applicable to all existing sites of the tenant and to sites that the tenant might add later.

You can modify the following PKI settings from the Edit Certificate Settings for Tenant page (**Administration > Certificate Management > VPN Authentication > Change**).

- **Certificate Authority (CA) Server Parameters**

The CA server manages the lifecycle of a certificate and publishes revoked certificates to the Certificate Revocation List (CRL) server. To obtain trusted CA certificates, CSO communicates with the CA server using the Simple Certificate Enrollment Protocol (SCEP).

To change the CA server parameters (URL, Password, and CRL URL):

1. On the Edit Certificate Settings for Tenant page, enter the new CA server URL and password in the **CA Server URL** and **Password** fields, respectively. Enter the new CRL server URL associated with the CA server in the **CRL Server** field.

2. Click **OK** to save your changes.

   You are returned to the VPN Authentication page, where a confirmation message appears indicating that a job is triggered to automatically renew certificates for all sites in the tenant.

   You can click the job link in the message to view the job details, or view the details on the Jobs (**Monitor > Jobs**) page.

   > **NOTE:** The CA server parameters are not updated if the PKI server is unreachable at the time that the job is triggered.

   After the job is completed successfully, a confirmation message appears indicating that the settings are updated. CSO also downloads the latest list of revoked certificates from the CA server.

- **CRL Server URL**

  You can choose to update only the CRL server URL associated with the CA server.

  To update the CRL server URL:

  1. On the Edit Certificate Settings for Tenant page, specify the new CRL server URL in the **CRL Server** field.

  2. Click **OK** to save your changes.

     You are returned to the VPN Authentication page, where a confirmation message appears indicating that a job is triggered.

     You can click the job link in the message to view the job details or view the details on the Jobs (**Monitor > Jobs**) page.

     After the job is completed successfully, a confirmation message appears indicating that the settings are updated. CSO downloads the latest list of revoked certificates from the CA server.

- **Certificate Renewal Method**

To change the certificate renewal method:

1. On the Edit Certificate Settings for Tenant page, click the Auto Renew Certificate toggle button to enable or disable the automatic renewal of certificates.

2. If you enable the Auto Renew Certificate toggle button, the **Renew Before Expiry** list appears.

   From the list, select the period before the expiry date on which the certificates should be automatically renewed:

   - 3 Days
   - 1 Week
   - 2 Weeks (default)
   - 1 Month

   If you disable the Auto Renew Certificate toggle button, the certificates should be manually renewed for each site before they expire. See "Modify PKI Settings for Selected Sites" on page 361 for the procedure to manually renew certificates for sites.

3. Click **OK** to save your changes.

   If you enabled the automatic renewal of certificates, CSO schedules a job to check the expiration date of certificates for all sites of the tenant (every 24 hours). Based on the expiration date that you've configured, CSO triggers a job to automatically renew the certificates.

   > **NOTE:** The certificate renewal job is not executed for sites that are down or that do not have connectivity to CSO at the time that the job is triggered.

   You are returned to the VPN Authentication page where a confirmation message appears indicating that the settings are updated.

RELATED DOCUMENTATION

About the VPN Authentication Page | 355

View and Edit Tenant Settings | 20

Modify PKI Settings for Selected Sites

The VPN authentication settings for a tenant are configured when the tenant is onboarded. If PKI Certificate is configured as the authentication type, tenant administrators can modify the PKI settings even after adding sites for the tenant. The changed settings are applicable to all existing sites of the tenant and to sites that the tenant might add later. To change the PKI settings for all sites in the tenant, see "Modify PKI Settings for All Sites" on page 358.

You can perform the following actions on the selected sites:

- Change the method of renewing PKI certificates:

  > **NOTE:** You can change the renewal method of PKI certificates for sites in a tenant only if you set the certificate renewal method for the tenant to automatic (that is, if you enable the Auto Renew Certificate toggle button).

  Do the following:

  1. Select **Administration > Certificate Management > VPN Authentication**.

     The VPN Authentication page appears.

  2. Select one or more sites from the list of available sites and click **Change Renewal Method**.

     A drop-down list appears.

  3. From the list, choose the renewal method (**Set Auto Renew** or **Set Manual Renew**).

     The **Edit Certificate Renewal Method** page appears asking you to confirm the renewal method.

  4. Click **Yes** to change the renewal method.

     You are returned to the VPN Authentication page, where a confirmation message appears indicating that the certificate renewal method is updated. The Renewal method column on the VPN Authentication page displays the updated renewal method for the selected sites.

- Manually renew certificates:

  1. Select **Administration > Certificate Management > VPN Authentication**.

     The VPN Authentication page appears.

  2. Select one or more sites from the list of available sites and click **Renew Certificate**.

     The Confirm Renew Certificate page appears.

3. Click **Yes** to manually renew the certificates.

   You are returned to the VPN Authentication page, where a confirmation message appears indicating that a certificate renewal job is triggered.

   You can click the job link in the message to view the job details, or view the details on the Jobs (**Monitor > Jobs**) page.

   > **NOTE:** The certificate renewal job is not executed for sites that are down or that do not have connectivity to CSO at the time that the job is triggered.

   If the job is completed successfully, a confirmation message appears on the VPN Authentication page.

RELATED DOCUMENTATION

- About the VPN Authentication Page | 355
- View and Edit Tenant Settings | 20

CHAPTER 14

Managing Juniper Identity Management Service

IN THIS CHAPTER

- Juniper Identity Management Service Overview | 363
- About the Identity Management Page | 366
- Configuring CSO and JIMS Connection | 367
- Configuring JIMS for an SRX Device | 369

Juniper Identity Management Service Overview

IN THIS SECTION

- Access Token Query | 364
- Batch or Periodic Query | 364
- IP Address Query | 365
- User Mapping Query | 365

Juniper Identity Management Service (JIMS) provides a robust and scalable user identification and IP address mapping implementation that includes endpoint context and machine ID. JIMS collects user identity information from different authentication sources, for SRX Series devices.

JIMS collects user identity information from a configured Active Directory and makes it available to SRX Series devices or vSRX instances. You can download and install Juniper Identity Management Service (JIMS), configure the CSO client on JIMS to obtain user identity information from the configured Active Directory, and use CSO and JIMS to manage user-based firewall policy intents on SRX Series devices and vSRX instances.

The SRX Series devices communicate with JIMS through HTTP or HTTPS connection. Use HTTP connection for debugging and HTTPS for deployments. SRX Series devices consist of primary and secondary JIMS configurations. These devices must always query the primary JIMS. The secondary JIMS is available as a

fall back option with limited resources. The secondary JIMS must be used when the HTTP GET query or a number of queries to the primary JIMS fails. SRX Series devices constantly scrutinize the failed primary JIMS and revert to the primary JIMS, once it is up and running.

When you request a JIMS report, the SRX Series device specifies the timestamp. JIMS forms an HTTPS response from the earliest known report since the requested timestamp. SRX Series devices request for the maximum number of reports to include in the response from JIMS. Along with the requested reports, JIMS always returns a cookie. In the subsequent requests to JIMs, SRX Series devices include cookies instead of timestamp to indicate the same context, same beginning timestamp, and to resume the same response from where it has stopped the previous time.

NOTE:

- IP and user mapping information might be inaccurate, if the user identities in JIMS are cleared, delayed, or missing.
- SRX firewall authentication can also push the authentication entries to JIMS.

The SRX Series device communicates with JIMS through HTTP or HTTPS messages to obtain the access token and query for user identities. The following different query modes are available and all queries can happen simultaneously.

Access Token Query

JIMS requires OAuth 2.0 protocol to authenticate or authorize. The SRX Series device user query function requires an access token to query the JIMS server. The SRX Series device uses the client credentials such as client ID and client secret to obtain an access token. These parameters must be consistent with the API client configured on JIMS.

Batch or Periodic Query

At the beginning, SRX Series device sends the batch queries to JIMS sequentially to obtain all the expected user identities. When there are no more entries in JIMS, SRX Series device periodically queries for the newly generated reports with the configured interval.

The timestamp is mentioned in the query to restart the response. The timestamp is expected in the query under the following circumstances:

- SRX Series device queries the JIMS server for the first time
- SRX Series device switches over to the secondary JIMS
- SRX Series device does the error recovery because of an internal error or upon receiving error response from JIMS

For all the other cases, SRX Series device provides the received cookie information in the query instead of a timestamp.

IP Address Query

SRX Series device can provide another query to JIMS specifying the IP address, if it has missed the data for the existing IP address flow. If there are many IP address queries in the queue, SRX Series device can keep multiple concurrent HTTP or HTTPS connections with JIMS to increase the throughput. However, the number of concurrent connections are restricted to less than or equal to 20 connections to reduce the load on JIMS.

User Mapping Query

SRX Series device can engage Captive Portal to obtain the user ID to authenticate the user. Once the user is authenticated, SRX Series device can issue another query to JIMS specifying the user ID and IP address to obtain user information. The firewall authentication uses the *https://<JIMS>/<query-api>/user/ip=<ip>&id=<id>&domain=<domain>* API to push an authentication success entry to JIMS with the user IP, user ID, and the domain. JIMS responds with the user information.

The difference between the IP address query and user query is that the IP address query does not have the user ID. Both these queries insert the user information to the internal cache of JIMS , and all SRX devices are updated with user information.

RELATED DOCUMENTATION

- About the Identity Management Page | **366**
- Configuring CSO and JIMS Connection | **367**
- Configuring JIMS for an SRX Device | **369**

About the Identity Management Page

To access this page, select **Administration** > **Identity Management**.

NOTE:

- For information on system requirements for installing JIMS, see System Requirements for Installing Juniper Identity Management Service
- For information on installing JIMS on your Windows server, see Installing Juniper Identity Management Service.

Use the **Identity Management** page to download and install JIMS, interface JIMS with CSO to obtain advanced user identity an active directory, and use CSO to push the JIMS configuration to SRX Series devices.

Figure 10 on page 366 illustrates the connectivity between, CSO, JIMS, and an SRX Series device.

Figure 10: CSO-JIMS-SRX Connectivity Configuration

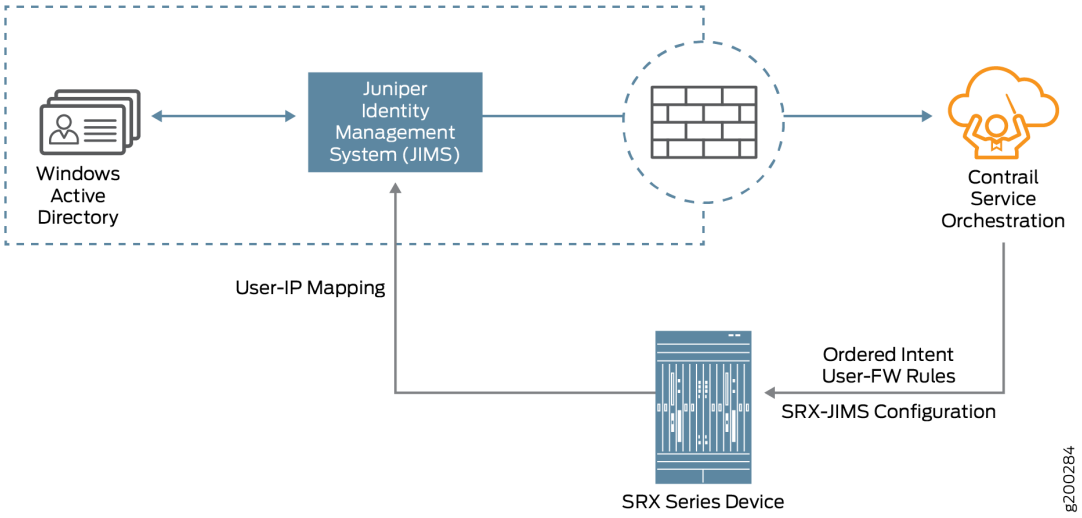

Tasks You Can Perform

You can perform the following tasks from this page:

- Download the JIMS executable to your Windows server using **Download JIMS**. Run the JIMS executable to install JIMS on your Windows server machine. See System Requirements for Installing Juniper Identity Management Service and Installing Juniper Identity Management Service.

After you have successfully installed JIMS, you can login into JIMS using your Windows user ID and password.

- Configure the connection between CSO and JIMS to import user and group lists from an Active Directory (AD) of your choice, using **JIMS to CSO**. See "Configuring CSO and JIMS Connection" on page 367.
- Configure the connection between JIMS and an SRX Series device. See "Configuring JIMS for an SRX Device" on page 369.

RELATED DOCUMENTATION

Juniper Identity Management Service Overview | 363
Configuring CSO and JIMS Connection | 367
Configuring JIMS for an SRX Device | 369
Preparing CSO Identity Management
JIMS v1.1 Feature Guide

Configuring CSO and JIMS Connection

Before you begin to configure the connection between CSO and JIMS, ensure that you have downloaded and installed JIMS. See System Requirements for Installing Juniper Identity Management Service and Installing Juniper Identity Management Service.

To configure a connection between CSO and JIMS:

1. Select **Administration** > **Identity Management**.

   The **Identity Management** page appears.

2. Click **JIMS-to-CSO Configuration** or the greater-than (>) symbol beside it.

   The **JIMS-to-CSO Configuration** panel expands. The panel displays a system-generated user name which cannot be changed, the last updated time of the user identity information from Active Directory and the connection status of the JIMS server(s).

   NOTE: If you have already configured a JIMS user account in CSO, the details of this connection is displayed in the **JIMS-to-CSO Configuration** panel.

3. The **Username** is auto-generated for each tenant. You will not be able to change it. Enter a password of your choice for your JIMS-to-CSO connection in the **Password** field.

   NOTE: The password must contain a number, an upper-case letter, and a special character.

   NOTE: The password you entered will appear encrypted. If you want to see the password that you entered as plain text, select **Show Password**.

4. Click **Save** to save your changes. The JIMS user credentials are saved.

   If you do not want to save your changes, click **Cancel**.

5. CSO and JIMS need to be connected in order for JIMS to push data to CSO. To set up this connection, you must configure the CSO client on JIMS, using the username and password that you created in the **JIMS-to-CSO Configuration** panel. For more information on configuring the CSO client on JIMS, see Configuring the Connection to a CSO Client.

6. Configure an Active Directory (AD) as a data source in JIMS, see Configuring the Connection to an Active Directory.

   NOTE: After your JIMS user credentials are saved, the password field changes to the **Change Password** link.

   If you want to change your password, click **Change Password**.

   The **Change Password** page appears.

   - Enter your new password in the **New Password** field and re-enter the same password in the **Confirm Password** field.
   - Click **OK** to save the new password. The updated password is saved.

     If you do not want to save your new password, click **Cancel** instead.

RELATED DOCUMENTATION

Juniper Identity Management Service Overview | 363

About the Identity Management Page | 366
Configuring JIMS for an SRX Device | 369

Configuring JIMS for an SRX Device

Configuring the connection between SRX Series devices to JIMS allows the JIMS server to send the IP address, username, and group relationship information to SRX Series devices through CSO. You can also configure a set of optional advanced settings for authentication timeout, domain filters, and choose to include or exclude user identity information in the communication between the JIMS server and the SRX Series device.

For every SRX Series device, you can configure the primary and secondary JIMS servers. The SRX Series device always queries the primary JIMS server. The secondary JIMS server is available as a fallback option with limited resources. The secondary JIMS server is used when a number of queries to the primary JIMS server fails. The SRX Series device constantly scrutinizes the failed primary JIMS server and reverts to the primary JIMS server, once it is up and running.

Before you begin, you need the following information:

- The IP address of the primary and secondary (optional) JIMS server.
- The client ID to obtain an OAuth token from the JIMS server for user queries.
- The client secret to obtain an OAuth token from the JIMS server for user queries.

To configure a connection between an SRX Series device and JIMS:

1. Select **Administration** > **Identity Management**.

   The **Identity Management** page appears.

2. Click **SRX-to-JIMS Configuration** or the greater-than (>) symbol beside it.

   The **SRX-to-JIMS Configuration** panel expands.

   > **NOTE:** If you have already configured JIMS for an SRX Series device, the details of this configuration is displayed in the **SRX-to-JIMS Configuration** panel.

3. Complete the configuration according to the guidelines provided in Table 104 on page 370.

4. Click **Save** to save the changes. JIMS is now configured for an SRX device.

   If you want to discard your changes, click **Cancel** instead.

Table 104 on page 370 provides guidelines on using the fields on the **SRX-to-JIMS Configuration** panel.

Table 104: Fields on the SRX-to-JIMS Configuration Panel

Field	Description
Identity	
IP Address	Enter a valid IPv4 or IPv6 address of the primary JIMS server. SRX Series devices always query the primary JIMS to obtain the user identities.
Secondary Identity	Enable this option to use the secondary JIMS server as a fallback when the primary JIMS server fails. By default, this option is disabled.
Secondary IP Address	Enter a valid IPv4 or IPv6 address of the secondary JIMS server. The secondary JIMS is available as a fall back option with limited resources. Use the secondary JIMS when the HTTP GET query or number of queries to the primary JIMS fails.
Client Credentials	
Client ID	Enter the client ID that the SRX Series device provides to JIMS server as part of its authentication. The SRX Series device must authenticate itself with the JIMS server to obtain an access token that allows the it to query the JIMS server for user identity information. The client ID must be consistent with the CSO client ID or username configured on the JIMS server.
Client Secret	Enter the client secret that the SRX Series device provides to the JIMS server as part of its authentication. The client secret must be consistent with the CSO client secret or password configured on the JIMS server.
Advanced Settings	
Authentication Entry Timeout	Enter the timeout interval (in minutes) after which, the idle entries in the JIMS authentication table expire. The timeout interval begins from when the user authentication entry is added to the authentication table. This value can be between 10 and 1440 minutes, where a value of 0 means no timeout. The default value is 69 minutes.
Include IP Address(es)	The SRX Series device sends a query to JIMS for the user identity information only for the IP addresses present in the selected address group; JIMS responds with the requested user identity information. Click **Add New Address** to create a new IP address group, see “Creating Addresses or Address Groups” on page 737.

Table 104: Fields on the SRX-to-JIMS Configuration Panel *(continued)*

Field	Description
Exclude IP Address(es)	The SRX Series device does not query JIMS for the user identity information for the excluded IP addresses present in the selected address group. Click **Add New Address** to create a new IP address group, see "Creating Addresses or Address Groups" on page 737.
Filter Domain(s)	The SRX Series device sends a query to JIMS for the user identity information within the specified domains. Enter a comma-separated list of up to 25 domain names. A domain name can be an alphanumeric string of up to 64 characters that can also contain dashes, underscores, and dots. Example: example.net

RELATED DOCUMENTATION

- Juniper Identity Management Service Overview | **363**
- About the Identity Management Page | **366**
- Configuring CSO and JIMS Connection | **367**

4

PART

Managing Policies, Profiles, and Proxies

Managing Firewall Policies | **373**

Managing UTM Profiles | **450**

Managing SLA Profiles and SD-WAN Policies | **499**

Managing NAT Policies | **554**

Managing IPS Signatures and Profiles | **591**

Managing SSL Proxies | **636**

Deploying Policies | **664**

Configuring Policies for SD-LAN | **670**

CHAPTER 15

Managing Firewall Policies

IN THIS CHAPTER

- Firewall Policy Overview | **374**
- About the Firewall Policy List Page | **376**
- About the Firewall Policy Name Page | **377**
- Adding a Firewall Policy | **378**
- Editing and Deleting Firewall Policies | **380**
- Adding Firewall Policy Intents | **381**
- Editing, Cloning, and Deleting Firewall Policy Intents | **387**
- Selecting Firewall Source | **389**
- Selecting Firewall Destination | **393**
- Firewall Policy Examples | **396**
- Firewall Policy Schedules Overview | **436**
- About the Firewall Policy Schedules Page | **437**
- Creating Schedules | **438**
- Editing, Cloning, and Deleting Schedules | **440**
- Deploying Firewall Policies | **441**
- About the Default Profiles for Unified Firewall Policy Page | **442**
- Editing Default Settings for the Unified Firewall Policy | **444**
- Importing Policies Overview | **446**
- Importing Firewall Policies | **448**

Firewall Policy Overview

Contrail Service Orchestration (CSO) provides the ability to create, modify, and delete firewall policy intents associated with a firewall policy. Firewall policies are presented as *intent-based policies*. A firewall policy intent controls transit traffic within a context that is derived out of the end-points defined in the intent. Intent-based firewall policies can incorporate both transport layer (Layer 4) and application layer (Layer 7) firewall constructs in a single intent. The underlying system, automatically analyzes the intent, translates them into the set of rules the devices understand. The choice of sequence and the assignment happens implicitly based on the endpoints in the intent definition. The intent consist of source and destination endpoints. Endpoints could be applications (L7), sites or site groups, IP address/address-groups, services, or departments.

NOTE:

- Intent based policies are not applicable for Hybrid WAN deployments.
- Starting from CSO Release 5.0.1, if a device (CPE or next-generation firewall) is running Junos OS Release 18.2R1 or later, a firewall policy acts as a unified firewall policy. In a unified firewall policy, dynamic application can be used as a match condition along with the existing match conditions. Therefore, a separate application firewall is not configured on the device to allow or block traffic to an application.

  However, If the device is running a version earlier than Junos OS Release 18.2R1, the firewall policy does not act as a unified firewall policy and application firewalls continue to be configured on the device.

  See Unified Security Policies for information about unified firewall policies.

Firewall policies provide security functionality by enforcing intents on traffic that passes through a device. Traffic is permitted or denied based on the action defined as the firewall policy intent.

A firewall policy provides the following features:

- Permits, rejects, or denies traffic based on the application in use.
- Identifies not only HTTP but also any application running on top of it, enabling you to properly enforce policies. For example, an application firewall intent could block HTTP traffic from Facebook but allow Web access to HTTP traffic from Microsoft Outlook.
- Provides the ability to enable advanced security protection by specifying one or more of the following:
  - Unified threat management (UTM) profile
  - SSL proxy profile
  - Intrusion prevention system (IPS) profile

In CSO, intents are categorized as zone-based intents and enterprise-based intents.

- Zone-based-intents are intents with zones as source and destination endpoints. The policies with zone-based intents can be applied to SD-WAN sites , hybrid WAN sites, and next-generation firewall sites. The parameters that you can define for zone-based intents are listed in Table 105 on page 375.

Table 105: Zone-based intents

Source End Points	Destination End points	Advanced Security Options	Supported Options
Zones	Zones	SSL Proxy Profile	Scheduler
Address	Address	UTM Profile	Logging
Users	Service (L4 port/protocol)	IPS Profile	
	Applications (Dynamic Applications)		

> **NOTE:** You cannot select a department or site as an endpoint in zone-based intents. The sites assigned to the policy are applicable for zone-based intents and are automatically considered for deployment.

- Enterprise-based intents are intents that contain sites, site-groups, departments, addresses as source and destination endpoints. Firewall policies with enterprise-based intents can be applied only to SD-WAN sites. The parameters that you can define for enterprise-based intents are listed in Table 106 on page 375.

Table 106: Enterprise-based intents

Source Endpoints	Destination Endpoints	Advanced Security Options	Supported Options
Sites	Sites	UTM Profile	Scheduler
Site-groups	Site-groups	IPS Profile	Logging
Departments	Departments		
Addresses	Addresses		
Users	Users		
	Service/Applications		

> **NOTE:**
> - Zones cannot be selected as source or destination endpoints for enterprise-based intents.
> - Intents added in CSO Release 4.1 and earlier are now called enterprise-based-intents.

RELATED DOCUMENTATION

About the Firewall Policy Name Page | 377
Firewall Policy Examples | 396
Adding Firewall Policy Intents | 381
Editing, Cloning, and Deleting Firewall Policy Intents | 387

About the Firewall Policy List Page

To access this page, select **Configuration** > **Firewall** > **Firewall Policy**.

Use this page to view and manage firewall policies associated with your site or site groups.

Tasks You Can Perform

You can perform the following tasks from this page:

- Add a firewall policy. See "Adding a Firewall Policy" on page 378.
- Edit or delete a firewall policy. See "Editing and Deleting Firewall Policies" on page 380.
- Import a firewall policy. See "Importing Firewall Policies" on page 448.
- Deploy a firewall policy. See "Deploying Firewall Policies" on page 441.
- Search for a firewall policy. Click the Search icon in the top right corner of the page to search for a firewall policy.

  You can enter partial text or full text of the keyword in the text box and press Enter.

  The search results are displayed on the same page.

Field Descriptions

Table 107 on page 376 provides guidelines on using the fields on the **Policy List** page.

Table 107: Fields on the Policy List Page

Field	Description
Policy Name	Name of the firewall policy.
Number of intents	Number of intents associated with the firewall policy.
Applied to	Sites to which the firewall policies are applied.

Table 107: Fields on the Policy List Page *(continued)*

Field	Description
Status	Status of firewall policy deployment.

RELATED DOCUMENTATION

Firewall Policy Overview | 374

About the Firewall Policy Name Page

To access this page, select **Configuration** > **Firewall** > **Firewall Policy** and click on a firewall policy.

Use this page to view and manage policy intents associated with your site or site groups. You can filter and sort this information to get a better understanding of what you want to configure.

Tasks You Can Perform

You can perform the following tasks from this page:

- Create a firewall policy intent. See "Adding Firewall Policy Intents" on page 381.
- Modify, clone or delete firewall policy intents. See "Editing, Cloning, and Deleting Firewall Policy Intents" on page 387.
- Deploy a firewall policy. See "Deploying Policies" on page 668.

  NOTE: An orange line is displayed against all undeployed firewall policy intents.

- Search for a firewall policy intent. Click the Search icon in the top right corner of the page to search for a firewall policy intent.

  You can enter partial text or full text of the keyword in the text box and press Enter.

  The search results are displayed on the same page.
- Show or hide columns. Click the **Show Hide Columns** icon at the top right corner of the page.
- View undeployed intents. Click the **Show Hide Columns** icon at the top right corner of the page and select **Undeployed Intent** under **Quick Filters**.

Field Descriptions

Table 108 on page 378 provides guidelines on using the fields on the **Firewall Policy** page.

Table 108: Fields on the Firewall Policy Page

Field	Description
Source	Source endpoint to which a firewall policy intent applies. A source endpoint can be addresses, sites, site groups, departments, users, or Internet (all in-bound traffic).
Destination	Destination endpoint to which a firewall policy intent applies. A destination endpoint can be addresses, services, sites, application signatures and groups, services and groups, or departments.
Options	Displays whether scheduling, logging, and UTM options are enabled for the firewall policy intent.
Total	Number of intents associated with the firewall policy.
Undeployed	Number of intents associated with the firewall policy that are either created new or updated, but are not yet deployed.

RELATED DOCUMENTATION

- Firewall Policy Overview | 374
- Adding Firewall Policy Intents | 381
- Firewall Policy Examples | 396
- Editing, Cloning, and Deleting Firewall Policy Intents | 387
- About the Deployments Page | 665
- Deploying Policies | 668

Adding a Firewall Policy

A firewall policy enforces rules for transit traffic, in terms of what traffic can pass through the firewall, and the actions that need to take place on traffic as it passes through the firewall.

Use this page to add a firewall policy and assign it to one or more sites.

NOTE: A single policy can have both enterprise based intents and zone based intents for SD-WAN sites, hybrid WAN sites, and next generation firewall sites.

To add a firewall policy:

1. Select **Configuration > Firewall > Firewall Policy**,

   The Firewall Policy page appears.

2. Click the plus icon (+).

   The Add Firewall Policy page appears.

3. Complete the configuration settings according to the guidelines provided inTable 109 on page 379.

   NOTE: Fields marked with an asterisk (*) are mandatory.

4. Click **OK**.

   The new firewall policy is created and and a confirmation message is displayed.

Table 109: Fields on the Add Firewall Policy Page

Field	Description
Name	Enter a unique string of alphanumeric characters that can include spaces and some special characters. The maximum length is 255 characters.
Description	Enter a description for the policy; the maximum length is 255 characters.
All Sites	Select the check box to apply the firewall policy to all sites.
Select Sites	Select one or more sites to which the policy must be applied. Select the sites from the **Available** column and click the right-arrow to move the sites to the **Selected** column.

Editing and Deleting Firewall Policies

IN THIS SECTION

- Editing Firewall Policies | 380
- Deleting Firewall Policies | 380

You can edit and delete firewall policies from the **Firewall Policy** page.

Editing Firewall Policies

> **NOTE:** You cannot modify the firewall policy name.

To modify the parameters configured for a firewall policy:

1. Select **Configuration > Firewall > Firewall Policy**.

   The **Firewall Policy** page appears, displaying the list of firewall policies.

2. Hover over the firewall policy that you want to edit, and then click the **...** icon that appears on the right side of the page.

3. Click **Edit**.

   The Edit Firewall Policy page appears displaying the same options that you entered while creating the firewall policy.

4. Modify the parameters following the guidelines provided in "Adding a Firewall Policy" on page 378.

5. Click **OK** to save the changes. If you want to discard your changes, click **Cancel** instead.

If you click **Save**, the modified policy appears on the **Firewall Policy** page.

Deleting Firewall Policies

To delete a firewall policy:

1. Select **Configuration > Firewall > Firewall Policy**.

The **Firewall Policy** page appears, displaying the list of firewall policies.

2. Select the firewall policy that you want to delete and then click the **...** icon that appears on the right side of the policy and click **Remove**.

   Alternatively, you can select the firewall policy and click the Delete icon.

   A message requesting confirmation for the deletion appears.

3. Click **Yes** to delete the selected firewall policy. If you want to discard your changes, click **Cancel** instead.

If you click **Yes**, the selected policy is deleted from the Firewall Policy page.

RELATED DOCUMENTATION

Adding a Firewall Policy | 378

Adding Firewall Policy Intents

Use this page to add a firewall intent that controls transit traffic within a context. The traffic is classified by matching its source and destination zones, the source and destination addresses, and the application that the traffic carries in its protocol headers with the policy database.

You can also enable advanced security protection by specifying one or more of the following:

- Unified threat management (UTM) profile
- SSL proxy profile
- Intrusion prevention system (IPS) profile

To configure a firewall policy intent:

1. Select **Configuration > Firewall > Firewall Policy**.

   The Firewall Policy page appears.

2. Click the firewall policy to which you want to add the intent.

   The *Firewall-Policy-Name* page appears.

3. Click the add icon (+).

   The option to create firewall policy intent appears inline on the *Firewall-Policy-Name* page.

4. Complete the configuration according to the guidelines provided in Table 110 on page 382.

5. Click **Save** to save the changes. If you want to discard your changes, click **Cancel** instead.

If you click **Save**, a new firewall policy intent with the provided configuration is saved and a confirmation message is displayed. Based on the source and destination end points, the intents are categorized as zone-based intents and enterprise-based intents.

> **NOTE:** After the policy intent is created, you must deploy the policy to ensure that the changes take effect on the applicable sites, departments, or applications. When a firewall policy intent is created, the Undeployed field is incremented by one indicating that intents are pending deployment.

Table 110: Fields on the <Firewall-Policy-Name> Page

Field	Description
General Information	
Name	Enter a unique string of alphanumeric characters, colons, periods, dashes, and underscores. No spaces are allowed and the maximum length is 255 characters. If you do not enter a name, the intent is saved with a default name assigned by CSO.
Description	Enter a description for the policy intent; maximum length is 1024 characters.

Table 110: Fields on the <Firewall-Policy-Name> Page *(continued)*

Field	Description
Select Schedule	Policy schedules enable you to define when a policy is active, and thus are an implicit match criterion. You can define the day of the week and the time of the day when the policy is active. For instance, you can define a security policy that opens or closes access based on business hours. Select a pre-saved schedule and the schedule options are populated with the selected schedule's data. You can add a schedule from the **End Points** panel, by selecting the schedule and clicking on the check mark icon (**√**). You can also create new schedules and then associate the schedule to your firewall policy. To create a new schedule and then add it to a firewall policy: 1. Click **Select Schedule**. 2. Click **Add schedule**. The **Create Schedules** page appears. 3. Create a new schedule. See "Creating Schedules" on page 438. The new schedule appears in the **End Points** tab, under **Schedules**. 4. Select the schedule and click on the add icon (+) to add it to the firewall policy.
Logging	Click the toggle button to enable logging; by default, logging is disabled. You can see the logged firewall events in the **Firewall Events** page by using **Monitor** > **Security Events** > **Firewall Events**. For more information, see "About the Firewall Events Page" on page 818.
Identify the traffic that the intent applies to	
Source	Click the add icon (**+**) to select the source end points on which the firewall policy intent applies, from the displayed list of addresses, departments, sites, site groups, users, zones, or the Internet. You can also select a source end point using the methods described in "Selecting Firewall Source" on page 389.

Table 110: Fields on the <Firewall-Policy-Name> Page *(continued)*

Field	Description
Destination	Click the add icon (+) to select the destination end points on which the firewall policy intent applies, from the displayed list of addresses, applications, application groups, departments, services, sites, site groups, zones or the Internet. You can also select a destination end point using the methods described in "Selecting Firewall Destination" on page 393.
Select Action	Click the add icon (+) to choose whether you want to permit, deny, or reject traffic between the source and destination. • **Allow**—Device permits traffic using the type of firewall authentication you applied to the policy. • **Deny**—Device silently drops all packets for the session and does not send any active control messages such as TCP Resets or ICMP unreachable. • **Reject**—Device sends a TCP reset if the protocol is TCP, and device sends an ICMP reset if the protocols are UDP, ICMP, or any other IP protocol. This option is useful when dealing with trusted resources so that applications do not waste time waiting for timeouts and instead get the active message.
Advanced Security	NOTE: This field is enabled only if you either select **Allow** for the action or if you select a zone as a source and destination. • UTM Profile—When you set the action to **Allow**, you can specify a UTM profile by selecting a profile from the list (under **UTM Profiles [UTM]**). You specify a UTM profile for protection against multiple threat types including spam and malware, and control access to unapproved websites and content. You can add a new UTM profile by clicking **+** in the End Points pane and selecting **UTM Profiles**. See "Creating UTM Profiles" on page 457. • IPS Profile—When you set the action to **Allow**, you can specify an IPS profile by selecting a profile from the list (under **IPS Profiles [IPS]**). You specify an IPS profile to monitor and prevent intrusions. • SSL Proxy Profile—When you configure a zone as part of the source and the destination, you can specify an SSL proxy profile by selecting a profile from the list (under **SSL Profiles [SSLP]**). You add an SSL proxy profile to ensure the secure transmission of data between a client and a server through a combination of privacy, authentication, confidentiality, and data integrity. You can also add a new SSL proxy profile by clicking **+** in the End Points pane and selecting **SSL Proxy Profile**. See "Creating SSL Forward Proxy Profiles" on page 655.
Add source and destination end points	

Table 110: Fields on the <Firewall-Policy-Name> Page *(continued)*

Field	Description
End Points	

Table 110: Fields on the <Firewall-Policy-Name> Page *(continued)*

Field	Description
	To add an end point to the source or destination: 1. Click on **Select Source** or **Select Destination** text box and then click the lesser-than icon on the right side of the page to open the **End Points** panel. The **End Points** panel displays the end points relevant to the source or destination based on your selection. • End points from addresses, departments, users, zones, and sites are displayed for source. **NOTE:** If JIMS is not configured for CSO, users will not be listed in the **End Points** panel. Instead you will be provided with an option to import users through the **Administration > Identity Management** page. To import users, click **Set Up** and follow the steps provided in "About the Identity Management Page" on page 366. • End points from addresses, applications, departments, services, zones, and sites are displayed for destination. **NOTE:** You can also search for a specific end point using the search option. 2. (Optional) Click on the edit icon (pencil symbol) to modify an end point. 3. (Optional) Click on the details icon on the right of the end point, to view more information about a source or destination end point. 4. Select the end point you want to add and click on the check mark icon (**√**) to add it the source or destination. The selected end point is added to the source or destination. To add new source and destination end points: 1. Click the less-than icon (<) on the right side of the page, to open the **End Points** panel. 2. Click on the add icon (+) on the top right of the **End Points** panel. A list of end points that you can add is displayed. 3. Select the end point you want to add. You can add the following end points: • Address or address group. See "Creating Addresses or Address Groups" on

Table 110: Fields on the <Firewall-Policy-Name> Page *(continued)*

Field	Description
	page 737. • Site or site group. See "Creating Site Groups" on page 197. • Department. See "Adding a Department" on page 767. • Service or service group. See "Creating Services and Service Groups" on page 744. • Application signature or application signature group. See "Adding Application Signatures" on page 757, and "Adding Application Signature Groups" on page 764. • Create a schedule. See "Creating Schedules" on page 438. 4. Click **Save** to add the new end point. The created end point is listed in the **End Points** panel. 5. Select the end point you want to add to the source or destination, and click on the check mark icon (√). The end point is added to the source or destination.

RELATED DOCUMENTATION

Firewall Policy Overview | 374

About the Firewall Policy Name Page | 377

Firewall Policy Examples | 396

Editing, Cloning, and Deleting Firewall Policy Intents | 387

Editing, Cloning, and Deleting Firewall Policy Intents

IN THIS SECTION

- Editing Firewall Policy Intents | 388
- Cloning Firewall Policy Intents | 388
- Deleting Firewall Policy Intents | 389

You can edit, clone, and delete firewall policy intents from the **Firewall Policy** page.

Editing Firewall Policy Intents

To modify the parameters configured for a firewall policy intent:

1. Select **Configuration > Firewall > Firewall Policy**.

   The **Firewall Policy** page appears, displaying the list of firewall policies.

2. Click the firewall policy for which you want to edit the firewall policy intents.

   The firewall policy intents are displayed in the Firewall Policy page.

3. Hover over the firewall policy intent that you want to edit, and then click the **...** icon that appears on the right side of the intent. Click **Edit**.

   The **Firewall Policy** page displays the same options as those that appear when you create a new firewall policy intent.

4. Modify the parameters following the guidelines provided in "Adding Firewall Policy Intents" on page 381.

5. Click **Save** to save the changes. If you want to discard your changes, click **Cancel** instead.

If you click **Save**, the modified intent appears on the **Firewall Policy** page.

Cloning Firewall Policy Intents

To clone a firewall policy intent:

1. Select **Configuration > Firewall > Firewall Policy**.

   The **Firewall Policy** page appears, displaying the intents associated with the policy.

2. Click the firewall policy for which you want to clone the firewall policy intents.

   The firewall policy intents are displayed in the Firewall Policy page.

3. Hover over the firewall policy intent that you want to clone, and then click the **...** icon that appears on the right side of the intent. Click **Clone**.

   The **Firewall Policy** page displays the same options as those that appear when you create a new firewall policy intent. Update the cloned intent as required.

4. Click **Save** to save the changes. If you want to discard your changes, click **Cancel** instead.

If you click **Save**, the cloned intent is added to the firewall policy and appears on the **Firewall Policy** page.

Deleting Firewall Policy Intents

To delete a firewall policy intent:

1. Select **Configuration > Firewall > Firewall Policy**.

   The **Firewall Policy** page appears, displaying the intents associated with the policy.

2. Click the firewall policy for which you want to delete the firewall policy intent.

   The firewall policy intents are displayed in the Firewall Policy page.

3. Select the firewall policy intent you want to delete, and then click the **...** icon that appears on the right side of the intent. Click **Delete**.

   An alert message appears, verifying that you want to delete the selected intent.

4. Click **Yes** to delete the selected intent. If you do not want to delete, click **Cancel** instead.

If you click **Yes**, the selected intent is deleted from the policy.

RELATED DOCUMENTATION

- Firewall Policy Overview | 374
- About the Firewall Policy Name Page | 377
- Firewall Policy Examples | 396
- Adding Firewall Policy Intents | 381

Selecting Firewall Source

IN THIS SECTION

- Adding an End Point as Firewall Source | 390
- Selecting Firewall Source Using Abbreviations | 391
- Selecting a Firewall Source from the End Points Panel | 391

Creating and Selecting a Firewall Source from the End Points Panel | 392
Creating Addresses from Source | 392

The following procedures provides various methods using which you can choose a firewall source end point:

Adding an End Point as Firewall Source

View and select the source end point from the complete list of addresses, sites, site groups, zones, or departments. You can also select the **Internet** option which denotes all in-coming traffic from outside your network.

NOTE: When you select **Any** address as a source, it implies traffic originating within the network.

NOTE:
The following conditions apply when you select **Internet** as a source end point:

- When **Internet** is not chosen as a source end point, it is implied that the traffic is originating within the network.
- If you chose **Internet** as a source, you cannot add other sites, site groups or departments as a source end point along with **Internet**.
- If you chose **Internet** as a source, the destination end point must be a site, site group, or department.

1. Click the **Source** field. A list of relevant endpoints are displayed.
2. Click on **View more results** link provided at the bottom of the source end points. The complete list of addresses, departments, users, sites, site groups, and zones is displayed in the **End Points** panel on the right.
3. (Optional) Click the edit icon to edit the address, users, department, or site group end point. You cannot edit a site end point.
4. Click check mark icon (√) to select the end point as a source.

Selecting Firewall Source Using Abbreviations

Enter an abbreviation in the **Source** field to select the source end point from a filtered list of source endpoints.

- To view a filtered list of addresses, enter **ADDR** or **addr**.
- To view a filtered list of departments, enter **DEPT** or **dept**.
- To view a filtered list of sites, enter **SITE** or **site**.
- To view a filtered list of site groups, enter **STGP** or **stgp**.
- To view a filtered list of user ids, enter **USER** or **user**.
- To view a filtered list of zones, enter **ZONE** or **zone**.

Click the endpoints in the filtered list to select them. You can also select the end point from the complete list of addresses, departments, users, sites, and site groups. See "Adding an End Point as Firewall Source" on page 390.

Selecting a Firewall Source from the End Points Panel

You can select a firewall source end point from the **End Points** panel. Alternately, you can create a new firewall source end point from the **End Points** panel, see "Creating and Selecting a Firewall Source from the End Points Panel" on page 392

To select an firewall source end point from the from the **End Points** panel:

1. Click on the **Source** field.

2. Click the lesser-than icon (**<**) on the right.

   The **End Points** panel appears, displaying the list of available addresses, departments, users, sites, site groups, and zones.

3. (Optional) To view more information about a source end point, click the details icon on the right of the end point. To edit the source end point, click the edit icon (pencil symbol) on the right of the end point.

   > **NOTE:** You can only edit or view details of a source end point if these options appear on right side of the end point when you hover over it. Not all endpoints provide these options.

4. Click the check mark icon (**√**) to add the end point as a source.

Creating and Selecting a Firewall Source from the End Points Panel

To create an new source end point from the **End Points** panel:

1. Click the add icon (**+**) on the top right of the panel and select the type of end point you want to create, among the options provided.

   Based on the option you select, the respective page appears. Fill in the required details to create a new end point.

   - To create a new address, see "Creating Addresses or Address Groups" on page 737.
   - To create a site or site group, see "Creating Site Groups" on page 197.

   After the end point is created, it appears in the **End Points** panel.

2. Click the check mark icon (**√**) to add the new end point as a source.

Creating Addresses from Source

You can use one of the following ways to create a new address from the **Source** field and use the newly created address as a source end point:

- Type the address directly in the **Source** field. If the address is valid, it is created immediately and added as a source end point.
- Create an address from the **Source** field, using the following steps:

  1. In the **Source** field, type **addr**. The **Add new address** link appears at the bottom of the list of addresses.
  2. Click **Add new address** to create a new address.

     The **Create Addresses** page appears.
  3. Configure the new address. See "Creating Addresses or Address Groups" on page 737.
  4. Click **Save** to save the new address.

     The new address is created, and will be listed as an option for the source. Select the new address to add it to the source.

RELATED DOCUMENTATION

Selecting Firewall Destination | 393

Adding Firewall Policy Intents | 381

Firewall Policy Overview | 374
About the Firewall Policy Name Page | 377
Editing, Cloning, and Deleting Firewall Policy Intents | 387

Selecting Firewall Destination

IN THIS SECTION

- Adding an End Point as Firewall Destination | 393
- Selecting Firewall Destination Using Abbreviations | 394
- Selecting a Firewall Destination from the End Points Panel | 394
- Creating and Selecting a Firewall Destination from the End Points Panel | 395
- Creating Addresses from Destination | 395

The following procedures provides various methods using which you can choose a firewall destination end point:

Adding an End Point as Firewall Destination

View and select the end point from the complete list of addresses, applications, application groups, departments, services, sites, site groups, or zones.

NOTE:

- When you choose **Any** address or service as the destination, it implies that traffic is flowing outside the network unless a site or department is mentioned explicitly.
- Unless you choose a site, site group, or department as a destination end point, it is implied the traffic will flow outside your network.

1. Click on **Destination**. A list of relevant end points are displayed.

2. Click on **View more results** link provided at the bottom of the destination end points. The complete list of addresses, departments, sites, and site groups is displayed in the **End Points** panel on the right.

3. (Optional) Click the edit icon to edit the address, department, or site group end point. You cannot edit a site end point.

4. Click check mark icon (√) to select the end point as a destination.

Selecting Firewall Destination Using Abbreviations

Enter an abbreviation in the **Destination** field to select the destination end point from a filtered list of destination endpoints.

- To view a filtered list of addresses, enter **ADDR** or **addr**.
- To view a filtered list of applications or application groups, enter **APPS** or **apps**.
- To view a filtered list of departments, enter **DEPT** or **dept**.
- To view a filtered list of services, enter **SVCS** or **svcs**.
- To view a filtered list of sites, enter **SITE** or **site**.
- To view a filtered list of site groups, enter **STGP** or **stgp**.
- To view a filtered list of zones, enter **ZONE** or **zone**.

Click the endpoints in the filtered list to select them. You can also select the end point from the complete list of addresses, departments, sites, and site groups. See “Adding an End Point as Firewall Destination” on page 393.

Selecting a Firewall Destination from the End Points Panel

You can select a firewall destination end point from the **End Points** panel. Alternately, you can create a new firewall destination end point from the **End Points** panel, see “Creating and Selecting a Firewall Destination from the End Points Panel” on page 395.

To select an firewall destination end point from the from the **End Points** panel:

1. Click on the **Destination** field.

2. Click the lesser-than icon (**<**) on the right.

   The **End Points** panel appears, displaying the list of available addresses, applications, application groups, departments, services, sites, site groups, or zones.

3. (Optional) To view more information about a destination end point, click the details icon on the right of the end point. To edit the destination end point, click the edit icon (pencil symbol) on the right of the end point.

> **NOTE:** You can only edit or view details of a destination end point if these options appear on right side of the end point when you hover over it. Not all endpoints provide these options.

4. Click the check mark icon (√) to add the end point as a destination.

Creating and Selecting a Firewall Destination from the End Points Panel

To create an new destination end point from the **End Points** panel:

1. Click the add icon (+) on the top right of the panel and select the type of end point you want to create, among the options provided.

   Based on the option you select, the respective page appears. Fill in the required details to add a new end point.

   - To add a new address, see “Creating Addresses or Address Groups” on page 737.
   - To add a site or site group department, see “Creating Site Groups” on page 197.
   - To add an application or application group, see “Adding Application Signatures” on page 757 and “Adding Application Signature Groups” on page 764.
   - To add a new service, see “Creating Services and Service Groups” on page 744.
   - To add an SSL proxy profile, see “Creating SSL Forward Proxy Profiles” on page 655.
   - To add an UTM Profile, see “Creating UTM Profiles” on page 457.

   After the end point is created, it appears in the **Endpoints** panel.

2. Click the check mark icon (√) to add the new end point as a destination.

Creating Addresses from Destination

You can use one of the following ways to create a new address from the **Destination** and use the newly created address as a destination end point:

- Type the address directly in the **Destination** field. If the address is valid, it is created immediately and added as a destination end point.

- Create an address from the **Destination** field, using the following steps:

  1. In the **Destination** field, type **addr**. The **Add new address** link appears at the bottom of the list of addresses.

  2. Click **Add new address** to create a new address.

     The **Create Addresses** page appears.

  3. Configure the new address. See "Creating Addresses or Address Groups" on page 737.

  4. Click **Save** to save the new address.

     The new address is created, and will be listed as an option for the destination. Select the new address to add it to the destination.

RELATED DOCUMENTATION

Adding Firewall Policy Intents | **381**

Firewall Policy Overview | **374**

About the Firewall Policy Name Page | **377**

Editing, Cloning, and Deleting Firewall Policy Intents | **387**

Firewall Policy Examples

IN THIS SECTION

- Example 1: Firewall Policy that Permits Traffic from Departments in Site A to the Departments in Site B | **398**
- Example 2: Firewall Policy that Permits Internet Access for all Departments in Site A and Site B | **400**
- Example 3: Firewall Policy that Permits Any Public Internet Address to Access the Sales Department in Site B | **403**
- Example 4: Firewall Policy that Permits Social Media Access to all Departments in Site A | **404**
- Example 5: Firewall Policy that Controls Access to Specific Applications for Various Departments | **406**
- Example 6: Firewall Policy that Denies Access to Social Networking Sites | **414**
- Example 7: Firewall Policy that Controls Access to an Address over the Internet (HTTP) | **417**
- Example 8: Firewall Policy that Permits or Denies the Use of HTTP or FTP as a Service | **423**

- Example 9: Firewall Policy that Denies Access to BitTorrent to the Finance Departments across both Site A and Site B | **425**
- Example 10: Firewall Policy that Allows Access to Facebook for Users in User Group A | **428**
- Example 11: Firewall Policy that Permits User B in Site A Access to YouTube with UTM Enabled | **432**
- Example 12: Firewall Policy that blocks access to Internet and allow access to Google Drive. | **435**

This topic provides information on how firewall policy intents that you define as part of your firewall policy is handled by Contrail Service Orchestration (CSO), using various examples. Each of the examples provide detailed explanation about how a firewall policy intent defined through the CSO GUI resolves into configuration in the system.

> **NOTE:** For more information, see "Firewall Policy Overview" on page 374 and "Adding Firewall Policy Intents" on page 381.

For easier understanding, all the examples have been defined to use the topology in illustrated in Figure 11 on page 397. In this topology, there are two sites—site A and site B. Each site has two departments defined as follows:

- Site A - IT (LAN segment LS1) and Finance (LAN segment LS2).
- Site B - Finance (LAN segment LS3) and Sales (LAN segment LS4).

Figure 11: Topology Diagram

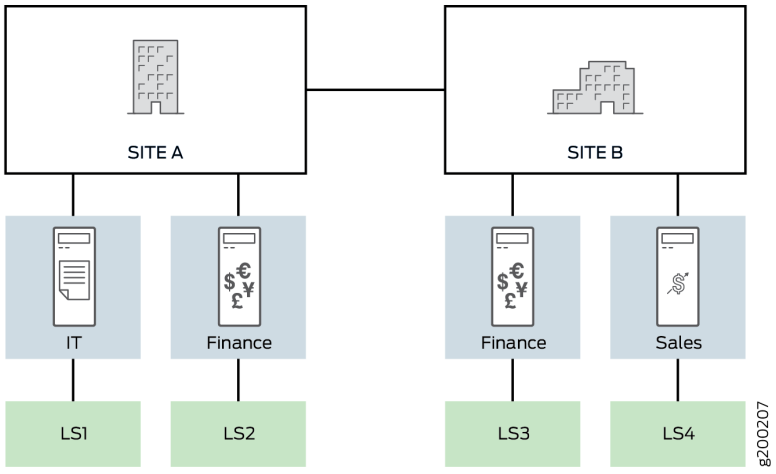

The following definitions are applicable to all the examples:

- While creating a site, you can designate some of the WAN interfaces to be breakout interfaces. These WAN interfaces can carry both site-to-site traffic (through the trust zone) and breakout traffic (through the untrust zone). The WAN interfaces can also be designated exclusively for carrying breakout traffic.
- A trust zone refers to the overlay interface that contains all the GRE tunnel interfaces, such as gr-0/0/0.1, gr-0/0/0.2, and IPSec interfaces, such as st0.1, st0.2 created between the sites.
- An untrust zone refers to the underlay interfaces (underlying physical interfaces) such as ge-0/0/0, ge-0/0/1.
- If you select an address or a service as a destination endpoint, CSO considers it as an address or service hosted on the Internet, unless the selected address or service is associated with a site.
- Table 111 on page 398 captures the addresses associated with the LAN segments used in the topology illustrated in Figure 11 on page 397.

Table 111: LAN Segments Definition

Site	Department	LAN Segment	LAN Segment Address
site A	IT	LS1	192.0.2.0/24
site A	Finance	LS2	192.168.1.0/24
site B	Finance	LS3	198.51.100.0/24
site B	Sales	LS4	203.0.113.0/24

The following examples help you understand the creation of intent-based firewall policies for various traffic scenarios across sources and destinations.

Example 1: Firewall Policy that Permits Traffic from Departments in Site A to the Departments in Site B

Define a firewall policy that permits traffic from the departments in site A to the departments in site B.

Table 112 on page 398 shows the firewall policy intent that is defined:

Table 112: Firewall Policy Intent Definition for Example - 1

Source	Destination	Action
site A	site B	Permit

Table 113 on page 399 shows how this firewall policy intent is resolved:

Table 113: Firewall Policy Intent Resolution for Example - 1

Site	Source Department	Source Address	Zone	Destination Address	Service	Intent Created
site A	Finance	[LS2]	Trust	[LS3, LS4]	Any	Intent 1__0
	IT	[LS1]	Trust	[LS3, LS4]	Any	Intent 1__1
site B	Trust	[LS3, LS4]	Sales	[LS2]	Any	Intent 1__0
	Trust	[LS3, LS4]	Finance	[LS1]	Any	Intent 1__1

Configuration Output Sample

Sample of configuration that permits traffic from departments in site A to the departments in site B.

The hierarchy level for the following configuration sample is **[edit security policies]**.

```
from-zone FINANCE to-zone trust {
       policy Intent_1__0 {
        match {
            source-address ls-192.168.1.0/24-SP50-L2;
            destination-address [ls-198.51.100.0/24-SP50-L3,
     ls-203.0.113.0/24-SP50-L4];
            application any;
      }
 then {
     permit;
      }
         }
     }
from-zone IT to-zone trust {
       policy Intent_1__1 {
        match {
            source-address ls-192.0.2.0/24-S42-L1;
            destination-address [ls-198.51.100.0/24-SP50-L3,
                        ls-203.0.113.0/24-SP50-L4];
            application any;
             }
        then {
       permit;
```

```
        }
           }
      }
```

Sample of configuration that permits traffic from departments in site B to the departments in site A.

The hierarchy level for the following configuration sample is **[edit security policies]**.

```
from-zone trust to-zone SALES {
       policy Intent_1__0 {
           match {
               source-address [ls-198.51.100.0/24-SP50-L3,
            ls-203.0.113.0/24-SP50-L4];
         destination-address ls-192.0.2.0/24-S42-L1;
               application any;
                }
           then {
               permit;
  }
           }
      }
from-zone trust to-zone FINANCE {
      policy Intent_1__1 {
          match {
              source-address [ls-198.51.100.0/24-SP50-L3,
         ls-203.0.113.0/24-SP50-L4];
              destination-address ls-192.168.1.0/24-SP50-L2;
              application any;
              }
          then {
              permit;
               }
   }
     }
```

Example 2: Firewall Policy that Permits Internet Access for all Departments in Site A and Site B

Define a firewall policy that permits all the department in site A and site B access to Internet.

Table 114 on page 401 shows the firewall policy intent that is defined:

Table 114: Firewall Policy Intent Definition for Example - 2

Source	Destination	Action
site A	http, https, icmp-ping, dns	Permit
site B	http, https, icmp-ping, dns	Permit

Table 115 on page 401 shows how this firewall policy intent is resolved:

Table 115: Firewall Policy Intent Resolution for Example - 2

Site	Source Department	Source Address	Zone	Destination Address	Service	Intent Created
site A	Finance	[LS2]	Untrust	Any	http, https, icmp-ping, dns	Intent 1__0
	IT	[LSI]	Untrust	Any	http, https, icmp-ping, dns	Intent 1__1
site B	Sales	[LS4]	Untrust	Any	http, https, icmp-ping, dns	Intent 1__0
	Finance	[LS3]	Untrust	Any	http, https, icmp-ping, dns	Intent 1__1

Configuration Output Sample

Sample of configuration that permits Internet access to all departments in site A.

The hierarchy level for the following configuration sample is **[edit security policies]**.

```
from-zone Finance to-zone untrust {
     policy Intent_1__0 {
        match {
               source-address ls-192.168.1.0/24-SP50-L2;
               destination-address any;
```

```
                application [junos-http junos-dns-tcp junos-https
                             junos-icmp-ping];
                }
            then {
                  permit;
                }
              }
        }
  from-zone IT to-zone untrust {
       policy Intent_1__1 {
          match {
                source-address ls-192.0.2.0/24-S42-L1;
                destination-address any;
                application [junos-http junos-dns-tcp junos-https
                             junos-icmp-ping];
                }
          then {
                permit;
              }
           }
       }
   policy-rematch;
```

Sample of configuration that permits Internet access to all departments in site B.

The hierarchy level for the following configuration sample is **[edit security policies]**.

```
from-zone Sales to-zone untrust {
     policy Intent_1__0 {
        match {
               source-address ls-203.0.113.0/24-SP50-L4;
               destination-address any;
               application [junos-http junos-dns-tcp junos-https
                            junos-icmp-ping];
              }
        then {
              permit;
             }
         }
     }
 from-zone Finance1 to-zone untrust {
```

```
      policy Intent_1__1 {
         match {
                source-address ls-198.51.100.0/24-SP50-L3;
                destination-address any;
                application [junos-http junos-dns-tcp junos-https
                             junos-icmp-ping];
                }
          then {
                permit;
               }
           }
      }
   policy-rematch;
```

Example 3: Firewall Policy that Permits Any Public Internet Address to Access the Sales Department in Site B

Define a firewall policy that permits any public Internet address access to a sales application hosted by the Sales department in site B.

> **NOTE:** For this example, breakout is not enabled and MPLS link type is used.

Table 116 on page 403 shows the firewall policy intent that is defined:

Table 116: Firewall Policy Intent Definition for Example - 3

Source	Destination	Action
Internet	Sales, site B	Permit

Table 117 on page 403 shows how this firewall policy intent is resolved:

Table 117: Firewall Policy Intent Resolution for Example - 3

Source Address	Zone	Destination Address	Service	Intent Created
Any public Internet address	Trust to Sales (No breakout)	[LS4]	Any	Intent 1__0

Configuration Output Example

Sample of configuration that permits any public Internet address to access the Sales department in site B.

The hierarchy level for the following configuration sample is **[edit security policies]**.

```
from-zone untrust to-zone Sales {
     policy Intent_1__0 {
        match {
               source-address any;
               destination-address ls-203.0.113.0/24-SP50-L4;
               application any;
               }
        then {
               permit;
             }
         }
     }
```

Example 4: Firewall Policy that Permits Social Media Access to all Departments in Site A

Define a firewall policy that permits all departments in site A access to Facebook.

Table 118 on page 404 shows the firewall policy intent that is defined:

Table 118: Firewall Policy Intent Definition for Example - 4

Source	Destination	Action
site A	Facebook	Permit

Table 119 on page 404 shows how this firewall policy intent is resolved:

Table 119: Firewall Policy Intent Resolution for Example - 4

Site	Source Address	Zone	Destination Address	Service	Intent Created	Application Firewall Profile
site A	[LS2]	Untrust	Facebook	Any	Intent 1__0	AppFwProfile_0
site A	[LS1]	Untrust	Facebook	Any	Intent 1__1	AppFwProfile_0

Configuration Output Example

Sample of configuration that controls access to Facebook for site A.

The hierarchy level for the following configuration sample is **[edit security policies]**.

```
from-zone Finance to-zone untrust {
     policy Intent_1__0 {
        match {
               source-address ls-192.168.1.0/24-SP50-L2;
               destination-address any;
               application any;
               }
         then {
               permit {
                  application-services {
                  application-firewall {
                  rule-set AppFwProfile_0;
                  }
                }
              }
            }
          }
        }
 from-zone IT to-zone untrust {
      policy Intent_1__1 {
         match {
                source-address ls-192.0.2.0/24-S42-L1;
                destination-address any;
                application any;
                }
          then {
                permit {
                   application-services {
                   application-firewall {
                   rule-set AppFwProfile_0;
                   }
                 }
               }
             }
           }
         }
  policy-rematch;
```

The hierarchy level for the following configuration sample is **[edit security application-firewall]**.

```
application-firewall {
   rule-sets AppFwProfile_0 {
       rule rule-1 {
               match {
                   dynamic-application junos:FACEBOOK-APP;
                   ssl-encryption any;
                     }
                then {
                    permit;
                }
            }
         default-rule {
           deny;
             }
          }
      }
```

Example 5: Firewall Policy that Controls Access to Specific Applications for Various Departments

Define a firewall policy that controls access to specific applications from various departments, with the following intents:

- The finance departments located in site A and site B (which are in different geographical locations) are permitted to access the news applications BBC and CNN.
- The IT department located in site A is denied access to the news applications BBC and CNN.
- Access to Telnet and SSH applications is given only to the finance departments.
- Access to Telnet and SSH applications is denied to all departments, except for the finance department.

Table 120 on page 406 shows the firewall policy intents that are to fulfil this requirement:

Table 120: Firewall Policy Intent Definition for Example - 5

Source	Destination	Action
Finance department, site A and Finance department, site B	BBC and CNN	Permit
IT department, site A	BBC and CNN	Deny
Finance department, site A and Finance department, site B	Telnet and SSH	Permit

Table 120: Firewall Policy Intent Definition for Example - 5 *(continued)*

Source	Destination	Action
Any (All addresses except the finance department)	Telnet and SSH	Deny

NOTE: The number of intents depends on the number of source sites within the given department and the number of destination sites.

Table 121 on page 407 shows how this firewall policy intent is resolved:

Table 121: Firewall Policy Intent Resolution for Example - 5

Source Department	Source Address	Zone	Destination Address	Service	Application Firewall Profile
Finance	[LS2]	Trust/Untrust	Any	Any	AppFwProfile_1 Permit: CNN/BBC Def. Rule : Permit
Finance	[LS3]	Trust/Untrust	Any	Any	AppFwProfile_1 Permit: CNN/BBC Def. Rule : Permit
IT	[LS1]	Trust/Untrust	Any	Any	AppFwProfile_3 Deny: CNN/BBC Def. Rule : Deny
Finance department, site A and Finance department, site B	[LS2, LS3]	Trust/Untrust	Any	Telnet, SSH	AppFwProfile_1-1 Permit: Telnet/SSH Def. Rule : Deny

Table 121: Firewall Policy Intent Resolution for Example - 5 *(continued)*

Source Department	Source Address	Zone	Destination Address	Service	Application Firewall Profile
IT department, site A	[LS1]	Trust/Untrust	Any	Telnet, SSH	AppFwProfile_3-1 Deny: Telnet/SSH Def. Rule : Deny

Configuration Output Example

Sample of configuration that controls access to specific applications for various departments in site A.

The hierarchy level for the following configuration sample is **[edit security policies]**.

```
from-zone Finance to-zone trust {
      policy Intent_3 {
        match {
            source-address ls-192.168.1.0/24-SP50-L2;
            destination-address any;
            application [junos-telnet junos-ssh];
        }
        then {
            permit {
                application-services {
                    application-firewall {
                        rule-set AppFwProfile_1-1;
                    }
                }
            }
        }
      }
      policy Intent_1 {
        match {
            source-address ls-192.168.1.0/24-SP50-L2;
            destination-address any;
            application any;
        }
        then {
            permit {
```

```
                application-services {
                    application-firewall {
                        rule-set AppFwProfile_1;
                    }
                }
            }
        }
      }
       policy Intent_4__0 {
        match {
            source-address any;
            destination-address any;
            application [junos-telnet junos-ssh];
        }
        then {
            permit;
            }
          }
     }
from-zone IT to-zone trust {
     policy Intent_4__1-1 {
        match {
            source-address ls-192.0.2.0/24-S42-L1;
            destination-address any;
            application [junos-telnet junos-ssh];
          }
        then {
            permit {
                application-services {
                    application-firewall {
                        rule-set AppFwProfile_3-1;
                   }
                }
             }
          }
       }
     policy Intent_2 {
        match {
            source-address ls-192.0.2.0/24-S42-L1;
            destination-address any;
            application any;
      }
```

```
          then {
              permit {
                  application-services {
                      application-firewall {
                          rule-set AppFwProfile_3;
                      }
                   }
                }
              }
           }
         policy Intent_4__1 {
           match {
              source-address any;
              destination-address any;
              application [junos-telnet junos-ssh];
               }
          then {
              deny;
              }
          }
      }
```

The hierarchy level for the following configuration sample is **[edit security application-firewall]**.

```
rule-sets AppFwProfile_1-1 {
     rule rule-1 {
        match {
            dynamic-application [junos:BBC junos:CNN];
            ssl-encryption any;
             }
        then {
            permit;
             }
           }
        default-rule {
        deny;
         }
       }
rule-sets AppFwProfile_3 {
     rule rule-2 {
```

```
        match {
            dynamic-application [junos:BBC junos:CNN];
            ssl-encryption any;
              }
        then {
            deny;
           }
         }
        default-rule {
        deny;
         }
       }
rule-sets AppFwProfile_1 {
      rule rule-3 {
        match {
            dynamic-application [junos:BBC junos:CNN];
            ssl-encryption any;
              }
        then {
            permit;
            }
          }
        default-rule {
        deny;
         }
        }
rule-sets AppFwProfile_3-1 {
     rule rule-4 {
        match {
            dynamic-application [junos:BBC junos:CNN];
            ssl-encryption any;
             }
        then {
            deny;
             }
            }
      default-rule {
        deny;
          }
       }
```

Sample of configuration that controls access to specific applications for various departments in site B.

The hierarchy level for the following configuration sample is **[edit security policies]**.

```
from-zone Finance to-zone trust {
     policy appQoe-36600-Permit-rule {
        match {
            source-address any;
            destination-address any;
            application appQoe-36000;
        }
        then {
            permit;
        }
      }
     policy Intent_3 {
        match {
            source-address ls-198.51.100.0/24-SP50-L3;
            destination-address any;
            application [ junos-telnet junos-ssh ];
        }
        then {
            permit {
                application-services {
                    application-firewall {
                        rule-set AppFwProfile_1-1;
                    }
                 }
               }
           }
        }
      policy Intent_1 {
        match {
            source-address ls-198.51.100.0/24-SP50-L3;
            destination-address any;
            application any;
           }
         then {
           permit {
               application-services {
                  application-firewall {
                        rule-set AppFwProfile_1;
                    }
                  }
```

```
                }
             }
          }
      policy Intent_4__1 {
         match {
            source-address any;
            destination-address any;
            application [junos-telnet junos-ssh];
          }
         then {
            deny;
             }
           }
        }
from-zone Sales to-zone trust {
      policy Intent_4__0 {
        match {
            source-address any;
            destination-address any;
            application [junos-telnet junos-ssh];
             }
        then {
            deny;
            }
       }
    }
  policy-rematch;
```

The hierarchy level for the following configuration sample is **[edit security application-firewall]**.

```
rule-sets AppFwProfile_1-1 {
     rule rule-ca2354d6-a7ba-488e-8c5a-91cbddfb9583-appFwRule {
        match {
            dynamic-application [junos:BBC junos:CNN];
             ssl-encryption any;
            }
        then {
            permit;
             }
          }
         default-rule {
         deny;
           }
```

```
        }
rule-sets AppFwProfile_1 {
       rule rule-ca2354d6-a7ba-488e-8c5a-91cbddfb9583-appFwRule {
         match {
            dynamic-application [junos:BBC junos:CNN];
            ssl-encryption any;
            }
          then {
            permit;
              }
            }
         default-rule {
        deny;
     }
  }
```

Example 6: Firewall Policy that Denies Access to Social Networking Sites

Define a firewall policy that denies access to networking sites such as Facebook and Twitter (defined as application group Social Networking) to the IT and finance departments located in Site A.

Table 122 on page 414 shows the firewall policy intent that is needed to fulfil this requirement:

Table 122: Firewall Policy Intent Definition for Example - 6

Source	Destination	Action
IT and Finance, site A	Application group Social Networking (Facebook and Twitter)	Deny

NOTE: Add site A if the IT or finance departments are present in different sites, but you only want to apply this firewall policy intent to the IT or finance departments present in site A.

Table 123 on page 415 shows how this firewall policy intent is resolved:

Table 123: Firewall Policy Intent Resolution for Example - 6

Source Department	Source Address	Zone	Destination Address	Service	Application Firewall Profile
Finance	[LS2]	Trust/Untrust	Any	Any	AppFwProfile_0 Deny: Social Networking (Apps) Def. Rule : Deny
IT	[LS1]	Trust/Untrust	Any	Any	AppFwProfile_1 Deny: Social Networking (Apps) Def. Rule : Deny

Configuration Output Example

Sample of configuration that denies access to social networking sites for departments in site A.

The hierarchy level for the following configuration sample is **[edit security policies]**.

```
from-zone IT to-zone untrust {
        policy Intent_1__0 {
           match {
               source-address  ls-192.0.2.0/24-S42-L1;
               destination-address any;
               application any;
              }
              then {
                 permit {
                     application-services {
                     application-firewall {
                     rule-set AppFwProfile_0;
                        }
                     }
                 }
               }
```

```
            }
        }
  from-zone Finance to-zone untrust {
         policy Intent_1__1 {
                match {
                    source-address ls-192.168.1.0/24-SP50-L2;
                    destination-address any;
                    application any;
                }
                then {
                    permit {
                        application-services {
                            application-firewall {
                                rule-set AppFwProfile_0;
                            }
                        }
                    }
                }
            }
        }
```

The hierarchy level for the following configuration sample is **[edit security application-firewall]**.

```
application-firewall {
        rule-sets AppFwProfile_0 {
            rule rule-b7e4ed02-e196-400a-88bf-f1de8973d30c-appFwRule {
                match {
                    dynamic-application-group Socialnetwork;
                    ssl-encryption any;
                }
                then {
                    deny;
                }
            }
            default-rule {
                deny;
            }
        }
```

```
    }
```

Example 7: Firewall Policy that Controls Access to an Address over the Internet (HTTP)

Define a firewall policy that controls access to an address over the Internet (HTTP) for various sites or site groups with the following intents:

- IP address prefix of site A and site B are permitted to access example.com.
- IP address prefix of site group Q1 are denied access to example-one.com. Site group Q1 consists of site A and site B.

Table 124 on page 417 shows the firewall policy intents that are needed to fulfil this requirement:

Table 124: Firewall Policy Intent Definition for Example - 7

Source	Service	Destination	Action
IP address prefix, site A and IP-Prefix, site B	HTTP	www.example.com	Permit
IP address prefix, site group Q1	HTTP	www.example-one.com	Deny

Table 125 on page 417 shows how this firewall policy intent is resolved:

Table 125: Firewall Policy Intent Resolution for Example - 7

Source Department	Source Address	Zone	Destination Address	Service	Application Firewall Profile
IT, Finance departments in site A	[LS1, LS2]	Trust/Untrust	www.example.com	Any	AppFwProfile_0 Permit: HTTP Def. Rule : Deny
Finance, Sales departments in site B	[LS3, LS4]	Trust/Untrust	www.example.com	Any	AppFwProfile_1 Permit: HTTP Def. Rule : Deny

Table 125: Firewall Policy Intent Resolution for Example - 7 *(continued)*

Source Department	Source Address	Zone	Destination Address	Service	Application Firewall Profile
IT, Finance departments in site A	[LS1, LS2]	Trust/Untrust	www.example-one.com	Any	AppFwProfile_2 Deny: HTTP Def. Rule : Deny
Finance, Sales departments in site B	[LS3, LS4]	Trust/Untrust	www.example-one.com	Any	AppFwProfile_3 Deny: HTTP Def. Rule : Deny

Configuration Output Example

Sample of configuration that controls access to an address over the Internet (HTTP) for site A.

The hierarchy level for the following configuration sample is **[edit security policies]**.

```
from-zone Finance to-zone untrust {
     policy Intent_4__0 {
        match {
            source-address ls-192.168.1.0/24-SP50-L2;
            destination-address www.example.com;
            application junos-http;
        }
        then {
            permit {
                application-services {
                    application-firewall {
                        rule-set AppFwProfile_0;
                    }
                 }
              }
           }
       }
      policy Intent_1__0 {
        match {
            source-address ls-192.168.1.0/24-SP50-L2;
```

```
            destination-address addr2;
            application junos-http;
        }
        then {
            permit {
                application-services {
                    application-firewall {
                        rule-set AppFwProfile_1;
                    }
                 }
              }
           }
        }
     }
from-zone IT to-zone untrust {
     policy Intent_4__1 {
        match {
            source-address ls-192.0.2.0/24-S42-L1;
            destination-address addr2;
            application junos-http;
        }
        then {
            permit {
                application-services {
                    application-firewall {
                        rule-set AppFwProfile_0;
                    }
                 }
              }
           }
        }
      policy Intent_1__1 {
         match {
            source-address ls-192.0.2.0/24-S42-L1;
            destination-address addr2;
            application junos-http;
         }
         then {
            permit {
                application-services {
                    application-firewall {
                        rule-set AppFwProfile_1;
```

```
                    }
                  }
                }
              }
            }
          }
       policy-rematch;
```

The hierarchy level for the following configuration sample is **[edit security application-firewall]**.

```
rule-sets AppFwProfile_1 {
     rule rule-ca2354d6-a7ba-488e-8c5a-91cbddfb9583-appFwRule {
        match {
            dynamic-application junos:YOUTUBE;
            ssl-encryption any;
        }
        then {
            deny;
          }
        }
        default-rule {
        deny;
        }
       }
rule-sets AppFwProfile_0 {
    rule rule-00f3879c-f3d7-4cb3-89b6-78328e3bff38-appFwRule {
        match {
            dynamic-application junos:CNN;
            ssl-encryption any;
        }
        then {
            permit;
         }
       }
rule rule-ca2354d6-a7ba-488e-8c5a-91cbddfb9583-appFwRule {
        match {
            dynamic-application junos:YOUTUBE;
            ssl-encryption any;
          }
        then {
```

```
            deny;
          }
        }
        default-rule {
        deny;
        }
      }
```

Sample of configuration that controls access to an address over the Internet (HTTP) for site B.

The hierarchy level for the following configuration sample is **[edit security policies]**.

```
from-zone Finance to-zone untrust {
       policy Intent_4__1 {
        match {
            source-address ls-198.51.100.0/24-SP50-L3;
            destination-address addr2;
            application junos-http;
        }
        then {
            permit {
                application-services {
                    application-firewall {
                        rule-set AppFwProfile_0;
                    }
                 }
              }
           }
        }
        policy Intent_1__1 {
         match {
            source-address ls-198.51.100.0/24-SP50-L3;
            destination-address addr2;
            application junos-http;
          }
           then {
            permit {
                application-services {
                    application-firewall {
                        rule-set AppFwProfile_1;
                    }
                  }
               }
```

```
            }
          }
       }
from-zone Sales to-zone untrust {
       policy Intent_4__0 {
           match {
               source-address ls-203.0.113.0/24-SP50-L4;
              destination-address addr2;
                 application junos-http;
            }
           then {
            permit {
                application-services {
                    application-firewall {
                        rule-set AppFwProfile_0;
                    }
                  }
                }
              }
          }
       policy Intent_1__0 {
           match {
               source-address ls-203.0.113.0/24-SP50-L4;
               destination-address addr2;
              application junos-http;
             }
        then {
              permit {
                  application-services {
                     application-firewall {
                     rule-set AppFwProfile_1;
                    }
                  }
                }
              }
            }
          }
   policy-rematch;
```

The hierarchy level for the following configuration sample is **[edit security application-firewall]**.

```
rule-sets AppFwProfile_1 {
     rule rule-ca2354d6-a7ba-488e-8c5a-91cbddfb9583-appFwRule {
```

```
        match {
            dynamic-application junos:YOUTUBE;
            ssl-encryption any;
         }
         then {
            deny;
           }
          }
         default-rule {
         deny;
        }
       }
rule-sets AppFwProfile_0 {
   rule rule-00f3879c-f3d7-4cb3-89b6-78328e3bff38-appFwRule {
        match {
            dynamic-application junos:CNN;
            ssl-encryption any;
        }
        then {
            permit;
          }
        }
      rule rule-ca2354d6-a7ba-488e-8c5a-91cbddfb9583-appFwRule {
        match {
            dynamic-application junos:YOUTUBE;
            ssl-encryption any;
             }
         then {
            deny;
             }
           }
          default-rule {
          deny;
         }
       }
```

Example 8: Firewall Policy that Permits or Denies the Use of HTTP or FTP as a Service

Define a firewall policy where a specific IP address that belongs to the IT department is permitted or denied the use of HTTP or FTP as a service.

Table 126 on page 424 shows the firewall policy intents that are needed to fulfil this requirement:

Table 126: Firewall Policy Intent Definition for Example - 8

Source	Service	Destination	Action
192.0.2.0	HTTP	example.com	Permit
192.0.2.0	FTP	example.com	Deny

Table 127 on page 424 shows how this firewall policy intent is resolved:

Table 127: Firewall Policy Intent Resolution for Example - 8

Source Department	Source Address	Zone	Destination Address	Service
IT, site A	192.0.2.0	Trust/Untrust	example.com	FTP
IT, site A	192.0.2.0	Trust/Untrust	example.com	HTTP

Configuration Output Example

Sample of configuration that allows access to HTTP

The hierarchy level for the following configuration sample is **[edit security policies]**.

```
from-zone IT to-zone trust {
      policy Intent_1__1 {
        match {
            source-address 192.0.2.0;
            destination-address example.com;
            application junos-ftp;
        }
        then {
            deny;
        }
      }
      policy Intent_4__1 {
        match {
            source-address 192.0.2.0;
            destination-address example.com;
            application junos-http;
          }
```

```
        then {
          permit;
        }
      }
    }
  policy-rematch;
```

Example 9: Firewall Policy that Denies Access to BitTorrent to the Finance Departments across both Site A and Site B

Define a firewall policy that denies access to BitTorrent for the Finance departments in site A and Site B.

Table 128 on page 425 shows the firewall policy intents that are needed to fulfil this requirement:

Table 128: Firewall Policy Intent Definition for Example - 9

Source	Destination	Action
site A, Finance department	BitTorrent	Deny
site B, Finance department	BitTorrent	Deny

Table 129 on page 425 shows how this firewall policy intent is resolved:

Table 129: Firewall Policy Intent Resolution for Example - 9

Site	Source Address	Zone	Destination Application	Service	Application Firewall Profile
Finance department, site A	[LS2]	Trust/Untrust	BitTorrent	Any	AppFwProfile_0 Deny: BitTorrent Def. Rule : Deny
Finance department, site B	[LS3]	Trust/Untrust	BitTorrent	Any	AppFwProfile_0 Deny: BitTorrent Def. Rule : Deny

Configuration Output Example

Sample of configuration that allows site A access to BitTorrent.

The hierarchy level for the following configuration sample is **[edit security policies]**.

```
from-zone Finance to-zone untrust {
      policy appQoe-36600-Permit-rule {
        match {
            source-address any;
            destination-address any;
            application appQoe-36000;
            }
        then {
            permit;
          }
       }
      policy Intent_1 {
        match {
            source-address ls-192.168.1.0/24-SP50-L2;
            destination-address any;
            application any;
             }
        then {
            permit {
                application-services {
                    application-firewall {
                        rule-set AppFwProfile_0;
                    }
                }
            }
            log {
                session-init;
                session-close;
                }
              }
           }
        }
       policy-rematch;
```

The hierarchy level for the following configuration sample is **[edit security application-firewall]**.

```
rule-sets AppFwProfile_0 {
     rule rule-2226740d-03a9-483c-b315-eddc9ae8619a-appFwRule {
        match {
            dynamic-application junos:BITTORRENT;
            ssl-encryption any;
        }
        then {
            deny;
          }
       }
         default-rule {
         deny;
        }
      }
```

Sample of configuration that allows site B to access to BitTorrent.

The hierarchy level for the following configuration sample is **[edit security policies]**.

```
from-zone Finance1 to-zone untrust {
       policy appQoe-36600-Permit-rule {
          match {
            source-address any;
            destination-address any;
            application appQoe-36000;
           }
           then {
              permit;
           }
        }
      policy Intent_4 {
          match {
            source-address ls-198.51.100.0/24-SP50-L3;
            destination-address any;
            application any;
             }
          then {
            permit {
                application-services {
                    application-firewall {
```

```
                    rule-set AppFwProfile_0;
                }
             }
          }
        log {
            session-init;
            session-close;
           }
        }
     }
   }
   policy-rematch;
```

The hierarchy level for the following configuration sample is **[edit security application-firewall]**.

```
rule-sets AppFwProfile_0 {
     rule rule-00f3879c-f3d7-4cb3-89b6-78328e3bff38-appFwRule {
        match {
            dynamic-application junos:BITTORRENT;
            ssl-encryption any;
        }
        then {
            deny;
        }
     }
     default-rule {
        deny;
     }
    }
```

Example 10: Firewall Policy that Allows Access to Facebook for Users in User Group A

Define a firewall policy where the users that are a part of user group A are provided access only to Facebook, and no other applications. User group A consists of users located in site A.

Table 130 on page 428 shows the firewall policy intent that is needed to fulfil this requirement:

Table 130: Firewall Policy Intent Definition for Example - 10

Source	Destination	Action
user group A, site A	Facebook	Permit

Table 131 on page 429 shows how this firewall policy intent is resolved:

Table 131: Firewall Policy Intent Resolution for Example - 10

Site	User/User Group	Source Address Range	Destination Address	Application
site A	user group A	192.0.2.0 to 192.0.2.20	Any	Facebook

Configuration Output Example

Sample of configuration that allows users in user group A access to Facebook.

The hierarchy level for the following configuration sample is **[edit security policies]**.

```
from-zone Finance to-zone untrust {
      policy appQoe-36600-Permit-rule {
        match {
            source-address any;
            destination-address any;
            application appQoe-36000;
         }
         then {
            permit;
            }
        }
      policy Intent_4__0 {
        match {
            source-address ls-192.168.1.0/24-SP50-L2;
            destination-address any;
            application any;
            source-identity "USERFW.LOCAL\Cert Publishers";
            }
        then {
            permit {
                application-services {
                    application-firewall {
                        rule-set AppFwProfile_0;
                    }
                }
            }
```

```
            log {
                session-init;
                session-close;
               }
            }
         }
      }
from-zone IT to-zone untrust {
      policy appQoe-36600-Permit-rule {
        match {
            source-address any;
            destination-address any;
            application appQoe-36000;
           }
        then {
            permit;
           }
        }
       policy Intent_4__1 {
         match {
            source-address ls-192.0.2.0/24-S42-L1;
            destination-address any;
            application any;
            source-identity "USERFW.LOCAL\Cert Publishers";
          }
          then {
            permit {
                application-services {
                    application-firewall {
                        rule-set AppFwProfile_0;
                    }
                }
            }
            log {
                session-init;
                session-close;
                }
             }
          }
       }
     policy-rematch;
```

The hierarchy level for the following configuration sample is **[edit security application-firewall]**.

```
rule-sets AppFwProfile_0 {
     rule rule-00f3879c-f3d7-4cb3-89b6-78328e3bff38-appFwRule {
        match {
            dynamic-application junos:FACEBOOK-APP;
            ssl-encryption any;
        }
        then {
            permit;
        }
      }
      default-rule {
        deny;
      }
    }
```

The hierarchy level for the following configuration sample is **[edit services user-identification identity-management]**.

```
connection {
       connect-method https;
       port 443;
       primary {
           address 10.213.50.50;
           client-id 1234;
           client-secret "$ABC123"; ## SECRET-DATA
           }
        token-api oauth_token/oauth;
        query-api user_query/v2;
        }
        batch-query {
        items-per-batch 200;
        query-interval 5;
      }
      ip-query {
      query-delay-time 15;
     }
```

Example 11: Firewall Policy that Permits User B in Site A Access to YouTube with UTM Enabled

Define a firewall policy where the User B located in Site A is provided access only to YouTube with UTM enabled. The user does not have permission to access any other applications.

Table 132 on page 432 shows the firewall policy intent that is needed to fulfil this requirement:

Table 132: Firewall Policy Intent Definition for Example - 11

Source	Destination	Action
user B, site A	YouTube	Permit

Table 133 on page 432 shows how this firewall policy intent is resolved:

Table 133: Firewall Policy Intent Resolution for Example - 11

Site	Source Address	User/User Group	Destination Address	UTM	Application
site A	192.0.2.22	user B	Any	Enabled	Facebook

Configuration Output Example

Sample of configuration that allows user B in site A access to YouTube, with UTM enabled.

The hierarchy level for the following configuration sample is **[edit security policies]**.

```
from-zone Finance to-zone untrust {
      policy Intent_4__0 {
         match {
            source-address ls-192.168.1.0/24-SP50-L2;
            destination-address any;
            application any;
            source-identity "userfw.local\CSO1";
          }
         then {
            permit {
                application-services {
                    utm-policy testUTM;
                    application-firewall {
                        rule-set AppFwProfile_0;
                    }
```

```
                }
            }
            log {
                session-init;
                session-close;
             }
            }
          }
       }
  from-zone IT to-zone untrust {
       policy Intent_4__1 {
         match {
             source-address ls-192.0.2.0/24-S42-L1;
             destination-address any;
             application any;
             source-identity "userfw.local\CSO1";
         }
         then {
             permit {
                 application-services {
                     utm-policy testUTM;
                     application-firewall {
                         rule-set AppFwProfile_0;
                     }
                 }
             }
             log {
                 session-init;
                 session-close;
               }
            }
         }
      }
     policy-rematch;
```

The hierarchy level for the following configuration sample is **[edit security utm]**.

```
feature-profile {
       web-filtering {
           type juniper-local;
```

```
          }
         }
       utm-policy testUTM {
       web-filtering {
       http-profile junos-wf-local-default;
        }
       anti-spam {
        smtp-profile junos-as-defaults;
        }
    traffic-options {
        sessions-per-client {
            over-limit log-and-permit;
           }
        }
     }
```

The hierarchy level for the following configuration sample is **[edit security application-firewall]**.

```
rule-sets AppFwProfile_0 {
     rule rule-00f3879c-f3d7-4cb3-89b6-78328e3bff38-appFwRule {
        match {
            dynamic-application junos:FACEBOOK-APP;
            ssl-encryption any;
        }
        then {
            permit;
        }
      }
      default-rule {
        deny;
       }
     }
```

The hierarchy level for the following configuration sample is **[edit services user-identification identity-management]**.

```
connection {
       connect-method https;
       port 443;
       primary {
          address 10.213.50.50;
```

```
        client-id 1234;
        client-secret "$ABC123"; ## SECRET-DATA
      }
     token-api oauth_token/oauth;
     query-api user_query/v2;
    }
   batch-query {
     items-per-batch 200;
     query-interval 5;
    }
   ip-query {
     query-delay-time 15;
  }
```

Example 12: Firewall Policy that blocks access to Internet and allow access to Google Drive.

The following section provides a sample firewall policy to block access to Internet and allow access to Google Drive. The firewall policy has one enterprise-based intent and one zone-based intent.

An enterprise-based intent to block access to Internet is provided in Table 134 on page 435.

Table 134: Sample Enterprise-based Intent

Rule Name	Source Endpoint	Destination Endpoint	Action
EnterpriseIntent_1	Engg (Department)	Internet	Deny

A zone-based intent to allow access to Google drive is provided in Table 135 on page 435.

Table 135: Sample Zone based Intent

Rule Name	Source Endpoint	Destination Endpoint	Action
ZoneIntent_1	Engg (Zone)	untrust(zone), google-drive	Allow

The intents in Table 134 on page 435 and Table 135 on page 435 result in firewall rules order that is provided in Table 136 on page 435.

Table 136: Sample firewall rule

Rule Name	Rule Order	Source Endpoint	Destination Endpoint	Action
ZoneIntent_1	1	Engg (Zone)	untrust(zone), google-drive	Allow
EnterpriseIntent_1	2	Engg (Department)	Internet	Deny

RELATED DOCUMENTATION

Firewall Policy Overview | 374

Adding Firewall Policy Intents | 381

Firewall Policy Schedules Overview

A schedule allows a policy to be active for a specified duration. If you want a policy to be active during a scheduled time, you must first create a schedule for that policy or link the policy to an existing schedule. When a schedule timeout expires, the associated policy is deactivated and all sessions associated with the policy are also timed out.

If a policy contains a reference to a schedule, that schedule determines when the policy is active. When a policy is active, it can be used as a possible match for traffic. A schedule lets you restrict access to, or remove a restriction from a resource, for a period of time.

A schedule uses the following guidelines:

- A schedule can have multiple policies associated with it; however, a policy cannot be associated with multiple schedules.
- A policy remains active as long as the schedule it refers to is also active.

  A schedule can be active during a single time slot, as specified by a start date and time, and a stop date and time.
- A schedule can be active forever (recurrent), but only as specified by the daily schedule. The schedule on a specific day (time slot) takes priority over the daily schedule.
- A scheduler can be active during a time slot, as specified by the weekday schedule.
- A scheduler be active within two different time slots (daily or for a specified duration).

RELATED DOCUMENTATION

About the Firewall Policy Schedules Page | 437

Firewall Policy Examples | 396

Creating Schedules | 438

Editing, Cloning, and Deleting Schedules | 440

About the Firewall Policy Schedules Page

To access this page, select **Configuration** > **Firewall** > **Schedules**.

The **Firewall Policy Schedules** page enables you to create, modify, clone, and delete schedules. A schedule allows you to restrict access to a resource, or remove a restriction to a resource, for a specified period of time.

Tasks You Can Perform

You can perform the following tasks from this page:

- Create a firewall policy schedule. See "Creating Schedules" on page 438.
- Modify, clone, or delete a firewall policy schedule. See "Editing, Cloning, and Deleting Schedules" on page 440.
- View the configured parameters of a schedule. Click the details icon that appears when you hover over the name of an image or click **More > Detailed View**.
- Show or hide columns about the firewall policy schedule. Click the **Show Hide columns** icon in the top right corner of the page and select columns that you want to view on the page.
- Search for a specific firewall policy schedule. Click the Search icon in the top right corner of the page to search for a firewall policy schedule.

  You can enter partial text or full text of the keyword in the text box and press Enter.

  The search results are displayed on the same page.

Field Descriptions

Table 137 on page 437 provides guidelines on using the fields on the **Firewall Policy Schedules** page.

Table 137: Fields on the Firewall Policy Schedules Page

Field	Description
Name	Name of the schedule; maximum length is 63 characters.
Description	Description for the schedule; maximum length is 900 characters.
Start Date	The date and time from when the schedule comes into effect.
End Date	The date and time from when the schedule ends.
Second Start Date	The second date and time from when the schedule comes into effect.

Table 137: Fields on the Firewall Policy Schedules Page *(continued)*

Field	Description
Second End Date	The second date and time from when the schedule ends.

RELATED DOCUMENTATION

- Firewall Policy Schedules Overview | **436**
- Firewall Policy Examples | **396**
- Creating Schedules | **438**
- Editing, Cloning, and Deleting Schedules | **440**

Creating Schedules

Use the **Create Schedules** page to create schedules. A schedule allows you to restrict access to a resource, or remove a restriction to a resource, for a specified period of time.

To configure a schedule:

1. Select **Configuration > Firewall > Schedules**.

   The **Firewall Policy Schedules** page appears.

2. Click the add icon (+).

   The **Create Schedules** page appears.

3. Complete the configuration of the schedule according to the guidelines provided in Table 138 on page 439.

4. Click **OK** to save the changes. If you want to discard your changes, click **Cancel** instead.

A new schedule is created. You can use this schedule to activate firewall policies for the times and dates configured in your schedules.

Table 138 on page 439 provides guidelines on using the fields to create a schedule.

Table 138: Fields on the Create Schedules Page

Field	Description
General Information	
Name	Required. Enter a unique name for the service. It must begin with an alphanumeric character and cannot exceed 63 characters. Dashes and underscores are allowed.
Description	Enter a description for your service. You should make this description as useful as possible for all administrators.
Dates	
Date Range	Select **Ongoing** if you want your schedules to always be active. Select **Custom** to configure two sets of start and end dates for a single schedule. For the first set, enter dates in the **Start Date** and **End Date** fields. You must enter the days in MM/DD/YYYY format. For the second set of the schedule, enter the start date in the **Second Start Date** field and enter the end date in the **Second End Date** field.
Times	
Time Ranges	Create a schedule to be active daily or for any specific times of the day.
Daily Options	Select **Daily** to make the schedule applicable daily. Select **Custom** to enter specific days and times. Click on a specific day to specify time options for an entire day, to exclude a specific day, or to enter time ranges for the selected day. You must enter the time in HH:MM:SS format. For example, if you click on Monday, you get a dialog box that allows you to specify whether you want the schedule to be active all day Monday, exclude Monday from the schedule, or have the schedule be active at specific times. Select **Specify the same time for all days** to enter a date and time that is applicable for all days.

RELATED DOCUMENTATION

- Firewall Policy Schedules Overview | 436
- About the Firewall Policy Schedules Page | 437
- Firewall Policy Examples | 396

Editing, Cloning, and Deleting Schedules

IN THIS SECTION

- Editing Schedules | 440
- Cloning Schedules | 440
- Deleting Schedules | 441

You can edit, clone, and delete schedules from the **Firewall Policy Schedules** page.

Editing Schedules

To modify the parameters configured for a schedule:

1. Select **Configuration > Firewall > Schedules**.

   The **Firewall Policy Schedules** page appears.

2. Select the schedule that you want to edit, and then click on the edit icon (pencil symbol) on the right top corner of the table, or right-click and select **Edit Schedule**.

   The **Edit Schedules** page appears, showing the same options as when creating a new schedule.

3. Modify the parameters according to the guidelines provided in “Creating Schedules” on page 438.

4. Click **OK** to save the changes. If you want to discard your changes, click **Cancel** instead.

If you click **OK**, the modified schedule appears on the **Firewall Policy Schedules** page.

Cloning Schedules

To clone a schedule:

1. Select **Configuration > Firewall Policy > Schedules**.

   The **Firewall Policy Schedules** page appears.

2. Right-click on the schedule that you want to clone and then click **Clone**, or select **More > Clone**.

   The **Clone Schedules** page appears with editable fields. You can modify the parameters according to the guidelines provided in “Creating Schedules” on page 438.

3. Click **OK** to save the changes. If you want to discard your changes, click **Cancel** instead.

If you click **OK**, the cloned schedule appears under the scheduled it is cloned from, in the **Firewall Policy Schedules**.

Deleting Schedules

To delete a schedule:

1. Select **Configuration > Firewall Policy > Schedules**.

   The **Firewall Policy Schedules** page appears.

2. Select the schedule you want to delete and then click the delete icon **(X)** .

   An alert message appears, verifying that you want to delete the schedule.

3. Click **Yes** to delete the selection. If you do not want to delete, click **Cancel** instead.

If you click **Yes**, the selected schedule is deleted.

RELATED DOCUMENTATION

- Firewall Policy Schedules Overview | **436**
- About the Firewall Policy Schedules Page | **437**
- Creating Schedules | **438**
- Firewall Policy Examples | **396**

Deploying Firewall Policies

After adding the intents to the firewall policies, you can deploy the firewall policy by clicking the **Deploy** option that is above the **End Points** panel. You can also deploy one or more policies from the **Firewall Policy** page.

To deploy firewall policies:

1. Select **Configuration > Firewall > Firewall Policy**.

   The Firewall Policy page appears.

2. Select one or more policies and click **Deploy.**

   The Deploy page appears.

3. In **Choose Deployment Time** options, select **Run Now** to deploy the policy immediately. Select **Schedule at a later time** and specify the date and time at which the policy should be deployed.

4. Click **Deploy**.

   A job is created. Click the job ID to go to the Jobs page and view the status of the deploy operation.

   NOTE: During deployment, CSO ensures the order of the zone-based intents and enterprise-based intents within and across the policies.

About the Default Profiles for Unified Firewall Policy Page

To access this page, select **Configuration** > **Firewall** > **Default Settings**.

Use this page to view and edit the default settings for unified firewall policies. In a unified firewall policy, dynamic application is used as a match criteria and therefore a separate application firewall is not configured on a device (CPE or next-generation firewall) to allow or block traffic to an application.

The unified firewall takes some time to detect the application in a traffic and act upon it. The default profiles help in providing security during that time.

NOTE: The unified firewall policy settings are applied on a device only when Junos OS version 18.2R1 or later is installed on the device.

The default settings comprise the following:

- A UTM profile to define antispam, antivirus, content filtering and web filtering behavior.
- An SSL proxy profile to define the action to be taken when server certificates are not authenticated.

- An IPS profile to define the actions to be taken when the traffic matches the attack objects specified in the IPS profile.
- Reject Settings to define an action when the firewall blocks traffic for a particular application:
  - Take no action
  - Provide a redirect URL to redirect the traffic to another application or URL.
  - Provide a block message to display or log a message indicating that the traffic for the particular application is blocked by the firewall policy.

Tasks You Can Perform

You can perform the following tasks from this page:

- View the default unified firewall settings—See Table 139 on page 443 describes the fields on this page.
- Modify the default profiles for the unified firewall policy—See "Editing Default Settings for the Unified Firewall Policy" on page 444.

Field Descriptions

Table 139 on page 443 describes the fields on the Default Profiles for the Unified Firewall Policy page.

Table 139: Default Profiles for the Unified Firewall Policy Page

Setting	Guideline
Default UTM Policy	UTM profile assigned for the unified firewall policy, which is set the default UTM policy.
Default SSL Profle	SSL proxy profile assigned for the unified firewall policy, which is set as the default SSL proxy profile.
Default IPS Profile	IPS profile assigned for the unified firewall policy, which is set as the default IPS policy on the device.
Reject Settings	
Reject Action	The action assigned to the unified firewall policy when a firewall blocks application traffic: • **None**: No message or redirection is provided. • **Redirect URL**: The firewall redirects the traffic to the specified URL. • **Text**: The firewall displays or logs the message configured for this field.

RELATED DOCUMENTATION

About the Firewall Policy Name Page | 377

Editing Default Settings for the Unified Firewall Policy

Use the Default Profiles for Unified Firewall Policy page to configure the default profile, SSL proxy profile, IPS profile, and reject or redirect URL or message in the unified firewall policy for a tenant. If you enable a default SSL proxy profile for the tenant, CSO sets the default SSL proxy profile for the tenant as the the default SSL profile in the unified firewall policy.

The unified firewall takes some time to detect the application in a traffic and act upon it. The default profiles help in providing security during that time. The default settings are applicable to all the unified firewall policies belonging to a tenant and pushed to all those sites where a firewall policy is deployed.

To configure default settings for the unified firewall policy:

1. Select **Configuration > Firewall > Default Settings** in Customer Portal.

   The Default Profiles for Unified Firewall Policy Settings page appears.

2. Click the **Edit** button.

   The fields on the page can now be modified.

3. Complete the configuration according to the guidelines provided in Table 140 on page 445.

4. Do one of the following:

   - Click **Cancel** to cancel the changes.
   - Click **OK** to save the changes.

     The settings are saved and a confirmation message is displayed.

     You can deploy the changes by editing the unified firewall policy and then redeploying it.

Table 140: Default Profiles for the Unified Firewall Policy

Setting	Guideline
Default UTM Policy	Select a default UTM profile (policy) from the drop-down list. Alternatively, click the **Add UTM Profile** to add a UTM profile and use it as the default UTM profile. The Create UTM Profiles wizard appears. For information about creating an UTM policy, see "Creating UTM Profiles" on page 457.
Default SSL Profle	Select a default SSL proxy profile from the drop-down list. Alternatively, click **Add SSL Profile** to add a new SSL proxy profile and use it as the default SSL proxy profile. . The Create SSL Proxy Profiles page appears. For information about configuring SSL proxy profiles, See "Creating SSL Forward Proxy Profiles" on page 655.
Default IPS Profile	Select the IPS profile that you want to associate with the unified firewall policy as the default IPS profile.
Reject Settings	
Reject Action	When the action of the firewall is set to deny a particular application traffic, provide an alternative URL to redirect such traffic or a reason for blocking the traffic and an action that a user can perform. Select one of the following: • **None**: Do nothing when an application's traffic is blocked by the firewall. • **Redirect URL**: Redirect traffic to a specified URL when the firewall blocks the traffic. If you select this option, you must specify the URL to which traffic should be redirected (in the Redirect URL field). • **Text**: Block traffic and display a message. If you select this option, you must enter the message (in the Block Message field) to be displayed or logged when the firewall blocks the traffic.
Redirect URL	If you chose **Redirect URL** for Reject Action, enter the URL to which an application traffic must be redirected.
Text	If you choose **Text** for Reject Action, enter the reason for blocking the traffic and what a user can do subsequently. You can enter a maximum of 256 alphanumeric characters including spaces.

RELATED DOCUMENTATION

UTM Overview | 451

SSL Forward Proxy Overview | 636

Importing Policies Overview

CSO supports importing policy configurations from next-generation firewall devices. You can discover existing policy configuration while onboarding next-generation firewall device (without enabling ZTP) or import policy configurations from Firewall and NAT policy pages (after ZTP). For more information about overview and configuration of ZTP on SRX Series devices, see Zero Touch Provisioning on SRX Series Devices.

- To import policy configuration after ZTP , see “Importing Firewall Policies” on page 448, and “Importing NAT Policies” on page 588.
- To discover existing policy configuration while onboarding next-generation firewall device (without enabling ZTP), see “Adding a Standalone Next Generation Firewall Site” on page 177, and “Adding an On-Premise Spoke Site with Next Generation Firewall and LAN Capabilities” on page 151.

CSO uses object name as the unique identifier for an object (such as addresses, services, schedulers, SSL profiles, unified threat management (UTM), and Layer 7 applications). During policy import, all objects that are supported by CSO are imported and all objects names are compared between what is in CSO and what is on the next generation firewall device. A conflict occurs when the name of the object to be imported matches an existing object, but the value of the object does not match. The object conflict resolution (OCR) operation is triggered to resolve the object name conflicts.

- If the object name does not exist in CSO, the object is added to CSO.
- If the object name exists in CSO with the same content, the existing object in CSO is used.
- If the object name exists in CSO with different content, the object conflict resolution operation is triggered, providing users with the following conflict resolution options:
  - Rename object
    - This is the default option.
    - By default, "_1" is added to the object name, or users can specify a new unique name.
    - Deploying the policy will delete the original object and add the object with the new name.
    - There is no functional change to the firewall policy (labels only).
  - Overwrite with imported value
    - The object in CSO is replaced with the object from the import operation.
    - The change will be reflected for all other devices that use this object after the policy deployment.

- There is no functional change to the firewall policy.
- There may be possible traffic impact to all other devices that use this object the next time the other device is updated from CSO.

- Keep existing object
  - The object name in CSO is used instead of what is on the next generation firewall device.
  - Policy deployment for the imported firewall policy will show the modification.
  - There may be possible traffic impact to this firewall because the content is different in some way.

The following section provides an example for importing policies. Here we use Address as an object type and see how to resolve the object name conflicts.

The existing objects in CSO are listed inTable 141 on page 447.

Table 141: Existing address in CSO

Object Name	Existing Value
Address1	198.51.100.10
Address2	198.51.100.20
Address3	198.51.100.30

The existing objects in the next generation firewall device are listed inTable 142 on page 447.

Table 142: Existing address in next-generation firewall device

Object Name	Existing Value
Address1	203.0.113.10/32
Address2	203.0.113.20/32
Address3	203.0.113.30/32

During policy import, OCR is triggered and the object conflicts between next generation firewall device and CSO. The resolution that we have chosen is listed in Table 143 on page 447.

Table 143: OCR while importing policies to CSO

Object Name in CSO	Object Type in CSO	Existing Value in CSO	Imported Value to CSO	Conflict Resolution	New Object Name in CSO
Address1	Address	198.51.100.10	203.0.113.10	Keep Existing Object	Address1_1

Table 143: OCR while importing policies to CSO *(continued)*

Object Name in CSO	Object Type in CSO	Existing Value in CSO	Imported Value to CSO	Conflict Resolution	New Object Name in CSO
Address2	Address	198.51.100.20	203.0.113.20	Overwrite with Imported value	Address2_1
Address3	Address	198.51.100.30	203.0.113.30	Rename Object	Address3_1

The object values and the result after resolving conflicts are listed in Table 144 on page 448.

Table 144: After importing policies to CSO

Discovered Object Name in CSO	Discovered Value in CSO	Result
Address1	198.51.100.10	No change
Address2	203.0.113.20	Content changed
Address3	198.51.100.30	No change
Address3_1	203.0.113.30	Address3_1 created

RELATED DOCUMENTATION

Contrail Service Orchestration (CSO) Deployment Guide

Importing Firewall Policies

Use this page to manually import a firewall policy from the discovered or onboarded sites (next generation firewall sites).

To import a firewall policy:

1. Select **Configuration > Firewall > Firewall Policy**.

   The Firewall Policy page appears.

2. Click **Import**.

   The Import Firewall Policies page appears displaying a list of discovered devices (next generation firewall devices).

3. Select the devices from which you want to import the firewall policies and click **Next**.

   The Discovered Services tab appears.

4. Select the policies that you want to import and click **Next**.

   The Resolve Conflicts tab appears.

5. If there are any conflicts with the imported objects, object conflict resolution(OCR) operation is triggered. The Conflicts window displays all the conflicts between CSO and the next generation firewall device. Select an object from the Conflicts window and click on any of the below option to resolve the object conflict.

   The resolution options are:

   - Rename Object—Rename the imported object. By default, "_1" is added to the object name, or you can specify a new name.
   - Overwrite with imported value—The object in CSO is replaced with the object from the import operation.
   - Keep existing object—The object name in CSO is used instead of what is on the next generation firewall device.

6. Click **Finish**.

   A summary of the discovered services is listed.

7. Review the summary and click **OK** to import the firewall policies.

   The import policy job is created and the firewall policies are imported from next generation firewall device to CSO. You can view the imported policy from the Firewall Policy page.

WHAT'S NEXT

After importing the firewall policy successfully, you can edit and deploy the policy. See Editing and Deleting Firewall Policies | 380, Editing, Cloning, and Deleting Firewall Policy Intents | 387, and Deploying Firewall Policies | 441.

RELATED DOCUMENTATION

Importing Policies Overview | 446

CHAPTER 16

Managing UTM Profiles

IN THIS CHAPTER

- UTM Overview | **451**
- Configuring UTM Settings | **453**
- About the UTM Profiles Page | **455**
- Creating UTM Profiles | **457**
- Editing, Cloning, and Deleting UTM Profiles | **460**
- About the Web Filtering Profiles Page | **462**
- Creating Web Filtering Profiles | **464**
- Editing, Cloning, and Deleting Web Filtering Profiles | **468**
- About the Antivirus Profiles Page | **470**
- Creating Antivirus Profiles | **472**
- Editing, Cloning, and Deleting Antivirus Profiles | **475**
- About the Antispam Profiles Page | **477**
- Creating Antispam Profiles | **478**
- Editing, Cloning, and Deleting Antispam Profiles | **480**
- About the Content Filtering Profiles Page | **482**
- Creating Content Filtering Profiles | **484**
- Editing, Cloning, and Deleting Content Filtering Profiles | **488**
- About the URL Patterns Page | **490**
- Creating URL Patterns | **491**
- Editing, Cloning, and Deleting URL Patterns | **492**
- About the URL Categories Page | **494**
- Creating URL Categories | **495**
- Editing, Cloning, and Deleting URL Categories | **497**

UTM Overview

IN THIS SECTION

- UTM Licensing | 452
- UTM Components | 452

Unified threat management (UTM) is a term used to describe the consolidation of several security features to protect against multiple threat types. The advantage of UTM is a streamlined installation and management of multiple security capabilities.

The following security features are provided as part of the UTM solution:

- Antispam—This feature examines transmitted messages to identify e-mail spam. E-mail spam consists of unwanted messages usually sent by commercial, malicious, or fraudulent entities. When the device detects an e-mail message deemed to be spam, it either drops the message or tags the message header or subject field with a preprogrammed string. The antispam feature uses a constantly updated Spamhaus Block List (SBL). Sophos updates and maintains the IP-based SBL.
- Full file-based antivirus—A virus is an executable code that infects or attaches itself to other executable code to reproduce itself. Some malicious viruses erase files or lock up systems. Other viruses merely infect files and overwhelm the target host or network with bogus data. The full file-based antivirus feature provides file-based scanning on specific application layer traffic, checking for viruses against a virus signature database. The antivirus feature collects the received data packets until it has reconstructed the original application content, such as an e-mail file attachment, and then scans this content.
- Express antivirus—Express antivirus scanning is offered as a less CPU-intensive alternative to the full file-based antivirus feature. The express antivirus feature is similar to the antivirus feature in that it scans specific application layer traffic for viruses against a virus signature database. However, unlike full antivirus, express antivirus does not reconstruct the original application content. Rather, it just sends (streams) the received data packets, as is, to the scan engine. With express antivirus, the virus scanning is executed by a hardware pattern-matching engine. This improves performance while scanning is occurring, but the level of security provided is lessened. Juniper Networks provides the scan engine.
- Content filtering—Content filtering blocks or permits certain types of traffic based on the MIME type, file extension, protocol command, and embedded object type.
- Web filtering—Web filtering enables you to manage Internet usage by preventing access to inappropriate Web content. The following types of Web filtering solutions are available:

- Integrated Web filtering—Blocks or permits Web access after the device identifies the category for a URL either from user-defined categories or from a category server (Websense provides the SurfControl Content Portal Authority (CPA) server).
- Redirect Web filtering—Intercepts HTTP requests and forwards the server URL to an external URL filtering server to determine whether to block or permit the requested Web access. Websense provides the URL filtering server.

UTM Licensing

All UTM components require licenses with the exception of content filtering, which uses the parameters defined in the content filtering profile. This is because Juniper Networks leverages third-party technology that is constantly updated to provide the most up-to-date inspection capabilities.

UTM Components

UTM components include custom objects, feature profiles, and UTM profiles that can be configured on SRX Series devices. From a high level, feature profiles specify how a feature is configured and then applied to UTM profiles, which in turn is applied to firewall policies, as shown in Figure 12 on page 452.

Figure 12: UTM Components

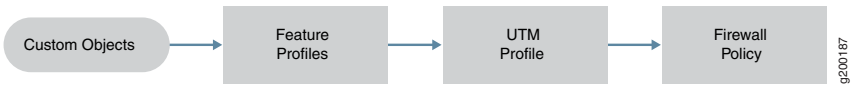

UTM profiles do not have their own seven-tuple rulebase; in a sense they inherit the rules from the firewall rule. The strength of the UTM feature comes from URL filtering, where you can have a separate configuration for different users or user groups.

- Custom objects—Although SRX Series devices support predefined feature profiles that can handle most typical use cases, there are some cases where you might need to define your own objects, specifically for URL filtering, antivirus filtering, and content filtering.
- Feature profiles—Feature profiles specify how components of each profile should function. You can configure multiple feature profiles that can be applied through different UTM profiles to firewall rules.
- UTM profiles—UTM profiles function as a logical container for individual feature profiles. UTM profiles are then applied to specific traffic flows based on the classification of rules in the firewall policy, thereby enabling you to define separate UTM profiles per firewall rule to differentiate the enforcement per

firewall rule. Essentially, the firewall rulebase acts as the match criteria, and the UTM profile is the action to be applied.

- Firewall policy—You can predefine feature profiles for the UTM profile that are then applied to the firewall rules. This gives you the advantage of using the predefined UTM profile for that one UTM technology (for example, antivirus or URL filtering), not both.

RELATED DOCUMENTATION

Configuring UTM Settings | 453

About the UTM Profiles Page | 455

Creating UTM Profiles | 457

Configuring UTM Settings

Use the Edit UTM Settings page to configure unified threat management (UTM) antispam, antivirus, and Web filtering settings for a tenant.

These settings are applicable to all the sites belonging to a tenant. The settings are pushed to all those sites where a firewall policy intent with UTM enabled is applicable.

To configure UTM settings:

1. Select **Configuration > Unified Threat Mgmt > UTM Settings** in Customer Portal.

   The Edit UTM Settings page appears.

2. Complete the configuration according to the guidelines provided in Table 145 on page 453.

3. Do one of the following:

   - Click **Reset** to reset the settings to the previously saved configured.
   - Click **OK** to save the settings.

     The settings are saved and a confirmation message is displayed.

Table 145: UTM Settings

Setting	Guideline
Antispam Settings	

Table 145: UTM Settings *(continued)*

Setting	Guideline
Address Whitelist	Select the URL pattern to be used as the antispam allowlist. Alternatively, click **Create a New Pattern** to create a new URL pattern to use as a allowlist. The Create URL Patterns page appears. For more information, see "Creating URL Patterns" on page 491 for an explanation of the fields on this page.
Address Blacklist	Select the URL pattern to be used as the antispam blocklist. Alternatively, click **Create a New Pattern** to create a new URL pattern to use as a blocklist.
Antivirus Settings	
MIME Whitelist	Enter one or more MIME types (separated by commas) to exclude from antivirus scanning.
Exception MIME Whitelist	Enter one or more MIME types (separated by commas) that are to be excluded from the list of MIME types specified as part of the MIME allowlist. This list is a subset of the MIME types that you specified in the MIME allowlist. For example, if you specify **video/** in the allowlist and **video/x-shockwave-flash** in the exception allowlist, all objects of MIME type **video/** except MIME type **video/x-shockwave-flash** are excluded from antivirus scanning.
URL Whitelist	Select the list of URLs that the antivirus settings can allow.
Web Filtering Settings	
URL Whitelist	Select the list of URLs that the Web filtering can allow; these URLs are excluded from Web filtering.
URL Blacklist	Select the list of URLs the Web filtering settings can block; these URLs are blocked from Web access.

RELATED DOCUMENTATION

About the UTM Profiles Page | 455

About the UTM Profiles Page

To access this page, select **Configuration > Unified Threat Mgmt > UTM Profiles** in Customer Portal.

Use this page to view and manage unified threat management (UTM) profiles. UTM profiles enable you to consolidate several security features into one system to protect against multiple threat types.

Tasks You Can Perform

You can perform the following tasks from this page:

- Create a UTM profile—See "Creating UTM Profiles" on page 457.
- Edit, clone, or delete a UTM profile—See "Editing, Cloning, and Deleting UTM Profiles" on page 460.
- Clear the selected UTM profiles—Click **Clear All Selections** to clear any UTM profiles that you might have selected.
- View the details of a UTM profile—Select the UTM profile for which you want to view the details and from the More or right-click menu, select **Detailed View**. The UTM Profile Details page appears. Table 147 on page 456 describes the fields on this page.
- Search for UTM profiles using keywords—Click the search icon and enter the search term in the text box and press Enter. The search results are displayed on the same page.

Field Descriptions

Table 146 on page 455 describes the fields on the UTM Profiles page.

Table 146: UTM Profiles Page Fields

Field	Description
Name	Name of the UTM profile.
Antispam	Information about the antispam profile associated with the UTM profile.
Antivirus	Information about the antivirus profiles associated with the UTM profile.
Content Filtering	Information about the content filtering profiles associated with the UTM profile.
Web Filtering	Information about the Web filtering profile associated with the UTM profile.
Description	Description of the UTM profile.

Table 147: UTM Profile Details Page Fields

Field	Description
General Information	
Name	Name of the UTM profile.
Description	Description of the UTM profile.
Traffic Options	
Action When Connection Limit Is Reached	Action to be taken when the configured connection limit per client is reached.
Web Filtering Profile	
HTTP	Web filtering profile to be used for HTTP traffic.
Antivirus Profile	
HTTP	Antivirus profile to be used for HTTP traffic.
FTP Upload	Antivirus profile to be used for FTP upload traffic.
FTP Download	Antivirus profile to be used for FTP download traffic.
IMAP	Antivirus profile to be used for IMAP traffic.
SMTP	Antivirus profile to be used for SMTP traffic.
POP3	Antivirus profile to be used for POP3 traffic.
Antispam Profile	
SMTP	Antispam profile to be used for SMTP traffic.

RELATED DOCUMENTATION

Creating UTM Profiles | 457

Creating UTM Profiles

Use the Create UTM Profiles page to configure UTM profiles. Unified threat management (UTM) consolidates several security features to protect against multiple threat types. The Create UTM Profiles wizard provides step-by-step procedures to create a UTM profile. You can configure antispam, antivirus, Web filtering, and content filtering profiles by launching the respective wizards from the wizard.

To create a UTM profile:

1. Select **Configuration > Unified Threat Mgmt > UTM Profiles** in Customer Portal.

   The UTM Profiles page appears.

2. Click the add icon (+) to create a new UTM profile.

   The Create UTM Profiles wizard appears, displaying brief instructions about creating a UTM profile.

3. Click **Next** to navigate to the next page.

4. Complete the configuration according to the guidelines provided in Table 148 on page 457.

   NOTE: Fields marked with * are mandatory.

5. Click **Finish**.

   A UTM profile is created. You are returned to the UTM Profiles page where a confirmation message is displayed. After you create a UTM profile, you can assign it to a firewall policy intent on the Firewall Policy page.

Table 148: UTM Profile Settings

Setting	Guideline
General	
Name	Enter a unique name for the UTM profile. The maximum length is 29 characters.
Description	Enter a description for the UTM profile. The maximum length is 255 characters.
Traffic Options **NOTE:** In an attempt to consume all available resources, a malicious user might generate a large amount of traffic all at once. To prevent such activity from succeeding, you can impose traffic options.	

Table 148: UTM Profile Settings *(continued)*

Setting	Guideline
Connection Limit per Client	Specify the connection limit per client for client connections on the device. The default is 2000 and a value of 0 means that there is no connection limit.
Action when connection limit is reached	Specify the action that must be taken when the connection limit is reached. The available actions are No action (default), Log and permit, and Block. Click **Next** to continue.
Web Filtering	
HTTP	Select the Web filtering profile to be applied for HTTP traffic. Alternatively, click **Create Another Profile** to create a Web filtering profile. The Create Web Filtering Profiles wizard appears. See "Creating Web Filtering Profiles" on page 464 for an explanation of the fields on this wizard. Click **Back** to go the preceding step or click **Next** to go to the next step.
Antivirus	
Apply to all protocols	Select this check box to apply a single antivirus profile to all traffic protocols. and then specify the profile in the **Default Profile** field. Clear the check box if you want to apply traffic-specific profiles.
Default Profile	Select the antivirus profile to be applied to all traffic protocols. Click **Back** to go the preceding step or click **Next** to go to the next step.
NOTE: Click **Create Another Profile** to create an antivirus profile that you can then assign. The Create Antivirus Profiles wizard appears. See "Creating Antivirus Profiles" on page 472 for an explanation of the fields on this wizard.	
HTTP	Select the antivirus profile to be applied to HTTP traffic.
FTP Upload	Select the antivirus profile to be applied to FTP upload traffic.
FTP Download	Select the antivirus profile to be applied to FTP download traffic.
IMAP	Select the antivirus profile to be applied to IMAP traffic.
SMTP	Select the antivirus profile to be applied to SMTP traffic.

Table 148: UTM Profile Settings *(continued)*

Setting	Guideline
POP3	Select the antivirus profile to be applied to POP3 traffic. Click **Back** to go the preceding step or click **Next** to go to the next step.
Antispam	
SMTP	Select the antispam profile to be applied for SMTP traffic. Alternatively, click **Create Another Profile** to create an antispam profile. The Create Antispam Profiles wizard appears. See "Creating Antispam Profiles" on page 478 for an explanation of the fields on this wizard. Click **Back** to go the preceding step or click **Next** to go to the next step.
Content Filtering	
Apply to all protocols	Select this check box to apply a single content filtering profile to all traffic protocols. and then specify the profile in the **Default Profile** field. Clear the check box if you want to apply traffic-specific profiles.
Default Profile	Select the content filtering profile to be applied to all traffic protocols. Click **Back** to go the preceding step or click **Next** to go to the next step.
NOTE: Click **Create Another Profile** to create a content filtering profile that you can then assign. The Create Content Filtering Profiles wizard appears. See "Creating Content Filtering Profiles" on page 484 for an explanation of the fields on this wizard.	
HTTP	Select the content filtering profile to be applied to HTTP traffic.
FTP Upload	Select the content filtering profile to be applied to FTP upload traffic.
FTP Download	Select the content filtering profile to be applied to FTP download traffic.
IMAP	Select the content filtering profile to be applied to IMAP traffic.
SMTP	Select the content filtering profile to be applied to SMTP traffic.
POP3	Select the content filtering profile to be applied to POP3 traffic. Click **Back** to go the preceding step.

RELATED DOCUMENTATION

About the UTM Profiles Page | 455
Configuring UTM Settings | 453

Editing, Cloning, and Deleting UTM Profiles

IN THIS SECTION

Editing UTM Profiles | 460
Cloning UTM Profiles | 461
Deleting UTM Profiles | 461

You can edit, clone, and delete UTM profiles from the UTM Profiles page. This topic has the following sections:

Editing UTM Profiles

To modify the parameters configured for a UTM profile:

> **NOTE:** You cannot modify the default profiles already present in the system.

1. Select **Configuration > Unified Threat Mgmt > UTM Profiles** in Customer Portal.

   The UTM Profiles page appears, displaying the existing UTM profiles.

2. Select the UTM profile that you want to edit and click the edit icon (pencil). Alternatively, right-click a profile and select **Edit Profile**.

   The Edit UTM Profiles page appears, displaying the same fields that are presented when you create a UTM profile.

3. Modify the UTM profile fields as needed.

4. Click **OK** to save your changes.

You are taken to the UTM Profiles page. A confirmation message appears indicating the status of the edit operation.

Cloning UTM Profiles

Cloning enables you to easily create a new UTM profile based on an existing one.

To clone a UTM profile:

1. Select **Configuration > Unified Threat Mgmt > UTM Profiles** in Customer Portal.

   The UTM Profiles page appears, displaying the existing UTM profiles.

2. Select the UTM profile that you want to clone and then select **More > Clone**. Alternatively, right-click a profile and select **Clone**.

   The Clone UTM Profiles page appears, displaying the same fields that are presented when you create a UTM profile.

3. Modify the UTM profile fields as needed.

4. Click **OK** to save your changes.

   You are taken to the UTM Profiles page. A confirmation message appears, indicating the status of the clone operation.

Deleting UTM Profiles

> **NOTE:** Before deleting a UTM profile, ensure that the profile is not used in a firewall policy intent. If you try to delete a profile that is used in a firewall policy intent, an error message is displayed.

To delete one or more UTM profiles:

1. Select **Configuration > Unified Threat Mgmt > UTM Profiles** in Customer Portal.

   The UTM Profiles page appears, displaying the existing UTM profiles.

2. Select one or more UTM profiles that you want to delete and click the delete icon (**X**). Alternatively, right-click a profile and select **Delete Profile**.

   An alert message appears, asking you to confirm the delete operation.

3. Click **Yes** to delete the selected UTM profiles.

   A confirmation message appears, indicating the status of the delete operation.

RELATED DOCUMENTATION

Creating UTM Profiles | 457

About the UTM Profiles Page | 455

About the Web Filtering Profiles Page

To access this page, select **Configuration > Unified Threat Mgmt > Web Filtering Profiles** in Customer Portal.

Use the Web Filtering Profiles page to view and manage Web filtering profiles. Web filtering profiles enable you to manage Internet usage by preventing access to inappropriate Web content over HTTP. Table 149 on page 462 lists the Web filtering solutions that are supported and the license requirements.

Table 149: Web Filtering Solutions Supported

Type	Description	License Requirement
Integrated Web Filtering	Blocks or permits Web access after the device identifies the category for a URL, either from user-defined categories or from a category server (SurfControl Content Portal Authority provided by Websense).	A separately licensed subscription service
Redirect Web Filtering	Intercepts HTTP requests and forwards the server URL to an external URL filtering server to determine whether to block or permit the requested Web access. Websense provides the URL filtering server.	Does not require a license.
Juniper Local Web Filtering	Intercepts every HTTP request in a TCP connection. In this case, the decision making is done on the device after it looks up a URL to determine whether it is in the allowlist or blocklist based on its user-defined category.	Does not require a license or a remote category server

Tasks You Can Perform

You can perform the following tasks from this page:

- Create a Web filtering profile—See "Creating Web Filtering Profiles" on page 464.
- Edit, clone, or delete a Web filtering profile—See "Editing, Cloning, and Deleting Web Filtering Profiles" on page 468.
- Clear the selected Web filtering profiles—Click **Clear All Selections** to clear any Web filtering profiles that you might have selected.
- View the details of a Web filtering profile—Select the Web filtering profile for which you want to view the details and from the More or right-click menu, select **Detailed View**. The Web Filtering Profile Details page appears. Table 151 on page 463 describes the fields on this page.
- Search for Web filtering profiles using keywords—Click the search icon and enter the search term in the text box and press Enter. The search results are displayed on the same page.

Field Descriptions

Table 150 on page 463 describes the fields on the Web Filtering Profiles page.

Table 150: Web Filtering Profiles Page Fields

Field	Description
Name	Name of the Web filtering profile.
Profile Type	Type of engine used for the profile: Juniper-enhanced or Websense redirect.
Default Action	Default action taken when the specified connection limit per client is reached.
Timeout	
Description	Description of the Web filtering profile.

Table 151: Web Filtering Profile Details Page Fields

Field	Description
General Information	
Name	Name of the Web filtering profile.
Description	Description of the Web filtering profile.
Engine Type	Type of engine used for the profile: Juniper-enhanced or Websense redirect.
Default Action	Default action taken when the specified connection limit per client is reached.

Table 151: Web Filtering Profile Details Page Fields *(continued)*

Field	Description
Fallback Options	
Default Action	Action taken for URL categories with no assigned action and for uncategorized URLs. This action is taken only if no reputation action is assigned.
Global Reputation Actions	Actions taken for the following site reputations: • **Very Safe** • **Moderately Safe** • **Fairly Safe** • **Suspicious** • **Harmful**
URL Categories	URL categories associated with the Web filtering profile.

RELATED DOCUMENTATION

Creating Web Filtering Profiles | 464
Editing, Cloning, and Deleting Web Filtering Profiles | 468

Creating Web Filtering Profiles

Web filtering profiles enable you to manage Internet usage by preventing access to inappropriate Web content over HTTP.

To create a Web filtering profile:

1. Select **Configuration > Unified Threat Mgmt > Web Filtering Profiles** in Customer Portal.

   The Web Filtering Profiles page appears.

2. Click the add icon (+) to create a new Web filtering profile.

   The Create Web Filtering Profiles wizard appears, displaying brief instructions about creating a Web filtering profile.

3. Click **Next** to navigate to the next page.

4. Complete the configuration according to the guidelines provided in Table 152 on page 465.

> **NOTE:** Fields marked with * are mandatory.

5. Click **Finish**.

   A Web filtering profile is created, which you can associate with a UTM profile. You are returned to the Web Filtering Profiles page where a confirmation message is displayed.

Table 152: Creating Web Filtering Profiles Settings

Setting	Guideline
General Information	
Name	Enter a unique name for the Web filtering profile. The maximum length is 29 characters.
Description	Enter a description for the Web filtering profile. The maximum length is 255 characters.
Timeout	Enter a timeout (in seconds) to wait for a response from the Websense server. The default is 15 seconds and the maximum is 1000 seconds.
Engine Type	Select an engine type for Web filtering: • (Default) Juniper Enhanced—UTM-enhanced Web filtering. • Websense Redirect—Redirect Web filtering profile.
Safe Search	Select the check box (default) to ensure that embedded objects, such as images on the URLs received from the search engines, are safe and that undesirable content is not returned to the client. Clear the check box to disable safe search redirects. **NOTE:** This option is available only for the Juniper Enhanced engine type. Save search redirect supports only HTTP and you cannot extract the URL for HTTPS. Therefore, it is not possible to generate a redirect response for HTTPS search URLs.
Custom Block Message/URL	Specify the redirect URL or a custom message to be sent when HTTP requests are blocked. The maximum length is 512 characters. **NOTE:** If a message begins with http: or https:, the message is considered a block message URL. Messages that begin with values other than http: or https: are considered custom block messages. Click **Back** to go the preceding step or click **Next** to go to the next step.

Table 152: Creating Web Filtering Profiles Settings *(continued)*

Setting	Guideline
Custom Quarantine Message	Define a custom message to allow or deny access to a blocked site based on a user's response to the message. The maximum length is 512 characters. The quarantine message contains the following information: • URL name • Quarantine name • Category (if available) • Site reputation (if available) For example, if you set the action for Enhanced_Search_Engines_and_Portals to quarantine, and you try to access www.search.yahoo.com, the quarantine message is as follows: ***The requested webpage is blocked by your organization's access policy***. Click **Back** to go the preceding step or click **Next** to go to the next step.
Account	Specify the user account associated with the Websense Web filtering profile.
Server	Specify the hostname or IP address for the Websense server.
Port	Specify the port number to use to communicate with the Websense server. The default port value is 15968. Click **Back** to go the preceding step or click **Next** to go to the next step.
Sockets	Enter the number of sockets used for communication between the client and the server. The default value is 8.
URL Categories	
Deny Action List	Click the **Add URL Categories** button to specify a list of URL categories that should be denied access. The Select URL Categories page appears. Complete the configuration according to the guidelines provided in Table 153 on page 468. The list of URL categories selected is displayed in a text box.
Log & Permit Action List	Specify a list of URL categories that are logged and then permitted. The Select URL Categories page appears. Complete the configuration according to the guidelines provided in Table 153 on page 468. The list of URL categories selected is displayed in a text box.

Table 152: Creating Web Filtering Profiles Settings *(continued)*

Setting	Guideline
Permit Action List	Specify a list of URL categories that should be permitted access. The Select URL Categories page appears. Complete the configuration according to the guidelines provided in Table 153 on page 468 The list of URL categories selected is displayed in a text box.
Quarantine Action List	Specify a list of URL categories that should be quarantined. The Select URL Categories page appears. Complete the configuration according to the guidelines provided in Table 153 on page 468. The list of URL categories selected is displayed in a text box. Click **Back** to go the preceding step or click **Next** to go to the next step.
Fallback Options	
Global Reputation Actions	Select this check box (default) if you want to apply global reputation actions. Enhanced Web filtering intercepts HTTP and HTTPS requests and sends the HTTP URL or the HTTPS source IP to the Websense ThreatSeeker Cloud (TSC). The TSC categorizes the URL into one of the predefined categories and also provides site reputation information for the URL to the device. The device determines if it can permit or block the request based on the information provided by the TSC. The URLs can be processed using their reputation score if there is no category available. Select the action that you want to take for the uncategorized URLs based on their reputation score: • **Very Safe**—Permit, log and permit, block, or quarantine a request if a site reputation of 90 through 100 is returned. By default, Permit is selected. • **Moderately Safe**—Permit, log and permit, block, or quarantine a request if a site reputation of 80 through 89 is returned. By default, Log and Permit is selected. • **Fairly Safe**—Permit, log and permit, block or quarantine a request if a site-reputation of 70 through 79 is returned. By default, Log and Permit is selected. • **Suspicious**—Permit, log and permit, block, or quarantine a request if a site reputation of 60 through 69 is returned. By default, Quarantine is selected. • **Harmful**—Permit, log and permit, block, or quarantine a request if a site reputation of zero through 59 is returned. By default, Block is selected.
Default Action	Choose the actions to be taken for URL categories with no assigned action and for uncategorized URLs. This is used only if no reputation action is assigned.

Table 152: Creating Web Filtering Profiles Settings *(continued)*

Setting	Guideline
Fallback Action	Select the fallback action, which is used when: • The ThreatSeeker Websense Cloud servers are unreachable. • A timeout occurs for requests to ThreatSeeker Cloud. • There are too many requests to be handled by the device.

Table 153: Select URL Categories Settings

Setting	Guideline
Show	Choose which URL categories should be displayed for selection: **All categories**, **Custom URL categories**, or **Websense URL categories**. The **Available** column of the **URL Categories** field displays URL categories based on your selection.
URL Categories	Select one or more URL categories in the **Available** column and click the forward arrow to confirm your selection. The selected URL categories are displayed in the **Selected** column. Alternatively, click **Create New URL Category** to create a URL category and assign it to the URL category. The Create URL Categories page appears; for more information, see "Creating URL Categories" on page 495. Click **OK** to confirm your selection. You are returned to the Create Web Filtering Profiles page.

RELATED DOCUMENTATION

Creating UTM Profiles | 457

Editing, Cloning, and Deleting Web Filtering Profiles

IN THIS SECTION

- Editing Web Filtering Profiles | 469
- Cloning Web Filtering Profiles | 469
- Deleting Web Filtering Profiles | 470

You can edit, clone, and delete Web filtering profiles from the Web Filtering Profiles page. This topic has the following sections:

Editing Web Filtering Profiles

To modify the parameters configured for a Web filtering profile:

> **NOTE:** You cannot modify the default profiles already present in the system.

1. Select **Configuration > Unified Threat Mgmt > Web Filtering Profiles** in Customer Portal.

   The Web Filtering Profiles page appears, displaying the existing Web filtering profiles.

2. Select the Web filtering profile that you want to edit and click the edit icon (pencil). Alternatively, right-click a profile and select **Edit Profile**.

   The Edit Web Filtering Profiles page appears, displaying the same fields that are presented when you create a Web filtering profile.

3. Modify the Web filtering profile fields as needed.

4. Click **OK** to save your changes.

   You are taken to the Web Filtering Profiles page. A confirmation message appears, indicating the status of the edit operation.

Cloning Web Filtering Profiles

Cloning enables you to easily create a new Web filtering profile based on an existing one.

To clone a Web filtering profile:

1. Select **Configuration > Unified Threat Mgmt > Web Filtering Profiles** in Customer Portal.

   The Web Filtering Profiles page appears, displaying the existing Web filtering profiles.

2. Select the Web filtering profile that you want to clone and then select **More > Clone**. Alternatively, right-click a profile and select **Clone**.

   The Clone Web Filtering Profiles page appears, displaying the same fields that are presented when you create a Web filtering profile.

3. Modify the Web filtering profile fields as needed.

4. Click **OK** to save your changes.

   You are taken to the Web Filtering Profiles page. A confirmation message appears, indicating the status of the clone operation.

Deleting Web Filtering Profiles

Before deleting a Web filtering profile, ensure that the profile is not used in a UTM profile that is, in turn, used in a firewall policy intent. If you try to delete a Web filtering profile that is used in a firewall policy intent, an error message is displayed.

To delete one or more Web filtering profiles:

1. Select **Configuration > Unified Threat Mgmt > Web Filtering Profiles** in Customer Portal.

   The Web Filtering Profiles page appears, displaying the existing Web filtering profiles.

2. Select one or more Web filtering profiles that you want to delete and click the delete icon (**X**). Alternatively, right-click a profile and select **Delete Profile**.

   An alert message appears, asking you to confirm the delete operation.

3. Click **Yes** to delete the selected Web filtering profiles.

   A confirmation message appears, indicating the status of the delete operation.

RELATED DOCUMENTATION

Creating Web Filtering Profiles | 464
About the Web Filtering Profiles Page | 462

About the Antivirus Profiles Page

To access this page, select **Configuration > Unified Threat Mgmt > Antivirus Profiles** in Customer Portal.

Use the Antivirus Profiles page to view and manage antivirus profiles. Antivirus profiles enable you to inspect files transmitted over several protocols (HTTP, FTP upload and download, IMAP, SMTP, and POP3) to determine whether the files exchanged are known malicious files, similar to how desktop antivirus software scans files for the same purpose.

Tasks You Can Perform

You can perform the following tasks from this page:

- Create an antivirus profile—See "Creating Antivirus Profiles" on page 472.
- Edit, clone, or delete an antivirus profile—See "Editing, Cloning, and Deleting Antivirus Profiles" on page 475.
- Clear the selected antivirus profiles—Click **Clear All Selections** to clear any antivirus profiles that you might have selected.
- View the details of an antivirus profile—Select the antivirus profile for which you want to view the details and from the More or right-click menu, select **Detailed View**. The Antivirus Profile Details page appears. Table 155 on page 471 describes the fields on this page.
- Search for antivirus profiles by using keywords—Click the search icon and enter the search term in the text box and press Enter. The search results are displayed on the same page.

Field Descriptions

Table 154 on page 471 describes the fields on the Antivirus Profiles page.

Table 154: Antivirus Profiles Page Fields

Field	Description
Name	Name of the antivirus profile.
Profile Type	Type of engine used for the profile.
Content Size Limit	Content size limit, in kilobytes, refers to accumulated TCP payload size.
Trickling Timeout	Number of seconds to wait for a response from the server.
Description	Description of the antivirus profile.

Table 155: Antivirus Profiles Details Page Fields

Field	Description
General Information	
Name	Name of the antivirus profile.
Description	Description of the antivirus profile.

Table 155: Antivirus Profiles Details Page Fields *(continued)*

Field	Description
Engine Type	Type of engine used for the profile.
Scan Options	
Content Size Limit	Content size limit, in kilobytes, refers to accumulated TCP payload size.
Fallback Options	
Default Action	Displays the default fallback action taken when the antivirus system encounters errors.
Content Size	Displays the actions taken if the content size exceeds a set limit.
Engine Error	Displays the action taken when an engine error occurs.

RELATED DOCUMENTATION

Creating UTM Profiles | 457

Creating Antivirus Profiles

Use the Create Antivirus Profiles page to configure antivirus profiles. The antivirus profile defines the content to scan for any malware and the action to be taken when malware is detected. After you create a profile, you can assign it to UTM profiles.

To create an antivirus profile:

1. Select **Configuration > Unified Threat Mgmt > Antivirus Profiles** in Customer Portal.

   The Antivirus Profiles page appears.

2. Click the add icon (+) to create a new antivirus profile.

   The Create Antivirus Profiles wizard appears, displaying brief instructions about creating an antivirus profile.

3. Click **Next** to navigate to the next page.

4. Complete the configuration according to the guidelines provided in Table 156 on page 473.

   NOTE: Fields marked with * are mandatory.

5. Click **Finish**.

   A summary page is displayed. Review the settings, and if you need to make any modifications, click the **Edit** link or the **Back** button.

6. Click **OK** to save the settings and create the profile.

   A message indicating the status of the create operation is displayed.

7. Click **Close**.

   You are returned to the Antivirus Profiles page.

Table 156: Antivirus Profile Settings

Setting	Guideline
General Information	
Name	Enter a unique name for the antivirus profile. The maximum length is 29 characters.
Description	Enter a description for the antivirus profile. The maximum length is 255 characters.
Engine Type	Displays the engine type used for scanning. Currently, **Sophos** is the only antivirus engine supported. Sophos antivirus is an in-the-cloud antivirus solution. The virus and malware database is located on external servers maintained by Sophos (Sophos Extensible List) servers, thus there is no need to download and maintain large pattern databases on the Juniper Networks device.
Fallback Options	

Table 156: Antivirus Profile Settings *(continued)*

Setting	Guideline
	Fallback options are used when the antivirus system experiences errors and must fall back to one of the previously configured actions to either deny (block) or permit the object. Specify the fallback options to use when there is a failure, or select the default action if no specific options are to be configured: • **Content Size**—Select an option to specify whether the content should be blocked (default) or logged and permitted if the content size the previously defined limit. • **Content Size Limit**—Enter the content size limit in kilobytes (KB) based on which action is taken. The range is 20 through 40,000 KB. The content size limit check occurs before the scan request is sent. The content size refers to accumulated TCP payload size. • **Engine Error**—Select the action to take (**Block** [default] or **Log and Permit**) when an engine error occurs. The term *engine error* refers all engine errors, including engine not ready, timeout, too many requests, password protected, corrupt file, decompress layer, and out of resources. • **Default Action**—Select the default action (**Block** [default] or **Log and Permit**) to take when an error occurs.
Notification Options	
	Use the notification options to configure a method of notifying the user when a fallback occurs or a virus is detected: • Fallback Deny—Select this option to notify mail senders that their messages were blocked. • Fallback Non-Deny—Select this option to warn mail recipients that they received unblocked messages despite problems. • Virus Detected—Select this option to notify mail recipients that their messages were blocked.

RELATED DOCUMENTATION

Creating UTM Profiles | 457

Editing, Cloning, and Deleting Antivirus Profiles

IN THIS SECTION

- Editing Antivirus Profiles | 475
- Cloning Antivirus Profiles | 475
- Deleting Antivirus Profiles | 476

You can edit, clone, and delete antivirus profiles from the Antivirus Profiles page. This topic has the following sections:

Editing Antivirus Profiles

To modify the parameters configured for an antivirus profile:

> **NOTE:** You cannot modify the default profiles already present in the system.

1. Select **Configuration > Unified Threat Mgmt > Antivirus Profiles** in Customer Portal.

   The Antivirus Profiles page appears, displaying the existing antivirus profiles.

2. Select the antivirus profile that you want to edit and then select the edit icon (pencil). Alternatively, right-click a profile and select **Edit Antivirus Profile**.

   The Edit Antivirus Profiles page appears, displaying the same fields that are presented when you create an antivirus profile.

3. Modify the antivirus profile fields as needed.

4. Click **OK** to save your changes.

   You are taken to the Antivirus Profiles page. A confirmation message appears, indicating the status of the edit operation.

Cloning Antivirus Profiles

Cloning enables you to easily create a new antivirus profile based on an existing one.

To clone an antivirus profile:

1. Select **Configuration > Unified Threat Mgmt > Antivirus Profiles** in Customer Portal.

   The Antivirus Profiles page appears, displaying the existing antivirus profiles.

2. Select the antivirus profile that you want to clone and then select **More > Clone**. Alternatively, right-click a profile and select **Clone**.

   The Clone Antivirus Profiles page appears, displaying the same fields that are presented when you create an antivirus profile.

3. Modify the antivirus profile fields as needed.

4. Click **OK** to save your changes.

   You are taken to the Antivirus Profiles page. A confirmation message appears, indicating the status of the clone operation.

Deleting Antivirus Profiles

Before deleting an antivirus profile, ensure that the profile is not used in a UTM profile that is, in turn, used in a firewall policy intent. If you try to delete an antivirus profile that is used in a firewall policy intent, an error message is displayed.

To delete one or more antivirus profiles:

1. Select **Configuration > Unified Threat Mgmt > Antivirus Profiles** in Customer Portal.

   The Antivirus Profiles page appears, displaying the existing antivirus profiles.

2. Select one or more antivirus profiles that you want to delete and then select the delete icon (**X**). Alternatively, right-click a profile and select **Delete Antivirus Profiles**.

   An alert message appears, asking you to confirm the delete operation.

3. Click **Yes** to delete the selected antivirus profiles.

   A confirmation message appears, indicating the status of the delete operation.

RELATED DOCUMENTATION

Creating Antivirus Profiles | 472

About the Antivirus Profiles Page | 470

About the Antispam Profiles Page

To access this page, select **Configuration > Unified Threat Mgmt > Antispam Profiles** in Customer Portal.

Use the Antispam Profiles page to view and manage antispam profiles, An antispam profile is used to examine transmitted e-mail messages to identify e-mail spam by using a constantly updated spam block list.

Tasks You Can Perform

You can perform the following tasks from this page:

- Create an antispam profile—See "Creating Antispam Profiles" on page 478.
- Edit, clone, or delete an antispam profile—See "Editing, Cloning, and Deleting Antispam Profiles" on page 480.
- Clear the selected antispam profiles—Click **Clear All Selections** to clear any antispam profiles that you might have selected.
- View the details of an antispam profile—Select the antispam profile for which you want to view the details and from the More or right-click menu, select **Detailed View**. The Antispam Profile Details page appears. Table 158 on page 478 describes the fields on this page.
- Search for antispam profiles using keywords—Click the search icon and enter the search term in the text box and press Enter. The search results are displayed on the same page.

Field Descriptions

Table 157 on page 477 describes the fields on the Antispam Profiles page.

Table 157: Antispam Profiles Page Fields

Field	Description
Name	Name of the antispam profile.
Blacklist	Indicates whether server-based spam filtering or local spam filtering is used.
Action	Action to be taken when spam is detected.
Custom Tag	Custom-defined tag that identifies an e-mail message as spam.
Description	Description of the antispam profile.

Table 158: Antispam Profile Details Page Fields

Field	Description
Name	Name of the antispam profile.
Description	Description of the antispam profile.
Sophos Blacklist	Indicates whether this field is enabled (server-based filtering) or disabled (local filtering).
Default Action	Action to be taken when spam is detected.
Custom Tag	Custom-defined tag that identifies an e-mail message as spam.

RELATED DOCUMENTATION

Creating UTM Profiles | 457

Creating Antispam Profiles

Use the Create Antispam Profiles page to configure antispam profiles.

E-mail spam consists of unwanted e-mail messages usually sent by commercial, malicious, or fraudulent entities. When the device detects an e-mail message deemed to be spam, it either blocks the message or tags the message header or subject field with a preprogrammed string. Antispam filtering allows you to use a third-party server-based spam block list (SBL) and to optionally create your own local allowlists (benign) and blocklists (malicious) for filtering against e-mail messages.

NOTE: Sophos updates and maintains the IP-based SBL. Antispam is a separately licensed subscription service.

After you create an antispam profile, you can assign it to UTM profiles.

To create an antispam profile:

1. Select **Configuration > Unified Threat Mgmt > Antispam Profiles** in Customer Portal.

   The Antispam Profiles page appears.

2. Click the add icon (+) to create a new antispam profile.

   The Create Antispam Profiles wizard appears, displaying brief instructions about creating an antispam profile.

3. Complete the configuration according to the guidelines provided in Table 159 on page 479.

   NOTE: Fields marked with * are mandatory.

4. Click **OK** save the settings and create the profile.

   A message indicating the status of the create operation is displayed. You are returned to the Antispam Profiles page.

Table 159: Antispam Profile Settings

Setting	Guideline
General Information	
Name	Enter a unique name for the antispam profile. The maximum length is 29 characters.
Description	Enter a description for the antispam profile. The maximum length is 255 characters.
Sophos Blacklist	Select this check box (the default) to use server-based spam filtering. If you clear the check box, local spam filtering is used. Server-based antispam filtering requires Internet connectivity with the spam block list (SBL) server. Domain Name Service (DNS) is required to access the SBL server. The firewall performs SBL lookups through the DNS protocol. **NOTE:** Server-based spam filtering supports only IP-based spam block list blocklist lookup. Sophos updates and maintains the IP-based spam block list. Server-based antispam filtering is a separately licensed subscription service.
Action	

Table 159: Antispam Profile Settings *(continued)*

Setting	Guideline
Default Action	Select the action to be taken when spam is detected: • **Tag Email Subject Line** • **Tag SMTP Header** • **Block Email** • **None**
Custom Tag	Enter a custom string for identifying a message as spam. The maximum length is 512 characters and the default is ***SPAM***.

RELATED DOCUMENTATION

Creating UTM Profiles | 457

Editing, Cloning, and Deleting Antispam Profiles

IN THIS SECTION

- Editing Antispam Profiles | 481
- Cloning Antispam Profiles | 481
- Deleting Antispam Profiles | 482

You can edit, clone, and delete antispam profiles from the Antispam Profiles page. This topic has the following sections:

Editing Antispam Profiles

To modify the parameters configured for an antispam profile:

> **NOTE:** You cannot modify the default profiles already present in the system.

1. Select **Configuration > Unified Threat Mgmt > Antispam Profiles** in Customer Portal.

   The Antispam Profiles page appears, displaying the existing antispam profiles.

2. Select the antispam profile that you want to edit and click the edit icon (pencil). Alternatively, right-click a profile and select **Edit Antispam Profile**.

   The Edit Antispam Profiles page appears, displaying the same fields that are presented when you create an antispam profile.

3. Modify the antispam profile fields as needed.

4. Click **OK** to save your changes.

   You are taken to the Antispam Profiles page. A confirmation message appears, indicating the status of the edit operation.

Cloning Antispam Profiles

Cloning enables you to easily create a new antispam profile based on an existing one.

To clone an antispam profile:

1. Select **Configuration > Unified Threat Mgmt > Antispam Profiles** in Customer Portal.

   The Antispam Profiles page appears displaying the existing antispam profiles.

2. Select the antispam profile that you want to clone and then select **More > Clone**. Alternatively, right-click a profile and select **Clone**.

   The Clone Antispam Profiles page appears, displaying the same fields that are presented when you create an antispam profile.

3. Modify the antispam profile fields as needed.

4. Click **OK** to save your changes.

You are taken to the Antispam Profiles page. A confirmation message appears, indicating the status of the clone operation.

Deleting Antispam Profiles

Before deleting an antispam profile, ensure that the profile is not used in a UTM profile that is, in turn, used in a firewall policy intent. If you try to delete an antispam profile that is used in a firewall policy intent, an error message is displayed.

To delete one or more antispam profiles:

1. Select **Configuration > Unified Threat Mgmt > Antispam Profiles** in Customer Portal.

   The Antispam Profiles page appears, displaying the existing antispam profiles.

2. Select one or more antispam profiles that you want to delete and click the delete icon (**X**). Alternatively, right-click a profile and select **Delete Antispam Profiles**.

   An alert message appears, asking you to confirm the delete operation.

3. Click **Yes** to delete the selected antispam profiles.

   A confirmation message appears, indicating the status of the delete operation.

RELATED DOCUMENTATION

About the Antispam Profiles Page | 477

About the Content Filtering Profiles Page

To access this page, select **Configuration > Unified Threat Mgmt > Content Filtering Profiles** in Customer Portal.

Use the Content Filtering Profiles page to view and manage content filtering profiles. Content filtering profiles enable you to block or permit certain types of traffic over several protocols (HTTP, FTP upload and download, IMAP, SMTP, and POP3) based on the MIME type, file extension, protocol command, and embedded object type.

Tasks You Can Perform

You can perform the following tasks from this page:

- Create a content filtering profile—See "Creating Content Filtering Profiles" on page 484.
- Edit, clone, or delete a content filtering profile—See "Editing, Cloning, and Deleting Content Filtering Profiles" on page 488.
- Clear the selected content filtering profiles—Click **Clear All Selections** to clear any content filtering profiles that you might have selected.
- View the details of a content filtering profile—Select the content filtering profile for which you want to view the details and from the More or right-click menu, select **Detailed View**. The Content Filtering Profile Details page appears. Table 161 on page 483 describes the fields on this page.
- Search for content filtering profiles by using keywords—Click the search icon, enter the search term in the text box, and press Enter. The search results are displayed on the same page.

Field Descriptions

Table 160 on page 483 describes the fields on the Content Filtering Profiles page.

Table 160: Content Filtering Profiles Page Fields

Field	Description
Name	Name of the content filtering profile.
Permit Command List	List of protocol commands permitted by the content filtering profile.
Block Command List	List of protocol commands blocked by the content filtering profile.
Notification Type	Type of notification that is sent when content is blocked.
Description	Description of the content filtering profile.

Table 161: Content Filtering Profiles Details Page Fields

Field	Description
General Information	
Name	Name of the content filtering profile.
Description	Description of the content filtering profile.
General Information	
Notify Mail Sender	Specifies whether the option to notify the e-mail sender is enabled or disabled.

Table 161: Content Filtering Profiles Details Page Fields *(continued)*

Field	Description
Notification Type	Type of notification that is sent when content is blocked.
Custom Notification Message	Custom notification message that is sent when content is blocked.
Protocol Commands	
Command Block List	List of protocol commands permitted by the content filtering profile.
Command Permit List	List of protocol commands blocked by the content filtering profile.
Content Types	
Block Content Types	List of harmful content types to be blocked.
File Extensions	
Extension Block List	File extensions to be blocked.
MIME	
MIME Block List	List of MIME types to be blocked.
MIME Permit List	List of MIME types to be permitted.

RELATED DOCUMENTATION

Creating UTM Profiles | 457

Creating Content Filtering Profiles

Use the Create Content Filtering Profiles page to configure content filtering profiles. Content filtering blocks or permits certain types of traffic based on the MIME type, file extension, and protocol command. The content filter controls file transfers across the device by checking traffic against configured filter lists. Table 162 on page 485 displays the types of content filters that you can configure as part of a content filtering profile.

NOTE: The content filtering profile evaluates traffic before all other UTM profiles. Therefore, if traffic meets criteria configured in the content filter, the content filter acts first upon this traffic.

Table 162: Supported Content Filter Types

Type	Description
MIME pattern filter	MIME patterns are used to identify the type of traffic in HTTP and MAIL protocols. There are two lists of MIME patterns that are used by the content filter to determine the action to be taken. The block MIME list contains a list of MIME type traffic that is to be blocked. The MIME exception list contains MIME patterns that are not to be blocked by the content filter and are generally subsets of items on the block list. **NOTE:** The exception list has a higher priority than the block list.
Block Extension List	Because the name of a file is available during the transfers, using file extensions is a highly practical way to block or allow file transfers. All protocols support the use of the block extension list.
Protocol Command Block and Permit Lists	Different protocols use different commands to communicate between servers and clients. By blocking or allowing certain commands, traffic can be controlled on the protocol command level. The block or permit command lists are intended to be used in combination, with the permit list acting as an exception list to the block list. **NOTE:** If a protocol command appears on both the permit list and the block list, the command is permitted.

To create a content filtering profile:

1. Select **Configuration > Unified Threat Mgmt > Content Filtering Profiles** in Customer Portal.

   The Content Filtering Profiles page appears.

2. Click the add icon (+) to create a new content filtering profile.

   The Create Content Filtering Profiles wizard appears, displaying brief instructions about creating a content filtering profile.

3. Click **Next** to navigate to the next page.

4. Complete the configuration according to the guidelines provided in Table 163 on page 486.

NOTE: Fields marked with * are mandatory.

5. Click **Finish**.

   A summary page is displayed. Review the settings and if you need to make any modifications click the **Edit** link or the **Back** button.

6. Click **OK** save the settings and create the profile.

   A message indicating the status of the create operation is displayed.

7. Click **Close**.

   You are returned to the Content Filtering Profiles page.

Table 163: Content Filtering Profile Settings

Setting	Guideline
General Information	
Name	Enter a unique name for the content filtering profile. The maximum length is 29 characters.
Description	Enter a description for the content filtering profile. The maximum length is 255 characters.
Notification Options	
Notify Mail Sender	Select this check box if you want to notify the sender when a failure occurs or a virus is detected. This check box is cleared by default.
Notification Type	Select the type of notification (**Protocol** or **Message**) from the drop-down list.
Custom Notification Message	Enter a custom notification message. The maximum length is 512 characters.
Protocol Commands	
Command Block List	Enter the protocol commands to be blocked for the HTTP, FTP, SMTP, IMAP, and POP3 protocols. Use commas to separate each command. Protocol commands allow you to control traffic at the protocol-command level.

Table 163: Content Filtering Profile Settings *(continued)*

Setting	Guideline
Command Permit List	Enter specific commands to be permitted for the HTTP, FTP, SMTP, IMAP, and POP3 protocols. Use commas to separate each command.
Content Types	
Block Content Type	Use the content filter to block other types of harmful files that the MIME type or the file extension cannot control. Select from the following types of content blocking (supported only for HTTP): • Active X • Windows executables (.exe) • HTTP cookie • Java applet • ZIP files
File Extensions	
Extension Block List	Use a file extension list to define a set of file extensions to block over HTTP, FTP, SMTP, IMAP, and POP3. Enter file extensions to block separated by commas. For example, exe, pdf, js, and so on.
MIME Types	
MIME Block List	Enter the MIME types you want to block over HTTP, FTP, SMTP, IMAP, and POP3 connections. Use commas to separate each MIME type.
MIME Permit List	Enter the MIME types you want to permit over HTTP, FTP, SMTP, IMAP, and POP3 connections. Use commas to separate each MIME type.

RELATED DOCUMENTATION

Creating UTM Profiles | 457

Editing, Cloning, and Deleting Content Filtering Profiles

IN THIS SECTION

- Editing Content Filtering Profiles | 488
- Cloning Content Filtering Profiles | 488
- Deleting Content Filtering Profiles | 489

You can edit, clone, and delete content filtering profiles from the Content Filtering Profiles page. This topic has the following sections:

Editing Content Filtering Profiles

To modify the parameters configured for a content filtering profile:

> **NOTE:** You cannot modify the default profiles already present in the system.

1. Select **Configuration > Unified Threat Mgmt > Content Filtering Profiles** in Customer Portal.

   The Content Filtering Profiles page appears, displaying the existing content filtering profiles.

2. Select the content filtering profile that you want to edit and click the edit icon (pencil). Alternatively, right-click a profile and select **Edit Profile**.

   The Edit Content Filtering Profiles page appears, displaying the same fields that are presented when you create a content filtering profile.

3. Modify the content filtering profile fields as needed.

4. Click **OK** to save your changes.

   You are taken to the Content Filtering Profiles page. A confirmation message appears, indicating the status of the edit operation.

Cloning Content Filtering Profiles

Cloning enables you to easily create a new content filtering profile based on an existing one.

To clone a content filtering profile:

1. Select **Configuration > Unified Threat Mgmt > Content Filtering Profiles** in Customer Portal.

   The Content Filtering Profiles page appears, displaying the existing content filtering profiles.

2. Select the content filtering profile that you want to clone and then select **More > Clone**. Alternatively, right-click a profile and select **Clone**.

   The Clone Content Filtering Profiles page appears, displaying the same fields that are presented when you create a content filtering profile.

3. Modify the content filtering profile fields as needed.

4. Click **OK** to save your changes.

   You are taken to the Content Filtering Profiles page. A confirmation message appears, indicating the status of the clone operation.

Deleting Content Filtering Profiles

Before deleting a content filtering profile, ensure that the profile is not used in a UTM profile that is, in turn, used in a firewall policy intent. If you try to delete a content filtering profile that is used in a firewall policy intent, an error message is displayed.

To delete one or more content filtering profiles:

1. Select **Configuration > Unified Threat Mgmt > Content Filtering Profiles** in Customer Portal.

   The Content Filtering Profiles page appears, displaying the existing content filtering profiles.

2. Select one or more content filtering profiles that you want to delete and click the delete icon (**X**). Alternatively, right-click a profile and select **Delete Profile**.

   An alert message appears, asking you to confirm the delete operation.

3. Click **Yes** to delete the selected content filtering profiles.

   A confirmation message appears, indicating the status of the delete operation.

RELATED DOCUMENTATION

Creating Content Filtering Profiles | 484

About the URL Patterns Page

To access this page, select **Configuration > Unified Threat Mgmt > URL Patterns** in Customer Portal.

Use this page to view, create, edit, clone, and delete URL patterns. A URL pattern contains a list of URLs.

Tasks You Can Perform

You can perform the following tasks from this page:

- Create a URL pattern—See "Creating URL Patterns" on page 491.
- Edit, clone, or delete a URL pattern—See "Editing, Cloning, and Deleting URL Patterns" on page 492.
- Clear the selected URL patterns—Click **Clear All Selections** to clear any URL patterns that you might have selected.
- View the details of a URL pattern—Select the URL pattern for which you want to view the details and from the More or right-click menu, select **Detailed View**. The URL Pattern Details page appears displaying the fields shown in Table 164 on page 490.
- Search for URL patterns using keywords—Click the search icon and enter the search term in the text box and press Enter. The search results are displayed on the same page.

Field Descriptions

Table 164 on page 490 describes the fields on the URL Patterns page.

Table 164: URL Patterns Page Fields

Field	Description
Name	Name of the URL pattern.
URLs	List of URLs in the URL pattern.
Description	Description of the URL pattern.

RELATED DOCUMENTATION

About the URL Categories Page | 494

Creating URL Patterns

Use this page to create URL patterns. You can also assign URL patterns to a URL category.

To create a URL pattern:

1. Select **Configuration > Unified Threat Mgmt > URL Patterns** in Customer Portal.

   The URL Patterns page appears.

2. Click the add icon (+) to create a URL pattern.

   The Create URL Patterns page is displayed.

3. Complete the configuration according to the guidelines provided in Table 165 on page 491.

   NOTE: Fields marked with * are mandatory.

4. Click **OK**.

   A new URL pattern is created and you are returned to the URL Patterns page.

Table 165: Create URL Patterns Settings

Settings	Guidelines
Name	Enter a unique name for the URL pattern. The name must begin with a letter or an underscore (_) and can contain alphanumeric characters and some special characters (_ -). The maximum length is 29 characters.
Description	Enter a description for the URL pattern. The maximum length is 255 characters.
URL Category	Select the URL category to which you want to assign the URL pattern. Alternatively, click **Create New URL Category** to create a URL category, enter the URL category name in the text box, and click **Save** to assign the URL pattern to the new category.

Table 165: Create URL Patterns Settings *(continued)*

Settings	Guidelines

Add URLs	Enter one or more URLs (separated by commas) in the text box, and click **Add**. The URLs are displayed in the URL List table.
	
	NOTE:
	• The following wildcard characters are supported:
	  ◦ asterisk (*)
	  ◦ period (.)
	  ◦ square brackets ([])
	  ◦ question mark (?)
	• Precede all wildcard characters with http://.
	• The asterisk (*) can only be used at the beginning of a URL and must be followed by a period (.).
	• The question mark (?) can only be used at the end of a URL.
	• The following are examples of wildcard syntaxes that are supported: http://*.example.net, http://www.example.ne?, and http://www.example.n??.
	• The following are examples of wildcard syntaxes that are not supported: *.example.???, http://*example.net, http://?, and www.example.ne?.

RELATED DOCUMENTATION

Creating URL Categories | 495

Editing, Cloning, and Deleting URL Patterns

IN THIS SECTION

Editing URL Patterns | 493
Cloning URL Patterns | 493
Deleting URL Patterns | 494

You can edit, clone, and delete URL patterns from the URL Patterns page. This topic has the following sections:

Editing URL Patterns

To modify the parameters configured for a URL pattern:

1. Select **Configuration > Unified Threat Mgmt > URL Patterns** in Customer Portal.

   The URL Patterns page appears, displaying the existing URL patterns.

2. Select the URL pattern that you want to edit and click the edit icon (pencil). Alternatively, right-click a pattern and select **Edit URL Patterns**.

   The Edit URL Patterns page appears, displaying the same fields that are presented when you create a URL pattern.

3. Modify the URL pattern fields as needed.

4. Click **OK** to save your changes.

   You are taken to the URL Patterns page. A confirmation message appears, indicating the status of the edit operation.

Cloning URL Patterns

Cloning enables you to easily create a new URL pattern based on an existing one.

To clone a URL pattern:

1. Select **Configuration > Unified Threat Mgmt > URL Patterns** in Customer Portal.

   The URL Patterns page appears, displaying the existing URL patterns.

2. Select the URL pattern that you want to clone and then select **More > Clone**. Alternatively, right-click a pattern and select **Clone**.

   The Clone URL Patterns page appears, displaying the same fields that are presented when you create a URL pattern.

3. Modify the URL pattern fields as needed.

4. Click **OK** to save your changes.

   You are taken to the URL Patterns page. A confirmation message appears, indicating the status of the clone operation.

Deleting URL Patterns

Before deleting a URL pattern, ensure that the URL pattern is not referenced in any UTM profiles that are, in turn, used in firewall policy intents or in URL categories referenced in the UTM settings. If you try to delete such a URL pattern, an error message is displayed.

To delete one or more URL patterns:

1. Select **Configuration > Unified Threat Mgmt > URL Patterns** in Customer Portal.

   The URL Patterns page appears, displaying the existing URL patterns.

2. Select one or more URL patterns that you want to delete and click the delete icon (**X**). Alternatively, right-click a pattern and select **Delete URL Pattern**.

   An alert message appears, asking you to confirm the delete operation.

3. Click **Yes** to delete the selected URL patterns.

   A confirmation message appears, indicating the status of the delete operation.

RELATED DOCUMENTATION

Creating URL Patterns | 491

About the URL Categories Page

To access this page, select **Configuration > Unified Threat Mgmt > URL Categories** in Customer Portal.

Use this page to view, create, edit, clone, and delete URL categories. A URL category is a list of URL patterns grouped under a single title.

Tasks You Can Perform

You can perform the following tasks from this page:

- Create a URL category—See "Creating URL Categories" on page 495.
- Edit, clone, or delete a URL category—See "Editing, Cloning, and Deleting URL Categories" on page 497.
- Clear the selected URL categories—Click **Clear All Selections** to clear any URL categories that you might have selected.

- View the details of a URL category—Select the URL category for which you want to view the details and from the More or right-click menu, select **Detailed View**. The URL Category Details page appears, displaying the details of the selected URL category; see Table 166 on page 495 for an explanation of the fields.
- Search for URL categories by using keywords—Click the search icon, enter the search term in the text box, and press Enter. The search results are displayed on the same page.

Field Descriptions

Table 166 on page 495 describes the fields on the URL Categories page.

Table 166: URL Categories Page Fields

Field	Description
Name	Name of the URL category.
URL Patterns	List of URL patterns in the URL category.
Definition Type	Indicates the type of URL category: • Predefined—URL categories that are loaded by default. • Custom—URL categories that are created by the user.
Description	Description of the URL category.

RELATED DOCUMENTATION

About the URL Patterns Page | 490

Creating URL Categories

Use this page to create URL categories. A URL category is a list of URL patterns grouped under a single title.

To create a URL category:

1. Select **Configuration > Unified Threat Mgmt > URL Categories** in Customer Portal.

   The URL Categories page appears.

2. Click the add icon (+) to create a URL category.

   The Create URL Categories page is displayed.

3. Complete the configuration according to the guidelines provided in Table 167 on page 496.

   NOTE: Fields marked with * are mandatory.

4. Click **OK**.

   A new URL category is created and you are returned to the URL Categories page.

Table 167: Create URL Categories Settings

Settings	Guidelines
Name	Enter a unique name for the URL category. The name must begin with a letter or an underscore (_) and can contain alphanumeric characters and some special characters (_ -). The maximum length is 59 characters.
Description	Enter a description for the URL pattern. The maximum length is 255 characters.
URL Patterns	Select one or more URL patterns in the **Available** column and click the forward arrow to confirm your selection. The selected URL patterns are displayed in the **Selected** column. Alternatively, click **Create a New Pattern** to create a URL pattern and assign it to the URL category. The Create URL Patterns page appears. For more information, see "Creating URL Patterns" on page 491 **NOTE:** You must select at least one URL pattern.

RELATED DOCUMENTATION

Editing, Cloning, and Deleting URL Categories | 497

Editing, Cloning, and Deleting URL Categories

IN THIS SECTION

- Editing URL Categories | 497
- Cloning URL Categories | 497
- Deleting URL Categories | 498

You can edit, clone, and delete URL categories from the URL Categories page. This topic has the following sections:

Editing URL Categories

To modify the parameters configured for a URL category:

1. Select **Configuration > Unified Threat Mgmt > URL Categories** in Customer Portal.

   The URL Categories page appears, displaying the existing URL categories.

2. Select the URL category that you want to edit and click the edit icon (pencil). Alternatively, right-click a category and select **Edit URL Categories**.

   The Edit URL Categories page appears, displaying the same fields that are presented when you create a URL category.

3. Modify the URL category fields as needed.

4. Click **OK** to save your changes.

   You are taken to the URL Categories page. A confirmation message appears, indicating the status of the edit operation.

Cloning URL Categories

Cloning enables you to easily create a new URL category based on an existing one.

To clone a URL category:

1. Select **Configuration > Unified Threat Mgmt > URL Categories** in Customer Portal.

   The URL Categories page appears, displaying the existing URL categories.

2. Select the URL category that you want to clone and then select **More > Clone**. Alternatively, right-click a category and select **Clone**.

   The Clone URL Categories page appears, displaying the same fields that are presented when you create a URL category.

3. Modify the URL category fields as needed.

4. Click **OK** to save your changes.

   You are taken to the URL Categories page. A confirmation message appears, indicating the status of the clone operation.

Deleting URL Categories

Before deleting a URL category, ensure that the URL category is not referenced in any UTM profiles that are, in turn, used in firewall policy intents or in the UTM settings. If you try to delete such a URL category, an error message is displayed.

To delete one or more URL categories:

1. Select **Configuration > Unified Threat Mgmt > URL Categories** in Customer Portal.

   The URL Categories page appears, displaying the existing URL categories.

2. Select one or more URL categories that you want to delete and click the delete icon (**X**). Alternatively, right-click a category and select **Delete URL Category**.

   An alert message appears, asking you to confirm the delete operation.

3. Click **Yes** to delete the selected URL categories.

   A confirmation message appears, indicating the status of the delete operation.

RELATED DOCUMENTATION

Creating URL Categories | 495

CHAPTER 17

Managing SLA Profiles and SD-WAN Policies

IN THIS CHAPTER

- SLA Profiles and SD-WAN Policies Overview | 499
- About the SD-WAN Policy Page | 503
- Creating SD-WAN Policy Intents | 505
- Editing and Deleting SD-WAN Policy Intents | 512
- About the SLA-Based Steering Profiles Page | 513
- Adding SLA-Based Steering Profiles | 517
- Editing and Deleting SLA-Based Steering Profiles | 524
- About the Path-Based Steering Profiles Page | 525
- Adding Path-Based Steering Profiles | 528
- Editing and Deleting Path-Based Steering Profiles | 530
- Breakout and Breakout Profiles Overview | 532
- About the Breakout Profiles Page | 535
- Adding Breakout Profiles | 540
- Adding Cloud Breakout Settings | 542
- Assigning Cloud Breakout Settings to Sites | 546
- Detaching Cloud Breakout Settings from Sites | 547
- Editing Breakout Profiles and Cloud Breakout Settings | 548
- Deleting Breakout Profiles and Cloud Breakout Settings | 550
- Configuring Breakout on SD-WAN Sites | 552

SLA Profiles and SD-WAN Policies Overview

IN THIS SECTION

- SLA Profiles | 500
- SD-WAN Policies | 501

Contrail Service Orchestration (CSO) enables you to create service-level agreement (SLA) profiles and map them to software-defined WAN (SD-WAN) policies for traffic management.

SLA Profiles

SLA profiles are created for applications or groups of applications for all tenants. An SLA profile consists of a set of configurable constraints that can be defined in the unified portal for both the Administration and Customer Portals. Table 168 on page 500 lists the categories of configurable constraints that are defined in an SLA profile.

Table 168: SLA Profile Categories

Category	Description

SLA profile parameters	You can define one or more than one of the following SLA profile parameters:
	• SLA Configuration—Whether to use recommended or custom values for the SLA threshold and SLA parameters.
	• SLA Threshold—Whether to use, liberal, baseline, or conservative settings for the threshold.
	• SLA parameters:
	  ◦ Packet loss—Percentage of data packets dropped by the network to manage congestion.
	  ◦ RTT—Target round-trip time (RTT) for the SLA profile.
	  ◦ Jitter—Difference between the maximum and minimum round-trip times (in ms) of a packet of data.

Path preference and failover	Paths are the WAN links to be used for the SLA profile. You can select MPLS, Internet, or any link as the preferred path. MPLS is more latency-sensitive than Internet. You can trigger the path failover criteria when any of the SLA parameters is violated.or when all the SLA parameters are violated.
Class of service	Class of service (CoS) provides different levels of service assurances to various forms of traffic. CoS enables you to divide traffic into classes and offer an assured service level for each class. The classes of service listed in increasing order of priority and sensitivity to latency are best effort, voice, interactive video, streaming audio or video, control, and business essential. The default CoS is voice.
Rate limiters	Rate limiters are defined for traffic shaping and efficient bandwidth utilization. You can define the following rate limiters: • Maximum upstream and downstream rates—The maximum upstream and downstream rate for all applications associated with the SLA profile. • Maximum upstream and downstream burst sizes—The maximum size of a steady stream of traffic sent at average rates that exceed the upstream and downstream rate limits for short periods.

NOTE: You must define at least one of the SLA parameters or path preference. You cannot leave both path preference and SLA parameters fields blank at the same time.

SD-WAN Policies

Applications are classified into the following categories:

- Cacheable applications, which refer to applications or application groups that are stored in the application cache when they are recognized by the device. After they are stored in the application cache, subsequent sessions are routed directly through the correct WAN link.
- Non-cacheable applications, which refer to applications or application groups that are not stored in the application cache and all sessions are first routed through the default path, and then routed to the correct WAN link based on the SD-WAN policy.

Policy intents consist of the following parameters:

- **Source**—A source endpoint that you can choose from a list of sites, site groups, and departments or a combination of all of these. The SD-WAN policy intent is applied to the selected source endpoint.
- **Destination**—A destination endpoint that you can choose from a list of applications and predefined or custom application groups. You can select a maximum of 32 applications or application groups as destination endpoints. The SD-WAN policy intent is applied to the selected destination endpoint.
- **Traffic Steering Profile**—Depending on whether you want to apply the policy intent to site-to-site traffic or breakout traffic, you can associate the traffic steering profile with the policy intent. The following options are available:
  - SLA-based steering profile— Applicable for site-to-site traffic
  - Path-based steering profile— Applicable for site-to-site traffic
  - Breakout profile—Applicable for breakout traffic (local, central, or cloud).
- **Intent name**—A unique name for the SD-WAN policy intent.

SD-WAN supports advanced policy-based routing (APBR). APBR enables you to dynamically define the routing behavior of the SD-WAN network based on applications. Dynamic application-based routing makes it possible to define policies and to switch WAN links on the fly based on the application's defined SLA parameters. The APBR mechanism classifies sessions based on applications and application signatures and uses policy intents to identify the best possible route for the application. When the best possible route does not meet the application's defined SLA requirements, the SD-WAN network finds the next best possible route to meet SLA requirements.

For example, consider an application in a site. If you want the application group to use custom throughput, latency, or jitter, you can create an SLA profile with these custom values. You can then create an intent

and configure the intent with the application and apply the custom SLA profile. When the intent is deployed, CSO determines the best suited WAN link to route traffic based in the application. If the WAN link fails to meet SLA requirements in runtime, the SD-WAN network switches WAN links to the next best suited path.

On the basis of the configured traffic-based steering profile constraints, you can categorize SD-WAN policies into three types:

- **Path-based steering policy**—If only the path preference is defined and none of the SLA parameters are defined in the SLA profile, then the policy is called a path-based steering policy. In path-based steering profile, you can define the path (MPLS or Internet) that must be used for a given traffic type profile, You cannot configure SLA parameters or path failover criteria for a path-based steering profile. The traffic type profile must be in enabled state in order to be used in any profile.
- **SLA-based steering policy**—If one or more SLA parameters in the SLA profile are defined, then the policy is called an SLA-based steering policy. In an SLA-based steering profile, each profile is associated with a traffic type profile and tracks the SLA parameters such as packet loss, Jitter and RTT. The traffic type profile must be in enabled state in order to be used in any profile. Based on your requirements, you can choose the recommended SLA threshold or enter custom SLA threshold for the traffic type profile. You can even set the path preference (Any, MPLS, or Internet) to switch traffic from one WAN interface to another based on the path failover criteria.

  When an intent is deployed on a site, if the WAN link chosen by the SD-WAN network does not meet the SLA requirements and the network performance deteriorates, then the site switches WAN links to meet the SLA requirements. The link switching is recorded as an SD-WAN event and displayed in the SD-WAN Events page in the customer portal and the *Tenant_name* SLA Performance pages in the administration and customer portals.
- **Breakout policy**—If local breakout, central breakout, or cloud breakout parameters are defined, then the policy is called a breakout policy.

RELATED DOCUMENTATION

About the SD-WAN Policy Page | **503**

Breakout and Breakout Profiles Overview | **532**

SD-WAN Events Overview | **845**

About the SD-WAN Policy Page

To access this page, select **Configuration > SD-WAN > SD-WAN Policy** in the Customer Portal.

SD-WAN policy intents help in optimum utilization of the WAN links and efficient load distribution of traffic. SD-WAN policy intents are applied to source endpoints (such as sites and departments) and destination endpoints (applications or application groups) and can be defined for site-to-site traffic (by using SLA profiles) or for breakout traffic (by using breakout profiles).

You can use the SD-WAN Policy page to view, create, edit, and deploy SD-WAN policy intents. SD-WAN policy intents use SLA profiles for traffic management. SD-WAN policies help in optimum utilization of the WAN links and efficient distribution of traffic. Every tenant has an SD-WAN policy and intents are created in the SD-WAN policy.

Tasks You Can Perform

You can perform the following tasks from this page:

- View existing SD-WAN policy intents. CSO provides pre-defined SD-WAN policy intents for the tenants. See Table 170 on page 504.

  NOTE: The pre-defined SD-WAN policy intents are available for the tenants only if the SP administrator has the downloaded the signature database prior to creating the tenants.

- Create SD-WAN policy intents. See “Creating SD-WAN Policy Intents” on page 505.
- Edit or delete SD-WAN policy intents. See “Editing and Deleting SD-WAN Policy Intents” on page 512.
- Deploy SD-WAN policy intents. See “Deploying Policies” on page 668.
- View the number of undeployed SD-WAN policy intents.
- Search for SD-WAN policy intents using keywords. Click the search icon. Enter partial text or full text of the keyword in the search bar and press Enter. The search results are displayed.

Field Descriptions

Table 169 on page 503 describes the fields on the SD-WAN Policy page.

Table 169: Fields on the SD-WAN Policy Page

Field	Description
Name	Displays the name of the SD-WAN policy intent.

Table 169: Fields on the SD-WAN Policy Page *(continued)*

Field	Description
Source	Displays the source endpoints that are configured for the policy intents. A source endpoint is chosen from sites, site groups, and departments or a combination of all of these to which the policy intent is applied.
Application	Displays the application destination endpoints that are configured for the policy intents. An application destination endpoint is chosen from a list of applications and predefined or custom application groups to which the policy intent is applied.
Traffic Steering Profile	Displays the breakout profile or the SLA profile associated with the policy intent.

Table 170: Pre-defined SD-WAN Policies

SD-WAN Policy Name	Applicable Sites	Application	Traffic Steering Profile
System-1	All Sites	CSO-Collaboration	CSO-AV
System-2	All Sites	CSO-Security	CSO-Sec
System-3	All Sites	CSO-Collaboration	CSO-Email
System-4	All Sites	CSO-Productivity	CSO-Productive
System-5	All Sites	CSO-File-Share	CSO-FileShare

RELATED DOCUMENTATION

- SLA Profiles and SD-WAN Policies Overview | **499**
- Creating SD-WAN Policy Intents | **505**
- Editing and Deleting SD-WAN Policy Intents | **512**

Creating SD-WAN Policy Intents

You can create policy intents for SD-WAN policies from the **SD-WAN Policy** page.

To create an SD-WAN policy intent:

1. Select **Configuration > SD-WAN > SD-WAN Policy** in Customer Portal.

   The SD-WAN Policy page appears.

2. Click the add icon (**+**).

   The options to create policy intents appear inline on the SD-WAN Policy page.

3. Enter the policy intent information according to the guidelines provided in Table 171 on page 506.

4. Click **Save** to create the policy intent.

   The SD-WAN policy intent is saved and a confirmation message is displayed.

> **NOTE:** After the policy intent is created, you must deploy the policy to ensure that the changes take effect on the applicable sites, departments, or applications. When an SD-WAN policy intent is created, the Undeployed field is incremented by one indicating that intents are pending deployment.

Table 171: Create SD-WAN Policy Intent Settings

Field	Guidelines
Source	

Table 171: Create SD-WAN Policy Intent Settings *(continued)*

Field	Guidelines
	You can select the source endpoints in one of the following ways: • Select source endpoints from the displayed list of departments, sites, or site groups, or a combination of these. Click the source endpoints to select them. • Select the source endpoints from the complete list of departments, sites, and site groups. To view the complete list of departments, sites, and site groups. 1. Click **View more results**. The complete list of departments, sites, and site groups is displayed in the **End Points** pane on the right. 2. (Optional) Hover over a department or site group and click the edit icon to edit the department or site group. You cannot edit a site. 3. Click the add icon (**+**) to select the endpoint. • Enter an abbreviation in the **Source** field to select the endpoint from a filtered list of departments, sites, or site groups. To view a filtered list of departments, sites, or site groups, enter DEPT, SITE, or STGP, respectively. The abbreviation is not case-sensitive. You can select the source endpoint in one of the following ways: ◦ Click the endpoints in the filtered list to select them. ◦ Click **View more results** to select the endpoint from the complete list of departments, sites, and site groups. ◦ Click **Add new department** or **Add new sitegroup** to create new departments or site groups and select them. The Create Site Group page or Add Department page appears based on your selection. See "Adding a Department" on page 767 and "Creating Site Groups" on page 197 for information about creating site groups and departments. • Create site groups or departments to select the source endpoint from the newly created site group or department. To create site groups or departments: 1. Click anywhere within the **Source** field. 2. Click the lesser-than icon (**<**) on the right. The list of available departments, sites, and site groups is displayed in the **End Points** pane on the right. 3. (Optional) To view more information about a source endpoint, hover over the endpoint click the details icon. 4. Click the add icon (**+**) on the top right of the pane.

Table 171: Create SD-WAN Policy Intent Settings *(continued)*

Field	Guidelines
	5. Click **Department** or **Site Group** as needed. The Add Department page or Create Site Group page appears based on your selection. See "Adding a Department" on page 767 and "Creating Site Groups" on page 197 for information about creating departments and site groups. 6. Click the check mark icon (√) if you want to save the department or site group to the policy intent. Alternatively, if you want to discard your updates, click **Cancel** instead.

Table 171: Create SD-WAN Policy Intent Settings *(continued)*

Field	Guidelines
Application	

Table 171: Create SD-WAN Policy Intent Settings *(continued)*

Field	Guidelines
	You can select the application endpoints in one of the following ways: **NOTE:** You can select **Any** as an application endpoint only if you are creating an SD-WAN policy intent that references a breakout profile. • Select application endpoints from the displayed list of applications and application groups. Click the endpoints to select them. • Select the application endpoints from the complete list of applications and application groups. To view the complete list of applications and applications groups. 1. Click **View more results**. The complete list of applications and applications groups is displayed in the **End Points** pane on the right. 2. (Optional) Hover over an application group and click the edit icon to edit the application group. 3. (Optional) Hover over an application and click the details icon to view details about the application. 4. Click the add icon (**+**) to select the endpoint. • Enter an abbreviation in the **Application** field to select the endpoint from a filtered list of applications and application groups. To view a filtered list of applications and application groups, enter apps or APPS. You can select the application endpoint in one of the following ways: ◦ Click the endpoints in the filtered list to select them. ◦ Click **View more results** to select the endpoint from the complete list of applications and applications groups. ◦ Click **Add new application** to create a new application group and select the application group. The Create Application Signature Group page appears. See "Adding Application Signature Groups" on page 764 for information about creating application groups. • Create custom application groups to select the application endpoint from the newly created application group. To create an application group: 1. Click anywhere within the **Application** field. 2. Click the lesser-than icon (**<**) on the right. The list of available applications, departments, sites, and site groups is displayed in the **End Points** pane on the right. 3. Click the add icon (**+**) on the top right of the pane.

Table 171: Create SD-WAN Policy Intent Settings *(continued)*

Field	Guidelines
	4. Click **Application**. The Create Application Signature Group page appears. See "Adding Application Signature Groups" on page 764 for information about creating application groups. 5. Click the check mark icon (√) if you want to save the application signature group to the policy intent. Alternatively, if you want to discard your updates, click **Cancel** instead.
Traffic Steering Profile	Select a breakout profile, SLA-based profile, or a path-based profile to apply to the source and application endpoints. You can select the profile in one of the following ways: • Select the breakout profile, SLA-based profile, or the path-based profile from the displayed list of profiles. Click the profile to select it. • Select the profile from the complete list of breakout, SLA-based, or path-based profiles: To view the complete list of breakout, SLA-based, or path-based profiles. 1. Click **View more results**. The complete list of profiles is displayed in the **End Points** pane on the right. 2. Click the add icon (**+**) to select the profile. • Select the profile by creating a custom SLA-based, path-based, or breakout profile: ◦ To create a breakout profile: 1. Click anywhere within the Profile field. 2. Click the lesser-than icon (**<**) on the right. The list of breakout profiles is displayed in the **End Points** pane on the right. 3. Click the add icon (**+**) on the top right of the pane and select **BRKT**. The Add Breakout Profile page appears. For more information, see "Adding Breakout Profiles" on page 540
Options	
Name	Enter a name for the policy intent.
Description	Enter a description for the policy intent.

RELATED DOCUMENTATION

SLA Profiles and SD-WAN Policies Overview | 499
About the SD-WAN Policy Page | 503
Editing and Deleting SD-WAN Policy Intents | 512
Deploying Policies | 668

Editing and Deleting SD-WAN Policy Intents

IN THIS SECTION

Editing SD-WAN Policy Intents | 512
Deleting SD-WAN Policy Intents | 513

You can edit or delete SD-WAN policy intents from the SD-WAN Policy page.

Editing SD-WAN Policy Intents

You can edit SD-WAN policy intents from the SD-WAN Policy page.

To edit an SD-WAN policy intent:

1. Hover over the SD-WAN policy intent that you want to edit, and then click the edit icon that appears on the right side of the policy intent.

   The options to create policy intents appear within the SD-WAN Policy page showing the same options that you see when you create a new SD-WAN policy intent.

2. Modify the parameters according to the guidelines provided in “Creating SD-WAN Policy Intents” on page 505.

3. Click **Save** to save your changes.

   Alternatively, click **Cancel** to discard your changes.

   NOTE: After you modify an SD-WAN policy intent, you must redeploy the policy to ensure that the changes take effect on the applicable sites, departments, or applications.

Deleting SD-WAN Policy Intents

If an SD-WAN intent is no longer needed, you can delete the SD-WAN policy intent from the SD-WAN Policy page.

To delete one or more SD-WAN policy intents:

1. Select one or more policy intents that you want to delete and click the delete icon (trash can).

   A page appears asking you to confirm the delete operation.

2. Click **Yes** to confirm that you want to delete the selected policy intents.

   A confirmation message appears indicating the status of the delete operation.

> **NOTE:** After you delete one or more SD-WAN policy intents, you must redeploy the policy to ensure that the changes take effect on the applicable sites, departments, or applications.

RELATED DOCUMENTATION

SLA Profiles and SD-WAN Policies Overview | 499

About the SD-WAN Policy Page | 503

Creating SD-WAN Policy Intents | 505

About the SLA-Based Steering Profiles Page

To access this page, select **Configuration > SD-WAN > SLA-Based Steering Profiles** in the Customer Portal.

You can use the SLA-Based Steering Profiles page to view information about service-level agreement (SLA)-based steering for the tenant profile in which you are logged in.

Tasks You Can Perform

You can perform the following tasks from this page:

- View details of SLA profiles for the tenants.
- Add an SLA-based steering profile for the tenant. See “Adding SLA-Based Steering Profiles” on page 517.

- Edit or delete an SLA-based steering profile. See “Editing and Deleting SLA-Based Steering Profiles” on page 524.
- Show or hide columns that contain information about SLA-based steering profiles. Click the **Show Hide columns** icon in the top right corner of the page and select columns that you want to view on the page.
- Search for SLA-based steering profiles using keywords. Click the search icon. Enter partial text or full text of the keyword in the search bar and press Enter. The search results are displayed.

Field Descriptions

Table 172 on page 514 shows the descriptions of the fields on the SLA-Based Steering Profiles page.

Table 172: Fields on the SLA-Based Steering Profiles Page

Field	Description	Displayed On
Name	Name of the SLA-based steering profile.	SLA-Based Steering Profiles page (SLA Profiles List tab) Detail for *SLA-Profile-Name* pane
Priority	Priority of the SLA-based steering profile. A value zero (0) indicates lower priority and one (1) indicates highest priority.	Detail for *SLA-Profile-Name* pane
Traffic Type Profile	Indicates the traffic type profile associated with the SLA-based steering profile. • VOICE-VIDEO • HIGH_PRIORITY_VIDEO • HOSTED_AV • PREMIUM_INTERNET • INTERNET	SLA-Based Steering Profiles page (SLA Profiles List tab) Detail for *SLA-Profile-Name* pane
Packet Loss (%)	Target packet loss for the SLA profile.	SLA-Based Steering Profiles page (SLA Profiles List tab) Detail for *SLA-Profile-Name* pane
Jitter (ms)	Target jitter for the SLA profile.	SLA-Based Steering Profiles page (SLA Profiles List tab) Detail for *SLA-Profile-Name* pane

Table 172: Fields on the SLA-Based Steering Profiles Page *(continued)*

Field	Description	Displayed On
RTT	Target round-trip time (RTT) for the SLA profile.	SLA-Based Steering Profiles page (SLA Profiles List tab) Detail for *SLA-Profile-Name* pane
SLA Probe Match	Indicates whether the profile requires the SLA probe to match all SLA criteria (All) or not (Any) .	Detail for *SLA-Profile-Name* pane
Created By	Name of the user who created the SLA-based steering profile.	SLA-Based Steering Profiles page (SLA Profiles List tab)
Path Preference	The preferred path for the SLA profile. The available options are: • MPLS • Internet • Any (default)	Detail for *SLA-Profile-Name* pane
Session-sampling %	Indicates the matching percentage of sessions for which you want to run the passive probes.	Detail for *SLA-Profile-Name* pane
SLA Violation Counts	Indicates the number of SLA violations after which you want CSO to switch paths.	Detail for *SLA-Profile-Name* pane
Sampling Period	The sampling period, in milliseconds, for which the SLA violations are counted.	Detail for *SLA-Profile-Name* pane
Switch Cool-off Period	The waiting period, in milliseconds, only after which you want the link switch to happen if an active link comes back online. This parameter helps prevent frequent switching of traffic between active and backup links.	Detail for *SLA-Profile-Name* pane
Path Failover Criteria	Indicates the path failover criteria for link switching. Path failover occurs when any (Any)of the SLA parameters is violated or when all (All) the SLA parameters are violated.	Detail for *SLA-Profile-Name* pane
Maximum Upstream Rate	The maximum upstream rate (in Kbps) for all applications associated with the SLA-based steering profile.	Detail for *SLA-Profile-Name* pane

Table 172: Fields on the SLA-Based Steering Profiles Page *(continued)*

Field	Description	Displayed On
Maximum Upstream Burst Size	The maximum upstream burst size (in bytes).	Detail for *SLA-Profile-Name* pane
Maximum Downstream Rate	The maximum downstream rate (in Kbps) for all applications associated with the SLA-based-steering profile.	Detail for *SLA-Profile-Name* pane
Maximum Downstream Burst Size	The maximum downstream burst size (in bytes).	Detail for *SLA-Profile-Name* pane

RELATED DOCUMENTATION

SLA Profiles and SD-WAN Policies Overview | 499

Adding SLA-Based Steering Profiles

You can use the Add SLA Profile page to add a new service-level agreement (SLA)-based steering profile, specify the traffic type profile, SLA configuration, SLA threshold, SLA parameters, path selection criteria, and rate limiting parameters for the profile. Table 173 on page 518 lists the SLA-based steering profiles that are tuned for specific application categories and traffic types.

Table 173: Predefined SLA-Based Steering Profiles

SLA-Based Steering Profiles	Traffic Type	Application Group	Applications Supported
CSO-AV	VOICE-VIDEO	CSO_Collaboration_AV	Skype for Business Zoom Video GotoMeeting Jive Jabber Citrix Online WebEx Zoho Meeting Google Hangout Adobe Connect

Table 173: Predefined SLA-Based Steering Profiles *(continued)*

SLA-Based Steering Profiles	Traffic Type	Application Group	Applications Supported
CSO-Productivity	PREMIUM-INTERNET	CSO_Productivity	ERP: Salesforce, Oracle, SAP Office365 (including SharePoint) Zendesk HRPayroll Zoho Office Suite Slack Square Concur Adobe Quickbooks Freshbooks Workday Project Management-MS PJ Basecamp Asana
CSO-Security	INTERNET	CSO_Security	Symantec McAfee Sophos Zonealarm Lookout

Table 173: Predefined SLA-Based Steering Profiles *(continued)*

SLA-Based Steering Profiles	Traffic Type	Application Group	Applications Supported
CSO-Email	PREMIUM-INTERNET	CSO_Collaboration_Email	MS Exchange IMAP POP3 Gmail OWA Yahoo
CSO-FileShare	INTERNET	CSO_File_Share	Box Dropbox Gsuite OneDrive Skype for Business-File Transfer Zoho Share

To add an SLA-based steering profile to the tenant:

1. Select **Configuration > SD-WAN > SLA-Based Steering Profiles**.

   The SLA-Based Steering Profiles page appears.

2. Click the add icon (**+**).

   The Add SLA Profile page appears.

3. Enter the SLA profile information according to the guidelines provided in Table 174 on page 521.

   NOTE: Fields marked with an asterisk (*) are mandatory.

4. Click **OK** to add the SLA profile.

   The SLA-Based Steering Profiles page appears with the new SLA profile information. You are returned to the SLA-Based Steering Profiles page and a confirmation message indicating that the SLA-based

steering profile was added is displayed. The page refreshes to display the SLA-based steering profile that you added.

Alternatively, if you want to discard your updates, click **Cancel** instead.

> **NOTE:** After you add an SLA-based steering profile, you must add an SD-WAN policy intent that references the SLA-based steering profile in order to enable site-to-site traffic.

Table 174: Fields on the Add SLA Profile page

Field	Guidelines
General	
Name	Enter a unique string that can contain alphanumeric characters and hyphens (-); the maximum length is 15 characters.
Traffic Type Profile	Choose a traffic type profile to apply the class-of-service configuration and priority to the SLA profile. You can select a traffic type profile only when it is in the **Enabled** state.
SLA Configuration	Choose one of the following options: • **Use Recommended**: To use the default SLA threshold and SLA parameters for the SLA-based steering profile. • **Enter Custom**: To specify customized values for SLA configuration and SLA parameters for the SLA-based steering profiles.
SLA Threshold	Choose one of the following options: • **Liberal**—To use a relaxed SLA threshold. • **Baseline**—To use the default SLA threshold. • **Conservative**—To use a strict SLA threshold.
SLA Parameters	
Packet Loss	Enter the target packet loss (in %) for the SLA-based steering profile. Packet loss is the percentage of data packets dropped by the network to manage congestion.
RTT	Enter the target round-trip time (RTT) for the SLA-based steering profile.
Jitter	Enter the target jitter (in ms) for the SLA-based steering profile. Jitter is the difference between the maximum and minimum round-trip times of a packet of data.

Table 174: Fields on the Add SLA Profile page *(continued)*

Field	Guidelines
Path Selection Criteria	
Path Preference	Select the preferred WAN link type to associate with the SLA profile. The options are Any, MPLS, and Internet. Any is the default value. Select the preferred path (MPLS, Internet, or Any) to be used for site-to-site traffic. If a WAN link type that matches the preferred path is enabled for site-to-site traffic, then that WAN link type is used for site-to-site traffic. If you specify that any path can be used, then there is no preference and all site-to-site-traffic-enabled links are used in a load-balancing mode.
Path Failover Criteria	Specify the failover criteria to determine how links are switched when the active links fail to meet the SLA criteria. In such cases, the traffic is routed to links that meet SLA criteria. Failover is supported only for MPLS or Internet links. **NOTE:** Path failover is supported only for bandwidth-optimized SD-WAN networks. Choose one of the following options: • **Does not meet one or more SLA parameters**—This triggers the path failover if any of the SLA parameters is violated. • **Does not meet all SLA parameters**—This triggers the path failover only when all the SLA parameters are violated.
Advanced Configuration-	
Rate Limiting	
Maximum Upstream Rate	Enter the maximum upstream rate (in Kbps) for all applications associated with the SLA profile. Range: 64 through 10,485,760 Kbps
Maximum Upstream Burst Size	Enter the maximum upstream burst size (in bytes). Range: 1 through 1,342,177,280 bytes
Maximum Downstream Rate	Enter the maximum downstream rate (in Kbps) for all applications associated with the SLA profile. Range: 64 through 10,485,760 Kbps

Table 174: Fields on the Add SLA Profile page *(continued)*

Field	Guidelines
Maximum Downstream Burst Size	Enter the maximum downstream burst size (in bytes). Range: 1 through 1,342,177,280
Loss Priority	Select a loss priority based on which packets can be dropped or retained when network congestion occurs. The chances of a packet getting dropped is the highest when the loss priority is set to **High**. Other available values are **Medium High**, **Medium Low**, and **Low**.
Real Time Optimized Mode Setting **NOTE:** The following fields are applicable only for sites configured with the real-time-optimized SD-WAN mode.	
SLA Sampling	
Session-sampling %	Enter the matching percentage of sessions for which you want to run the passive probes.
SLA-violation-count	Enter the number of SLA violations after which you want CSO to switch paths. The range is 1 through 32.
Sampling-period	Enter the sampling period, in seconds, for which the SLA violations are counted. The range is 2 through 60.
Switch-cool-off-period	Enter the waiting period, in seconds, only after which you want the link switch to happen if an active link comes back online. This parameter helps prevent frequent switching of traffic between active and backup links. The range is 5 through 300.

RELATED DOCUMENTATION

- SLA Profiles and SD-WAN Policies Overview | **499**
- About the SLA-Based Steering Profiles Page | **513**
- Editing and Deleting SLA-Based Steering Profiles | **524**

Editing and Deleting SLA-Based Steering Profiles

IN THIS SECTION

- Editing an SLA-Based Steering Profile | 524
- Deleting SLA-Based Steering Profiles | 525

You can use the SLA-Based Steering Profiles page to edit and delete SLA profiles.

> **NOTE:** You cannot edit the predefined SLA-Based steering profiles that are automatically created by Contrail Service Orchestration (CSO).

Editing an SLA-Based Steering Profile

To edit an SLA-based steering profile:

> **NOTE:** If you edit an SLA-based steering profile that is used in an SD-WAN policy intent, then that SD-WAN policy is marked for redeployment.

1. Select **Configuration > SD-WAN > SLA-Based Steering Profiles**.

   The SLA-Based Steering Profiles page appears.

2. Select the SLA-based steering profile that you want to edit, and click the Edit (pencil)icon.

   The Edit SLA Profile page appears displaying the same fields that are presented when you add a SLA-based steering profile. For more information, see "Adding SLA-Based Steering Profiles" on page 517.

3. Modify the fields as needed.

   > **NOTE:** You cannot edit the SLA-based steering profile name.

4. Click **OK**.

You are returned to the SLA-Based Steering Profiles page. The modifications that you made are saved and a confirmation message is displayed.

Deleting SLA-Based Steering Profiles

You can delete the SLA-based steering profile if they are no longer needed. To delete one or more SLA-based steering profile:

> **NOTE:** You cannot delete an SLA-based steering profile if it is referenced by one or more SD-WAN policy intents.

1. Select **Configuration > SD-WAN > SLA Based Steering Profiles**.

   The SLA-Based Steering Profiles page appears.

2. Select the SLA-based steering profiles that you want to delete and click the delete (trash can) icon.

   A popup dialog appears asking you to confirm the deletion.

3. Click **Yes**.

   You are returned to the SLA-Based Steering Profiles page. The selected SLA-based steering profile is deleted and a confirmation message is displayed.

RELATED DOCUMENTATION

SLA Profiles and SD-WAN Policies Overview | 499

About the SLA-Based Steering Profiles Page | 513

Adding SLA-Based Steering Profiles | 517

About the Path-Based Steering Profiles Page

To access this page, select **Configuration > SD-WAN > Path-Based Steering Profiles** in the Customer Portal.

You can use the Path-Based Steering Profiles page to view information about path profiles for the tenant profile in which you are logged in.

Tasks You Can Perform

You can perform the following tasks from this page:

- View details of path-based steering profiles for the tenant.
- Add path-based steering profiles for for the tenant. See "Adding Path-Based Steering Profiles" on page 528.
- Edit or delete a path-based steering profile. See "Editing and Deleting Path-Based Steering Profiles" on page 530.
- Show or hide columns that contain information about path-based steering profile. Click the **Show Hide columns** icon in the top right corner of the page and select columns that you want to view on the page.
- Search for path-based steering profiles using keywords. Click the search icon. Enter partial text or full text of the keyword in the search bar and press Enter. The search results are displayed.

Field Descriptions

Table 175 on page 526 shows the descriptions of the fields on the Path-Based Steering Profiles page.

Table 175: Fields on the Path-Based Steering Profiles Page

Field	Description	Displayed on
Name	Name of the path-based-steering profile.	Path-Based Steering Profiles Page (Path Profiles List tab) Detail for *Path-Profile-Name* pane
Traffic Type Profile	Indicates the traffic type profile associated with the path-based-steering profile. • VOICE-VIDEO • HIGH_PRIORITY_VIDEO • HOSTED_AV • PREMIUM_INTERNET • INTERNET	Path-Based Steering Profiles Page (Path Profiles List tab) Detail for *Path-Profile-Name* pane
Path Preference	The preferred path for the SLA profile. The available options are: • MPLS • Internet	Path-Based Steering Profiles Page (Path Profiles List tab) Detail for *Path-Profile-Name* pane
Created by	The name of the user who created the path profile.	Path-Based Steering Profiles Page (Path Profiles List tab)

Table 175: Fields on the Path-Based Steering Profiles Page *(continued)*

Field	Description	Displayed on
Priority	Priority of the path-based steering profile. A value zero (0) indicates lower priority and one (1) indicates highest priority.	Detail for *Path-Profile-Name* pane
Packet Loss	Target packet loss for the SLA profile.	Detail for *Path-Profile-Name* pane
RTT	Target round-trip time (RTT) for the SLA profile.	Detail for *Path-Profile-Name* pane
Jitter	Target jitter for the SLA profile.	Detail for *Path-Profile-Name* pane
SLA Probe Match	Indicates whether the profile requires the SLA probe to match all SLA criteria (All) or not (Any) .	Detail for *Path-Profile-Name* pane
Session-sampling %	Indicates the matching percentage of sessions for which you want to run the passive probes.	Detail for *Path-Profile-Name* pane
SLA Violation Counts	Indicates the number of SLA violations after which you want CSO to switch paths.	Detail for *Path-Profile-Name* pane
Sampling Period	The sampling period, in milliseconds, for which the path-based steering profile violations are counted.	Detail for *Path-Profile-Name* pane
Switch Cool-off Period	The waiting period, in milliseconds, only after which you want the link switch to happen if an active link comes back online. This parameter helps prevent frequent switching of traffic between active and backup links.	Detail for *Path-Profile-Name* pane
Path Failover Criteria	Indicates the path failover criteria for link switching. Path failover occurs when any (Any)of the path-based steering profile parameters is violated or when all (All) the path-based steering profile parameters are violated.	Detail for *Path-Profile-Name* pane
Maximum Upstream Rate	The maximum upstream rate (in Kbps) for all applications associated with the path-based steering profile.	Detail for *Path-Profile-Name* pane
Maximum Upstream Burst Size	The maximum upstream burst size (in bytes).	Detail for *Path-Profile-Name* pane

Table 175: Fields on the Path-Based Steering Profiles Page *(continued)*

Field	Description	Displayed on
Maximum Downstream Rate	The maximum downstream rate (in Kbps) for all applications associated with the path-based-steering profile.	Detail for *Path-Profile-Name* pane
Maximum Downstream Burst Size	The maximum downstream burst size (in bytes).	Detail for *Path-Profile-Name* pane

RELATED DOCUMENTATION

SLA Profiles and SD-WAN Policies Overview | 499

Adding Path-Based Steering Profiles

You can use the Add Path Profile page to add a new path-based steering profile, and specify the traffic type profile, path preference, and advanced configuration for the profile.

To add a path-based steering profile to the tenant:

1. Select **Configuration > SD-WAN > Path-Based Steering Profiles**.

   The Path-Based Steering Profiles page appears.

2. Click the add (+) icon.

   The Add Path Profile page appears.

3. Enter the path-based steering profile information according to the guidelines provided in Table 176 on page 529.

   NOTE: Fields marked with an asterisk (*) are mandatory.

4. Click **OK** .

You are returned to the Path-Based Steering Profiles page and a confirmation message indicating that the path-based steering profile was added is displayed. The page refreshes to display the path-based steering profile that you added.

> **NOTE:** After you add a path-based steering profile, you must add an SD-WAN policy intent that references the path-based steering profile in order to enable site-to-site traffic.

Table 176: Fields on the Add Path Profile page

Field	Guidelines
Name	Enter a unique string that can contain alphanumeric characters and hyphens (-); the maximum length is 15 characters.
Traffic Type Profile	Choose a traffic type profile to apply the class-of-service configuration and priority to the SLA profile. You can select a traffic type profile only when it is in the **Enabled** state.
Path Preference	Select the preferred WAN link type to associate with the SLA profile. The options are MPLS, and Internet.
Advanced Configuration	
Maximum Upstream Rate	Enter the maximum upstream rate (in Kbps) for all applications associated with the SLA profile. Range: 64 through 10,485,760 Kbps
Maximum Upstream Burst Size	Enter the maximum burst size (in bytes). Range: 1 through 1,342,177,280 bytes
Maximum Downstream Rate	Enter the maximum downstream rate (in Kbps) for all applications associated with the SLA profile. Range: 64 through 10,485,760 Kbps
Maximum Downstream Burst Size	Enter the maximum burst size (in bytes). Range: 1 through 1,342,177,280 bytes

Table 176: Fields on the Add Path Profile page *(continued)*

Field	Guidelines
Loss Priority	Select a loss priority based on which packets can be dropped or retained when network congestion occurs. The chances of a packet getting dropped is the highest when the loss priority is set to **High**. Other available values are **Medium High**, **Medium Low**, and **Low**.

RELATED DOCUMENTATION

- SLA Profiles and SD-WAN Policies Overview | 499
- About the Path-Based Steering Profiles Page | 525
- Editing and Deleting Path-Based Steering Profiles | 530

Editing and Deleting Path-Based Steering Profiles

IN THIS SECTION

- Editing a Path-Based Steering Profile | 531
- Deleting a Path-Based Steering Profile | 532

You can use the Path-Based Steering Profiles page to edit and delete path-based steering profiles.

Editing a Path-Based Steering Profile

To edit a path-based steering profile:

> **NOTE:** If you edit a path-based steering profile that is used in an SD-WAN policy intent, then that SD-WAN policy is marked for redeployment.

1. Select **Configuration > SD-WAN > Path-Based Steering Profiles**.

   The Path-Based Steering Profiles page appears.

2. On the Path Profiles tab, select the path-based steering profile that you want to edit.

3. Click the edit (pencil) icon.

   The Edit Path Profile page appears displaying the same fields that are presented when you add a path-based steering profile. For more information, see "Adding Path-Based Steering Profiles" on page 528.

4. Modify the fields as needed.

   > **NOTE:** You cannot edit the path profile name.

5. Click **OK**.

   You are returned to the Path-Based Steering Profiles page. The modifications that you made are saved and a confirmation message is displayed.

Deleting a Path-Based Steering Profile

You can delete path-based steering profiles if they are no longer needed. To delete one or more path-based steering profiles:

> **NOTE:** You cannot delete a path-based steering profile if it is referenced by one or more SD-WAN policy intents.

1. Select **Configuration > SD-WAN > Path-Based Steering Profiles**.

   The Path-Based Steering Profiles page appears.

2. On the Path Profiles List tab, select the path profiles that you want to delete.

3. Click the delete (trash can) icon.

   A popup dialog appears asking you to confirm the deletion.

4. Click **Yes**.

   You are returned to the Path-Based Steering Profiles page. The selected path-based steering profiles are deleted and a confirmation message is displayed.

RELATED DOCUMENTATION

- SLA Profiles and SD-WAN Policies Overview | **499**
- About the Path-Based Steering Profiles Page | **525**
- Adding Path-Based Steering Profiles | **528**

Breakout and Breakout Profiles Overview

IN THIS SECTION

- Cloud Breakout | **534**
- Breakout Profiles | **534**

SD-WAN Policy Intents for Breakout | 534
Benefits of Breakout Profiles | 535

Site-to-site traffic between spoke sites of a tenant is sent (on overlay tunnels) directly from one site to another depending on the tenant topology or through the hub or enterprise hub. However, for Internet-bound or Software as a Service (SaaS) traffic, you can break out the traffic in different ways:

- Local breakout—The traffic exits the VPN directly at the site and goes to the destination.

  NOTE: If underlay BGP is enabled for a WAN link, then the routes learnt from BGP are installed for local breakout; CSO does not generate the static default route.

- Backhaul or central breakout—The traffic exits the VPN at the provider hub or at the enterprise hub (if a enterprise hub is associated with the spoke site) and then goes to the destination.
- Cloud breakout—The traffic is sent from the site to a designated cloud-based security platform instead of traffic being sent over an underlay.

  NOTE: From CSO Release 4.1.0 onwards, Zscaler is the only cloud-based security platform supported.

In CSO Release 4.0, only local breakout and central breakout (backhaul) are supported and the breakout option is enabled only at the site level. However, from CSO Release 4.1.0 onward, breakout is supported at the site, department, and application (cacheable only) levels by using breakout profiles that are applied using SD-WAN policy intents. Non-cacheable applications follow the site-specific or department-specific behavior as configured in the SD-WAN policy intent.

NOTE: For sites added in CSO Release 4.1.0 onward, you cannot configure breakout *directly* at the site level and must use breakout profiles referenced in SD-WAN policy intents for this purpose.

Cloud Breakout

In releases before CSO Release 5.1.0, as part of providing the tunneled breakout to Zscaler, the tunnel source public IP address was obtained only from the WAN interface. With pool-based NAT supported from Release 5.1.0 onward, the tunnel creation to Zscaler (when pool-based NAT is configured) obtains the source address from the WAN link's NAT pool.

When multiple Zscaler tunnels are needed on a WAN interface (for example, when primary and secondary cloud breakout nodes are configured), the pool IP address must be large enough to accommodate these tunnels. In the case of multiple Zscaler tunnels, no two Zscaler tunnels will have the same source IP address. However, the IP address that is used as Zscaler tunnel's source address, can also be used in the NAT pools.

Breakout Profiles

The following three types of breakout profiles are supported in CSO:

- Local breakout (underlay)
- Backhaul (central breakout)
- Cloud breakout

After you add a breakout profile, you must create an SD-WAN policy intent specifying the source (site, site group, or department) and application and the applicable breakout profile.

SD-WAN Policy Intents for Breakout

For SD-WAN policy intents configured at different source endpoints, the following is applicable:

- Site—A policy intent configured at the site level applies to all the departments within the site. In addition, by default, the site-level configuration is also applicable to all applications because the default configuration for applications is **Any**.
- Department—A policy intent configured at the department level (for tenants with network segmentation enabled) overrides the policy intent configured at the site level. Similar to the behavior for the site-level policy intent, by default, a department-level policy intent is also applicable to all applications because the default configuration for applications is **Any**.
- Application (cacheable only)—A policy intent (at the application level) where you specify one or more cacheable applications overrides the policy intent specified at either the department level or the site level *only* for the specified applications.

Benefits of Breakout Profiles

- Breakout profiles used in intent-based Internet breakout policies (through SD-WAN policy intents) give users granular control over the Internet breakout behavior for specific applications.

RELATED DOCUMENTATION

Adding Breakout Profiles | 540

Adding Cloud Breakout Settings | 542

Creating SD-WAN Policy Intents | 505

About the Breakout Profiles Page

IN THIS SECTION

- Tasks You Can Perform | 535
- Breakout Profiles Field Descriptions | 536
- Cloud Breakout Settings Field Descriptions | 538

To access this page, click **Configuration > SD-WAN > Breakout Profiles**.

You can use the Breakout Profiles page to view existing breakout profiles, add local, backhaul, and cloud breakout profiles, edit breakout profiles, and delete breakout profiles. You can also add settings for cloud breakout, edit cloud breakout settings, assign the settings to one or more sites, detach the settings from one or more sites, and delete the settings.

The breakout profiles are displayed on the Breakout Profiles tab and the cloud breakout settings are displayed on the Cloud Breakout Settings tab.

Tasks You Can Perform

You can perform the following tasks from this page:

- View existing breakout profiles—See Table 177 on page 536 for a description of the fields.
- View the details of a breakout profile—On the Breakout Profiles tab, select a breakout profile and from the More menu, select **Detail View**. The Detail for *Breakout-Profile-Name* pane appears on the right-hand side of the page. See Table 177 on page 536 for a description of the fields on this pane.
- View existing cloud breakout settings—See Table 178 on page 538 for a description of the fields.
- View the details of cloud breakout settings—On the Cloud Breakout Settings tab, select a cloud breakout setting and from the More menu, select **Detail View**. The Detail for *Cloud-Breakout-Setting-Name* pane appears on the right-hand side of the page. See Table 178 on page 538 for a description of the fields on this pane.
- Add a breakout profile—See "Adding Breakout Profiles" on page 540.
- Edit a breakout profile—See "Editing Breakout Profiles and Cloud Breakout Settings" on page 548.
- Delete a breakout profile—See "Deleting Breakout Profiles and Cloud Breakout Settings" on page 550.
- Add cloud breakout settings—See "Adding Cloud Breakout Settings" on page 542.
- Edit cloud breakout settings—See "Editing Breakout Profiles and Cloud Breakout Settings" on page 548.
- Delete cloud breakout settings—See "Deleting Breakout Profiles and Cloud Breakout Settings" on page 550.
- Assign cloud breakout settings to one or more sites—See "Assigning Cloud Breakout Settings to Sites" on page 546.
- Detach cloud breakout settings from one or more sites—See "Detaching Cloud Breakout Settings from Sites" on page 547.

Breakout Profiles Field Descriptions

Table 177: Breakout Profiles Field Descriptions

Field	Description	Displayed On
Name	Name of the breakout profile.	Breakout Profiles page (Breakout Profiles tab) Detail for *Breakout-Profile-Name* pane
Type	Indicates whether the breakout profile is for local breakout (underlay) or backhaul (central breakout) or cloud breakout.	Breakout Profiles page (Breakout Profiles tab) Detail for *Breakout-Profile-Name* pane
Description	Description of the breakout profile.	Breakout Profiles page (Breakout Profiles tab)

Table 177: Breakout Profiles Field Descriptions *(continued)*

Field	Description	Displayed On
Path Preference	Indicates the preferred path to be used for breakout traffic: • MPLS • Internet • Any, which indicates no preference.	Breakout Profiles page (Breakout Profiles tab)
Added by	Username of the user who added the breakout profile.	Breakout Profiles page (Breakout Profiles tab)
FqName	Internal name of the breakout profile.	Breakout Profiles page (Breakout Profiles tab)
Rate Limiting	Indicates whether rate limiting is enabled or disabled for the breakout profile.	Detail for *Breakout-Profile-Name* pane
Downstream Rate	Indicates the maximum downstream rate (in Kbps) for all cacheable applications associated with the breakout profile.	Detail for *Breakout-Profile-Name* pane
Downstream Burst Size	Indicates the maximum downstream burst size (in bytes) for all cacheable applications associated with the breakout profile.	Detail for *Breakout-Profile-Name* pane
Upstream Rate	Indicates the maximum upstream rate (in Kbps) for all cacheable applications associated with the breakout profile.	Detail for *Breakout-Profile-Name* pane
Upstream Burst Size	Indicates the maximum upstream burst size (in bytes) for all cacheable applications associated with the breakout profile.	Detail for *Breakout-Profile-Name* pane
Loss Priority	Indicates the loss priority associated with the breakout profile. The loss priority determines which packets are dropped or retained when network congestion occurs.	Detail for *Breakout-Profile-Name* pane

Cloud Breakout Settings Field Descriptions

Table 178: Cloud Breakout Settings Field Descriptions

Field	Description	Displayed On
General		
Name or Profile-Name	Name of the cloud breakout setting.	Breakout Profiles page (Cloud Breakout Settings tab) Detail for *Cloud-Breakout-Setting-Name* pane
Tunnel Type	Overlay tunnel type (IPSEC or GRE) used to break out traffic to the cloud breakout node.	Breakout Profiles page (Cloud Breakout Settings tab) Detail for *Cloud-Breakout-Setting-Name* pane
Primary Gateway	IPv4 address of the primary cloud breakout node.	Breakout Profiles page (Cloud Breakout Settings tab) Detail for *Cloud-Breakout-Setting-Name* pane
Primary Link Type	Preferred type of WAN link to be used for breaking out the traffic to the primary cloud breakout node. If a WAN link type that matches the preferred path is enabled for breakout, then that WAN link type is used for breakout traffic.	Breakout Profiles page (Cloud Breakout Settings tab) Detail for *Cloud-Breakout-Setting-Name* pane
Secondary Gateway	IPv4 address of the secondary cloud breakout node.	Breakout Profiles page (Cloud Breakout Settings tab) Detail for *Cloud-Breakout-Setting-Name* pane
Secondary Link Type	Preferred type of WAN link to be used for breaking out the traffic to the secondary cloud breakout node.	Breakout Profiles page (Cloud Breakout Settings tab) Detail for *Cloud-Breakout-Setting-Name* pane

Table 178: Cloud Breakout Settings Field Descriptions *(continued)*

Field	Description	Displayed On
Sites	If the cloud breakout settings are assigned to one or more sites, the names of the sites are displayed; if not, this field is blank.	Breakout Profiles page (Cloud Breakout Settings tab)
Provider	Name of the cloud service provider.	Detail for *Cloud-Breakout-Setting-Name* pane
IPSEC Configuration Parameters		
Domain Name	The domain name to generate the fully qualified domain name (FQDN) that is used by the cloud security providers to identify the IPsec tunnel end points. The domain name is populated based on the customer domain name that you provided while onboarding the tenant (**Administration Portal > Tenants > Add Tenant > Tenant Properties > Cloud Breakout Settings**).	Detail for *Cloud-Breakout-Setting-Name* pane
Phase 1 Parameters		
Encryption Type	The encryption type for IPsec proposals.	Detail for *Cloud-Breakout-Setting-Name* pane
Authentication Type	The IPsec authentication algorithm for security association.	Detail for *Cloud-Breakout-Setting-Name* pane
DH Group	The Diffie-Hellman (DH) group.	Detail for *Cloud-Breakout-Setting-Name* pane
Phase 2 Parameters		
Encryption Type	The encryption type for IPsec proposals.	Detail for *Cloud-Breakout-Setting-Name* pane
Authentication Type	The IPsec authentication algorithm for security association.	Detail for *Cloud-Breakout-Setting-Name* pane

RELATED DOCUMENTATION

Breakout and Breakout Profiles Overview | 532

Adding Breakout Profiles

You use the Add Breakout Profile page to add a local breakout (underlay), backhaul, or a cloud breakout profile. A cloud breakout profile is added by Contrail Service Orchestration (CSO) by default.

To add a breakout profile:

1. Select **Configuration > SD-WAN > Breakout Profiles**.

   The Breakout Profiles page appears.

2. On the **Breakout Profiles** tab, click the add icon (**+**).

   The Add Breakout Profile page appears.

3. Complete the configuration according to the guidelines provided in Table 179 on page 541.

   NOTE: Fields marked with an asterisk (*) are mandatory.

4. Click **OK**.

   You are returned to the Breakout Profiles page (Breakout Profiles tab) and a confirmation message indicating that the breakout profile was added is displayed. The page refreshes to display the breakout profile that you added.

   NOTE: After you add a breakout profile, you must add an SD-WAN policy intent that references the breakout profile in order to enable breakout traffic.

Table 179: Fields on the Add Breakout Profile Page

Field	Description
Type	Select the type of breakout profile that you want to add: • **Local Breakout (Underlay)**—Select this option if you want traffic to break out locally (on the underlay) from the site. • **Backhaul**—Select this option if you want traffic to break out through a hub or a enterprise hub (if configured). • **Local Breakout (Cloud)**—Select to break out traffic through a cloud-based security platform. Currently, Zscaler is the only cloud-based security platform supported.
Name	Enter a unique name for the breakout profile. You can use alphanumeric characters and hyphens (-); the maximum length is 15 characters.
Description	Enter a description for the breakout profile.
Traffic Type Profile	Select a traffic type profile to apply class of service parameters to the breakout traffic. You can select only a traffic type profile that is enabled.
Preferred Path	Select the preferred path (MPLS, Internet, or Any) to be used for breaking out the traffic. If a WAN link type that matches the preferred path is enabled for breakout, then that WAN link type is used for breakout traffic. If you specify that any path can be used, then there is no preference and all breakout-enabled links are used in a load-balancing mode.
Advanced Configuration	
Rate Limiting	Click the toggle button to enable rate limiting of breakout traffic for cacheable applications. By default, rate limiting is disabled. If you enable rate limiting, you must specify the upstream and downstream parameters, and the loss priority.
Upstream Rate	Specify the maximum upstream rate (in Kbps) for all cacheable applications associated with the breakout profile.
Upstream Burst Size	Specify the maximum size (in bytes) of a steady stream of traffic sent at average rates that exceed the upstream rate limit for short periods.
Downstream Rate	Specify the maximum downstream rate (in Kbps) for all cacheable applications associated with the breakout profile.

Table 179: Fields on the Add Breakout Profile Page *(continued)*

Field	Description
Downstream Burst Size	Specify the maximum size (in bytes) of a steady stream of traffic sent at average rates that exceed the downstream rate limit for short periods.
Loss Priority	Select a loss priority based on which packets are dropped or retained when network congestion occurs. Packet drops are most likely when the loss priority is High and least likely when the loss priority is Low.

RELATED DOCUMENTATION

Breakout and Breakout Profiles Overview | 532

Creating SD-WAN Policy Intents | 505

Adding Cloud Breakout Settings

You use the Add Cloud Breakout Settings page to add cloud breakout settings that you can then apply to sites.

To add cloud breakout settings:

1. Select **Configuration > SD-WAN > Breakout Profiles**.

   The Breakout Profiles page appears.

2. On the **Cloud Breakout Settings** tab, click the add icon (**+**).

   The Add Cloud Breakout Settings page appears.

3. Complete the configuration according to the guidelines provided in Table 180 on page 543.

   NOTE: Fields marked with an asterisk (*) are mandatory.

4. Click **OK**.

> **NOTE:** If the gateway is unreachable, an error message **Gateway is unreachable. Do you want to proceed with profile creation?** is displayed. If you want to continue with the cloud breakout profile creation, click **Yes**, else click **Cancel**.

You are returned to the Breakout Profiles page (Cloud Breakout Settings tab) and a confirmation message indicating that the breakout settings are added is displayed.

After you add cloud breakout settings, you can assign the settings to one or more sites. Assigning cloud breakout settings to sites provisions the cloud breakout node (Zscaler) overlay. For traffic to flow, you must reference the cloud breakout profile in an SD-WAN policy intent.

Table 180: Fields on the Add Cloud Breakout Settings Page

Field	Description
Name	Enter a unique name for the cloud breakout settings. You can use alphanumeric characters and hyphens (-); the maximum length is 15 characters.
Tunnel Type	Select the type of overlay tunnel (IPSEC or GRE) used to break out the traffic to the cloud breakout node.
IPsec Configuration Parameters	
Domain Name	Displays the domain name that is used to generate the fully qualified domain name (FQDN) for SD-WAN policies. The FQDN is used by the cloud security providers to identify the IPsec tunnels. The domain name is populated based on the customer domain name that you provided while onboarding the tenant (**Administration Portal > Tenants > Add Tenant > Tenant Properties > Cloud Breakout Settings**). Though the domain name is populated automatically, you can modify the domain name.
Phase 1	In Phase 1, the SD-WAN spoke site and the cloud breakout node establish a secure tunnel to negotiate the IPsec security associations (SAs).
Encryption Type	Select an encryption type for IPsec proposals: • **AES-256-CBC (default)**—Advanced Encryption Standard (AES) 256-bit encryption algorithm in Cipher Block Chaining (CBC) mode. • **AES-192-CBC**—AES 192-bit encryption algorithm. • **AES-128-CBC**—AES 128-bit encryption algorithm. • **3DES-CBC**—Triple Data Encryption Algorithm (3DES) in CBC mode. Has a block size of 24 bytes; the key size is 192 bits long.

Table 180: Fields on the Add Cloud Breakout Settings Page *(continued)*

Field	Description
Authentication Type	Select an IPsec authentication algorithm for security association: • **SHA-256 (default)**—Secure Hash Algorithm (SHA) that converts a text of any length into a string of 256 bits. • **SHA-384**—Produces a 384-bit string. • **SHA1**—Produces a 160-bit string.
DH Group	Specify the Diffie-Hellman (DH) group to match the IPsec encryption algorithm: • **GROUP2 (default)**—1024-bit Modular Exponential (MODP) algorithm. • **GROUP5**—1536-bit MODP algorithm. • **GROUP14**—2048-bit MODP algorithm.
Phase 2	In Phase 2, the SD-WAN spoke site and the cloud breakout node negotiate the IPsec SAs for encrypting and authenticating the exchange of data.
Encryption Type	Select an encryption type for IPsec proposals. • **NULL (default)**—No encryption. This is the default. • **AES-256-CBC**—AES 256-bit encryption algorithm. • **AES-192-CBC**—AES 192-bit encryption algorithm. • **AES-128-CBC**—AES 128-bit encryption algorithm.
Authentication Type	Select an IPsec authentication algorithm for security association. • **HMAC-MD5-96 (default)**—Produces a 128-bit digest. This is the default. • **HMAC-SHA-256-128**—Produces a 256-bit digest, truncated to 128 bits. • **HMAC-SHA1-96**—Produces a 160-bit digest.
Protocol	Displays the protocol as ESP (default). Encapsulating Security Payload (ESP) protocol provides a means to ensure privacy (encryption), source authentication and content integrity (authentication). **NOTE:** You cannot edit the protocol.
Primary Gateway	Configuration for the primary cloud breakout node.
Link Type	Select the preferred type of WAN link (MPLS or Internet) to be used for breaking out the traffic to the primary cloud breakout node. If a WAN link type that matches the preferred path is enabled for breakout, then that WAN link type is used for breakout traffic.

Table 180: Fields on the Add Cloud Breakout Settings Page *(continued)*

Field	Description
IP Address/Hostname	Enter the IPv4 address or host name of the primary cloud breakout node. Currently, Zscaler is the only cloud-based security platform supported. The IP address or hostname, is validated. If the IP address or host name is not reachable, the **Host Unreachable** message is displayed.
Preshared Key	Enter the preshared key used for IKE authentication with the primary cloud breakout node. The preshared key is provided by the Zscaler. The key that you enter is masked.
Confirm Preshared Key	Reenter the preshared key for confirmation.
Secondary Gateway	Configuration for the secondary cloud breakout node.
Link Type	Select the preferred type of WAN link (MPLS or Internet) to be used for breaking out the traffic to the secondary cloud breakout node. If a WAN link type that matches the preferred path is enabled for breakout, then that WAN link type is used for breakout traffic.
IP Address/Hostname	Enter the IPv4 address or host name of the secondary cloud breakout node. Currently, Zscaler is the only cloud-based security platform supported. The IP address or hostname, is validated. If the IP address or host name is not reachable, the **Host Unreachable** message is displayed.
Preshared Key	Enter the preshared key used for IKE authentication with the secondary cloud breakout node. The preshared key is provided by the Zscaler. The key that you enter is masked.
Confirm Preshared Key	Reenter the preshared key for confirmation.

RELATED DOCUMENTATION

Breakout and Breakout Profiles Overview | 532

Creating SD-WAN Policy Intents | 505

Assigning Cloud Breakout Settings to Sites

You use the Assign Cloud Breakout Settings to Sites page to assign cloud breakout settings to one or more sites. You assign cloud breakout settings to one or more sites to provision tunnels from the sites to the cloud breakout node. For breakout traffic from the site, the cloud breakout profile must be referenced in an SD-WAN policy intent.

NOTE:

- If you want a site to have cloud breakout enabled, you must assign cloud breakout settings for that site.
- A site can have only one cloud breakout setting associated with it at any given time.

To assign one or more sites to a cloud breakout profile:

1. Select **Configuration > SD-WAN > Breakout Profiles**.

   The Breakout Profiles page appears.

2. On the **Cloud Breakout Settings** tab, select a cloud breakout profile and click **Assign Sites**.

   The Assign Cloud Breakout Settings to Sites page appears displaying the name of the cloud breakout setting and the existing sites to which you can assign the setting. All SD-WAN sites that have local breakout enabled will be displayed in the Available sites column.

3. In the Sites field, select one or more sites in the Available column and click the right arrow icon to move the selected sites to the Selected column. You can also use the search icon on the top right of each column to search for sites names.

   Alternatively, if you want to remove sites that you previously selected for assignment, select one or more sites in the Selected column and click the left arrow icon to move the selected sites back to the Available column.

   NOTE: You must select at least one site before proceeding.

4. Click **Next**.

   The Edit Site Tunnels tab is displayed.

5. Review the configuration and modify the settings, if needed.

   - For IPsec Tunnels, ensure that the format for the FQDN is as follows:

- *Site-name*.primary_link.primary_gateway.1@*Customer-Domain-Name* for the primary gateway primary link
- *Site-name*.backup_link.primary_gateway.1@*Customer-Domain-Name* for the primary gateway backup link
- *Site-name*.primary_link.backup_gateway.1@*Customer-Domain-Name* for the secondary gateway primary link
- *Site-name*.backup_link.backup_gateway.1@*Customer-Domain-Name* for the secondary gateway backup link

Where *Site-Name* is the name of the site (in CSO) for which the breakout is configured and *Customer-Domain-Name* is the name of the customer domain (in CSO) that you added while onboarding the tenant (**Administration Portal > Tenants > Add Tenant > Tenant Properties > Cloud Breakout Settings**).

- For GRE tunnels, ensure that the primary and secondary gateway internal IP prefix is same as provided by the Zscaler.

6. Select the local links (WAN links) to create the tunnel.

7. Select the link mode as Active-Active or Active-Backup. The primary link is always set to active mode and is used to send the traffic. If secondary link is set to active, the CPE device will load balance the traffic on both primary and secondary links. If the secondary link is set to backup, then secondary link will not be used to send traffic unless the primary link fails.

8. Click **OK**.

   A Job is created and you are returned to the Breakout Profiles page (Cloud Breakout Settings tab). After successful completion of the job, the names of the sites to which the settings are assigned are displayed in the Sites column.

RELATED DOCUMENTATION

Breakout and Breakout Profiles Overview | 532
Detaching Cloud Breakout Settings from Sites | 547

Detaching Cloud Breakout Settings from Sites

You must detach the cloud breakout settings from one or more sites (by using the Detach Cloud Breakout From Sites page) before editing or deleting the cloud breakout settings.

To detach one or more sites from a cloud breakout profile:

1. Select **Configuration > SD-WAN > Breakout Profiles**.

   The Breakout Profiles page appears.

2. On the **Cloud Breakout Settings** tab, select a cloud breakout profile and click **Detach Sites**.

   The Detach Cloud Breakout From Sites page appears displaying the name of the cloud breakout setting and the existing sites to which the setting was assigned.

3. In the **Sites** field, select one or more sites in the Available column and click the right arrow icon to move the selected sites to the Selected column. You can also use the search icon on the top right of each column to search for sites.

   Alternatively, if you want to remove the selected sites, select one or more sites in the Selected column and click the left arrow icon to move the selected sites back to the Available column.

   > **NOTE:** You must select at least one site before proceeding to the next step.

4. Click **Save**.

   A job is created and you are returned to the Breakout Profiles page (Cloud Breakout Settings tab). After successful completion of the job, the names of the sites to which the settings are detached are removed from the Sites column of the cloud breakout settings tab.

RELATED DOCUMENTATION

Breakout and Breakout Profiles Overview | 532

Assigning Cloud Breakout Settings to Sites | 546

Editing Breakout Profiles and Cloud Breakout Settings

IN THIS SECTION

- Editing Breakout Profiles | 549
- Editing Cloud Breakout Settings | 549

On the Breakout Profiles page, you can edit breakout profiles and cloud breakout settings that are not assigned to sites.

> **NOTE:** You cannot edit the cloud breakout profile that is automatically created by Contrail Service Orchestration (CSO).

Editing Breakout Profiles

To edit a breakout profile:

> **NOTE:** If you edit a breakout policy that is used in an SD-WAN policy intent, then that SD-WAN policy is marked for redeployment.

1. Select **Configuration > SD-WAN > Breakout Profiles**.

   The Breakout Profiles page appears.

2. On the **Breakout Profiles** tab, select the breakout profile that you want to edit.

3. Click the edit (pencil) icon.

   The Edit Breakout Profile page appears displaying the same fields that are presented when you add a breakout profile. For more information, see “Adding Breakout Profiles” on page 540.

4. Modify the fields as needed.

   > **NOTE:** You can modify only some fields when you are editing a breakout profile

5. Click **OK**.

   You are returned to the Breakout Profiles page. The modifications that you made are saved and a confirmation message is displayed.

Editing Cloud Breakout Settings

Before editing a cloud breakout setting, ensure that the setting is detached from the site. The edit (pencil) icon is disabled for cloud breakout settings that are assigned to sites.

To edit cloud breakout settings that are not assigned to sites:

1. Select **Configuration > SD-WAN > Breakout Profiles**.

   The Breakout Profiles page appears.

2. On the **Cloud Breakout Settings** tab, select the cloud breakout setting that you want to edit.

3. Click the edit (pencil) icon.

   The Edit Cloud Breakout page appears displaying the same fields that are presented when you add cloud breakout settings. For more information, see "Adding Cloud Breakout Settings" on page 542.

4. Modify the fields as needed.

   NOTE: You can modify only some fields when you are editing a breakout profile

5. Click **OK**.

   You are returned to the Breakout Profiles page. The modifications that you made are saved and a confirmation message is displayed.

RELATED DOCUMENTATION

Breakout and Breakout Profiles Overview | 532

About the Breakout Profiles Page | 535

Detaching Cloud Breakout Settings from Sites | 547

Deleting Breakout Profiles and Cloud Breakout Settings

IN THIS SECTION

Deleting Breakout Profiles | 551

Deleting Cloud Breakout Settings | 551

On the Breakout Profiles page, you can delete breakout profiles that are not used in SD-WAN policy intents and cloud breakout settings that are not assigned to sites.

Deleting Breakout Profiles

To delete a breakout profile that is not used in an SD-WAN policy intent:

1. Select **Configuration > SD-WAN > Breakout Profiles**.

   The Breakout Profiles page appears.

2. On the **Breakout Profiles** tab, select the breakout profile that you want to delete.

3. Click the delete (trash can) icon.

   A popup dialog appears asking you to confirm the deletion.

4. Click **Yes**.

   You are returned to the Breakout Profiles page. The selected breakout profile is deleted and a confirmation message is displayed.

Deleting Cloud Breakout Settings

Before deleting a cloud breakout setting, ensure that the setting is detached from the site. The delete (trash can) icon is disabled for cloud breakout settings that are assigned to sites.

To delete cloud breakout settings that are not assigned to sites:

1. Select **Configuration > SD-WAN > Breakout Profiles**.

   The Breakout Profiles page appears.

2. On the **Cloud Breakout Settings** tab, select the cloud breakout setting that you want to delete.

3. Click the delete (trash can) icon.

   A popup dialog appears asking you to confirm the deletion.

4. Click **Yes**.

   You are returned to the Breakout Profiles page. The selected cloud breakout setting is deleted and a confirmation message is displayed.

RELATED DOCUMENTATION

Breakout and Breakout Profiles Overview | 532

About the Breakout Profiles Page | 535

Detaching Cloud Breakout Settings from Sites | 547

Configuring Breakout on SD-WAN Sites

The following is the workflow for configuring breakout (local breakout [underlay], backhaul [central breakout], or cloud breakout):

1. *Before* configuring breakout, ensure that you complete the following tasks:

   a. If you are using enterprise hub sites, add, configure, and activate one or more enterprise hub sites. See "Adding Enterprise Hubs with SD-WAN Capability or SD-WAN and LAN Capabilities" on page 61.

   b. Add, configure, and activate one or more on-premise spoke sites with SD-WAN capability. See "Adding an On-Premise Spoke Site with SD-WAN Capability" on page 100.

      NOTE: You must attach an on-premise spoke site with SDWAN capability to a provider hub site or an enterprise hub site, or to both hub sites.

   c. (Optional) If you are using application-based breakout, ensure that you install the application ID license (if it is required for the device) and signatures on the devices (associated with the sites).

2. Depending on the type of breakout you are configuring, add one or more breakout profiles for the following types of breakout:

   - Local breakout (underlay)
   - Backhaul (central breakout)
   - Cloud breakout

   See "Adding Breakout Profiles" on page 540.

3. For cloud breakout, add cloud breakout settings and then assign the cloud breakout settings to one or more on-premise spoke or enterprise hub sites. See "Adding Cloud Breakout Settings" on page 542 and "Assigning Cloud Breakout Settings to Sites" on page 546.

4. Add one or more SD-WAN policy intents in which you reference the previously-added breakout profiles. See “Creating SD-WAN Policy Intents” on page 505.

5. Deploy the SD-WAN policy. See “Deploying Policies” on page 668.

6. Configure firewall policy intents to allow Internet-bound traffic from the sites or departments for which you configured breakout (through the SD-WAN policy intent). See “Adding Firewall Policy Intents” on page 381.

7. Deploy the firewall policy. See “Deploying Policies” on page 668.

8. For cloud breakout using Zscaler, ensure that the user IDs in the Zscaler account are configured as follows:

   - *Site-Name*.primary.1@*Tenant-Name*.com for the primary tunnel
   - *Site-Name*.backup.1@*Tenant-Name*.com for the secondary tunnel

   Where *Site-Name* is the name of the site (in CSO) for which the breakout is configured and *Tenant-Name* is the name of the tenant (in CSO) to which the site belongs.

RELATED DOCUMENTATION

Breakout and Breakout Profiles Overview | 532

CHAPTER 18

Managing NAT Policies

IN THIS CHAPTER

- NAT Policies Overview | **555**
- About the NAT Policies Page | **558**
- Creating NAT Policies | **559**
- Editing and Deleting NAT Policies | **561**
- About the Single NAT Policy Page | **562**
- Creating NAT Policy Rules | **564**
- Editing, Cloning, and Deleting NAT Policy Rules | **571**
- Deploying NAT Policy Rules | **573**
- Selecting NAT Source | **574**
- Selecting NAT Destination | **578**
- NAT Pools Overview | **582**
- About the NAT Pools Page | **582**
- Creating NAT Pools | **584**
- Editing, Cloning, and Deleting NAT Pools | **586**
- Deploying NAT Policies | **588**
- Importing NAT Policies | **588**

NAT Policies Overview

Network Address Translation (NAT) is a form of network masquerading where you can hide devices or sites between zones or interfaces. A trusted zone is a segment of a network on which security measures are applied. It is usually assigned to the internal LAN. An example of an untrusted zone is the internet. NAT modifies the IP addresses of the packets moving between the trusted and untrusted zones.

Whenever a packet exits a NAT device (when traversing from the internal LAN to the external WAN), the device performs a translation on the packet's IP address by rewriting it with an IP address that was specified for external use. After translation, the packet appears to have originated from the gateway rather than from the original device within the network. This process hides your internal IP addresses from the other networks and keeps your network secure.

Using NAT also enables you to use more internal IP addresses. As these IP addresses are hidden, there is no risk of conflict with an IP address from a different network. This helps you conserve IP addresses.

CSO supports three types of NAT:

- Source NAT— Translates the source IP address of a packet leaving a trust zone (outbound traffic). It translates the traffic originating from the device in the trust zone. The source IP address of the traffic (which is a private IP address), is translated to a public IP address that can be accessed by the destination device specified in the NAT rule. The destination IP address is not translated.

  The following uses cases show the support for source NAT translation between IPv6 and IPv4 address domains:

  - Translation from one IPv6 subnet to another IPv6 subnet without Network Address Port Translation (NAPT), also known as Port Address Translation (PAT).
  - Translation from IPv4 addresses to IPv6 prefixes along with IPv4 address translation.
  - Translation from IPv6 hosts to IPv6 hosts with or without NAPT.
  - Translation from IPv6 hosts to IPv4 hosts with or without NAPT.
  - Translation from IPv4 hosts to IPv6 hosts with or without NAPT.

- Destination NAT—Translates the destination IP address of a packet. Using destination NAT, an external device can send packets to a hidden internal device. As an example, consider the case of a webserver behind a NAT device. Traffic to the WAN-facing public IP address (the destination IP address) is translated to the internal webserver private IP address.

  The following uses cases show the support for destination NAT translation between IPv6 and IPv4 address domains:

  - Mapping of one IPv6 subnet to another IPv6 subnet
  - Mapping between one IPv6 host and another IPv6 host

  - Mapping of one IPv6 host (and optional port number) to another special IPv6 host (and optional port number)
  - Mapping of one IPv6 host (and optional port number) to another special IPv4 host (and optional port number)
  - Mapping of one IPv4 host (and optional port number) to another special IPv6 host (and optional port number)

- Static NAT— Always translates a private IP address to the same public IP address. It translates traffic from both sides of the network (both source and destination). For example, a web-server with a private IP address can access the Internet using a static, one-to-one address translation. In this case, outgoing traffic from the web-server undergoes source NAT translation, and incoming traffic to the web-server undergoes destination NAT translation.

  The following uses cases show the support for static NAT translation between IPv6 and IPv4 address domains:

  - Mapping of one IPv6 subnet to another IPv6 subnet.
  - Mapping between one IPv6 host and another IPv6 host.
  - Mapping between IPv4 address *a.b.c.d* and IPv6 address *Prefix::a.b.c.d*.
  - Mapping between IPv4 hosts and IPv6 hosts.
  - Mapping between IPv6 hosts and IPv4 hosts.

CSO also supports persistent NAT where address translations are maintained in the database for a configurable amount of time after a session ends.

Table 181 on page 556 shows the persistent NAT support for different source NAT and destination NAT addresses.

Table 181: Persistent NAT Support

Source NAT Address	Translated Address	Destination NAT Address	Persistent NAT
IPv4	IPv6	IPv4	No
IPv4	IPv6	IPv6	No
IPv6	IPv4	IPv4	Yes
IPv6	IPv6	IPv6	No

Table 182 on page 557 and Table 183 on page 557 show the translated address pool selection for source NAT, destination NAT, and static NAT addresses.

Table 182: Translated Address Pool Selection for Source NAT

Source NAT Address	Destination Address	Pool Address
IPv4	IPv4	IPv4
IPv4	IPv6 - Subnet must be greater than 96	IPv6
IPv6	IPv4	IPv4
IPv6	IPv6	IPv6

Table 183: Translated Address Pool Selection for Destination NAT And Static NAT

Source NAT Address	Destination Address	Pool Address
IPv4	IPv4	IPv4 or IPv6
IPv4	IPv6 - Subnet must be greater than 96	IPv4 or IPv6
IPv6	IPv4	IPv4
IPv6	IPv6	IPv4 or IPv6

NOTE:

- For source NAT, the proxy Neighbor Discovery Protocol (NDP) is available for NAT pool addresses. For destination NAT and static NAT, the proxy NDP is available for destination NAT addresses.
- A NAT pool can have a single IPv6 subnet or multiple IPv6 hosts.
- You cannot configure the overflow pool if the address type is IPv6.
- NAT pools permit address entries of only one version type: IPv4 or IPv6.

RELATED DOCUMENTATION

About the NAT Policies Page | **558**

Creating NAT Policies | **559**

Editing and Deleting NAT Policies | **561**

Editing, Cloning, and Deleting NAT Policy Rules | **571**

About the NAT Policies Page

To access this page, select **Configuration** > **NAT** > **NAT Policies**.

Use the **NAT Policies** page to create, modify, clone, and delete NAT policies and policy rules. You can filter and sort this information to get a better understanding of what you want to configure.

Tasks You Can Perform

You can perform the following tasks from this page:

- Create a NAT policy. See "Creating NAT Policies" on page 559.
- Modify or delete a NAT policy. See "Editing and Deleting NAT Policies" on page 561.
- Create, modify, clone, and delete NAT policy rules. See "About the Single NAT Policy Page" on page 562.
- Search for a specific NAT policy. Click the Search icon in the top right corner of the page to search for a NAT policy.

  You can enter partial text or full text of the keyword in the text box and press Enter.

  The search results are displayed on the same page.
- Show or hide columns. Click the **Show Hide Columns** icon in the top right corner of the page.

Field Descriptions

Table 184 on page 558 provides guidelines on using the fields on the **NAT Policies** page.

Table 184: Fields on the NAT Policies Page

Field	Description
Name	Displays the name of the NAT policy.
Installed On	Displays the sites on which the NAT policy is assigned.
Rules	Number of rules assigned to the NAT policy.
Undeployed	Number of undeployed rules associated with the NAT policy.

RELATED DOCUMENTATION

NAT Policies Overview | 555

Creating NAT Policies | 559
Editing and Deleting NAT Policies | 561
About the Single NAT Policy Page | 562

Creating NAT Policies

Use the **Create NAT Policy** page to create NAT policies.

To create a NAT policy:

1. Select **Configuration > NAT > NAT Policies**.

   The **NAT Policies** page appears.

2. Click the add icon (+).

   The **Create NAT Policy** page displays fields required for creating and configuring a NAT policy.

3. Complete the configuration according to the guidelines provided in Table 185 on page 559.

   NOTE: You can associate only a single device or a device cluster with a site.

4. Click **OK** to save the changes.

   A NAT policy with the configuration you provided is created.

Table 185: Fields on the Create NAT Policy Page

Field	Description
Name	Enter a unique string of alphanumeric characters, colons, periods, dashes, and underscores. No spaces are allowed and the maximum length is 255 characters.
Description	Enter a description for the policy; the maximum length is 1024 characters.

Table 185: Fields on the Create NAT Policy Page *(continued)*

Field	Description
Manage Auto-Proxy ARP	The Address Resolution Protocol (ARP) protocol translates IPv4 addresses to MAC addresses. Typically, an interface responds with its MAC address only when an ARP request for its IP address is received. A proxy ARP implies that the same interface will proxy for other IP addresses (that is, respond to ARP requests for other IP addresses). Managing a proxy ARP automatically enables the selection of an appropriate interface for any address (as part of a NAT rule) that is not an actual interface address. Proxy ARP management applies to translated addresses in a source NAT rule or to a destination address in a destination NAT rule. **NOTE:** When creating a source NAT rule with pool translation, the address pool assigned must be in the same subnet as the outgoing interface selected. **NOTE:** When creating a destination NAT rule, the external WAN interface can be a proxy for another IP address in the same subnet as the original IP address of the interface.
Sites Applied On	Select the sites on which you want to apply the policy in the **Available** column and move them to the **Selected** column by clicking the greater-than icon (>). **NOTE:** The **Available** column lists only those sites that do not have a NAT policy associated with them.
Sequence No.	Click **Select Policy Sequence**. The **Select Policy Sequence** page appears, displaying all NAT policies. Select the policy you want to reorder and select **Move Policy Up** or **Move Policy Down** to reorder your NAT policy among the existing policies.

RELATED DOCUMENTATION

- NAT Policies Overview | 555
- About the NAT Policies Page | 558
- Editing and Deleting NAT Policies | 561
- About the Single NAT Policy Page | 562
- Creating NAT Policy Rules | 564
- Editing, Cloning, and Deleting NAT Policy Rules | 571

Editing and Deleting NAT Policies

IN THIS SECTION

- Editing NAT Policies | 561
- Deleting NAT Policies | 561

You can edit or delete a NAT policy from the **NAT Policies** page.

Editing NAT Policies

To modify the parameters configured for a NAT Policy:

1. Select **Configuration > NAT > NAT Policies**.

   The **NAT Policies** page appears.

2. Hover over the NAT policy you want to edit, and then click on the edit icon (pencil symbol) on the right side of the table.

   The **Edit NAT Policy** page appears, showing the same fields as those seen when you create a new NAT policy.

3. Modify the parameters according to the guidelines provided in "Creating NAT Policies" on page 559.

4. Click **OK** to save the changes. If you want to discard your changes, click **Cancel** instead.

If you click **OK**, you will see the modified NAT policy in the **NAT Policies** page.

Deleting NAT Policies

To delete a NAT policy:

1. Select **Configuration > NAT > NAT Policies**.

   The **NAT Policies** page appears.

2. Hover over the NAT policy you want to delete and then click the delete icon **(X)**.

   An alert message appears, verifying that you want to delete your selection.

3. Click **Yes** to delete the selection. If you do not want to delete, click **Cancel** instead.

If you click **Yes**, the NAT policy is deleted.

> **NOTE:** When the NAT policy is deleted, the NAT rules associated with the policy are deleted from device.

RELATED DOCUMENTATION

- NAT Policies Overview | 555
- About the NAT Policies Page | 558
- Creating NAT Policies | 559
- Editing, Cloning, and Deleting NAT Policy Rules | 571

About the Single NAT Policy Page

To access this page, select **Configuration** > **NAT** > **NAT Policies**. The **NAT Policies** page appears displaying all existing NAT policies. Click on a NAT policy to view the rules associated with it.

The *Single NAT Policy* page displays the NAT rules associated with the NAT policy, and keep track of the number and order of rules for each policy. You can also create a new NAT rule, modify the configured parameters of existing NAT rules, clone, and delete NAT rules, using this page.

Tasks You Can Perform

You can perform the following tasks from this page:

- Create a NAT rule. See "Creating NAT Policy Rules" on page 564.
- Update the sequence of the NAT rules using the up and down arrows that appear when you hover over the NAT rule.
- Modify, clone, and delete NAT rules. See "Editing, Cloning, and Deleting NAT Policy Rules" on page 571.
- Deploy a NAT rule. See "Deploying NAT Policy Rules" on page 573.
- Search for a specific NAT rule. Click the Search icon in the top right corner of the page to search for a NAT rule.
- Show or hide columns. Click the **Show Hide Columns** icon in the top right corner of the page.

Field Descriptions

Table 186 on page 563 provides information on the fields in the NAT rules contained within this NAT policy.

Table 186: Fields on the Single NAT Policy Page

Field	Description
Source	Displays the source endpoint on which the NAT policy applies. A source endpoint can be an address, protocol, interface, routing instance, zone, or port.
Destination	Displays the destination endpoint on which the NAT policy applies. A destination endpoint can be an address, interface, service, routing instance, zone, or port.
Translation	Displays the translation type applied on the incoming or outgoing traffic.
Details	Displays the type of NAT rule. A NAT rule can be of type source, static, or destination.

The **Total Rules** field on the top right corner of the page displays the total number of rules associated with the NAT policy. The **Undeployed** field displays the number of undeployed rules associated with the NAT policy. To deploy undeployed rules, click **Deploy**. See "Deploying NAT Policy Rules" on page 573.

RELATED DOCUMENTATION

- NAT Policies Overview | 555
- About the NAT Policies Page | 558
- Creating NAT Policies | 559
- Editing and Deleting NAT Policies | 561
- Creating NAT Policy Rules | 564
- Editing, Cloning, and Deleting NAT Policy Rules | 571
- Deploying NAT Policy Rules | 573

Creating NAT Policy Rules

NAT processing centers on the evaluation of NAT rule sets and rules. A rule set determines the overall direction of the traffic to be processed. After a rule set that matches the traffic is found, each rule in the rule set is evaluated for a match. NAT rules can match on the following packet information:

- Source and destination address
- Source port (for source and static NAT only)
- Destination port

The first rule in the rule set that matches the traffic is used. If a packet matches a rule in a rule set during session establishment, traffic is processed according to the action specified by that rule.

To create a new NAT rule, click the NAT policy name. The *Single NAT Policy* page appears, providing your with options to configure NAT rules. Alternately, you can click on the rule number listed under **Rules** against the policy, to create a new rule. You can configure the following types of NAT rules:

- Static—To add a static NAT rule, click **Add Static NAT Rule** or click **Create** on the top right corner and select **Static**.
- Source—To add a source NAT rule, click **Add Source NAT Rule** or click **Create** on the top right corner and select **Source**.
- Destination—To add a destination NAT rule, click **Add Destination NAT Rule** or click **Create** on the top right corner and select **Destination**.

Depending on the type of rule you have chosen, some fields in the rule will not be applicable. In addition to defining rules between zones and interfaces, you can define NAT rules with virtual routers defined on the device. These rules can be successfully published and updated on the device.

To create a NAT policy rule:

1. Select **Configuration > NAT > NAT Policies**.

   The **NAT Policies** page appears, displayed the existing NAT policies.

2. Click the name of the NAT policy for which you want to create rules. Alternately, you can click on the number listed under **Rules** against a NAT policy.

   The ***Single NAT Policy*** page appears.

3. Click **Create** and select either **Source**, **Static**, or **Destination**. The page displays fields for creating a NAT rule.

4. Complete the configuration according to the guidelines provided in Table 187 on page 565.

5. Click **OK** to save the changes. If you want to discard your changes, click **Cancel** instead.

A NAT rule with the configuration you provided is created.

Table 187 on page 565 provides guidelines on using the fields on the **Single NAT Policy** page.

Table 187: Fields on the Single NAT Policy Page for Creating NAT Rules

Field	Description

Source	Click the add icon (+) to select the source endpoints on which the NAT policy rule applies, from the displayed list of addresses, protocols, interfaces, routing instances, zones, or ports.
	
	The possible endpoints for source differ based on whether the NAT rule is a source, destination, or static NAT rule.
	
	• The possible endpoints for source for a source NAT rule are:
	  ◦ Addresses
	  ◦ Routing instances, interfaces, or zones
	  ◦ Protocols
	  ◦ Ports
	• The possible endpoints for source for a destination NAT rule are:
	  ◦ Addresses
	  ◦ Routing instances, interfaces, or zones
	  ◦ Protocols
	• The possible endpoints for source for a static NAT rule are:
	  ◦ Addresses
	  ◦ Routing instances, interfaces, or zones
	  ◦ Ports
	
	You can also select a source endpoint by using the methods described in "Selecting NAT Source" on page 574.

Table 187: Fields on the Single NAT Policy Page for Creating NAT Rules *(continued)*

Field	Description

Destination	Click the add icon (+) to select the destination endpoints on which the NAT policy rule applies, from the displayed list of addresses, interfaces, services, routing instances, zones, or ports.
	
	The possible endpoints for destination differ based on whether the NAT rule is a source, destination, or static NAT rule.
	
	• The possible endpoints for destination for a source NAT rule are:
	  ◦ Addresses
	  ◦ Routing instances, interfaces, or zones
	  ◦ Services
	  ◦ Ports
	• The possible endpoints for destination for a destination NAT rule are:
	  ◦ Addresses
	  ◦ Services
	  ◦ Ports
	• The possible endpoints for destination for a static NAT rule are:
	  ◦ Addresses
	  ◦ Ports
	
	You can select a destination endpoint by using the methods described in "Selecting NAT Destination" on page 578.
	
	NOTE: When you create a destination NAT rule for traffic arriving on an interface that terminates a VPN link, the translation process may break the VPN link. This will happen if the destination address in a destination NAT rule is specified only as the WAN-facing IP address of that interface. For example, in the following NAT rule, any traffic destined to Wan.IP will get translated to the destination pool and will break functionality of the VPN link packets terminating on this interface.
	
	[Any.Address] --> [Wan.IP] :: [Dest-Pool-1]
	
	Therefore, the recommendation in such cases is to use a destination NAT rule with destination field as **[Address + Port]**. For example:
	
	[Any.Address] --> [Wan.IP + Port] :: [Dest-Pool-1]

Translation	

Table 187: Fields on the Single NAT Policy Page for Creating NAT Rules *(continued)*

Field	Description
Translation Type	Specify the translation type for the incoming traffic. The translation options vary based on whether you are creating a source, static, or destination NAT rule. Chose one among the following translation types for a source NAT rule: • None—No translation is required for the incoming traffic. • Interface—Performs interface-based translations on the source or destination packet. • Pool—Performs pool-based translations on the source or destination packet. Click on the add icon (+) in the **Select Pool** field to choose the translation pool. You can also create a new pool by clicking **Add new pool**. See "Creating NAT Pools" on page 584. Chose one among the following translation types for a static NAT rule: • Address—Performs address-based translations on the source or destination packet. Click on the add icon (+) in the **Select Address** field to choose the translation address. You can also create a new address by clicking **Add new address**. See "Creating Addresses or Address Groups" on page 737. **NOTE:** In an SD-WAN environment, it is mandatory that you select the routing instance corresponding to the translation address. You can select the routing instance for a translation address using the **Advanced Settings** page. For more information on **Advanced Settings**, see Table 189 on page 570. • Corresponding IPv4—Uses the corresponding IPv4 address to perform translations on the source or destination packet. Chose one among the following translation types for a destination NAT rule: • None—No translation is required for the incoming traffic. • Pool—Performs pool-based translations on the source or destination packet. Click on the add icon (+) in the **Select Pool** field to choose the translation pool. You can also create a new pool by clicking **Add new pool**. See "Creating NAT Pools" on page 584. **NOTE:** In an SD-WAN environment, the destination NAT pool selected should be configured with a site and a routing instance corresponding to the pool address. For example, a webserver with IP address (IP1) is running in the HR department. To create a destination NAT pool corresponding to this webserver IP address, you must specify the following mandatory fields while creating the NAT pool: **Address - IP1** **Site - the site hosting the webserver** **Routing instance - natVR_HR**

Table 187: Fields on the Single NAT Policy Page for Creating NAT Rules *(continued)*

Field	Description
Advanced Settings (Optional)	Click **Configure** to configure advance settings for a source or static NAT rule. For more information about advanced settings for the translation types **Interface** and **Pool** for a source NAT rule, see Table 188 on page 568. For more information about advanced settings for the translation types **Interface** and **Pool** for a static NAT rule, see Table 189 on page 570
Details	
Name	Enter a unique string of alphanumeric characters, colons, periods, dashes, and underscores. No spaces are allowed and the maximum length is 255 characters.
Description	Enter a description for the policy intent; maximum length is 1024 characters.
End Points	Create source and destination endpoints such as addresses and services. • To create an address, click the add icon (+) and select **Address**. See "Creating Addresses or Address Groups" on page 737 to configure the parameters of the address. • To create a service, click the add icon (+) and select **Service**. See "Creating Services and Service Groups" on page 744 to configure the parameters of the service. To edit the configured parameters of an address or service, hover over it and click on the edit icon (pencil symbol).

Table 188 on page 568 provides guidelines on using the fields on the **Advanced Settings** page for a source NAT rule.

Table 188: Fields on the Advanced Settings Page for Source NAT Rule

Field	Description

Persistent	Enable the check box to ensure that all requests from the same internal transport address are mapped to the same reflexive transport address.
	NOTE: For persistence to be applicable for the NAT policy, ensure that port overloading is turned off for the device to which the NAT policy is applicable. Use the following command to turn off port overloading for a device:
	`[Edit mode]`
	`set security nat source interface port-overloading off`

Table 188: Fields on the Advanced Settings Page for Source NAT Rule *(continued)*

Field	Description
Persistent NAT Type	Configure persistent NAT mappings. • Permit any remote host—All requests from a specific internal IP address and port are mapped to the same reflexive transport address. (The reflexive transport address is the public IP address and port created by the NAT device closest to the STUN server.) Any external host can send a packet to the internal host by sending the packet to the reflexive transport address. • Permit target host—All requests from a specific internal IP address and port are mapped to the same reflexive transport address. An external host can send a packet to an internal host by sending the packet to the reflexive transport address. The internal host must have previously sent a packet to the external host's IP address. • Permit target host port—All requests from a specific internal IP address and port are mapped to the same reflexive transport address. An external host can send a packet to an internal host by sending the packet to the reflexive transport address. The internal host must have previously sent a packet to the external host's IP address and port.
Inactivity Timeout	The amount of time, in seconds, that the persistent NAT binding remains in the site's memory when all the sessions of the binding entry have ended. When the configured timeout is reached, the binding is removed from memory. The value of the inactivity timeout can range from 60 through 7200 seconds. The default value of the inactivity timeout is 60 seconds.
Maximum Session Number	Maximum session number—The maximum number of sessions with which a persistent NAT binding can be associated. For example, if the maximum session number of the persistent NAT rule is 65,536, then a 65,537th session cannot be established if that session uses the persistent NAT binding created from the persistent NAT rule. The range is 8 through 65,536. The default is 30 sessions.
Address Mapping	Select an address from the available list.
Pool Address	Displays the NAT pool address.
Host Address Base	Displays the base address of the original source IP address range. The host address base is used for IP address shifting.
Port Translation	Displays whether port translation is enabled or disabled for this NAT rule.
Overflow Pool Type	Displays the source pool to be used when the current address pool is exhausted.
Overflow Pool Name	Displays the name of the overflow pool.

Table 188: Fields on the Advanced Settings Page for Source NAT Rule *(continued)*

Field	Description
Mapped Port Type	Specify the type of port mapping: • Port—Enter a value for **Port**, ranging from 0 through 65,535. • Range—Enter the port range values in the **Start** and **End** fields, ranging from 0 through 65,535.

Table 189 on page 570 provides guidelines on using the fields on the **Advanced Settings** page for a static NAT rule.

Table 189: Fields on the Advanced Settings Page for Static NAT Rule

Field	Description
Mapped Port Type	Specify the type of port mapping: • Port—Enter a value for **Port**, ranging from 0 through 65,535. • Range—Enter the port range values in the **Start** and **End** fields, ranging from 0 through 65,535.
Routing Instance	Select the routing instance for the static NAT rule.

RELATED DOCUMENTATION

- About the Single NAT Policy Page | **562**
- Editing, Cloning, and Deleting NAT Policy Rules | **571**
- Deploying NAT Policy Rules | **573**
- NAT Policies Overview | **555**
- About the NAT Policies Page | **558**
- Creating NAT Policies | **559**
- Editing and Deleting NAT Policies | **561**

Editing, Cloning, and Deleting NAT Policy Rules

IN THIS SECTION

- Editing NAT Policy Rules | 571
- Cloning NAT Policy Rules | 571
- Deleting NAT Policy Rules | 572

You can edit, clone, or delete a NAT policy rule from the ***NAT Policy*** page.

Editing NAT Policy Rules

To modify the parameters configured for an NAT policy rule:

1. Select **Configuration > NAT > NAT Policies**.

   The **NAT Policies** page appears, displaying the NAT policies.

2. Select the NAT policy whose rules you want to edit.

   The selected ***NAT Policy*** appears, displaying the rules associated with the NAT policy.

3. Hover over the NAT policy rule that you want to modify and click on the edit icon (pencil symbol) that appears on the right side of the NAT policy rule. The page changes to display the same fields that you use to create a NAT policy rule.

4. Complete the configuration according to the guidelines provided in "Creating NAT Policy Rules" on page 564.

5. Click **OK** to save the changes. If you want to discard your changes, click **Cancel** instead.

The modified NAT policy rule appears on the ***NAT Policy*** page.

Cloning NAT Policy Rules

To clone a NAT policy rule:

1. Select **Configuration > NAT > NAT Policies**.

   The **NAT Policies** page appears, displaying the NAT policies.

2. Select the NAT policy whose rule you want to clone.

   The selected ***NAT Policy*** appears, displaying the rules associated with the NAT policy.

3. Hover over the NAT policy rule that you want to clone and click on the clone icon that appears on the right side of the NAT policy rule.

   The cloned NAT policy rule appears below the current rule.

You can modify the parameters configured for the cloned NAT policy rule or rename it as required.

Deleting NAT Policy Rules

To delete a NAT policy rule:

1. Select **Configuration > NAT > NAT Policies**.

   The **NAT Policies** page appears, displaying the NAT policies.

2. Select the NAT policy whose rule you want to delete.

   The selected ***NAT Policy*** appears, displaying the rules associated with the NAT policy.

3. Hover over the NAT policy rule you want to delete and then click the delete icon **(X)** .

   An alert message appears, verifying that you want to delete your selection.

4. Click **Yes** to delete the selection. If you do not want to delete, click **Cancel** instead.

If you click **Yes**, the selected NAT policy rule is deleted.

RELATED DOCUMENTATION

- About the Single NAT Policy Page | **562**
- Creating NAT Policy Rules | **564**
- Deploying NAT Policy Rules | **573**
- NAT Policies Overview | **555**
- About the NAT Policies Page | **558**
- Creating NAT Policies | **559**
- Editing and Deleting NAT Policies | **561**

Deploying NAT Policy Rules

To deploy an NAT policy rule:

1. Select **Configuration > NAT Policy > Policies**.

2. Click on the name of the NAT policy rules displayed.

   The NAT policy rule page appears.

3. Click **Deploy**.

   The **Deploy** page appears.

4. Configure your deployment as required. See "Deploying Policies" on page 668.

   All the NAT policy rules associated with the NAT policy are deployed. That is, the entire NAT policy is deployed.

   > **NOTE:** By default, all the NAT policy rules associated with the NAT policy (the entire NAT policy) are deployed when you click **Deploy**. Suppose you select a particular NAT policy rule and click **Deploy**, even then, all the NAT policy rules associated with that NAT policy are deployed.

RELATED DOCUMENTATION

- About the Single NAT Policy Page | 562
- Creating NAT Policy Rules | 564
- Editing, Cloning, and Deleting NAT Policy Rules | 571
- NAT Policies Overview | 555
- About the NAT Policies Page | 558
- Creating NAT Policies | 559
- Editing and Deleting NAT Policies | 561

Selecting NAT Source

IN THIS SECTION

- Adding an Endpoint as NAT Source | 574
- Selecting Interfaces when GWR Resides Inside an NFX Box | 574
- Selecting NAT Source Using Abbreviations | 575
- Selecting a NAT Source from the End Points Panel | 576
- Creating and Selecting a NAT Source from the End Points Panel | 576
- Creating Addresses from Source Field | 577

The following procedures provides various methods using which you can choose an endpoint as a NAT source:

Adding an Endpoint as NAT Source

View and select the source endpoint from the complete list of addresses, protocols, interfaces, zones, routing instances, or ports.

1. Click the **Source** field. A list of relevant endpoints are displayed.

2. Click the **View more results** link provided at the bottom of the source endpoints. The complete list of addresses, protocols, interfaces, and ports is displayed in the **End Points** panel on the right.

3. (Optional) Click the edit icon to edit the address, protocol, interface, zones, routing instances, or port endpoint.

4. Click check mark icon (√) to select the endpoint as a source.

Selecting Interfaces when GWR Resides Inside an NFX Box

The physical interfaces of an NFX box are mapped to the virtual interfaces of the Gateway Router (GWR) (vSRX) as given in Table 190 on page 575. These are the default mappings provided by CSO. You may change these interface mappings based on your requirements.

Table 190: NFX and GWR Interface Mapping

NFX Physical Interface	GWR Virtual Interface
WAN 0 (ge-0/0/10)	ge-0/0/2
WAN 1 (ge-0/0/11)	ge-0/0/3
WAN 2 (xe-0/0/12)	ge-0/0/7
WAN 3 (xe-0/0/13)	ge-0/0/8
LAN-X (ge-0/0/X)	Ge-0/0/06.<vlan-id-for-X>

When you create a new NAT rule and an NFX physical interface is intended as the source endpoint, select the respective mapped GWR interface.

Selecting NAT Source Using Abbreviations

Enter an abbreviation in the **Source** field to select the source endpoint from a filtered list of source endpoints.

- To view a filtered list of addresses, enter **ADDR** or **addr**.
- To view a filtered list of protocols, enter **PROT** or **prot**.
- To view a filtered list of interfaces, enter **INTR** or **intr**.
- To view a filtered list of zones, enter **ZONE** or **zone**.
- To view a filtered list of routing instances, enter **ROUT** or **rout**.

Click the endpoints in the filtered list to select them.

You can add a port number as a source endpoint. To do so:

1. Type **PORT** or **port** in the **Source** field.
2. Press Tab.
3. Enter the port number and press Enter.

   You can also enter a range of ports by using the separator **-**. For example, you can enter **10-20**.

The entered port value is selected as a source endpoint.

You can also select the endpoint from the complete list of addresses, protocols, interfaces, zones, and routing instances. See “Adding an Endpoint as NAT Source” on page 574.

Selecting a NAT Source from the End Points Panel

You can select a NAT source endpoint from the **End Points** panel. Alternately, you can create a new NAT source endpoint from the **End Points** panel, see "Creating and Selecting a NAT Source from the End Points Panel" on page 576.

To select an NAT source endpoint from the **End Points** panel:

1. Click the **Source** field.

2. Click the lesser-than icon (**<**) on the right.

   The **End Points** panel appears, displaying the list of available addresses, interfaces, protocols, zones, and routing instances.

3. (Optional) To view more information about a source endpoint, click the details icon on the right of the endpoint. To edit the source endpoint, click the edit icon (pencil symbol) on the right of the endpoint.

   > **NOTE:** You can only edit or view details of a source endpoint if these options appear on right side of the endpoint when you hover over it. Not all endpoints provide these options.

4. Click the check mark icon (**√**) to add the endpoint as a source.

Creating and Selecting a NAT Source from the End Points Panel

To create a new source endpoint from the **End Points** panel:

1. Click the add icon (**+**) on the top right of the panel and select the type of endpoint you want to create, among the options provided.

   Based on the option you select, the respective page appears. Fill in the required details to create a new endpoint.

   - To create a new address, see "Creating Addresses or Address Groups" on page 737.
   - To create a new service, see "Creating Services and Service Groups" on page 744.
   - To create a new NAT pool, see "Creating NAT Pools" on page 584.

   After the endpoint is created, it appears in the **Endpoints** panel.

2. Click the check mark icon (**√**) to add the new endpoint as a source.

Creating Addresses from Source Field

You can use one of the following ways to create a new address from the **Source** field and use the newly created address as a source endpoint:

- Type the address directly in the **Source** field. If the address is valid, it is created immediately and added as a source endpoint.
- Create an address from the **Source** field, using the following steps:

  1. In the **Source** field, type **addr**. The **Add new address** link appears at the bottom of the list of addresses.

  2. Click **Add new address** to create a new address.

     The **Create Addresses** page appears.

  3. Configure the new address. See "Creating Addresses or Address Groups" on page 737.

  4. Click **Save** to save the new address.

     The new address is created, and will be listed as an option for the source. Select the new address to add it to the source.

RELATED DOCUMENTATION

Selecting NAT Destination | 578

Creating NAT Policy Rules | 564

Editing, Cloning, and Deleting NAT Policy Rules | 571

Deploying NAT Policy Rules | 573

About the Single NAT Policy Page | 562

NAT Policies Overview | 555

About the NAT Policies Page | 558

Creating NAT Policies | 559

Editing and Deleting NAT Policies | 561

Selecting NAT Destination

IN THIS SECTION

- Adding an Endpoint as NAT Destination | 578
- Selecting Interfaces when GWR Resides Inside an NFX Box | 578
- Selecting NAT Destination Using Abbreviations | 579
- Selecting a NAT Destination from the End Points Panel | 580
- Creating and Selecting a NAT Destination from the End Points Panel | 580
- Creating Addresses from Destination Field | 581
- Creating Services from Destination Field | 581

The following procedures provides various methods that you can use to choose an endpoint as a NAT destination:

Adding an Endpoint as NAT Destination

View and select the destination endpoint from the complete list of addresses, interfaces, services, zones, routing instances, or ports.

1. Click the **Destination** field. A list of relevant endpoints are displayed.

2. Click the **View more results** link provided at the bottom of the destination endpoints. The complete list of addresses, interfaces, services, zones, and routing instances, is displayed in the **End Points** panel on the right.

3. (Optional) Click the edit icon to edit the address, service, or port endpoint.

4. Click check mark icon (√) to select the endpoint as a destination.

Selecting Interfaces when GWR Resides Inside an NFX Box

The physical interfaces of an NFX box are mapped to the virtual interfaces of the Gateway Router (GWR) (vSRX) as given in Table 191 on page 579. These are the default mappings provided by CSO. You may change these interface mappings based on your requirements.

Table 191: NFX and GWR Interface Mapping

NFX Physical Interface	GWR Virtual Interface
WAN 0 (ge-0/0/10)	ge-0/0/2
WAN 1 (ge-0/0/11)	ge-0/0/3
WAN 2 (xe-0/0/12)	ge-0/0/7
WAN 3 (xe-0/0/13)	ge-0/0/8
LAN-X (ge-0/0/X)	Ge-0/0/06.<vlan-id-for-X>

When you create a new NAT rule and an NFX physical interface is intended as the destination endpoint, select the respective mapped GWR interface.

Selecting NAT Destination Using Abbreviations

Enter an abbreviation in the **Destination** field to select the destination endpoint from a filtered list of destination endpoints.

- To view a filtered list of addresses, enter **ADDR** or **addr**.
- To view a filtered list of interfaces, enter **INTR** or **intr**.
- To view a filtered list of services, enter **SVCS** or **svcs**.
- To view a filtered list of zones, enter **ZONE** or **zone**.
- To view a filtered list of routing instances, enter **ROUT** or **rout**.

Click the endpoints in the filtered list to select them.

You can add a port number as a destination endpoint. To do so:

1. Enter **PORT** or **port** in **Destination**.
2. Press Tab.
3. Enter the port number and press Enter.

   You can also enter a range of ports by using the separator **-**. For example, you can enter **10-20**.

The entered port value is selected as a destination endpoint.

You can also select the endpoint from the complete list of addresses, interfaces, services, zones, and routing instances. See “Adding an Endpoint as NAT Destination” on page 578.

Selecting a NAT Destination from the End Points Panel

You can select a NAT destination endpoint from the **End Points** panel. Alternately, you can create a new NAT destination endpoint from the **End Points** panel, see "Creating and Selecting a NAT Destination from the End Points Panel" on page 580.

To select a NAT destination endpoint from the **End Points** panel:

1. Click the **Destination** field.

2. Click the lesser-than icon (**<**) on the right.

   The **End Points** panel appears, displaying the list of available addresses, interfaces, services, zones, and routing instances.

3. (Optional) To view more information about a destination endpoint, click the details icon on the right of the endpoint. To edit the destination endpoint, click the edit icon (pencil symbol) on the right of the endpoint.

   > **NOTE:** You can only edit or view details of a destination endpoint if these options appear on right side of the endpoint when you hover over it. Not all endpoints provide these options.

4. Click the check mark icon (**√**) to add the endpoint as a destination.

Creating and Selecting a NAT Destination from the End Points Panel

To create a new destination endpoint from the **End Points** panel:

1. Click the add icon (**+**) on the top right of the panel and select the type of endpoint you want to create, among the options provided.

   Based on the option you select, the respective page appears. Fill in the required details to create a new endpoint.

   - To create a new address, see "Creating Addresses or Address Groups" on page 737.
   - To create a new service, see "Creating Services and Service Groups" on page 744.

   After the endpoint is created, it appears in the **Endpoints** panel.

2. Click the check mark icon (**√**) to add the new endpoint as a destination.

Creating Addresses from Destination Field

You can use one of the following ways to create a new address from the **Destination** and use the newly created address as a destination endpoint:

- Type the address directly in the **Destination** field. If the address is valid, it is created immediately and added as a destination endpoint.
- Create an address from the **Destination** field, using the following steps:

  1. In the **Destination** field, type **addr**. The **Add new address** link appears at the bottom of the list of addresses.

  2. Click **Add new address** to create a new address.

     The **Create Addresses** page appears.

  3. Configure the new address. See “Creating Addresses or Address Groups” on page 737.

  4. Click **Save** to save the new address.

     The new address is created, and will be listed as an option for the destination. Select the new address to add it to the destination.

Creating Services from Destination Field

To create a new service from the **Destination** field and use the newly created service as a destination endpoint:

1. In the **Destination** link, type **svcs**. The **Add new service** link appears at the bottom of the list of services.

2. Click **Add new service** to create a new service.

   The **Create Services** page appears.

3. Configure the new service. See “Creating Services and Service Groups” on page 744.

4. Click **Save** to save the new service.

   The new service is created, and will be listed as an option for the destination. Select the new service to add it to the destination.

RELATED DOCUMENTATION

About the Single NAT Policy Page | 562
Editing, Cloning, and Deleting NAT Policy Rules | 571
Creating NAT Policy Rules | 564
Deploying NAT Policy Rules | 573
NAT Policies Overview | 555
About the NAT Policies Page | 558
Creating NAT Policies | 559
Editing and Deleting NAT Policies | 561

NAT Pools Overview

A NAT pool is a set of IP addresses that you can define and use for address translation. NAT policies perform address translation by translating internal IP addresses to the addresses in these pools. Unlike static NAT, where there is a one-to-one mapping that includes destination IP address translation in one direction and source IP address translation in the reverse direction, with source NAT, you translate the original source IP address to an IP address in the address pool. With destination NAT, you translate the original destination address to an IP address in the address pool.

RELATED DOCUMENTATION

NAT Policies Overview | 555
About the NAT Pools Page | 582
Creating NAT Pools | 584
Editing, Cloning, and Deleting NAT Pools | 586

About the NAT Pools Page

To access this page, select **Configuration** > **NAT** > **Pools**.

Use the **NAT Pools** page to create, modify, clone, and delete NAT pools. You can filter and sort this information to get a better understanding of what you want to configure.

Tasks You Can Perform

You can perform the following tasks from this page:

- Create a NAT pool. See "Creating NAT Pools" on page 584.
- Modify, clone, or delete a NAT pool. See "Editing, Cloning, and Deleting NAT Pools" on page 586.
- View unused NAT pools by selecting **More > Show Unused**. Delete unused NAT pools by selecting **More > Delete Unused Items**.
- View duplicate NAT pools. Select **More > Show Duplicates**. The **Show Duplicates** page appears, displaying duplicate NAT pools. To delete a duplicate NAT pool, select it and click the delete icon (**X**).
- View the details of a NAT pool by selecting **More > Detailed View**, or by right-clicking a NAT pool and select **Detailed View**.
- Search for a specific NAT pool. Click the Search icon in the top right corner of the page to search for a NAT pool.

  You can enter partial text or full text of the keyword in the text box and press Enter.

  The search results are displayed on the same page.
- Show or hide columns. Click the **Show Hide Columns** icon at the top right corner of the page.

Table 192 on page 583 provides description of the fields on the **NAT Pools** page.

Table 192: Fields on the NAT Pools Page

Field	Description
Name	Displays the name of the NAT pool.
Pool Address	Displays the IP address of the NAT pool.
Description	Displays the description provided about the NAT pool when it was created.
Pool Type	Displays the NAT pool type. A NAT pool can be of type **Source** or **Destination**.

RELATED DOCUMENTATION

NAT Pools Overview | 582

Creating NAT Pools | 584

Editing, Cloning, and Deleting NAT Pools | 586

Creating NAT Pools

Use the **Create NAT Pools** page to create NAT pools.

To create a NAT pool:

1. Select **Configuration > NAT > Pools**.

   The **NAT Pools** page appears.

2. Click the add icon (**+**).

   The **Create NAT Pool** page displays fields required for creating and configuring a NAT pool.

3. Complete the configuration according to the guidelines provided in Table 193 on page 584.

4. Click **OK** to save the changes. A NAT pool with the configuration you provided is created.

   If you want to discard your changes, click **Cancel** instead.

Table 193 on page 584 provides guidelines on using the fields on the **Create NAT Pool** page.

Table 193: Fields on the Create NAT Pool Page

Field	Description
General Information	
Name	Enter a unique string of alphanumeric characters, dashes, and underscores. Colons, and periods are not allowed, and the maximum length is 31 characters.
Description	Enter a description for the new NAT pool; maximum length is 1024 characters.
Pool Type	Select a NAT pool type to configure: • Source • Destination
Pool Address	Select a NAT pool address or click **Add new address** to create a new NAT pool address.
Routing Instance	
Site	Select the site to which the NAT pool is applicable. **NOTE:** In a hub and spoke topology, both hub and spoke sites are listed in the **Site** drop-down. Ensure that you select only a spoke site, when you are creating a destination NAT pool.

Table 193: Fields on the Create NAT Pool Page *(continued)*

Field	Description
Routing Instance	Select the required routing instance from the list of available routing instances for the selected site.
Advanced	
Host Address Base	Enter the base address of the original source IP address range. The **Host Address Base** is used for IP address shifting.
Translation	Select the translation type for the incoming traffic: • No Translation—There is no translation required for the incoming traffic. • Port/Range—Set the global default single port range for source NAT pools with port translation. • Overload—Multiple source addresses are translated to pool addresses. If you set **Overload** as the translation type, the value of the **Pool Address** field cannot be an IP range or subnet, but it will be a single address.
Address Pooling	Select a NAT address pooling behavior: • Paired—Use this option for applications that require all sessions associated with one internal IP address to be translated to the same external IP address for multiple sessions. • Non-Paired—Use this option for applications that can be assigned IP addresses in a round-robin fashion.
Port	Enter the port number for the destination NAT pool type.
Start	Enter the start port range for the source NAT pools, if the translation type is Port/Range. The value of the port range can be any value between 1024 to 65535.
End	Enter the end port range. The value of the port range can be any value between 1024 to 65535.
Port Overloading Factor	Configure the port overloading capacity for a source NAT pool. If the factor is set to *x*, each translated IP address has *x* times the maximum number of ports available. The value of the port overloading factor can range between 2 and 32.
Address Sharing	Enable address sharing so that multiple internal IP addresses can be mapped to the same external IP address. Select this option only when the source NAT pool is configured with no port translation. When a source NAT pool has only one or a few external IP addresses available, the address sharing option with a many-to-one address mapping increases NAT resources and improves traffic.

Table 193: Fields on the Create NAT Pool Page *(continued)*

Field	Description
Overflow Pool Type	Select a source pool to use when the current address pool is exhausted. • Interface—Allow the egress interface IP address to support overflow. • Pool—Name of the source address pool. ◦ Overflow Pool—When addresses from the original source NAT pool are exhausted, IP addresses and port numbers are allocated from the overflow pool. A user-defined source NAT pool or an egress interface can be used as the overflow pool. (When the overflow pool is used, the pool ID is returned with the address.)

RELATED DOCUMENTATION

NAT Pools Overview | 582

About the NAT Pools Page | 582

Editing, Cloning, and Deleting NAT Pools | 586

Editing, Cloning, and Deleting NAT Pools

IN THIS SECTION

- Editing NAT Pools | 586
- Cloning NAT Pools | 587
- Deleting NAT Pools | 587

Editing NAT Pools

To modify the parameters configured for a NAT pool:

1. Select **Configuration > NAT > Pools**.

   The **NAT Pools** page appears.

2. Select the NAT pool that you want to edit, and click the edit icon (pencil symbol) at the top right corner of the table, or right-click and select **Edit NAT Pool**.

The **Edit NAT Pool** page appears, displaying the same options that are displayed when creating a new NAT pool.

3. Modify the parameters according to the guidelines provided in "Creating NAT Pools" on page 584.

4. Click **OK** to save the changes. If you click **OK**, you see the modified NAT pool in the **NAT Pools** page.

   If you want to discard your changes, click **Cancel** instead.

Cloning NAT Pools

To clone a NAT pool:

1. Select **Configuration > NAT > Pools**.

   The **NAT Pools** page appears.

2. Right-click the NAT pool that you want to clone and then click **Clone**, or select **More > Clone**.

   The **Clone NAT Pool** page appears with editable fields. Modify the parameters of the cloned NAT pool as per your requirements.

3. Click **OK** to save the changes. If you click **OK**, the cloned NAT pool appears at the end of the NAT pools list in the **NAT Pools** page.

   If you want to discard your changes, click **Cancel** instead.

Deleting NAT Pools

To delete a NAT pool:

1. Select **Configuration > NAT > Pools**.

   The **NAT Pools** page appears.

2. Select the NAT pool you want to delete and then click the delete icon **(X)** .

   An alert message appears, verifying that you want to delete the NAT pool.

3. Click **Yes** to delete the NAT pool. If you click **Yes**, the selected NAT pool is deleted.

   If you do not want to delete, click **Cancel** instead.

RELATED DOCUMENTATION

NAT Pools Overview | 582

About the NAT Pools Page | 582

Creating NAT Pools | 584

Deploying NAT Policies

After adding the intents to the NAT policies, you can deploy the NAT policy by clicking the **Deploy** option that is above the **End Points** panel. You can also deploy one or more policies from the **NAT Policies** page.

To deploy NAT policies:

1. Select **Configuration > NAT > NAT Policies**.

   The NAT Policies page appears.

2. Select one or more policies and click **Deploy.**

   The Deploy page appears.

3. In **Choose Deployment Time** options, select **Run Now** to deploy the policy immediately. Select **Schedule at a later time** and specify the date and time at which the policy should be deployed.

4. Click **Deploy**.

   A job is created. Click the job ID to go to the Jobs page and view the status of the deploy operation.

Importing NAT Policies

Use this page to manually import a firewall policy from the discovered or onboarded sites (next generation firewall sites).

To import a NAT policy:

1. Select **Configuration > NAT > NAT Policies**.

   The NAT Policy page appears.

2. Click **Import**.

   The Import NAT Policies page appears displaying a list of discovered devices (next generation firewall devices).

3. Select the devices from which you want to import the NAT policies and click **Next**.

   The Discovered Services tab appears.

4. Select the NAT policies that you want to import and click **Next**.

   The Resolve Conflicts tab appears.

5. If there are any conflicts with the imported objects, object conflict resolution(OCR) operation is triggered. The Conflicts window displays all the conflicts between CSO and the next generation firewall device. Select an object from the Conflicts window and click on any of the below option to resolve the object conflict.

   The resolution options are:

   - Rename Object—Rename the imported object. By default, "_1" is added to the object name, or you can specify a new name.
   - Overwrite with imported value—The object in CSO is replaced with the object from the import operation.
   - Keep existing object—The object name in CSO is used instead of what is on the next generation firewall device.

6. Click **Finish**.

   A summary of the discovered services is listed.

7. Review the summary and click **OK** to import the NAT policies.

   The import policy job is created and the NAT policies are imported from next generation firewall device to CSO. You can view the imported policy from the NAT Policies page.

WHAT'S NEXT

> After importing the NAT policy successfully, you can edit and deploy the policy. See Editing and Deleting NAT Policies | **561**, Editing, Cloning, and Deleting NAT Pools | **586**, and Deploying NAT Policies | **588**.

RELATED DOCUMENTATION

- Importing Policies Overview | 446

CHAPTER 19

Managing IPS Signatures and Profiles

IN THIS CHAPTER

- About the IPS Signatures Page | **591**
- Create IPS Signatures | **596**
- Create IPS Signature Static Groups | **604**
- Create IPS Signature Dynamic Groups | **605**
- Edit, Clone, and Delete IPS Signatures | **611**
- Edit, Clone, and Delete IPS Signature Static Groups | **613**
- Edit, Clone, and Delete IPS Signature Dynamic Groups | **616**
- About the IPS Profiles Page | **618**
- Create IPS Profiles | **620**
- Edit, Clone, and Delete IPS Profiles | **621**
- About the <IPS-Profile-Name> / Rules Page | **623**
- Create IPS or Exempt Rules | **625**
- Edit, Clone, and Delete IPS or Exempt Rules | **633**

About the IPS Signatures Page

IN THIS SECTION

- Tasks You Can Perform | **592**
- Field Descriptions | **592**

To access this page, select **Configure** > **IPS** > **IPS Signature**.

Use intrusion prevention system (IPS) signatures to monitor and prevent intrusions. IPS compares traffic against signatures of known threats and blocks traffic when a threat is detected.

Tasks You Can Perform

- View the details of an IPS signature—Select an IPS signature and click **More > Details**, or mouse over the IPS signature and click the **Detailed View** icon. The IPS Signature Details View page appears. See Table 195 on page 594 for an explanation of fields on this page.
- View the details of an IPS signature static group—Select an IPS signature static group and click **More > Details**, or mouse over the IPS signature static group and click the **Detailed View** icon. The IPS Static Group Details page appears. See Table 196 on page 595 for an explanation of fields on this page.
- View the details of an IPS signature dynamic group—Select an IPS signature dynamic group and click **More > Details**, or mouse over the IPS signature dynamic group and click the **Detailed View** icon. The IPS Signature Dynamic Details View page appears. See Table 197 on page 595 for an explanation of fields on this page.
- Create an IPS signature—See "Create IPS Signatures" on page 596.
- Create an IPS signature static group—See "Create IPS Signature Static Groups" on page 604.
- Create an IPS signature dynamic group—See "Create IPS Signature Dynamic Groups" on page 605.
- Edit, clone, or delete an IPS signature—See "Edit, Clone, and Delete IPS Signatures" on page 611.
- Edit, clone, or delete an IPS signature static group—See "Edit, Clone, and Delete IPS Signature Static Groups" on page 613.
- Edit, clone, or delete an IPS signature dynamic group—See "Edit, Clone, and Delete IPS Signature Dynamic Groups" on page 616.
- Search for IPS signatures, static groups or dynamic groups by using keywords—Click the search icon and enter the search term in the text box and press Enter. The search results are displayed on the same page.
- Filter IPS signatures, static groups or dynamic groups—Click the filter icon (funnel) and specify one or more filtering criteria. The filtered results are displayed on the same page.
- Sort IPS signatures, static groups or dynamic groups—Click a column name to sort the data in the grid (table) based on the column name.

  NOTE: Sorting is applicable only to some fields.

- Show or hide columns—Click the **Show Hide Columns** icon at the top right corner of the page and select the columns that you want displayed on the page.

Field Descriptions

Table 194 on page 593 describes the field on the IPS Signatures page.

Table 194: Fields on the IPS Signatures Page

Field	Description
Name	Name of the IPS signature, IPS signature static group, or IPS signature dynamic group.
Severity	Severity level of the attack that the signature will report.
Category	Category of the attack object.
Object Type	Displays the type of attack object: • Static Group • Dynamic Group • Signature • Protocol Anomaly • Compound Attack
Recommended	Indicates whether the attack objects are recommended by Juniper (True) or not (False).
Action	Action taken when the monitored traffic matches the attack objects specified in the IPS rules.
Definition Type	Displays whether the IPS signature, static group, or dynamic group was created by CSO (Predefined) or user-created (Custom).
CVE	Displays the Common Vulnerabilities and Exposures (CVE) identifier or name associated with the threat.
CERT	Displays the computer emergency response team (CERT) advisory number associated with the threat.
BUG	Displays the list of bugs that are related to the signature attack.
False Positives	Displays the frequency with which the attack produces a false positive on your network.
Service	Protocol or service that the attack uses to enter your network.
Performance Impact	Performance impact of the IPS signature.
Direction	Direction of the traffic for which the attack is detected; for example, client to server.

Table 195: Fields on the IPS Signature Details View Page

Field	Description
Name	Name of the IPS signature.
Description	Description of the IPS signature.
URL(s)	Displays the URLs that have the details about the signature attack. For example, http://www.faqs.org/rfcs/rfc2865.html.
Category	See Table 194 on page 593.
Recommended	See Table 194 on page 593.
Action	See Table 194 on page 593.
Keywords	Keywords associated with the IPS signature.
Severity	See Table 194 on page 593.
BUGS	See Table 194 on page 593.
CERT	See Table 194 on page 593.
CVE	See Table 194 on page 593.
Signature Details	
Binding	Protocol or service that the attack uses to enter your network.
Service	For service binding, displays the service the attack uses to enter your network.
Time Count	Number of time that IPS detects the attack in a specified time scope.

Table 195: Fields on the IPS Signature Details View Page *(continued)*

Field	Description
Signature	Displays (in a table) the signature attack objects configured as part of the IPS signature. For each row, the following fields are displayed: • No.—Unique identifier for the signature attack object. • Context—Attack context, which defines the location of the signature where IPS should look for the attack. • Direction—Connection direction of the attack. • Pattern—Signature pattern (in Juniper's proprietary regular expression syntax) of the attack to be detected. • Regex—Regular expression to match malicious or unwanted behavior over the network. • Negated—Indicates whether the pattern should be excluded from being matched (true) or not (false).
Anomaly	Displays (in a table) the protocol anomaly attack objects configured as part of the IPS signature. For each row, the following fields are displayed: • No.—Unique identifier for the anomaly. • Anomaly—Protocol or service for which the anomaly is defined. • Direction—Connection direction of the attack.

Table 196: Fields on the IPS Static Group Details Page

Field	Description
Name	Name of the IPS signature static group.
Description	Description of the IPS signature static group.
Group Members	Displays the IPS signatures or IPS signature dynamic groups that are part of the IPS static group. See Table 194 on page 593 for an explanation of the fields in the table. To view the details, select a row, click **More > Details**, or mouse over a row and click the **Detailed View** icon. Depending on the object type, the IPS Signature Details View page or IPS Signature Dynamic Details View page appears. See Table 195 on page 594 and Table 197 on page 595 for an explanation of the fields on these pages.

Table 197: Fields on the IPS Signature Dynamic Details View Page

Field	Description
Name	Name of the IPS signature dynamic group.

Table 197: Fields on the IPS Signature Dynamic Details View Page *(continued)*

Field	Description
Severity	Severity filters used for the dynamic group.
Service	Services filters used for the dynamic group.
Category	Category filters used for the dynamic group.
Recommended	Indicates whether predefined attack objects recommended by Juniper are added to the dynamic group (true) or not (false).
Direction	Traffic direction (for which the attack is detected) filters used for the dynamic group.
Performance Impact	Performance impact filter used for the dynamic group.
False Positive	False positive filter used for the dynamic group.
Age of Attack	Age of the attack (in years) used as a filter for the dynamic group.
CVSS Score	Common Vulnerability Scoring System (CVSS) score used as a filter for the dynamic group.
File Type	File type of the attack used as a filter for the dynamic group.
Vulnerability Type	Vulnerability type of the attack used as a filter for the dynamic group.
Object Type	Type of object (anomaly or signature) used as a filter for the dynamic group.
Vendor Description	Vendor or product that the attack belongs to.

RELATED DOCUMENTATION

About the IPS Profiles Page | 618

Create IPS Signatures

The signature database in Contrail Service Orchestration (CSO) contains predefined intrusion prevention system (IPS) signatures that you can use. From the Create IPS Signature page, users with the tenant

administrator role or a custom role with appropriate IPS tasks can also create customized IPS signatures to block newer attacks or unknown attacks.

To create a customized IPS signature:

1. Select **Configuration > IPS > IPS Signatures**.

   The IPS Signatures page appears.

2. Select **Create > IPS Signature**.

   The Create IPS Signature page appears.

3. Complete the configuration according to the guidelines in Table 198 on page 597.

   NOTE: Fields marked with an asterisk (*) are mandatory.

4. Click **OK**.

   You are returned to the IPS Signatures page and a message indicating that the signature is created is displayed.

After you create an IPS signature, you can use the signature in an IPS or an exempt rule and reference the IPS profile (containing the rule) in a firewall policy that you can then deploy on the device.

Table 198: Create IPS Signature Settings

Setting	Guideline
Name	Enter a unique name for the IPS signature that is a string of alphanumeric characters and some special characters (colon, hyphen, period, and underscore). No spaces are allowed and the maximum length is 255 characters.
Description	Enter a description for the IPS signature; the maximum length is 1024 characters.
Category	Enter a predefined category or a new category. The category can contain alphanumeric characters and special characters (hyphen and underscore) and must begin with an alphanumeric character. No spaces are allowed and the maximum length is 63 characters. You use categories to group attack objects and then within each category, you can assign severity levels to the attack objects.

Table 198: Create IPS Signature Settings *(continued)*

Setting	Guideline
Action	Select the action to take when the monitored traffic matches the attack objects specified in the IPS rule: • None—No action is taken. Use this action to only generate logs for some traffic. • Close Client & Server—Closes the connection and sends a TCP reset (RST) packet to both the client and the server. • Close Client—Closes the connection and sends an RST packet to the client, but not to the server. • Close Server—Closes the connection and sends an RST packet to the server, but not to the client. • Ignore—Stops scanning traffic for the rest of the connection if an attack match is found. IPS disables the rulebase for the specific connection. • Drop—Drops all packets associated with the connection, preventing traffic for the connection from reaching its destination. Use this action to drop connections for traffic that is not prone to spoofing. • Drop Packet—Drops a matching packet before it can reach its destination but does not close the connection. Use this action to drop packets for attacks in traffic that is prone to spoofing, such as UDP traffic. Dropping a connection for such traffic could result in a denial of service that prevents you from receiving traffic from a legitimate source-IP address.
Keywords	Enter unique identifiers that can be used to search and sort signatures. Keywords should relate to the attack and the attack object. For example, Amanda Amindexd Remote Overflow.
Severity	Select a severity level for the attack that the signature will report: • Critical—Contains attack objects matching exploits that attempt to evade detection, cause a network device to crash, or gain system-level privileges. • Major—Contains attack objects matching exploits that attempt to disrupt a service, gain user-level access to a network device, or activate a Trojan horse previously loaded on a device. • Minor—Contains attack objects matching exploits that detect reconnaissance efforts attempting to access vital information through directory traversal or information leaks. • Warning—Contains attack objects matching exploits that attempt to obtain noncritical information or scan a network with a scanning tool. • Info—Contains attack objects matching normal, harmless traffic containing URLs, DNS lookup failures, SNMP public community strings, and peer-to-peer (P2P) parameters. You can use informational attack objects to obtain information about your network.
Signature Details	

Table 198: Create IPS Signature Settings *(continued)*

Setting	Guideline
Binding	Select the protocol or service that the attack uses to enter your network: • IP—Match the attack for a specified protocol type number, which you must specify in the Protocol field. • IPv6—Match the attack for a specified protocol type number (for the header following the IPv6 header), which you must specify in the Next Header field • ICMP—Match the attack for ICMP packets. • IPv6—Match the attack for ICMPv6 packets. • TCP—Match the attack for specified TCP ports or port ranges, which you must specify in the Port Range(s) field. • UDP—Match the attack for specified UDP ports or port ranges. • RPC—Match the attack for a specified remote procedure call (RPC) program number, which you must specify in the Program Number field. • Service—Match the attack for a specified service, which you must choose from the Service field.
Protocol	For IP binding, specify the transport layer protocol number that you want matched to the attack. Range: 1 through 139 excluding 1, 6, and 17.
Next Header	For IPv6 binding, specify the transport layer protocol number for the next header following the IPv6 header with which to match the attack. Range: 1 through 139 excluding 6, 17, and 58.
Port Range(s)	For TCP or UDP binding, specify a port number or a port range (*min-port-no-max-port-no* format) that you want matched to the attack.
Program Number	For RPC binding, specify the RPC program number (ID) that you want matched to the attack.
Service	For service binding, select the service that you want matched to the attack.
Time Count	Specify the number of times that IPS detects the attack within the specified time scope before triggering an event.

Table 198: Create IPS Signature Settings *(continued)*

Setting	Guideline
Time Scope	Specify the scope within which the counting of the attack occurs: • Source IP—Detect attacks from the source IP address for the specified time count regardless of the destination IP address. • Dest IP—Detect attacks from the destination IP address for the specified time count regardless of the source IP address. • Peer—Detect attacks between source and destination IP addresses of the sessions for the specified time count.
Match Assurance	Specify a false positives filter to track attack objects based on the frequency that the attack produces a false positive on your network: • High—Provides information on the frequently tracked false positive occurrences. • Medium—Provides information on the occasionally tracked false positive occurrences. • Low—Provides information on the rarely tracked false positive occurrences.
Performance Impact	Specify this filter to select only the appropriate attacks based on performance impact; for example to filter out slow-performing attack objects: • High—Add high performance impact attack objects that are vulnerable to an attack. The performance impact of signatures is high7 to high9, where the application identification is slow. • Medium—Add medium performance impact attack objects that are vulnerable to an attack. The performance impact of signatures is medium4 to medium6, where the application identification is normal. • Low—Add low performance impact attack objects that are vulnerable to an attack. The performance impact of signatures is low1 to low3, where the application identification is faster. • Unknown—Add attack objects whose performance impact is unknown.

Table 198: Create IPS Signature Settings *(continued)*

Setting	Guideline
Add Signature	You can specify one or more signature attack objects that use a stateful attack signature (a pattern that always exists within a specific section of the attack) to detect known attacks. **NOTE:** For a customized IPS signature, you must specify at least one signature attack object or anomaly. • To add a signature attack object: 1. Click the add (**+**) icon. The Add Signature page appears. 2. Complete the configuration according to the guidelines in Table 199 on page 603. 3. Click **OK**. You are returned to the previous page and the signature attack object is displayed in the table. • To modify a signature attack object that you added: 1. Select an attack object and click the edit (pencil) icon. The Edit Signature page appears, displaying the same fields that appear when you add a signature attack object. 2. Modify the fields as needed. See Table 199 on page 603. 3. Click **OK**. Your modifications are saved and you are returned to the previous page. • To delete a signature attack object that you added: 1. Select an attack object and click the delete (trash can) icon. A popup appears asking you to confirm the delete operation. 2. Click **Yes**. The signature attack object is deleted and you are returned to the previous page.

Table 198: Create IPS Signature Settings *(continued)*

Setting	Guideline
Add Anomaly	**NOTE:** • The Add Anomaly field is displayed only if you specify a service binding. • For a customized IPS signature, you must specify at least one signature attack object or anomaly. Protocol anomaly attack objects detect abnormal or ambiguous messages within a connection according to the set of rules for the particular protocol being used. You can add, modify, or delete anomaly attack objects: • To add an anomaly: 1. Click the add (+) icon. The Add Anomaly page appears. 2. Complete the configuration according to the guidelines in Table 200 on page 603. 3. Click **OK**. You are returned to the previous page and the anomaly is displayed in the table. • To modify an anomaly that you added: 1. Select an anomaly and click the edit (pencil) icon. The Edit Anomaly page appears, displaying the same fields that appear when you add an anomaly. 2. Modify the fields as needed. See Table 200 on page 603. 3. Click **OK**. Your modifications are saved and you are returned to the previous page. • To delete an anomaly that you added: 1. Select an anomaly and click the delete (trash can) icon. A popup appears asking you to confirm the delete operation. 2. Click **Yes**. The signature anomaly is deleted and you are returned to the previous page.

Table 199: Add Signature Settings

Setting	Guideline
Signature No.	Displays the system-generated signature number; you cannot modify this field.
Context	Select the attack context, which defines the location of the signature where IPS should look for the attack in a specific Application Layer protocol.
Direction	Select the connection direction of the attack: • Any—Detect the attack for traffic in either direction. • Client to-Server—Detect the attack only in client-to-server traffic. • Server to Client—Detect the attack only in server to client traffic.
Pattern	Enter the signature pattern (in Juniper Networks proprietary regular expression syntax) of the attack you want to detect. An attack pattern can be a segment of code, a URL, or a value in a packet header and the signature pattern is the syntactical expression that represents that attack pattern. For example, use **\[<character-set>\]** for case-insensitive matches.
Regex	Enter a regular expression to define rules to match malicious or unwanted behavior over the network. For example: For the syntax \[hello\], the expected pattern is hello, which is case sensitive. The example matches can be: hElLo, HEllO, and heLLO.
Negated	Select this check box to exclude the specified pattern from being matched. When you negate a pattern, the attack is considered matched if the pattern defined in the attack does not match the specified pattern.

Table 200: Add Anomaly Settings

Setting	Guideline
Anomaly No.	Displays the system-generated anomaly number; you cannot modify this field.
Anomaly	Select the protocol (service) whose anomaly is being defined in the attack.
Direction	Select the connection direction of the attack: • **Any**—Detect the attack for traffic in either direction. • **Client to-Server**—Detect the attack only in client-to-server traffic. • **Server to Client**—Detect the attack only in server to client traffic.

RELATED DOCUMENTATION

Create IPS Profiles | 620

Create IPS Signature Static Groups

The signature database in Contrail Service Orchestration (CSO) contains predefined intrusion prevention system (IPS) signature static groups that you can use. Users with the tenant administrator role or a custom role with appropriate IPS tasks can also create customized IPS signature static groups from the Create IPS Signature Static Group page. Static groups enable better manageability because you can group different types of signatures into one entity.

To create a customized IPS signature static group:

1. Select **Configuration > IPS > IPS Signatures**.

   The IPS Signatures page appears.

2. Select **Create > Static Group**.

   The Create IPS Signature Static Group page appears.

3. Complete the configuration according to the guidelines in Table 201 on page 604.

   NOTE: Fields marked with an asterisk (*) are mandatory.

4. Click **OK**.

   You are returned to the IPS Signatures page and a message that the static group was successfully created is displayed.

After you create an IPS signature static group, you can use the static group in an IPS or an exempt rule and reference the IPS profile (containing the rule) in a firewall policy that you can then deploy on the device.

Table 201: Create IPS Signature Static Group Settings

Setting	Guideline
Name	Enter a unique name for the IPS signature static group that is a string of alphanumeric characters, colons, periods, dashes, and underscores. No spaces are allowed and the maximum length is 255 characters.

Table 201: Create IPS Signature Static Group Settings *(continued)*

Setting	Guideline
Description	Enter a description for the IPS signature static group; the maximum length is 1024 characters.
Group Members	You can add one or more IPS signatures, static groups, or dynamic groups to be members of the static group that you are creating. In addition, you can delete group members after adding them. **NOTE:** You must add at least one IPS signature, static group, or dynamic group to proceed. • To add one or more group members: 1. Click the add (+) icon. The Add IPS Signatures page appears displaying the existing predefined and custom IPS signatures, static groups, and dynamic groups in a table. 2. Select one or more group members by clicking the check boxes corresponding to the rows. 3. Click **OK**. You are returned to the previous page and the group members that you added are displayed in the table. • To delete one or more group members that you added: 1. Select the group members that you want do delete and click the delete (trash can) icon. A warning message appears asking you to confirm the deletion. 2. Click **Yes**. The group members are deleted.

RELATED DOCUMENTATION

Create IPS Profiles | 620

Create IPS Signature Dynamic Groups

The signature database in Contrail Service Orchestration (CSO) contains predefined intrusion prevention system (IPS) signature dynamic groups that you can use. Users with the tenant administrator role or a

custom role with appropriate IPS tasks can also create customized IPS signature dynamic groups (based on a specified filter criteria) from the Create IPS Signature Dynamic Group page.

The filter criteria that you specify are matched only to predefined or customized IPS signatures, and not to IPS static groups dynamic groups. When a new signature database is used, the dynamic group membership is automatically updated based on the filter criteria for that group.

To create a customized IPS signature dynamic group:

1. Select **Configuration > IPS > IPS Signatures**.

   The IPS Signatures page appears.

2. Select **Create > Dynamic Group**.

   The Create IPS Signature Static Group page appears.

3. Complete the configuration according to the guidelines in Table 202 on page 606.

   NOTE: Fields marked with an asterisk (*) are mandatory.

4. (Optional) Click **Preview Filtered Signatures** to check if the signatures that match the dynamic group are consistent with the filter criteria that you specified.

   The IPS Signatures page appears displaying the list of IPS signatures matching the filters. If the signatures do not match, you can tweak the filter criteria as needed. Click **Close** to go back to the previous page.

5. Click **OK**.

   You are returned to the IPS Signatures page and a message indicating that the dynamic group was successfully created is displayed.

After you create an IPS signature dynamic group, you can use the dynamic group in an IPS or an exempt rule and reference the IPS profile (containing the rule) in a firewall policy that you can then deploy on the device.

Table 202: Create IPS Signature Dynamic Group Settings

Setting	Guideline
Name	Enter a unique name for the IPS signature dynamic group that is a string of alphanumeric characters, colons, periods, hyphens, and underscores. No spaces are allowed and the maximum length is 255 characters.

Table 202: Create IPS Signature Dynamic Group Settings *(continued)*

Setting	Guideline
Filter Criteria	You select one or more filters to define the attributes of IPS signatures that will be added to the IPS signature dynamic group that you are creating. Filters apply to existing signatures (already downloaded in CSO) and to new signatures when they are downloaded. IPS signatures that match any of the filters that you configure are included as part of the signature group.
Severity	
Info	Select the **Enable** check box to include IPS signatures with the severity level Info.
Warning	Select the **Enable** check box to include IPS signatures with the severity level Warning.
Minor	Select the **Enable** check box to include IPS signatures with the severity level Minor.
Major	Select the **Enable** check box to include IPS signatures with the severity level Major.
Critical	Select the **Enable** check box to include IPS signatures with the severity level Critical.
Service	
Service	Specify the services that you want to use to filter for IPS signatures that should be included as part of the dynamic group. Select one or more services listed in the **Available** column and click the forward arrow to confirm your selection. The selected services are displayed in the Selected column.
Category	
Category	Specify the categories that you want to use to filter for IPS signatures that should be included as part of the dynamic group. Select one or more categories listed in the **Available** column and click the forward arrow to confirm your selection. The selected categories are displayed in the Selected column.
Recommended	

Table 202: Create IPS Signature Dynamic Group Settings *(continued)*

Setting	Guideline
Recommended	This filter is based on attack objects recommended by Juniper Networks. Select one of the following: • **None**—Don't use this filter. • **Yes**—Add predefined attacks recommended by Juniper Networks to the dynamic group. • **No**—Add predefined attacks that are not recommended by Juniper Networks to the dynamic group.
Direction	You use this filter to add IPS signatures to the dynamic group based on the traffic direction of the attacks. If you specify more than one traffic direction (Any, Client-to-Server, and Server-to-Client), you must select a value in the Expression field.
Any	Select one of the following: • **None** (default): Do not use this filter. • **Yes**: Include IPS signatures that track traffic from client to server or server to client. • **No**: Do not include IPS signatures that track traffic from client to server or server to client.
Client-to-Server	Select one of the following: • **None** (default): Do not use this filter. • **Yes**: Include IPS signatures that track traffic from client to server. • **No**: Do not include IPS signatures that track traffic from client to server.
Server-to-Client	Select one of the following: • **None** (default): Do not use this filter • **Yes**: Include IPS signatures that track traffic from server to client. • **No**: Do not include IPS signatures that track traffic from server to client.
Expression	If you specified more than one direction filter, you must specify how the signatures should be matched: • **OR**—Include signatures that match any of the specified traffic directions. • **AND**—Include signatures that match all of the specified traffic directions.
Performance Impact	

Table 202: Create IPS Signature Dynamic Group Settings *(continued)*

Setting	Guideline
Unknown	Select the **Enable** check box to include IPS signatures with the performance impact Unknown.
Low	Select the **Enable** check box to include IPS signatures with the performance impact Low.
Medium	Select the **Enable** check box to include IPS signatures with the performance impact Medium.
High	Select the **Enable** check box to include IPS signatures with the performance impact High.
False Positives	
Unknown	Select the **Enable** check box to include IPS signatures with the match assurance Unknown.
Low	Select the **Enable** check box to include IPS signatures with the match assurance Low.
Medium	Select the **Enable** check box to include IPS signatures with the match assurance Medium.
High	Select the **Enable** check box to include IPS signatures with the match assurance High.
Age of Attack	
Age of Attack	Enter the age of the attack (in years) to be used as a filter criteria to include IPS signatures as part of the dynamic group. Range: 1 through 100.
Expression	Select whether the IPS signatures should be filtered based on whether the age of attack in the signature is greater than (default) or less than the value that you specified.
CVSS Score	

Table 202: Create IPS Signature Dynamic Group Settings *(continued)*

Setting	Guideline
CVSS Score	Specify the Common Vulnerability Scoring System (CVSS) to be used as a filter criteria to include IPS signatures as part of the dynamic group. Range: Decimal number between 0 and 10.
Expression	Select whether the IPS signatures should be filtered based on whether the CVSS score of the attack is greater than (default) or less than the value that you specified.
Other Filters	
Excluded	Select one of the following: • **None** (default): Do not use this filter • **Yes**: Include excluded attack objects as part of the dynamic group. • **No**: Do not include excluded attack objects as part of the dynamic group.
File Type	Select the file type of the attack to be used as a filter criteria; for example, flash.
Vulnerability Type	Select the vulnerability type of the attack to be used as a filter criteria; for example, overflow.
Object Type	Specify this filter to group attack objects by type (anomaly or signature).
Signature	Select the **Enable** check box to add signatures based on stateful signature attack objects specified in the signature. A stateful attack signature is a pattern that always exists within a specific section of the attack. Stateful signature attack objects also include the protocol or service used to perpetrate the attack and the context in which the attack occurs.
Protocol Anomaly	Select the **Enable** check box to add signatures of attacks that violate protocol specifications (RFCs and common RFC extensions).
Vendor Description	
Product Type	Specify this filter to include signatures belonging to the selected product type.
Vendor Name	Specify this filter to include signatures belonging to the selected vendor.
Title	Specify this filter to include signatures belonging to the selected product name. The product names are populated only when you select a product type and a vendor.

RELATED DOCUMENTATION

Create IPS Profiles | 620

Edit, Clone, and Delete IPS Signatures

IN THIS SECTION

- Edit IPS Signatures | 611
- Clone IPS Signatures | 612
- Delete IPS Signatures | 612

Users with the tenant administrator role or a custom role with appropriate IPS tasks can edit, clone, or delete IPS signatures.

Edit IPS Signatures

You can edit only customized IPS signatures and not predefined (system-generated) signatures.

To edit a customized IPS signature:

1. Select **Configuration > IPS > IPS Signatures**.

   The IPS Signatures page appears.

2. Select a customized IPS signature and click the edit (pencil) icon.

   The Edit IPS Signature page appears, displaying the same fields that are presented when you create an IPS signature.

3. Modify the IPS signature fields as needed. See “Create IPS Signatures” on page 596.

   NOTE: You can modify all fields except the name.

4. Click **OK** to save your changes.

You are returned to the IPS Signatures page and a message that the IPS signature was successfully updated is displayed.

If the IPS signature was used in an IPS or exempt rule that is deployed on the device (through the firewall policy), then the firewall policy is marked for deployment. You must deploy the firewall policy for the changes to take effect on the device.

Clone IPS Signatures

Cloning enables you to easily create a new IPS signature based on an existing one. You can clone predefined or customized IPS signatures and modify the parameters as needed.

To clone an IPS signature:

1. Select **Configuration > IPS > IPS Signatures**.

   The IPS Signatures page appears.

2. Select an IPS signature and select **More > Clone**.

   The Clone IPS Signature page appears, displaying the same fields that are presented when you create an IPS signature.

3. Modify the IPS signature fields as needed. See "Create IPS Signatures" on page 596.

4. Click **OK** to save your changes.

   You are returned to the IPS Signatures page and a message that the IPS signature was successfully created is displayed.

After you clone an IPS signature, you can use the signature in an IPS or an exempt rule and reference the IPS profile (containing the rule) in a firewall policy that you can then deploy on the device.

Delete IPS Signatures

NOTE:

- You can delete only customized (user-created) IPS signatures that are not used in an IPS or exempt rule.
- You cannot delete predefined (system-generated) IPS signatures.

To delete one or more customized IPS signatures:

1. Select **Configuration > IPS > IPS Signatures**.

The IPS Signatures page appears.

2. Select one or more customized IPS signatures and click the delete (trash can) icon

   A warning message appears asking you to confirm the deletion.

3. Click **Yes** to proceed with the deletion.

   You are returned to the IPS Signatures page and a message indicating the status of the delete operation is displayed.

RELATED DOCUMENTATION

Create IPS Profiles | 620

Edit, Clone, and Delete IPS Signature Static Groups

IN THIS SECTION

- Edit IPS Signature Static Groups | 613
- Clone IPS Signature Static Groups | 614
- Delete IPS Signature Static Groups | 615

Users with the tenant administrator role or a custom role with appropriate IPS tasks can edit, clone, or delete IPS signature static groups.

Edit IPS Signature Static Groups

You can edit only customized IPS signature static groups and not predefined (system-generated) static groups.

To edit a customized IPS signature static group:

1. Select **Configuration > IPS > IPS Signatures**.

   The IPS Signatures page appears.

2. Select a customized IPS signature static group and click the edit (pencil) icon.

   The Edit IPS Signature Static Group page appears, displaying the same fields that are presented when you create an IPS signature static group.

3. Modify the IPS signature static group fields as needed. See "Create IPS Signature Static Groups" on page 604.

   NOTE: You can modify all fields except the name.

4. Click **OK** to save your changes.

   You are returned to the IPS Signatures page and a message that the IPS signature static group was successfully updated is displayed.

If the IPS signature static group was used in an IPS or exempt rule that is deployed on the device (through the firewall policy), then the firewall policy is marked for deployment. You must deploy the firewall policy for the changes to take effect on the device.

Clone IPS Signature Static Groups

Cloning enables you to easily create a new IPS signature static group based on an existing one. You can clone predefined or customized IPS signature static groups and modify the parameters as needed.

To clone an IPS signature static group:

1. Select **Configuration > IPS > IPS Signatures**.

   The IPS Signatures page appears.

2. Select an IPS signature static group and select **More > Clone**.

   The Clone IPS Signature Static Group page appears, displaying the same fields that are presented when you create an IPS signature static group.

3. Modify the IPS signature static group fields as needed. See "Create IPS Signature Static Groups" on page 604.

4. Click **OK** to save your changes.

   You are returned to the IPS Signatures page and a message that the IPS signature static group was successfully created is displayed.

After you clone an IPS signature static group, you can use the static group in an IPS or an exempt rule and reference the IPS profile (containing the rule) in a firewall policy that you can then deploy on the device.

Delete IPS Signature Static Groups

NOTE:

- You can delete only customized (user-created) IPS signature static groups that are not used in an IPS or exempt rule.
- You cannot delete predefined (system-generated) IPS signature static groups.

To delete one or more customized IPS signature static groups:

1. Select **Configuration > IPS > IPS Signatures**.

   The IPS Signatures page appears.

2. Select one or more customized IPS signature static groups and click the delete (trash can) icon

   A warning message appears asking you to confirm the deletion.

3. Click **Yes** to proceed with the deletion.

   You are returned to the IPS Signatures page and a message indicating the status of the delete operation is displayed.

RELATED DOCUMENTATION

Create IPS Profiles | 620

Edit, Clone, and Delete IPS Signature Dynamic Groups

IN THIS SECTION

- Edit IPS Signature Dynamic Groups | 616
- Clone IPS Signature Dynamic Groups | 617
- Delete IPS Signature Dynamic Groups | 618

Users with the tenant administrator role or a custom role with appropriate IPS tasks can edit, clone, or delete IPS signature dynamic groups.

Edit IPS Signature Dynamic Groups

You can edit only customized IPS signature dynamic groups and not predefined (system-generated) dynamic groups.

To edit a customized IPS signature dynamic group:

1. Select **Configuration > IPS > IPS Signatures**.

   The IPS Signatures page appears.

2. Select a customized IPS signature dynamic group and click the edit (pencil) icon.

   The Edit IPS Signature Dynamic Group page appears, displaying the same fields that are presented when you create an IPS signature dynamic group.

3. Modify the IPS signature dynamic group fields as needed. See "Create IPS Signature Dynamic Groups" on page 605.

   NOTE: You can modify all fields except the name.

4. (Optional) Click **Preview Filtered Signatures** to check if the signatures that match the dynamic group are consistent with the filter criteria that you specified.

The IPS Signatures page appears displaying the list of IPS signatures matching the filters. If the signatures do not match, you can tweak the filter criteria as needed. Click **Close** to go back to the previous page.

5. Click **OK** to save your changes.

   You are returned to the IPS Signatures page and a message indicating that the IPS signature dynamic group was successfully updated is displayed.

If the IPS signature dynamic group was used in an IPS or exempt rule that is deployed on the device (through the firewall policy), then the firewall policy is marked for deployment. You must deploy the firewall policy for the changes to take effect on the device.

Clone IPS Signature Dynamic Groups

Cloning enables you to easily create a new IPS signature dynamic group based on an existing one. You can clone predefined or customized IPS signature dynamic groups and modify the parameters as needed.

To clone an IPS signature dynamic group:

1. Select **Configuration > IPS > IPS Signatures**.

   The IPS Signatures page appears.

2. Select an IPS signature dynamic group and select **More > Clone**.

   The Clone IPS Signature Dynamic Group page appears, displaying the same fields that are presented when you create an IPS signature dynamic group.

3. Modify the IPS signature dynamic group fields as needed. See "Create IPS Signature Dynamic Groups" on page 605.

4. (Optional) Click **Preview Filtered Signatures** to check if the signatures that match the dynamic group are consistent with the filter criteria that you specified.

   The IPS Signatures page appears displaying the list of IPS signatures matching the filters. If the signatures do not match, you can tweak the filter criteria as needed. Click **Close** to go back to the previous page.

5. Click **OK** to save your changes.

   You are returned to the IPS Signatures page and a message that the IPS signature dynamic group was successfully created is displayed.

After you clone an IPS signature dynamic group, you can use the dynamic group in an IPS or an exempt rule and reference the IPS profile (containing the rule) in a firewall policy that you can then deploy on the device.

Delete IPS Signature Dynamic Groups

> **NOTE:**
> - You can delete only customized (user-created) IPS signature dynamic groups that are not used in an IPS or exempt rule.
> - You cannot delete predefined (system-generated) IPS signature dynamic groups.

To delete one or more customized IPS signature dynamic groups:

1. Select **Configuration > IPS > IPS Signatures**.

   The IPS Signatures page appears.

2. Select one or more customized IPS signature dynamic groups and click the delete (trash can) icon

   A warning message appears asking you to confirm the deletion.

3. Click **Yes** to proceed with the deletion.

   You are returned to the IPS Signatures page and a message indicating the status of the delete operation is displayed.

RELATED DOCUMENTATION

Create IPS Profiles | 620

About the IPS Profiles Page

IN THIS SECTION

- Tasks You Can Perform | 619
- Field Descriptions | 619

To access this page, select **Configure** > **IPS** > **IPS Profiles**.

Use intrusion prevention system (IPS) IPS Profiles page to manage IPS profiles. IPS profiles can be associated with IPS or exempt rules and deployed on a device by associating a profile with a firewall intent and deploying the firewall policy on the device.

Tasks You Can Perform

- Create an IPS profile—See "Create IPS Profiles" on page 620.
- Edit, clone, or delete an IPS profile—See "Edit, Clone, and Delete IPS Profiles" on page 621.
- Manage the IPS rules associated with an IPS profile—Click the ***IPS-Profile-Name*** to manage the IPS rules associated with the IPS profile. See "About the <IPS-Profile-Name> / Rules Page" on page 623.
- Search for IPS profiles by using keywords—Click the search icon and enter the search term in the text box and press Enter. The search results are displayed on the same page.
- Sort IPS profiles—Click a column name to sort the data in the grid (table) based on the column name.

> **NOTE:** Sorting is applicable only to some fields.

Field Descriptions

Table 203 on page 619 describes the field on the IPS Profiles page.

Table 203: Fields on the IPS Profiles Page

Field	Description
Name	Name of the IPS profile. Click the ***IPS-Profile-Name*** to manage the IPS rules associated with the IPS profile. The *IPS-Profile-Name* / Rules page appears. See "About the <IPS-Profile-Name> / Rules Page" on page 623.
Definition Type	Indicates whether the IPS profile was system-generated (PREDEFINED) or created by a user (CUSTOM).
Description	Description of the IPS profile.

RELATED DOCUMENTATION

About the IPS Signatures Page | 591

Create IPS Profiles

Contrail Service Orchestration (CSO) contains predefined intrusion prevention system (IPS) profiles that you can use. You can create customized IPS profiles from the Create IPS Profile page.

To create a customized IPS profile:

1. Select **Configuration > IPS > IPS Profiles**.

   The IPS Profiles page appears.

2. Click the add (+) icon.

   The Create IPS Profile page appears.

3. Complete the configuration according to the guidelines in Table 204 on page 620.

   NOTE: Fields marked with an asterisk (*) are mandatory.

4. Click **OK**.

   You are returned to the IPS Profiles page and a confirmation message is displayed indicating that the IPS profile is created.

After you create an IPS profile, you can add one or more IPS or exempt rules to the profile, and use the IPS profile in a firewall policy intent.

Table 204: Create IPS Profile Settings

Setting	Guideline
Name	Enter a unique name for the IPS profile that is a string of alphanumeric characters and some special characters (colon, hyphen, period, and underscore). No spaces are allowed and the maximum length is 255 characters.
Description	Enter a description for the IPS profile; the maximum length is 255 characters.

RELATED DOCUMENTATION

Create IPS or Exempt Rules | 625
Adding Firewall Policy Intents | 381

Edit, Clone, and Delete IPS Profiles

IN THIS SECTION

- Edit IPS Profiles | 621
- Clone IPS Profiles | 622
- Delete IPS Profiles | 622

You can edit, clone, or delete IPS profiles.

Edit IPS Profiles

You can edit only customized IPS profiles and not predefined (system-generated) profiles.

To edit a customized IPS profile:

1. Select **Configuration > IPS > IPS Profiles**.

   The IPS Profiles page appears.

2. Select a customized IPS profile and click the edit (pencil) icon.

   The Edit IPS Profile page appears, displaying the same fields that are presented when you create an IPS profile.

3. Modify the IPS profile fields as needed. See “Create IPS Profiles” on page 620.

   NOTE: You can only modify the description and not the name.

4. Click **OK** to save your changes.

   You are returned to the IPS Profiles page and a message that the IPS profile was successfully updated is displayed.

If the IPS profile is referenced in a firewall policy intent, then the firewall policy is marked for deployment. You must deploy the firewall policy for the changes to take effect on the device.

Clone IPS Profiles

Cloning enables you to easily create a new IPS profile based on an existing one. You can clone predefined or customized IPS profiles and modify the parameters as needed.

To clone an IPS profile:

1. Select **Configuration > IPS > IPS Profiles**.

   The IPS Profiles page appears.

2. Select an IPS profile and select **More > Clone**.

   The Clone IPS Profile page appears, displaying the same fields that are presented when you create an IPS profile.

3. Modify the IPS profile fields as needed. See "Create IPS Profiles" on page 620.

4. Click **OK** to save your changes.

   You are returned to the IPS Profiles page and a message that the IPS profile was successfully created is displayed.

Delete IPS Profiles

NOTE:

- You can delete only customized IPS profiles that are not referenced in a firewall policy intent.
- You cannot delete predefined (system-generated) IPS profiles.

To delete one or more customized IPS profiles:

1. Select **Configuration > IPS > IPS Profiles**.

   The IPS Profiles page appears.

2. Select one or more customized IPS profiles and click the delete (trash can) icon

   A warning message appears asking you to confirm the deletion.

3. Click **Yes** to proceed with the deletion.

You are returned to the IPS Profiles page and a message indicating the status of the delete operation is displayed.

RELATED DOCUMENTATION

About the <IPS-Profile-Name> / Rules Page | 623

About the <IPS-Profile-Name> / Rules Page

IN THIS SECTION

Tasks You Can Perform | 623
Field Descriptions | 624

To access this page, select **Configure** > **IPS** > **IPS Profiles** > ***IPS-Profile-Name***.

Use the *IPS-Profile-Name* / Rules page to manage intrusion prevention system (IPS) rules and exempt rules. IPS profiles can be associated with IPS or exempt rules and deployed on a device by associating the IPS profile with a firewall policy intent and deploying the firewall policy on the device.

Tasks You Can Perform

- Create an IPS rule—See "Create IPS or Exempt Rules" on page 625.
- Create an exempt rule—See "Create IPS or Exempt Rules" on page 625.
- Edit, clone, or delete IPS or exempt rules—See "Edit, Clone, and Delete IPS or Exempt Rules" on page 633.
- Search for rules by using keywords—Click the search icon and enter the search term in the text box and press Enter. The search results are displayed on the same page.
- Filter rules—Click the filter icon (funnel) and specify one or more filtering criteria. The filtered results are displayed on the same page.

NOTE: Filtering is applicable only to some fields.

Field Descriptions

Table 205 on page 624 describes the field on the *IPS-Profile-Name* / Rules page.

Table 205: Fields on the <IPS-Profile-Name> / Rules Page

Field	Description
Name	Name of the IPS rule or exempt rule.
IPS Signatures	Displays the IPS signatures associated with the IPS rule or exempt rule. If there is more than one signature associated with the rule, the number of additional signatures is displayed. Mouse over the number to view the full list of signatures.
IPS Action	• For IPS rules, displays the action to be taken when the rule is matched. • For exempt rules, displays **Not Applicable** because exempt rules are not associated with an action.

Additional Actions	• For IPS rules, displays:
	  ◦ Configured, if additional actions (to be taken when the rule is matched) are configured. Mouse over the gear icon to view the additional actions configured.
	  ◦ Not Configured, if no additional actions are configured.
	• For exempt rules, displays **Not Applicable** because exempt rules are not associated with any actions.

Details	Displays whether the rule is an IPS rule or an exempt rule. Mouse over the **Details** field and then mouse over the ellipsis (**...**) displayed to access a menu to edit, clone, or delete the rule. See "Edit, Clone, and Delete IPS or Exempt Rules" on page 633.

RELATED DOCUMENTATION

About the IPS Signatures Page | **591**

About the IPS Profiles Page | **618**

Create IPS or Exempt Rules

IN THIS SECTION

- Create IPS Rules | 625
- Create Exempt Rules | 632

You can create intrusion prevention system (IPS) rules or exempt rules only for customized IPS profiles.

Create IPS Rules

To create an IPS rule:

1. Select **Configuration > IPS > IPS Profiles**.

   The IPS Profiles page appears.

2. Click ***IPS-Profile-Name*** for the profile for which you want to create a rule.

   The *IPS-Profile-Name* / Rules page appears.

3. Select **Create > IPS Rule**.

   The parameters for an IPS rule appear inline at the top of the page.

4. Complete the configuration according to the guidelines in Table 206 on page 626.

   NOTE: Fields marked with an asterisk (*) are mandatory.

5. Click **Save** to save your changes.

   The changes are saved and a confirmation message appears at the top of the page.

   You can use the IPS profile in a firewall policy intent and deploy the firewall policy on the device, which deploy the IPS and exempt rules associated with the profile.

Table 206: Create IPS Rule Settings

Setting	Guideline
Rule Name	CSO generates a unique rule name by default. You can modify the name if needed. The name must begin with an alphanumeric character and can contain alphanumeric characters and some special characters (colons, hyphens, forward slashes, periods, and underscores); 63-character maximum.
Description	Enter a description for the rule; the maximum length is 1024 characters.
IPS Signatures	You can add one or more IPS signatures and IPS signature static and dynamic groups to be associated with the rule: 1. Click inside the text box with the **+** icon. A list of IPS signatures and IPS signature static and dynamic groups appears. 2. (Optional) Enter a search term and press Enter to filter the list of items displayed. 3. Click a list item to add it to the IPS signatures and IPS signature static or dynamic groups associated with the rule. 4. (Optional) Repeat the preceding step to add more signatures, static groups, and dynamic groups. 5. Click the **View more results** link to view the full list of IPS signatures and IPS signature static and dynamic groups. The full list is displayed in the End Points panel on the right. To add one or more signatures, static groups, or dynamic groups: a. Mouse over a list item and select the check box that appears. b. Repeat the preceding step for the other signatures, static groups, or dynamic groups that you want to add. c. Click the check mark icon (✓) at the top of the End Points panel, and select **Signatures**. The signatures, static groups, or dynamic groups that you selected are added and displayed in the IPS Signatures field.

Table 206: Create IPS Rule Settings *(continued)*

Setting	Guideline
IPS Action	Select the action to be taken when the monitored traffic matches the attack objects specified in the rules: • None—No action is taken. Use this action to only generate logs for some traffic. • Ignore—Stops scanning traffic for the rest of the connection if an attack match is found. IPS disables the rulebase for the specific connection. • Close Client and Server—Closes the connection and sends a TCP reset (RST) packet to both the client and the server. • Close Client—Closes the connection and sends an RST packet to the client, but not to the server. • Close Server—Closes the connection and sends an RST packet to the server, but not to the client. • Drop Connection—Drops all packets associated with the connection, preventing traffic for the connection from reaching its destination. Use this action to drop connections for traffic that is not prone to spoofing. • Drop Packet—Drops a matching packet before it can reach its destination but does not close the connection. Use this action to drop packets for attacks in traffic that is prone to spoofing, such as UDP traffic. Dropping a connection for such traffic could result in a denial of service that prevents you from receiving traffic from a legitimate source-IP address. • Recommended (default)—Uses the action that Juniper Networks recommends when that attack is detected. All predefined attack objects have a default action associated with them. • Diffserv Marking—Assigns the specified differentiated services code point (DSCP) value to the packet in an attack and pass the packet on normally. When you select Diffserv Marking, you must enter a DSCP value: 1. Click the **Code Point: *Vaule*** hyperlink. The Code point for Diffserve Marking action popup appears. 2. In the **Code Point** field, enter a DSCP value from 0 through 63. 3. Click **OK**. You are returned to the previous page; the value that you entered is displayed

Table 206: Create IPS Rule Settings *(continued)*

Setting	Guideline
Additional Actions	

Table 206: Create IPS Rule Settings *(continued)*

Setting	Guideline
	In addition to the IPS action, you can configure one or more of the following additional actions: • Notifications—When attacks are detected, you can choose to log the attack and create log records with attack information and send that information to the log server. To configure notifications: 1. Click the **Notification** link. The Notification page appears. 2. Complete the configuration according to the guidelines shown in Table 207 on page 630. 3. Click **OK**. You are returned to the previous page. A gear icon next to the Notification link indicates that you have configured notification settings. • IP actions—When attacks are detected, you can configure actions that you want IPS to take against future connections that use the same IP address. To configure IP actions: 1. Click the **IP Action** link. The IP Action page appears. 2. Complete the configuration according to the guidelines shown in Table 208 on page 631. 3. Click **OK**. You are returned to the previous page. A gear icon next to the IP Action link indicates that you have configured IP action settings. • Additional actions—When attacks are detected, you can configure additional actions that you want CSO to take. To configure additional actions: 1. Click the **Additional** link. The Additional page appears. 2. Complete the configuration according to the guidelines shown in Table 209 on page 632. 3. Click **OK**.

Table 206: Create IPS Rule Settings *(continued)*

Setting	Guideline
	You are returned to the previous page. A gear icon next to the Additional link indicates that you have configured additional settings.

Table 207: Notification Settings

Setting	Guideline
Attack Logging	Select the **Enable** check box to log an attack when it is detected.
Alert Flag	Select the **Enable** check box to set the alert flag in the attack log.
Log Packets	Select the **Enable** check box to log packets when an attack is detected. In response to a rule match, you can capture the packets received before and after the attack for further offline analysis of attacker behavior. You can configure the number of pre-attack and post-attack packets to be captured for this attack, and limit the duration of post-attack packet capture by specifying a timeout value. You must specify at least one of the Packets Before, Packets After, or Post Window Timeout fields.
Packets Before	Specify the number of packets received before an attack that should be captured for further analysis of the behavior of the attack. Range: 1 through 255.
Packets After	Specify the number of packets received after an attack that should be captured for further analysis of attacker behavior. Range: 1 through 255.
Post Window Timeout	Specify a time limit (in seconds) for capturing packets received after an attack. No packets are captured after the specified timeout has elapsed. Range: 1 through 1800.

Table 208: IP Action Settings

Setting	Guideline
IP Action	Select the action to be taken on future connections that use the same IP address: **NOTE:** If there is an IP action match with more than one rule, then the most severe IP action of all the matched rules is applied. In decreasing order of severity, the actions are block, close, and notify. • **None** (default)—Do not take any action. This is similar to if you did not configure the IP action. • **IP Notify**—Don't take any action on future traffic but log the event. • **IP Close**—Close future connections of new sessions that match the IP address by sending RST packets to the client and server. • **IP Block**—Block future connections of any session that matches the IP address.
IP Target	Specify how the traffic should be matched for the configured IP actions: • **None**—Do not match any traffic. • **Destination Address**—Match traffic based on the destination IP address of the attack traffic. • **Service**—For TCP and UDP, matches traffic based on the source IP address, source port, destination IP address, and destination port of the attack traffic. • **Source Address**—Matches traffic based on the source IP address of the attack traffic. • **Source Zone**—Matches traffic based on the source zone of the attack traffic. • **Source Zone Address**—Matches traffic based on the source zone and source IP address of the attack traffic. • **Zone Service**—Matches traffic based on the source zone, destination IP address, destination port, and protocol of the attack traffic.
Refresh Timeout	Select the **Enable** check box to refresh the IP action timeout (that you specify in the Timeout Value field) if future traffic matches the IP actions configured.
Timeout Value	Configure the number of seconds that you want the IP action to remain in effect. For example, if you configure a timeout of 3600 seconds (1 hour) and traffic matches the IP actions configured, the IP action remains in effect for 1 hour. Range: 0 through 64,800 seconds.
Log Taken	Select the **Enable** check box to log the information about the IP action against the traffic that matches a rule.
Log Creation	Select the **Enable** check box generate an event when the IP action filter is triggered.

Table 209: Additional Settings

Setting	Guideline
Severity	Select a severity level to override the inherited attack severity in the rules. The most dangerous level is critical, which attempts to crash your server or gain control of your network. Informational is the least dangerous level and is used by network administrators to discover holes in their security systems.
Terminal	Select the **Enable** check box to mark the IPS rule as terminal. When a terminal rule is matched, the device stops matching for the rest of the rules in that IPS profile.

Create Exempt Rules

To create an exempt rule:

1. Select **Configuration > IPS > IPS Profiles**.

   The IPS Profiles page appears.

2. Click ***IPS-Profile-Name*** for the profile for which you want to create a rule.

   The *IPS-Profile-Name* / Rules page appears.

3. Select **Create > Exempt Rule**.

   The parameters for an exempt rule appear inline at the top of the page.

4. You can configure only the following fields:

   - Rule Name
   - Description
   - IPS Signatures

   See Table 206 on page 626 for an explanation of these fields.

5. Click **Save** to save your changes.

   The changes are saved and a confirmation message appears at the top of the page.

   You can use the IPS profile in a firewall policy intent and deploy the firewall policy on the device, which deploy the IPS and exempt rules associated with the profile.

RELATED DOCUMENTATION

Adding Firewall Policy Intents | 381

Edit, Clone, and Delete IPS or Exempt Rules

IN THIS SECTION

- Edit IPS or Exempt Rules | 633
- Clone IPS or Exempt Rules | 634
- Delete IPS or Exempt Rules | 634

You can edit, clone, or delete IPS or exempt rules.

Edit IPS or Exempt Rules

You can edit IPS and exempt rules associated only with customized IPS profiles and not rules associated with predefined (system-generated) profiles.

To edit an IPS or an exempt rule:

1. Select **Configuration > IPS > IPS Profiles**.

   The IPS Profiles page appears.

2. Click ***IPS-Profile-Name*** for the profile for which you want to edit a rule.

   The *IPS-Profile-Name* / Rules page appears.

3. Mouse over the **Details** field, then mouse over the ellipsis (...) that appears, and from the menu, select **Edit**.

   The rule that you selected for editing appears inline at the top of the page.

4. Modify the rule as needed. See “Create IPS or Exempt Rules” on page 625.

> **NOTE:** You can modify all fields except the name.

5. Click **Save** to save your changes.

   The changes are saved and a confirmation message appears at the top of the page.

If the IPS or exempt rule belongs to an IPS profile that is referenced in a firewall policy intent, then the firewall policy is marked for deployment. You must deploy the firewall policy for the changes to take effect on the device.

Clone IPS or Exempt Rules

Cloning enables you to easily create a new IPS or exempt rule based on an existing one. You can clone IPS and exempt rules associated only with customized IPS profiles and not rules associated with predefined (system-generated) profiles.

To clone an IPS or an exempt rule:

1. Select **Configuration > IPS > IPS Profiles**.

   The IPS Profiles page appears.

2. Click ***IPS-Profile-Name*** for the profile for which you want to clone a rule.

   The *IPS-Profile-Name* / Rules page appears.

3. Select a rule and select **More > Clone**. Alternatively, Mouse over the **Details** field, then mouse over the ellipsis (...) that appears, and from the menu, select **Clone**.

   The rule that you selected for cloning appears inline at the top of the page.

4. Modify the rule as needed. See “Create IPS or Exempt Rules” on page 625.

5. Click **Save** to save your changes.

   The new rule is created and a confirmation message appears at the top of the page.

Delete IPS or Exempt Rules

You can delete IPS and exempt rules associated only with customized IPS profiles and not rules associated with predefined (system-generated) profiles.

To delete one or more IPS or exempt rules:

1. Select **Configuration > IPS > IPS Profiles**.

   The IPS Profiles page appears.

2. Click ***IPS-Profile-Name*** for the profile for which you want to delete one or more rules.

   The *IPS-Profile-Name* / Rules page appears.

3. Select one or more rules and click the delete (trash can) icon

   A warning message appears asking you to confirm the deletion.

4. Click **Yes** to proceed with the deletion.

   A message indicating the status of the delete operation appears at the top of the page.

If the IPS or exempt rule that you deleted belongs to an IPS profile that is referenced in a firewall policy intent, then the firewall policy is marked for deployment. You must deploy the firewall policy for the changes to take effect on the device.

RELATED DOCUMENTATION

Adding Firewall Policy Intents | **381**

CHAPTER 20

Managing SSL Proxies

IN THIS CHAPTER

- SSL Forward Proxy Overview | 636
- About the SSL Proxy Policy Page | 642
- Creating SSL Proxy Policy Intents | 644
- Editing, Cloning, and Deleting SSL Proxy Policy Intents | 648
- Understanding How SSL Proxy Policy Intents Are Applied | 651
- About the SSL Proxy Profiles Page | 653
- Creating SSL Forward Proxy Profiles | 655
- Editing, Cloning, and Deleting SSL Forward Proxy Profiles | 659
- Configuring and Deploying an SSL Forward Proxy Policy | 662

SSL Forward Proxy Overview

IN THIS SECTION

- Supported Ciphers in Proxy Mode | 638
- Server Authentication | 639
- Root CA | 640
- Trusted CA List | 640
- Session Resumption | 641
- SSL Proxy Logs | 641

Secure Sockets Layer (SSL) is an application-level protocol that provides encryption technology for the Internet. SSL, also called *Transport Layer Security* (TLS), ensures the secure transmission of data between a client and a server through a combination of privacy, authentication, confidentiality, and data integrity. SSL relies on certificates and private–public key exchange pairs for this level of security.

Server authentication guards against fraudulent transmissions by enabling a Web browser to validate the identity of a Web server. Confidentiality mechanisms ensure that communications are private. SSL enforces confidentiality by encrypting data to prevent unauthorized users from eavesdropping on electronic communications. Finally, message integrity ensures that the contents of a communication have not been tampered with.

SSL forward proxy is a transparent proxy; that is, it performs SSL encryption and decryption between the client and the server, but neither the server nor the client can detect its presence. SSL forward proxy ensures that it has the keys to encrypt and decrypt the payload:

- For the server, SSL forward proxy acts as a client—Because SSL forward proxy generates the shared pre-master key, it determines the keys to encrypt and decrypt.
- For the client, SSL forward proxy acts as a server—SSL forward proxy first authenticates the original server and replaces the public key in the original server certificate with a key that is known to it. It then generates a new certificate by replacing the original issuer of the certificate with its own identity and signs this new certificate with its own public key (provided as a part of the proxy profile configuration). When the client accepts such a certificate, it sends a shared pre-master key encrypted with the public key on the certificate. Because SSL forward proxy replaced the original key with its own key, it is able to receive the shared pre-master key. Decryption and encryption take place in each direction (client and server), and the keys are different for both encryption and decryption.

Figure 13 on page 638 shows how SSL forward proxy works on an encrypted payload. When application firewall (AppFW) is configured, SSL forward proxy acts as an SSL server terminating the SSL session from the client and a new SSL session is established to the server. The device decrypts and then re-encrypts all SSL forward proxy traffic. SSL forward proxy uses the following services:

- SSL-T-SSL terminator on the client side.
- SSL-I-SSL initiator on the server side.
- Configured AppFW services use the decrypted SSL sessions.

Figure 13: SSL Forward Proxy on an Encrypted Payload

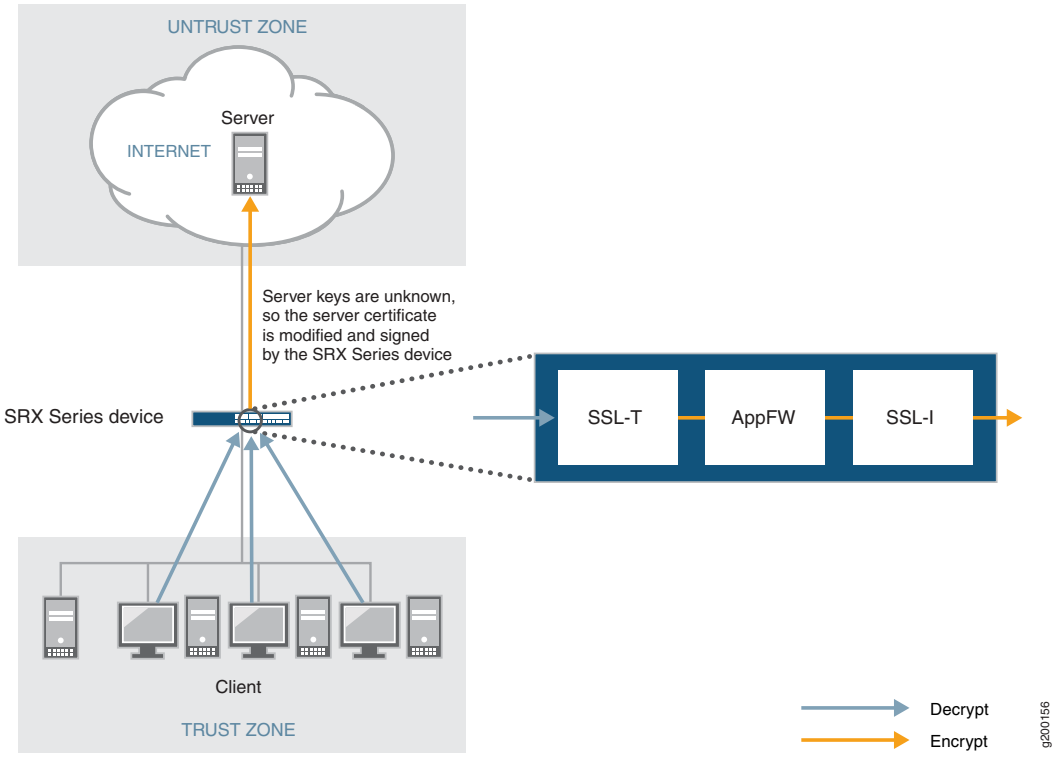

This topic has the following sections:

Supported Ciphers in Proxy Mode

An SSL cipher comprises encryption ciphers, authentication method, and compression. Table 210 on page 638 displays a list of supported ciphers. NULL ciphers are excluded.

The following SSL protocols are supported:

- SSLv3
- TLS1

Table 210: Supported Ciphers in Proxy Mode

SSL Cipher	Key Exchange Algorithm	Data Encryption	Message Integrity
RSA_WITH_RC4_128_MD5	RSA key exchange	128-bit RC4	Message Digest 5 (MD5) hash

Table 210: Supported Ciphers in Proxy Mode *(continued)*

SSL Cipher	Key Exchange Algorithm	Data Encryption	Message Integrity
RSA_WITH_RC4_128_SHA	RSA key exchange	128-bit RC4	Secure Hash Algorithm (SHA) hash
RSA_WITH_DES_CBC_SHA	RSA key exchange	DES CBC	SHA hash
RSA_WITH_3DES_EDE_CBC_SHA	RSA key exchange	3DES EDE/CBC	SHA hash
RSA_WITH_AES_128_CBC_SHA	RSA key exchange	128-bit AES/CBC	SHA hash
RSA_WITH_AES_256_CBC_SHA	RSA key exchange	256-bit AES/CBC	SHA hash
RSA_EXPORT_WITH_RC4_40_MD5	RSA-export	40-bit RC4	MD5 hash
RSA_EXPORT_WITH_DES40_CBC_SHA	RSA-export	40-bit DES/CBC	SHA hash
RSA_EXPORT1024_WITH_DES_CBC_SHA	RSA 1024 bit export	DES/CBC	SHA hash
RSA_EXPORT1024_WITH_RC4_56_MD5	RSA 1024 bit export	56-bit RC4	MD5 hash
RSA_EXPORT1024_WITH_RC4_56_SHA	RSA 1024 bit export	56-bit RC4	SHA hash
RSA-WITH-AES-256-GCM-SHA384	RSA key exchange	256-bit AES/GCM	SHA384 hash
RSA-WITH-AES-256-CBC-SHA256	RSA key exchange	256-bit AES/CBC	SHA256 hash
RSA-WITH-AES-128-GCM-SHA256	RSA key exchange	128-bit AES/GCM	SHA256 hash
RSA-WITH-AES-128-CBC-SHA256	RSA key exchange	128-bit AES/CBC	SHA256 hash

Server Authentication

Implicit trust between the client and the device (because the client accepts the certificate generated by the device) is an important aspect of SSL proxy. It is extremely important that server authentication is not compromised; however, in reality, self-signed certificates and certificates with anomalies are in abundance. Anomalies can include expired certificates, instances of common name not matching a domain name, and so forth.

You can specify that the SSL forward proxy should ignore server authentication completely. In this case, SSL forward proxy ignores errors encountered during the server certificate verification process (such as CA signature verification failure, self-signed certificates, and certificate expiry).

You can specify whether the SSL proxy should ignore server authentication errors or not during the creation of an SSL forward proxy profile.

- If you specify that server authentication errors should *not* be ignored, the following scenarios occur:
  - If authentication succeeds, a new certificate is generated by replacing the keys and changing the issuer name to the issuer name that is configured in the root CA certificate in the proxy profile.
  - If authentication fails, the connection is dropped.
- If you specify that server authentication errors should be ignored, the following scenarios occur:

  NOTE: We do not recommend that you configure this option for authentication because configuring it results in websites not being authenticated at all. However, you can use this option to effectively identify the root cause for dropped SSL sessions.

  - If the certificate is self-signed, a new certificate is generated by replacing the keys only. The issuer name is not changed. This ensures that the client browser displays a warning that the certificate is not valid.
  - If the certificate has expired or if the common name does not match the domain name, a new certificate is generated by replacing the keys and changing the issuer name to SSL-PROXY: DUMMY_CERT:GENERATED DUE TO SRVR AUTH FAILURE. This ensures that the client browser displays a warning that the certificate is not valid.

Root CA

In a public key infrastructure (PKI) hierarchy, the root CA is at the top of the trust path. The root CA identifies the server certificate as a trusted certificate.

Trusted CA List

SSL forward proxy ensures secure transmission of data between a client and a server. Before establishing a secure connection, SSL forward proxy checks certificate authority (CA) certificates to verify signatures on server certificates. For this reason, a reasonable list of trusted CA certificates is required to effectively authenticate servers.

Session Resumption

An SSL session refers to the set of parameters and encryption keys that are created when a full handshake is performed. A connection is the conversation or active data transfer that occurs within the session. The computational overhead of a complete SSL handshake and generation of master keys is considerable. In short-lived sessions, the time taken for the SSL handshake can be more than the time for data transfer. To improve throughput and still maintain an appropriate level of security, SSL session resumption provides a mechanism for caching sessions so that session information, such as the pre-master secret key and agreed-upon ciphers, can be cached for both the client and the server. The cached information is identified by a session ID. In subsequent connections, both parties agree to use the session ID to retrieve the information rather than create a new pre-master secret key. Session resumption shortens the handshake process and accelerates SSL transactions.

SSL Proxy Logs

When logging is enabled in an SSL proxy profile, the SSL proxy can generate the messages shown in Table 211 on page 641.

Table 211: SSL Proxy Logs

Log Type	Description
SSL_PROXY_SSL_SESSION_DROP	Logs generated when a session is dropped by SSL proxy.
SSL_PROXY_SSL_SESSION_ALLOW	Logs generated when a session is processed by SSL proxy even after encountering some minor errors.
SSL_PROXY_SESSION_IGNORE	Logs generated if non-SSL sessions are initially mistaken as SSL sessions.
SSL_PROXY_SESSION_WHITELIST	Logs generated when a session is allowed.
SSL_PROXY_ERROR	Logs used for reporting errors.
SSL_PROXY_WARNING	Logs used for reporting warnings.
SSL_PROXY_INFO	Logs used for reporting general information.

All logs contain similar information; the message field contains the reason for the log generation. One of three prefixes shown in Table 212 on page 642 identifies the source of the message. Other fields are descriptively labeled.

Table 212: SSL Proxy Log Prefixes

Prefix	Description
system	Logs generated because of errors related to the device or an action taken as part of the SSL proxy profile. Most logs fall into this category.
openssl error	Logs generated during the handshake process if an error is detected by the openssl library.
certificate error	Logs generated during the handshake process if an error is detected in the certificate (X.509 related errors).

RELATED DOCUMENTATION

- About the SSL Proxy Policy Page | 642
- About the SSL Proxy Profiles Page | 653
- Certificates Overview | 349

About the SSL Proxy Policy Page

To access this page, select **Configuration > SSL Proxy > Policy** in Customer Portal.

Use the SSL Proxy Policy page to view and manage SSL proxy policy intents. You can also deploy the SSL proxy policy immediately or schedule the deployment for later.

NOTE:

- When an SSL proxy intent is deployed, the corresponding certificates used in the SSL profile (associated with the SSL proxy intent) are automatically deployed to the applicable sites.
- If the application firewall (AppFW) service is not configured in the corresponding firewall policy intent, then the SSL forward proxy services are bypassed even if an SSL proxy profile is attached to a firewall policy. Therefore, ensure that AppFW is configured for the firewall policy intents that should go through SSL inspection. If AppFw is not included in the policy intent, this does not cause an error; however, the SSL proxy action does not take place even though sessions are matched.

Tasks You Can Perform

You can perform the following tasks from this page:

- Create SSL proxy policy intents—See "Creating SSL Proxy Policy Intents" on page 644.
- Edit, clone, or delete SSL proxy policy intents—See "Editing, Cloning, and Deleting SSL Proxy Policy Intents" on page 648.
- Search for SSL proxy policy intents by using keywords—Click the search icon and enter the search term in the text box and press Enter. The search results are displayed on the same page.
- Filter SSL proxy policy intents—Click the filter icon and select whether you want to show or hide column filters or apply a quick filter. Depending on your selection, you can filter the policy intents based on source, destination, or both, or view the filtered results. The filtered results are displayed on the same page.
- Deploy the SSL proxy policy—See "Deploying Policies" on page 668.

Field Descriptions

Table 213 on page 643 describes the fields on SSL Proxy Policy page.

Table 213: SSL Proxy Policy Page Fields

Field	Description
Total Intents	Total number of policy intents in the SSL proxy policy.
Undeployed	Number of SSL proxy policy intents that have not yet been deployed.
For each SSL proxy policy intent, the following information is displayed in a grid:	
Source	Source endpoints to which an SSL proxy policy intent applies.
Destination	Destination endpoints to which an SSL proxy policy intent applies.
SSL Proxy Profile	Name of the SSL proxy profile associated with the policy intent.
Options	Name and description of the SSL proxy policy intent.

RELATED DOCUMENTATION

SSL Forward Proxy Overview | 636

Creating SSL Proxy Policy Intents

You can configure an SSL proxy policy intent inline on the SSL Proxy Policy page. An SSL proxy policy intent enables you to configure an SSL proxy between source and destination endpoints by associating the latter with an SSL proxy profile.

To create an SSL proxy policy intent:

1. Select **Configuration > SSL Proxy > Policy** in Customer Portal.

   The SSL Proxy Policy page appears.

2. Click the add icon (**+**).

   The options to create policy intents appear inline on the SSL Proxy Policy page.

3. Enter the policy intent information according to the guidelines provided in Table 214 on page 645

4. Click **Save**.

   The SSL proxy policy intent is saved and a confirmation message is displayed.

> **NOTE:** After the policy intent is created, you must deploy the policy to ensure that the changes take effect on the applicable sites. When an SSL proxy policy intent is created, the **Undeployed** field is incremented by one indicating that intents are pending deployment.

Table 214: Create SSL Proxy Policy Intent Settings

Setting	Guideline
Source	A source endpoint can be an IP address, an IP address group, a site, a site group, or a department, or or a combination of these. **NOTE:** A source IP address value of **Any** signifies any IP address from any site. Specify one or more source endpoints in one of the following ways: • Click the add icon (**+**) and select the endpoints from the list of previously configured endpoints. • Filter the endpoints by entering a search term or one or more predefined keywords in the **Source** field and select one or more endpoints. Table 215 on page 647 displays the list of predefined keywords. • Click the **View more results** link to view additional configured endpoints. The list of endpoints is displayed in the **End Points** panel on the right. Do one of the following: ◦ To add one endpoint at a time, select an endpoint and click the check mark icon (✓) that appears when you hover over the endpoint. ◦ To add multiple endpoints, select one or more endpoints that you want to add, click the check mark icon (✓) at the top of the **End Points** panel, and select **Source**. ◦ Filter the endpoints by entering a search term or one or more predefined keywords in the **End Points** field and select one or more endpoints. Table 215 on page 647 displays the list of predefined keywords. **NOTE:** You can also create endpoints by clicking the add icon (**+**) in the End Points panel. Table 216 on page 648 displays the endpoints that can be created.

Table 214: Create SSL Proxy Policy Intent Settings *(continued)*

Setting	Guideline
Destination	A destination endpoint can be an IP address, an IP address group, a site, a site group, or a department, or or a combination of these. **NOTE:** A destination IP address value of **Any** signifies traffic going to the Internet (any address). Traffic within sites (internal traffic) is not covered by the destination IP address value of **Any**. If you want to cover traffic between two sites, ensure that the sites are included in both the source and destination endpoints. Specify one or more destination endpoints in one of the following ways: • Click the add icon (**+**) and select the endpoints from the list of previously configured endpoints. • Filter the endpoints by entering a search term or one or more predefined keywords in the **Destination** field and select one or more endpoints. Table 215 on page 647 displays the list of predefined keywords. • Click the **View more results** link to view additional configured endpoints. The list of endpoints is displayed in the **End Points** panel on the right. Do one of the following: ◦ To add one endpoint at a time, select an endpoint and click the check mark icon (✓) that appears when you hover over the endpoint. ◦ To add multiple endpoints, select one or more endpoints that you want to add, click the check mark icon (✓) at the top of the **End Points** panel, and select **Destination**. ◦ Filter the endpoints by entering a search term or one or more predefined keywords in the **End Points** field and select one or more endpoints. Table 215 on page 647 displays the list of predefined keywords. **NOTE:** You can also create endpoints by clicking the add icon (**+**) in the End Points panel. Table 216 on page 648 displays the endpoints that can be created.

Table 214: Create SSL Proxy Policy Intent Settings *(continued)*

Setting	Guideline
SSL Proxy Profile	Specify an SSL proxy profile to associate with the SSL proxy policy intent in one of the following ways: • Click the add icon (+) and select the SSL proxy profile from the list of previously configured profiles. • Filter the profiles by entering a search term in the **SSL Proxy Profile** field and select a profile. • Create a SSL proxy profile—Click the **Add New Profile** link. The Create SSL Proxy Profiles page appears. See"Creating SSL Forward Proxy Profiles" on page 655. **NOTE:** You can also create profiles by clicking the add icon (+) in the End Points panel and selecting **SSL Proxy Profiles**. • Click the **View more results** link to view additional configured profiles. The list of SSL proxy profiles is displayed in the **End Points** panel on the right. To add a profile, select it and click the check mark icon (✓) that appears when you hover over the profile.
Details	Enter the name of the SSL proxy policy intent in the first text box. If you do not enter a name, the system-generated name is used. The name that you enter must begin with an alphanumeric character and can contain alphanumeric characters and some special characters (- _). The maximum length is 63 characters. Enter the description of the SSL proxy policy intent in the second text box.

Table 215: Keywords for Filtering Endpoints

Endpoint	Keyword	Applicable to
Address or Address Group	**addr** or **ADDR**	Source Destination
Site	**site** or **SITE**	Source Destination
Site Group	**stgp** or **STGP**	Source Destination
Department	**dept** or **DEPT**	Source Destination

Table 216: Creating Endpoints

Endpoint	Procedure
Address or Address Group	Click the add icon (+) and select **Address**. The Create Addresses page appears. See "Creating Addresses or Address Groups" on page 737.
Site Group	Click the add icon (+) and select **Site Group**. The Create Site Group page appears. See "Creating Site Groups" on page 197.
Department	Click the add icon (+) and select **Department**. The Create Department page appears. See "Adding a Department" on page 767.

RELATED DOCUMENTATION

SSL Forward Proxy Overview | 636

Editing, Cloning, and Deleting SSL Proxy Policy Intents

IN THIS SECTION

- Editing SSL Proxy Policy Intents | 649
- Cloning SSL Proxy Policy Intents | 649
- Deleting SSL Proxy Policy Intents | 650

You can edit, clone, and delete SSL proxy policy intents from the SSL Proxy Policy page. This topic has the following sections:

Editing SSL Proxy Policy Intents

To modify the parameters configured for an SSL proxy policy intent:

1. Select **Configuration > SSL Proxy > Policy** in Customer Portal.

   The SSL Proxy Policy page appears, displaying the intents associated with the policy.

2. Hover over the SSL proxy policy intent that you want to edit, and then click the edit icon (pencil symbol) that appears on the right side of the intent.

   You can now modify the policy intent inline on the SSL Proxy Policy page.

3. Modify the parameters following the guidelines provided in "Creating SSL Proxy Policy Intents" on page 644.

4. Click **Save** to save your changes.

   The SSL proxy policy intent is saved and a confirmation message is displayed.

> **NOTE:** After a policy intent is modified, you must redeploy the policy to ensure that the changes take effect on the relevant sites. When an SSL proxy policy intent is modified, the **Undeployed** field is incremented by one indicating that intents are pending deployment.

Cloning SSL Proxy Policy Intents

Cloning enables you to easily create a new SSL proxy policy intent based on an existing one.

To clone an SSL proxy policy intent:

1. Select **Configuration > SSL Proxy > Policy** in Customer Portal.

   The **SSL Proxy Policy** page appears, displaying the intents associated with the policy.

2. Hover over the SSL proxy policy intent that you want to clone, and then click the clone icon that appears on the right side of the intent.

   You can modify the cloned policy intent inline on the SSL Proxy Policy page.

3. Modify the parameters following the guidelines provided in "Creating SSL Proxy Policy Intents" on page 644.

4. Click **Save** to save your changes.

The SSL proxy policy intent is cloned and a confirmation message is displayed.

> **NOTE:** After a policy intent is cloned, you must redeploy the policy to ensure that the changes take effect on the relevant sites. When an SSL proxy policy intent is cloned, the **Undeployed** field is incremented by one indicating that one or more intents are pending deployment.

Deleting SSL Proxy Policy Intents

To delete one or more SSL proxy policy intents:

1. Select **Configuration > SSL Proxy > Policy** in Customer Portal.

   The **SSL Proxy Policy** page appears, displaying the intents associated with the policy.

2. Select the SSL proxy policy intents that you want to delete and then click the delete icon (**X**).

   You are asked to confirm the delete operation.

3. Click **Yes** to delete the selected SSL proxy policy intents.

   A confirmation message appears indicating the status of the delete operation.

> **NOTE:** After one or more policy intents are deleted, you must redeploy the policy to ensure that the changes take effect on the applicable sites.

RELATED DOCUMENTATION

About the SSL Proxy Policy Page | 642

Understanding How SSL Proxy Policy Intents Are Applied

IN THIS SECTION

- Example 1: Firewall Policy Intent and SSL Proxy Policy Intent Match | 651
- Example 2: Firewall Policy Intent and SSL Proxy Policy Intent Do Not Match | 652
- Example 3: Applying SSL Proxy Policy Intents on Internal (Site-to-Site) Traffic | 652

When you deploy an SSL proxy policy, SSL proxy profiles are deployed to the applicable sites based on SSL proxy policy intents. The deployments of firewall and SSL polices are related in that firewall policy deployments take into account the last-deployed SSL snapshots and vice versa. Therefore, even if an SSL proxy profile is deployed to the applicable sites, it is *applied* only to traffic to which the firewall policy intent applies.

The decision regarding *which* SSL proxy profile is attached to a firewall policy intent is based on matching criteria between SSL proxy policy and firewall policy intents. In addition, if there is a match between the SSL proxy policy intent and the firewall policy intent, the SSL profile is applied *only* to the policy intents that are common between the firewall and the SSL proxy policies.

The following examples demonstrate the matching logic between SSL proxy policy and firewall policy intents.

Example 1: Firewall Policy Intent and SSL Proxy Policy Intent Match

Table 217 on page 651 shows an example of a firewall policy intent and an SSL proxy policy intent that match, which means that the SSL proxy profile attaches to the firewall policy intent. In this case, the firewall policy intent has a source and destination of **Any** IP address, which signifies traffic from any IP address from any site to any IP address on the Internet. The SSL proxy policy intent has a source of **Any** IP address, which signifies any IP address *from* any site, and a destination IP address of 198.51.100.0.

Therefore, there is a match between the firewall policy intent and the SSL proxy policy intent and the SSL proxy profile is applied *only* to traffic from any IP address of any site to the IP address 198.51.100.0.

Table 217: (Example) Match Between Firewall Policy Intent and SSL Proxy Policy Intent

Type	Source	Destination	Action or Profile
Firewall policy intent	IP address—Any	IP address—Any	Allow
SSL proxy policy intent	IP address—Any	IP address—198.51.100.0	SSL-Profile-1

Example 2: Firewall Policy Intent and SSL Proxy Policy Intent Do Not Match

Table 218 on page 652 shows an example of a firewall policy intent and an SSL proxy policy intent that do not match, which means that the SSL proxy profiles do not attach.

Although, at first glance, it *appears* that an SSL proxy policy intent with a source and destination IP address **Any** should match a firewall policy intent with a source IP address **Any** and destination department Finance, this is not the case because of what the IP address **Any** signifies in the destination.

For both firewall and SSL proxy policy intents:

- A source IP address value of **Any** signifies any IP address *from* any site.
- A destination IP address value of **Any** signifies traffic going *to* the Internet—that is, to any IP address on the Internet. Traffic *within* sites (internal traffic) is not covered by the destination IP address value of **Any**.

In this example, the firewall policy intent applies to traffic from any IP address (from any site) to the Finance department. However, the SSL proxy policy intent applies to traffic from any IP address (from any site) to any IP address on the Internet. This means that there is no match between the firewall policy intent and the SSL proxy policy intent and the SSL proxy profile does not attach.

Table 218: (Example) No Match Between Firewall Policy Intent and SSL Proxy Policy Intent

Type	Source	Destination	Action or Profile
Firewall policy intent	IP address—Any	Department—Finance	Allow
SSL proxy policy intent	IP address—Any	IP address—Any	SSL-Profile-2

Example 3: Applying SSL Proxy Policy Intents on Internal (Site-to-Site) Traffic

> **NOTE:** SSL forward proxy typically might not be used for site-to-site traffic, but this example is provided as an explanation of how an SSL proxy policy intent applies to site-to-site traffic.

Consider a scenario in which you have three sites (A, B, C) and you want to configure an SSL proxy for traffic between the sites. Table 219 on page 653 displays the firewall policy and SSL proxy policy intents that you can use for such a scenario.

Both the firewall policy intent and the SSL proxy policy intent use Site A, Site B, and Site C as the source and destination. Therefore, the firewall policy intent and the SSL proxy policy intent match, and the SSL proxy profile attaches to the firewall policy intent.

NOTE: The destination must be Site A, Site B, and Site C because the destination IP address **Any** signifies any IP address on the *Internet*.

Table 219: (Example) Firewall Policy and SSL Proxy Policy Intents for Site-to-Site Traffic

Type	Source	Destination	Action or Profile
Firewall Policy Intent	Site A, Site B, Site C	Site A, Site B, Site C	Allow
SSL Proxy Policy Intent	Site A, Site B, Site C	Site A, Site B, Site C	SSL-Profile-3

RELATED DOCUMENTATION

SSL Forward Proxy Overview | 636

Configuring and Deploying an SSL Forward Proxy Policy | 662

About the SSL Proxy Profiles Page

To access this page, click **Configuration > SSL Proxy > Profiles** in Customer Portal.

Use the SSL Proxy Profiles page to view and manage SSL proxy profiles.

Tasks You Can Perform

You can perform the following tasks from this page:

- Create an SSL proxy profile—See "Creating SSL Forward Proxy Profiles" on page 655.
- Edit, clone, or delete an SSL proxy profile—See "Editing, Cloning, and Deleting SSL Forward Proxy Profiles" on page 659.
- View the details of an SSL proxy profile—Select the SSL proxy profile for which you want to view the details and from the More or right-click menu, select **Detailed View**. The View SSL Proxy Profile Details page appears. Table 221 on page 654 describes the fields on this page.
- Search for SSL proxy profiles using keywords—Click the search icon and enter the search term in the text box and press Enter. The search results are displayed on the same page.

Widget Descriptions

Table 220 on page 654 describes the fields on the SSL Proxy Profiles page.

Table 220: Fields on the SSL Proxy Profiles Page

Field	Description
Name	Name of the SSL proxy profile.
Preferred Cipher	Preferred cipher associated with the profile.
Custom Ciphers	The set of ciphers, if the preferred cipher is **Custom**, which the SSH server uses to perform encryption and decryption functions.
Exempted Address	Addresses that can are exempted from SSL forward proxy processing.
Description	Description of the SSL proxy profile.
Root Certificate	Root certificate associated with the SSL proxy profile.

Table 221: View SSL Forward Proxy Profile Details Page Fields

Field	Description
General Information	
Name	Name of the SSL proxy profile.
Description	Description of the SSL proxy profile.
Preferred Cipher	Preferred cipher associated with the proxy profile.
Custom Ciphers	The set of ciphers, if the preferred cipher is **Custom**, which the SSH server uses to perform encryption and decryption functions.
Flow Trace Enabled	Indicates whether flow tracing is enabled or disabled.
Certificates	Displays the root certificate and the trusted certificate authorities associated with the root certificate.
Exempted Address	Addresses that can are exempted from SSL forward proxy processing.
Exempted URL Categories	URL categories that are exempted from SSL forward proxy processing.
Actions	

Table 221: View SSL Forward Proxy Profile Details Page Fields *(continued)*

Field	Description
Ignore	Indicates whether server authentication failure is ignored (**Enabled**) or not (**Disabled**).
Session Resumption	Indicates whether session information is cached to enable session resumption (**Enabled**) or not (**Disabled**).
Logging	If logging is enabled, indicates the type of events that are logged.
Renegotiation	Indicates the type of renegotiation required if there is a change in SSL parameters after a session is created and SSL tunnel transport is established.

RELATED DOCUMENTATION

About the SSL Proxy Policy Page | 642

Creating SSL Forward Proxy Profiles

Use this page to configure SSL forward proxy profiles. SSL proxy is enabled as an application service within a security policy. You specify the traffic that you want the SSL proxy enabled on as match criteria and then specify the SSL proxy profile to be applied to the traffic.

To create an SSL forward proxy profile:

> **NOTE:** Ensure that you have a root certificate imported for the tenant before you create an SSL forward proxy profile. You can import SSL certificates (root and trusted) from the Certificates page (**Administration > Certificates**) and associate the certificates with SSL forward proxy profiles.

1. Select **Configuration > SSL Proxy > Profiles** in Customer Portal.

   The SSL Proxy Profiles page appears.

2. Click the add icon (+) to create an SSL forward proxy profile.

   The Create SSL Proxy Profiles page appears.

3. Complete the configuration according to the guidelines provided in Table 222 on page 656.

> **NOTE:** Fields marked with an asterisk (*) are mandatory.

4. Click **OK**.

   An SSL forward proxy profile is created. You are returned to the SSL Proxy Profiles page where a confirmation message is displayed.

   The SSL forward proxy profile can be used in an SSL proxy policy intent (**Configuration > SSL Proxy > Policy**).

Table 222: Creating SSL Forward Proxy Profile Settings

Setting	Guideline
General Information	
Name	Enter a unique name for the profile, which is string of alphanumeric characters and some special characters (- _). No spaces are allowed and the maximum length is 63 characters.
Description	Enter a description for the profile. The maximum length is 255 characters.
Preferred Cipher	Select a preferred cipher. Preferred ciphers enable you to define an SSL cipher that can be used with acceptable key strength. You can select from the following categories: • **None** (Default)—Do not specify a preferred cipher. • **Medium**—Use ciphers with key strength of 128 bits or greater. • **Strong**—Use ciphers with key strength of 168 bits or greater. • **Weak**—Use ciphers with key strength of 40 bits or greater. • **Custom**—Configure a custom cipher suite.

Table 222: Creating SSL Forward Proxy Profile Settings *(continued)*

Setting	Guideline
Custom Ciphers	If you specified **Custom** as the preferred cipher, you can define a custom cipher list by selecting ciphers. Select the set of ciphers that the SSH server can use to perform encryption and decryption functions. The available custom ciphers are: • rsa-with-RC4-128-md5—RSA, 128- bit RC4, MD5 hash • rsa-with-RC4-128-sha—RSA, 128-bit RC4, SHA hash • rsa-with-des-cbc-sha—RSA, DES/CBC, SHA hash • rsa-with-3DES-ede-cbc-sha—RSA, 3DES EDE/CBC, SHA hash • rsa-with-aes-128-cbc-sha—RSA, 128-bit AES/CBC, SHA hash • rsa-with-aes-256-cbc-sha—RSA, 256 bit AES/CBC, SHA hash • rsa-export-with-rc4-40-md5—RSA-export, 40 bit RC4, MD5 hash • rsa-export-with-des40-cbc-sha—RSA-export, 40 bit DES/CBC, SHA hash • rsa-export1024-with-des-cbc-sha—RSA 1024 bit export, DES/CBC, SHA hash • rsa-export1024-with-rc4-56-md5—RSA 1024 bit export, 56 bit RC4, MD5 hash • rsa-export1024-with-rc4-56-sha—RSA 1024 bit export, 56 bit RC4, SHA hash • rsa-with-aes-256-gcm-sha384—RSA, 256 bit AES/GCM, SHA384 hash • rsa-with-aes-256-cbc-sha256—RSA, 256 bit AES/CBC, SHA256 hash • rsa-with-aes-128-gcm-sha256—RSA, 128 bit AES/GCM, SHA256 hash • rsa-with-aes-128-cbc-sha256—RSA, 256 bit AES/CBC, SHA256 hash • ecdhe-rsa-with-aes-256-gcm-sha384—ECDHE, RSA, 256 bit AES/GCM, SHA384 hash • ecdhe-rsa-with-aes-256-cbc-sha384—ECDHE, RSA, 256 bit AES/CBC, SHA384 hash • ecdhe-rsa-with-aes-256-cbc-sha—ECDHE, RSA, 256 bit AES/CBC, SHA hash • ecdhe-rsa-with-aes-3des-ede-cbc-sha—ECDHE, RSA, 3DES, EDE/CBC, SHA hash • ecdhe-rsa-with-aes-128-gcm-sha256—ECDHE, RSA, 128 bit AES/GCM, SHA256 hash • ecdhe-rsa-with-aes-128-cbc-sha256—ECDHE, RSA, 128 bit AES/CBC, SHA256 hash • ecdhe-rsa-with-aes-128-cbc-sha—ECDHE, RSA, 128 bit AES/CBC, SHA hash
Flow Trace	Select this option to enable flow tracing to enable the troubleshooting of policy-related issues.
Root Certificate	Select or add a root certificate. In a public key infrastructure (PKI) hierarchy, the root certificate authority (CA) is at the top of the trust path.

Table 222: Creating SSL Forward Proxy Profile Settings *(continued)*

Setting	Guideline
Trusted Certificate Authorities	Choose whether you want to add all trusted certificates present on the device (**All**) or select specific trusted certificates. Before establishing a secure connection, the SSL proxy checks CA certificates to verify signatures on server certificates. **NOTE:** • Specifying that all trusted certificates should be used means that all trusted certificates on a particular device (site) will be used during SSL policy deployment. • If you specify that all trusted certificates should be used in an SSL forward proxy profile, you must ensure that at least one trusted certificate is installed on the device.
Actions	
Exempted Addresses	Exempted addresses include addresses that you want to exempt from undergoing SSL proxy processing. To specify exempted addressees, select one or more addresses in the **Available** column and click the forward arrow to confirm your selection. The selected addresses are then displayed in the **Selected** column. These addresses are used to create allowlists that bypass SSL forward proxy processing. Because SSL encryption and decryption are complicated and expensive procedures, network administrators can selectively bypass SSL proxy processing for some sessions. Such sessions typically include connections and transactions with trusted servers or domains with which network administrators are very familiar. There are also legal requirements to exempt financial and banking sites. Such exemptions are achieved by configuring the IP addresses or domain names of the servers under allowlists. **NOTE:** You can also add addresses by clicking **Add New Address**. The Create Addresses page appears. See "Creating Addresses or Address Groups" on page 737.
Exempted URL Categories	Select the previously defined URL categories to create allowlists that bypass SSL forward proxy processing. The selected URL categories are exempted during SSL inspection.
Server Auth Failure	Select this check box to ignore errors encountered during the server certificate verification process (such as CA signature verification failure, self-signed certificates, and certificate expiry). This check box is cleared by default. We do not recommend this option for authentication, because configuring it results in websites not being authenticated at all. However, you can use this option to effectively identify the root cause for dropped SSL sessions.

Table 222: Creating SSL Forward Proxy Profile Settings *(continued)*

Setting	Guideline
Session Resumption	Select this check box to disable session resumption. This check box is cleared by default. To improve throughput and still maintain an appropriate level of security, SSL session resumption provides a session-caching mechanism so that session information, such as the pre-master secret key and agreed-upon ciphers, can be cached for both the client and server.
Logging	Select one or more events to be logged. You can choose to log all events, warnings, general information, errors, or different sessions (added to the allowlist, allowed, dropped, or ignored). Logging is disabled by default.
Renegotiation	Select one of the following options if a change in SSL parameters requires renegotiation: • **None** (default)—Indicates that renegotiation is not required. • **Allow**—Allow secure and nonsecure renegotiation. • **Allow-secure**—Allow secure negotiation only. • **Drop**—Drop session on renegotiation request. After a session is created and SSL tunnel transport has been established, a change in SSL parameters requires renegotiation. SSL forward proxy supports both secure (RFC 5746) and nonsecure (TLS v1.0 and SSL v3) renegotiation. When session resumption is enabled, session renegotiation is useful in the following situations: • Cipher keys need to be refreshed after a prolonged SSL session. • Stronger ciphers need to be applied for a more secure connection.

RELATED DOCUMENTATION

About the SSL Proxy Policy Page | 642

Editing, Cloning, and Deleting SSL Forward Proxy Profiles

IN THIS SECTION

- Editing SSL Forward Proxy Profiles | 660
- Cloning SSL Forward Proxy Profiles | 660
- Deleting SSL Forward Proxy Profiles | 661

You can edit, clone, and delete SSL forward proxy profiles from the SSL Proxy Profiles page. This topic has the following sections:

Editing SSL Forward Proxy Profiles

To modify the parameters configured for an SSL forward proxy profile:

> **NOTE:** If an SSL forward proxy profile is already used in an SSL proxy policy intent, we recommend that you do not modify the profile name. If you want to create a profile with a new name, clone the existing profile and modify the name.

1. Select **Configuration > SSL Proxy > Profiles**.

   The SSL Proxy Profiles page appears, displaying the existing SSL forward proxy profiles.

2. Select the SSL forward proxy profile that you want to edit and click the edit icon (pencil). Alternatively, right-click a profile and select **Edit Profile**.

   The Edit SSL Proxy Profile page appears showing the same fields that are presented when you create an SSL forward proxy profile.

3. Modify the SSL forward proxy profile fields as needed.

4. Click **OK** to save your changes.

   You are taken to the SSL Proxy Profiles page. A confirmation message appears, indicating the status of the edit operation.

   > **NOTE:** If an SSL forward proxy profile that is associated with an SSL proxy policy intent is modified, you must redeploy the SSL proxy policy to ensure that the changes take effect on the site.

Cloning SSL Forward Proxy Profiles

Cloning enables you to easily create a new SSL forward proxy profile based on an existing one.

To clone an SSL forward proxy profile:

1. Select **Configuration > SSL Proxy > Profiles**.

   The SSL Proxy Profiles page appears displaying the existing SSL forward proxy profiles.

2. Select the SSL forward proxy profile that you want to clone and select **More > Clone**. Alternatively, right-click a profile and select **Clone**.

   The Clone SSL Proxy Profile page appears, showing the same fields that are presented when you create an SSL forward proxy profile.

3. Modify the SSL forward proxy profile fields as needed.

4. Click **OK** to save your changes.

   You are taken to the SSL Proxy Profiles page. A confirmation message appears, indicating the status of the clone operation.

Deleting SSL Forward Proxy Profiles

To delete one or more SSL forward proxy profiles:

> **NOTE:** If you try to delete an SSL forward proxy profile that is associated with an SSL proxy policy intent, a message is displayed indicating that the profile cannot be deleted.

1. Select **Configuration > SSL Proxy > Profiles**.

   The SSL Proxy Profiles page appears, displaying the existing SSL forward proxy profiles.

2. Select one or more SSL forward proxy profiles that you want to delete and click the delete icon (**X**). Alternatively, right-click a profile and select **Delete SSL Proxy Profile**.

   An alert message appears asking you to confirm the delete operation.

3. Click **Yes** to delete the selected SSL forward proxy profiles.

   A confirmation message appears indicating the status of the delete operation.

> **NOTE:** If the deleted SSL forward proxy profile is associated with an SSL proxy policy intent, you must redeploy the SSL proxy policy to ensure that the changes take effect on the site.

RELATED DOCUMENTATION

Creating SSL Forward Proxy Profiles | 655

About the SSL Proxy Profiles Page | 653

Configuring and Deploying an SSL Forward Proxy Policy

The following is the workflow for configuring and deploying an intent-based SSL forward proxy policy in CSO:

1. Obtain the root certificate and private key from your trusted certificate authority (CA).

2. Combine the root certificate and private key into a single file.

3. Import the certificate and private key file (on the Import Certificate page); see "Importing a Certificate" on page 352.

4. (Optional) Install the imported certificate on one or more sites (on the Install Certificate page); see "Installing and Uninstalling Certificates" on page 354.

5. By default, Juniper Networks ships trusted certificates for sites that use HTTPS. These certificates are installed automatically by CSO when the site is successfully provisioned.

   If you want to use additional trusted certificates, import and install the certificates as explained in Step 3 and 4.

6. Create an SSL proxy profile (on the Create SSL Proxy Profiles) page; see "Creating SSL Forward Proxy Profiles" on page 655.

   NOTE:
   - Use the imported root certificate when you create the SSL proxy profile.
   - For trusted certificates, specify that all trusted certificates on the device are used (select **All** in the **Trusted Certificate Authorities** field).

7. Create an SSL proxy policy intent that uses the SSL proxy profile that you created (on the SSL Proxy Policy page); see "Creating SSL Proxy Policy Intents" on page 644.

8. Deploy the SSL proxy policy; see "Deploying Policies" on page 668.

NOTE:

- Ensure that the root and trusted certificates are imported into CSO before the policy is deployed.
- If you have not installed the certificates referenced in the SSL proxy profile, then they are automatically installed when the SSL proxy policy is deployed.

9. For Internet access from an SRX Series device by using the SSL proxy, ensure that you import the root certificate (obtained in Step 1) into the browsers of the clients accessing the Internet.

NOTE: If you do not import the certificate, the traffic does not go through for clients in the LAN segments.

RELATED DOCUMENTATION

SSL Forward Proxy Overview | **636**

Understanding How SSL Proxy Policy Intents Are Applied | **651**

CHAPTER 21

Deploying Policies

IN THIS CHAPTER

- Deploying Policies Overview | 664
- About the Deployments Page | 665
- Using the Deployment Icon to Deploy Policies | 667
- Deploying Policies | 668

Deploying Policies Overview

When you finish creating and verifying your security configurations, you can deploy these configurations and keep them ready to be pushed to the security devices. CSO enables you to push security configurations to the devices all at once by providing a single interface that is intuitive.

The deployment workflow provides the ability to save and publish different services to be updated at a later time to the appropriate firewalls (during downtime). This enables administrators to review their firewall and NAT policies before updating the device. Administrators also save troubleshooting time, avoid errors, and save costs associated with errors. Verify and tweak your security configurations before updating them to the device. This approach helps you keep the configurations ready and update these configurations to the devices during the maintenance window.

When you deploy policies, the process takes into account the priority and precedence values set on the policy and the order of rules on the device. Rules are published in the order of their priority groups.

If you change the priority or precedence of a published policy, the policy must be republished for the changes to take effect. Sometimes, changing priority or precedence in one policy can affect other policies in the same priority group. However, such dependent policies do not need to be republished in order for their changes in priority or precedence to take effect. It will be enough if the policy which is updated is republished.

There are three ways in which you can view and deploy your security configurations:

- Click on the deployment icon present in the CSO Customer Portal banner and use the deployment panel that appears, to deploy policies. See “Using the Deployment Icon to Deploy Policies” on page 667.

> **NOTE:** The deployment icon is highlighted in orange if there are undeployed configurations.

- Use the **Deployments** page. See "About the Deployments Page" on page 665.
- Select a firewall, NAT or SD-WAN policy from its respective landing pages and click **Deploy**. For more information, see "Deploying Policies" on page 668.

RELATED DOCUMENTATION

Using the Deployment Icon to Deploy Policies | 667

About the Deployments Page | 665

Deploying Policies | 668

About the Deployments Page

To access this page, click **Configuration** > **Deployments**.

Use this page to deploy or schedule the deployment of undeployed SD-WAN, NAT, and firewall policies. Undeployed policies refer to newly created firewall policy rules or NAT policies. These changes do not come into effect until the policies are deployed. The **Deploy** page provides scheduling options for you to deploy these policies.

Tasks You Can Perform

You can perform the following task from this page:

- Deploy a policy. See "Deploying Policies" on page 668.

Field Descriptions

Table 223 on page 666 provides guidelines on using the fields on the **Deployments** page.

Table 223: Fields on the Deployments Page

Field	Description
Awaiting Deployment	The **Awaiting Deployment** tab displays all the policies that are awaiting deployment. The following fields provide more information about the undeployed policies: • **Name**—Name of the policy that needs to be deployed. • **Deployment Type**—Type of the policy that needs to be deployed. • **Summary**—Description of the policy. • **Owner**—The tenant who has created the policy. • **Last updated**—The last time the policy was updated. If you want to deploy a policy, select the policy and click **Deploy**. The policy is deployed and will no longer appear in the **Awaiting Deployment** tab. If you want to refresh the **Awaiting Deployment** tab, click the refresh icon provided below the details table.
Scheduled	The **Scheduled** tab displays all the policies that have been scheduled for deployment on a certain date and time. The following fields provide more information about scheduled policies: • **Name**—Name of the policy. • **Deployment Type**—Type of the policy that needs to be deployed. • **Summary**—Description of the policy. • **Schedule**—The date and time at which the policy is scheduled to be deployed. • **Status**—Displays whether the scheduled policy has been deployed or not. • **Next Run**—Date and time when the scheduled deployments will be run. If you want to deploy a scheduled policy immediately, select the policy and click **Deploy Now**. If you want to modify the deployment schedule of a policy, select the policy and click the edit icon (pencil icon). The **Deploy** page appears displaying the current scheduling information. See "Deploying Policies" on page 668, to update the schedule.
History	The **History** tab displays all the policies that have been deployed. The following fields provide more information about deployed policies: • **Name**—Name of the deployed policy. • **Deployment Type**—Type of the deployed policy. • **Summary**—Description of the policy. • **Status**–Displays the status of the deployed policy. • **Job Details**—Details of the job. • **Deployed On**—Date and time the policy was deployed. If you want to redeploy a policy, select the policy and click **Re-Deploy**. The policy is redeployed and the **History** tab details changes to reflect this information.

RELATED DOCUMENTATION

Deploying Policies Overview | 664
Using the Deployment Icon to Deploy Policies | 667
Deploying Policies | 668

Using the Deployment Icon to Deploy Policies

CSO provides an option of viewing and deploying policies through the deployment panel, that appears when you click on the deployment icon. The deployment icon is highlighted in orange if there are undeployed policies.

To deploy policies through the deployment panel:

1. Click the deployment icon on the Customer Portal banner.

   The deployment panel appears. For information about the panel, see Table 224 on page 667.

2. Hover over the policy you want to deploy. The **Deploy** option appears on the right side of the policy.

3. Click **Deploy** to deploy the policy. For more information, see “Deploying Policies” on page 668.

Table 224 on page 667 provides guidelines on using the fields on the deployment panel.

Table 224: Fields on the Deployment Panel

Field	Description
Awaiting Deployment	The **Awaiting Deployment** tab displays all the policies that are awaiting deployment.
In Progress	The **In Progress** tab displays all the policies that are currently being deployed.

RELATED DOCUMENTATION

Deploying Policies Overview | 664
About the Deployments Page | 665
Deploying Policies | 668

Deploying Policies

You can deploy firewall, NAT, SD-WAN, and SSL proxy policies added by various services immediately or schedule the deployment for a later date and time.

To configure a deployment:

1. You can initiate the deployment of a policy in the following ways:
   - Select a policy from the **Awaiting Deployment** tab on the **Deployments** page and click **Deploy**.
   - Select a policy from the **Scheduled** tab on the **Deployments** page and click **Deploy**.
   - Select a policy from the **Scheduled** tab on the **History** page and click **Re-Deploy**.
   - Use the deployment icon on the Customer Portal banner. For more information about deploying policies using the deployment icon, see "Using the Deployment Icon to Deploy Policies" on page 667.

     NOTE: The deployment icon is highlighted in orange if there are undeployed policies.

   - Select **Configuration > Firewall > Firewall Policy**. The **Firewall Policy** page appears, displaying the intents associated with the policy. Click **Deploy**.
   - Select **Configuration > NAT > NAT Policies** and select the NAT policy you want to deploy. Click **Deploy**.
   - Select **Configuration > SSL Proxy > Policy**. The SSL Proxy Policy page appears, displaying the intents associated with the policy. Click **Deploy**.
   - Select an SD-WAN policy intent on the **SD-WAN Policy** page and click **Deploy**.
2. The **Deploy** page appears. In **Choose Deployment Time** options, select **Run Now** to deploy the policy immediately.

   Select **Schedule at a later time** to deploy the policy at a later date and time. For scheduling options, see Table 225 on page 668.

3. Click **Deploy**.

Table 225 on page 668 provides guidelines on using the fields on the **Deploy** page.

Table 225: Fields on the Deploy Page

Field	Description
Summary	
Policies	The summary of the policy that is to be deployed.

Table 225: Fields on the Deploy Page *(continued)*

Field	Description
Choose Deployment Time	

Type	• Select **Run now** if you want to deploy the policy immediately.
	• Select **Schedule at a later time** if you want to schedule the deployment for a later date and time.
	  ◦ Click on the calendar icon to choose the date for the deployment in MM/DD/YYYY format.
	  ◦ Enter the time for the deployment in HH:MM:SS format. You can choose the 12 hour (AM or PM) or 24 hour format to specify the time by selecting the option from the drop-down list provided beside the time field.

RELATED DOCUMENTATION

Deploying Policies Overview | 664

Using the Deployment Icon to Deploy Policies | 667

About the Deployments Page | 665

CHAPTER 22

Configuring Policies for SD-LAN

IN THIS CHAPTER

- NAC Overview | **671**
- About the Port Profiles Page | **673**
- Add Port Profiles | **674**
- Edit, Clone, and Delete Port Profiles | **679**
- About the Authentication Profiles Page | **682**
- Add Authentication Profiles | **684**
- Edit, Clone, and Delete an Authentication Profile | **690**
- About the Access Profiles Page | **692**
- Add Access Profiles | **693**
- Edit, Clone, and Delete Access Profiles | **696**
- About the RADIUS Server Profiles Page | **699**
- Add RADIUS Server Profiles | **700**
- Edit, Clone, and Delete RADIUS Server Profiles | **702**
- Firewall Filters Overview | **705**
- Configure a Firewall Filter for an EX Series Switch | **706**
- About the EX Firewall Filters Page | **706**
- Add Firewall Filters | **707**
- Delete Firewall Filters | **708**
- About the < Firewall-Filters-Name> / Terms Page | **709**
- Add Terms to Firewall Filters | **710**
- Edit, Clone, and Delete Terms | **713**
- Deploy or Redeploy a Port Profile | **715**
- Enable Ports | **716**
- Disable Ports | **717**
- Edit Configuration of Ports | **717**

NAC Overview

Network access control (NAC) is a set of policies, applied on a switch, to enforce security so that only trusted users and devices connected to the LAN of a campus or branch network is granted access to network resources. NAC also monitors and controls the activities of the users and devices after network access is granted.

You can implement NAC on a switch by configuring and deploying the port authentication profile on the switch ports.

An authentication profile defines the authentication method, fallback options, and other settings such as number of retries, maximum number of authentication requests that can be allowed for a supplicant, authentication server timeout, and so on, related to the communication between the switch and the supplicant (a user or device such as printer).

When you implement NAC in your network:

1. A supplicant sends a request for network access to the switch, which acts as the internet gateway for a campus or branch network. A supplicant can be a computer (desktop or laptop), a tablet, a phone, a headless device such as a printer, camera, or industrial controls, or a wireless access point.
2. The switch requests for credentials (username, password) of the supplicant, which the supplicant provides.
3. The switch validates the supplicant credentials by using the RADIUS server (authentication server).

   If the supplicant credentials are valid, the switch grants access to the campus or branch network through one of its ports. Based on the firewall filters configured, the switch enforces polices on the supplicant to restrict access to the network resources.

   If the switch is unable to validate the supplicant credentials, the supplicant is denied network access completely or is restricted to access only the internet.
4. Detailed session records including user and device details, session types, and service details are maintained in the RADIUS accounting server for troubleshooting, class-of-service (CoS) control, and billing purposes.

Types of Authentication

CSO allows you to configure the following types of authentication on the switch:

- 802.1x—In this type of authentication, the supplicant is authenticated based on the credentials provided.

  You can configure the following three modes for authenticating the supplicant in a LAN or WLAN when using the 802.1x authentication:

  - Single: The switch authenticates only the first supplicant requesting network access, in a LAN. All other supplicants in the LAN that connect later to the port are allowed full access without any further authentication, based on the first supplicant's authentication.
  - Single Secure: Allows only one supplicant in a LAN to connect to a port. No other supplicant in the LAN is allowed to connect until the first supplicant logs out.
  - Multiple: Allows multiple supplicants in a LAN to connect to a port at the same time. Each supplicant is authenticated individually.
- MAC address authentication—In this authentication, the switch sends the MAC address of the supplicant, instead of the credentials to the RADIUS server for authentication. This type of authentication is used when the switch does not receive the authentication request packet from the supplicant. This method is also used by users (contractors or guests) and devices such as printers, cameras, and industrial controls that do not have a user interface to provide credentials.

When a supplicant is not authenticated because the RADIUS server is inaccessible or the supplicant has provided incorrect credentials, you can configure the switch to:

- Allow network access.
- Deny network access
- Move the supplicant to a specific VLAN (server-fail VLAN when the RADIUS server is inaccessible and server-reject VLAN when the supplicant's credentials are rejected by the RADIUS server)

If the RADIUS server is unreachable when reauthenticating a supplicant during a session, the supplicant is allowed access based on prior authentication. However, a new supplicant requesting network access are denied network access.

You can configure a guest VLAN on the switch to provide limited network access (only to the Internet) for:

- Guests and contractors
- Devices that are not 802.1X enabled and on which the MAC RADIUS authentication is not configured or supported.

For detailed information about network access control on an EX Series switch, see User Access and Authentication Feature Guide.

RELATED DOCUMENTATION

Add Authentication Profiles | 684

Add Access Profiles | 693

Deploy or Redeploy a Port Profile | 247

About the Port Profiles Page

To access this page, select **Configuration > SD-LAN > Port Profiles** in Customer Portal.

Use this page to view, edit, clone, and delete port profiles. A port profile authenticates supplicants (users and devices such as a printer) trying to access network resources and configure parameters such as flow control, link mode, storm control, MAC limit, and so on, on the switch ports. If you include firewall filters for egress and ingress traffic in the port profile, you can permit or deny access to network resources granularly based on source and destination IP addresses, MAC addresses, ports, and protocols in the packets.

Tasks You Can Perform

You can perform the following tasks from this page:

- Add a port profile—See "Add Port Profiles" on page 674.
- Edit, clone, or delete a port profile—See "Edit, Clone, and Delete Port Profiles" on page 679.
- Clear the selected port profiles—Click **Clear All Selections** to clear any port profiles that you have selected.
- Search for port profiles using keywords—Click the search icon and enter the search term in the text box and press Enter. The search results are displayed on the same page.

Field Descriptions

Table 226 on page 673 describes the fields on the Port Profiles page.

Table 226: Fields on the Port Profiles Page

Field	Description
Profile Name	Name of the port profile.
Description	Description of the port profile.
Port Mode	Operating mode defined in the profile for a port—Trunk, Access.
Port Authentication Profile	Authentication profile assigned on the port.
Firewall Filter (Ingress)	Firewall filter assigned for the ingress traffic.
Firewall Filter (Egress)	Firewall filter assigned for the egress traffic.

RELATED DOCUMENTATION

Add Authentication Profiles | 684

Add Firewall Filters | 707

Deploy or Redeploy a Port Profile | 247

Add Port Profiles

Use the Add Port Profiles page in Customer Portal to add port profiles. You add a port profile by assigning an authentication profile, a firewall filter for the ingress traffic, a firewall filter for the egress traffic, and configuring port parameters.

To add a port profile:

1. Select **Configuration > SD-LAN > Port Profiles** in Customer Portal.

   The Port Profiles page appears.

2. Click the add icon (**+**) to add a new port profile.

   The Add Port Profile wizard appears. The wizard provides step-by-step procedures to add the port profile.

3. Complete the configuration according to the guidelines provided in Table 227 on page 674.

   NOTE: Fields marked with * are mandatory.

4. Click **OK**.

   The port profile is added to CSO. You are returned to the Port Profiles page where a confirmation message is displayed.

After you add a port profile, you can deploy the port profile on one or more ports of the switch. See "Deploy or Redeploy a Port Profile" on page 247.

Table 227: Port Profile Settings

Setting	Guideline
General	

Table 227: Port Profile Settings *(continued)*

Setting	Guideline
Profile Name	Enter a unique name for the port profile which can contain only alphanumeric characters and hyphen (-); 32-characters maximum.
Profile Description	Enter a description for the port profile.
Basic Settings	
Port Mode	Select whether the port should be configured as a trunk port or an access port: • Trunk—The port can be used to connect with other switches or routers. • Access—The port can be used to connect to access points and end-user devices such as laptops or printer.
Port Authentication Settings	
Authentication Profile	Select an authentication profile to be used in the port profile. **NOTE:** An authentication profile defines the authentication method and other parameters related to communication between the switch and the supplicant. You must configure the authentication profile before referencing it in the port profile. To add an authentication profile, see "Add Authentication Profiles" on page 684. If you select None (default), no authentication profile is associated with the port profile. You can use this option when you do not want 802.1x authentication to be configured on a port.
Firewall Filter	
Firewall Filter Profile (Ingress)	Select a firewall filter profile to be used for the ingress traffic. You must configure the firewall filter before referencing it in the port profile. .
Firewall Filter Profile (Egress)	Select a firewall filter profile to be used for the egress traffic. You must configure the firewall filter before referencing it in the port profile. .
Advanced Settings	

Table 227: Port Profile Settings *(continued)*

Setting	Guideline
Link Settings	Click the toggle button to enable or disable (default) link settings on a port. If you disable this setting, the port uses the default configurations for auto negotiation, flow control, MTU, speed, and link mode. Enable this option to modify the default configuration for auto negotiation, flow control, MTU, and so on.
Auto Negotiation	Click the toggle button to enable (default) or disable autonegotiation on the port. Auto negotiation enables a port to determine the data transmission speed and the duplex mode based on the speed and duplex mode of the peer port. If you have enabled autonegotiation and also configured link mode and speed, the ports use the configured values for link mode and speed. If you disable autonegotiation, you must configure values for link mode and speed.
Flow Control	Click the toggle button to enable (default) or disable flow control on a port. Flow control enables a port to regulate network traffic so that there is no data loss during congestion. If you disable flow control, you lose data during congestion.
MTU	Enter the size (in bytes) of the maximum transmission unit (MTU) that can be transmitted through a port. Range: 256 to 9,216 bytes Default: 1,514 bytes
Speed	Select the maximum transmission speed of a port (in GB or MB). If you enable auto-negotiation and select a transmission speed, the port uses the value configured here for transmission speed. Default: 1G

Table 227: Port Profile Settings *(continued)*

Setting	Guideline
Link Mode	Select the mode of the links configured on a port: • Automatic (default) —The port automatically selects the duplex mode based on the duplex mode of the peer port. • Full Duplex—The port allows data to be sent and received at the same time over a link. • Half Duplex—The port allows data to be either only received or sent at a given time over a link. If you enable autonegotiation and also select a value for link mode, the port considers the value configured here for the operating mode of the links established on the port.
Storm Control Settings	Click the toggle button to enable or disable (default) storm control settings on a port. If you disable this setting, the port uses the defalt value for storm control. Enable this option to modify the default storm control value.
Storm Control	Enter the bandwidth (in kbps) or the percentage of the bandwidth, beyond which a port can drop packets. Also, select whether the value you enter indicates the percentage or the bandwidth, from the drop-down list. The default unit is percentage. Range: For bandwidth, 100 through 100,000,000 kbps. For percentage, 1 through 100. Default: 80 percent.
Power over Ethernet (PoE) Settings	Click the toggle button to enable or disable (default) PoE on a port. If you disable this setting, the default PoE setting is configured on a port when you deploy the profile on the port. Enable this option to modify the default PoE settings.
Maximum Power	Enter the maximum power (in Watts) that a port can provide. Range: 1 through 90 Watts. Default: 30 Watts.

Table 227: Port Profile Settings *(continued)*

Setting	Guideline
Priority	Select a priority (Low or High) for a port to be used as a source for powering a device connected to the port. If power is insufficient for all PoE ports, the PoE power to low-priority ports is shut down before power to high-priority ports is shut down. Among ports that have the same assigned priority, the power priority is determined by port number, with lower-numbered ports having higher priority. Default: Low
Port Security Settings	Click the toggle button to enable or disable (default) security on a port. If you disable this setting, the default port security is configured on a port when you deploy the profile on the port. Enable this option to modify the default settings.
Trust DHCP	Click the toggle button to enable (default) or disable trusting traffic from a DHCP server. If you disable this option, the port drops packets sent to and received from a DHCP server.
MAC Limit	Click the toggle button to enable or disable (default) setting the maximum number of MAC addresses that can be stored in the MAC table for a port. If you enable MAC Limit, you must configure a value for MAC limit and MAC limit action. If you disable this option, you cannot limit the MAC addresses that are learnt within a VLAN and, therefore, enforce security against the flooding of the Ethernet switching table.
MAC Limit	Enter the maximum number of MAC addresses that a switch can store in the MAC table for a port. Range: 1 through 10,000. Default: 1.

Table 227: Port Profile Settings *(continued)*

Setting	Guideline
MAC Limit Action	Select the action that a port must take when the number of entries in the port MAC table exceeds the MAC limit value: • Drop (Default)—Drop the packet, but do not generate an alarm • Shutdown—Disable the port and generate an alarm, an SNMP trap, or a system log entry. • Drop and Log—Drop the packet and generate an alarm, an SNMP trap, or system log entry.

RELATED DOCUMENTATION

Add Access Profiles | 693

Edit, Clone, and Delete Port Profiles

IN THIS SECTION

- Edit a Port Profile | 680
- Clone a Port Profile | 680
- Delete a Port Profile | 681

You can edit, clone, and delete port profiles from the Port Profiles page.

Edit a Port Profile

NOTE:

- When you edit a port profile that is already deployed on a port, you must redeploy the profile on the port for the changes to take effect.
- When you edit an authentication profile or firewall filter that is used in a port profile and if the port profile is deployed on one or more ports, you must redeploy the port profile on the ports for the changes in the authentication profile and firewall filter to take effect.

To edit a port profile:

1. Select **Configuration > SD-LAN > Port Profiles** in Customer Portal.

   The Port Profiles page appears, displaying the configured port profiles.

2. Select the port profile that you want to edit and click the **edit** icon (pencil).

   The Edit Port Profiles page appears, displaying the same fields that were presented when you added a port profile.

3. Modify the port profile fields as needed. Refer to Table 227 on page 674 topic to modify the fields.

   NOTE: You cannot edit the port profile name.

4. Click **OK** to save your changes.

   You are taken to the Port Profiles page. A confirmation message appears indicating the status of the edit operation.

Clone a Port Profile

To clone a port profile:

1. Select **Configuration > SD-LAN > Port Profiles** in Customer Portal.

   The Port Profiles page appears, displaying the configured port profiles.

2. Select the port profile that you want to clone and click **Clone**.

   The Clone Port Profile page appears. A default name appears for the cloned port profile as CLONE-*Port Profile Name*.

3. (Optional) Edit the name for the cloned profile.

   The name for the port profile should be unique and contain only alphanumeric characters and hyphen (-); 32-characters maximum.

4. Click **OK** to create a new profile.

   The Port Profiles page appears. A confirmation message appears indicating the status of the clone operation.

5. After you clone the profile, select the profile and click **Edit** to modify the parameters as needed. See Table 227 on page 674 to modify the parameters.

After you clone and modify the port profile, you can deploy the profile on the switch ports. See "Deploy or Redeploy a Port Profile" on page 247 for details.

Delete a Port Profile

> **NOTE:** You cannot delete a port profile when the port profile is deployed on one or more ports. To delete the port profile, first disassociate the port profile from all the ports on which it is deployed and then attempt to delete the profile; see "Dissociate a Profile from a Port" on page 253 for details.

To delete a port profile:

1. Select **Configuration > SD-LAN > Port Profiles** in Customer Portal.

   The Port Profiles page appears, displaying the configured port profiles.

2. Select the port profile that you want to delete and click the **Delete** icon (dustbin).

   An alert message appears, asking you to confirm the delete operation.

3. Click **Yes** to delete the selected port profile.

   After the profile is deleted successfully, a message indicating that the port profile is deleted appears on top of the page and the port profile is removed from Port Profiles page.

RELATED DOCUMENTATION

About the Port Profiles Page | 673

Deploy or Redeploy a Port Profile | 247
Dissociate a Profile from a Port | 253

About the Authentication Profiles Page

To access this page, select **Configuration > SD-LAN > Authentication Profiles** in Customer Portal.

Use this page to view, clone, edit, and delete authentication profiles. An authentication profile enables you to define parameters to authenticate a user. You can define the following parameters in an authentication profile—the authentication method, fallback options, and other settings (for example, number of retries, maximum number of requests that can be allowed, and authentication server timeout) related to the communication between the switch and a supplicant.

Tasks You Can Perform

You can perform the following tasks from this page:

- Add an authentication profile—See "Add Authentication Profiles" on page 684.
- Edit, clone, or delete an authentication profile—See "Edit, Clone, and Delete an Authentication Profile" on page 690.
- Clear the selected authentication profiles—Click **Clear All Selections** to clear any authentication profiles that you might have selected.
- Search for authentication profiles using keywords—Click the search icon and enter the search term in the text box and press Enter. The search results are displayed on the same page.

Field Descriptions

Table 228 on page 682 describes the fields on the Authentication Profiles page.

Table 228: Authentication Profiles Page Fields

Field	Description
Profile Name	Name of the authentication profile.
Description	A description about the authentication profile.

Table 228: Authentication Profiles Page Fields *(continued)*

Field	Description
Supplicant Mode	The mode of authenticating supplicants: • Single—Authenticates only the first supplicant in a LAN. All other supplicants in the LAN that connect later to the port are allowed access without any further authentication, based on the first supplicant's authentication. • Single Secure—Allows only one supplicant in a LAN to connect to the port. No other supplicant in the LAN is allowed to connect until the first supplicant logs out. • Multiple—Allows multiple supplicants in a LAN to connect to the port. Each supplicant is authenticated individually.
Primary Authentication Method	The primary method for authenticating a supplicant: • dot1x—IEEE 802.1X standard for port-based network access control (PBNAC); protects Ethernet LANs from unauthorized user access. The dot1x method blocks all traffic to and from a supplicant at the port until the supplicant's credentials are presented and matched on the authentication server (a RADIUS server). When the supplicant is authenticated, the switch allows traffic from and to the supplicant to transmit through it. • MAC RADIUS—Used for network devices (such as a printer or a camera) connected in a LAN that needs to access network resources, but do not support the 802.1X standard. When a switch detects a supplicant that is not 802.1X-enabled on its port, the switch transmits the MAC address of the supplicant to the authentication server. The server then tries to match the MAC address with a list of MAC addresses in its database. If the MAC address matches an address in the list, the supplicant is authenticated.
Secondary Authentication Method	The secondary method for authenticating a supplicant when the switch is unable to validate a supplicant by using the primary method : • None • dot1x, when MAC RADIUS is set as the primary authentication method. • MAC RADIUS, when the dot1x method is set as the primary authentication method.

Table 228: Authentication Profiles Page Fields *(continued)*

Field	Description
Server Fail	The action that the switch takes when the RADIUS servers are unavailable for authenticating a supplicant: • None—No action is taken. If network access is already granted to a supplicant, the access is maintained. • Deny—Network access is denied to the supplicant. • Permit—Network access is permitted to the supplicant. If a RADIUS server timeout occurs during reauthentication, traffic is allowed from and to the supplicant as the supplicant is already authenticated. • Use Cache—Recognizes already connected supplicants and reauthenticates the supplicant when there is a RADIUS server timeout (new supplicants are denied access): • VLAN ID—Moves a supplicant to a specified VLAN (server fail VLAN) if a RADIUS server timeout occurs:
Server Reject	The action the switch takes when the switch is unable to validate a supplicant because of incorrect credentials provided by the supplicant: • None—No action is taken and the supplicant is denied network access. • VLAN ID—Moves the supplicant to a specified VLAN (server reject VLAN) with limited network access (Internet only)
Guest	The action the switch takes for temporary users such as guests or contractors: • None—No action is taken and the supplicant is denied network access. • VLAN ID—Moves the supplicants to a specified VLAN (guest VLAN) with limited network access (Internet only)

RELATED DOCUMENTATION

NAC Overview | **671**

Add Port Profiles | **674**

Add Access Profiles | **693**

Add Authentication Profiles

Use the Add Authentication Profiles page in Customer Portal to add authentication profiles.

Adding an authentication profile involves the following steps:

1. Define the primary and secondary methods for authenticating a supplicant—802.1x (dot1x), MAC RADIUS.
2. Define the action that the port must take when the RADIUS server is not reachable or a user is not authenticated (fallback options).
3. Define the authentication process parameters, such as the number of times that the switch can request for user authentication, whether a user must be reauthenticated at regular intervals, the number of times a switch can attempt to contact the RADIUS server for authenticating a user, and so on.

To add an authentication profile:

1. Select **Configuration > SD-LAN > Authentication Profiles** in Customer Portal.

   The Authentication Profiles page appears, displaying the existing authentication profiles.

2. Click the **Add** icon (**+**).

   The Add Authentication Profiles wizard appears.

3. Complete the configuration according to the guidelines provided in Table 229 on page 685.

   NOTE: Fields marked with * are mandatory.

4. Click **OK**.

   An authentication profile is added. You are returned to the Authentication Profiles page where a confirmation message is displayed.

   After you add an authentication profile, you can assign it to a port profile. See “Add Port Profiles” on page 674.

Table 229: Fields on the Add Authentication Profile Page

Setting	Guideline
General	
Profile Name	Enter a unique name for the authentication profile, which can only contain alphanumeric characters and hyphen (-); 15-character maximum.
Profile Description	Enter a description for the authentication profile.

Table 229: Fields on the Add Authentication Profile Page *(continued)*

Setting	Guideline
Supplicant Mode	Select a mode for authenticating the supplicant: • Single—Authenticates only the first supplicant in a LAN. All other supplicants in the LAN that connect to the port later are allowed or denied access without any authentication, based on the first supplicant's authentication. • Single Secure—Allows only one supplicant in a LAN to connect to the port. No other supplicant in the LAN is allowed to connect until the first supplicant logs out. • Multiple—Allows multiple supplicants in a LAN to connect to the port. Each supplicant is authenticated individually.
Authentication Method	
Primary Method	Select the primary method of authenticating a supplicant: • dot1x—IEEE 802.1X standard for port-based network access control (PBNAC); protects Ethernet LANs from unauthorized user access. The 802.1x method blocks all traffic to and from a supplicant at the port until the supplicant's credentials are presented and matched on the authentication server (a RADIUS server). When the supplicant is authenticated, the switch allows traffic from and to the supplicant to transmit through it. • MAC RADIUS—Used for supplicants, connected in a LAN that need to access network resources, such as printer or camera, but do not support the 802.1X standard. When a switch detects a supplicant that is not 802.1X-enabled, the switch transmits the MAC address of the supplicant to the authentication server. The server then tries to match the MAC address with a list of MAC addresses in its database. If the MAC address matches an address in the list, the supplicant is authenticated.
Secondary Method	The secondary method for authenticating a supplicant when the switch is unable to validate a supplicant by using the primary method: • None • dot1x, when MAC RADIUS is selected as the primary authentication method. • MAC RADIUS, when dot1x is selected as the primary authentication method.
Fallback Options	
You can configure authentication fallback options to specify how supplicants connected to a switch are supported if the RADIUS authentication server becomes unavailable or sends a RADIUS access-reject message.	

Table 229: Fields on the Add Authentication Profile Page *(continued)*

Setting	Guideline
Server Fail	Select an action that the switch applies to supplicants when the authentication servers are not reachable. The switch can accept or deny access to supplicants or maintain the access already granted to supplicants before the RADIUS timeout occurred. You can also configure the switch to move the supplicants to a specific VLAN. • None—No action is taken. If network access is already granted to a supplicant, the access is maintained. • Deny—Network access is denied to the supplicant. • Permit—Network access is permitted to the supplicant. If a RADIUS server timeout occurs during reauthentication, traffic is allowed from and to the supplicant because the supplicant is already authenticated. • Use Cache—Recognizes already connected supplicants and reauthenticates the supplicants when there is a RADIUS timeout; new supplicants are denied access. • VLAN ID—Moves a supplicant to a specified VLAN (server-fail VLAN) if a RADIUS server timeout occurs: If you select this option, enter the VLAN ID in the text box that appears below the Server Fail field. **NOTE:** The server-fail VLAN should be already configured on the site containing the switch.
VLAN ID	If you select VLAN ID for the Server Fail option, enter the VLAN ID of the VLAN to which the supplicant must be assigned.
Server Reject	The action the switch takes when the switch is unable to validate a supplicant because of incorrect credentials provided by the supplicant: • None—No action is taken and the supplicant is denied network access. • VLAN ID—Moves the supplicant to a specified VLAN (server-reject VLAN) with limited network access (Internet only). The server-reject VLAN is already configured on the switch. If you select this option, enter the VLAN ID in the text box that appears below the Server Reject field. **NOTE:** The server-reject VLAN should be already configured on the site containing the switch.
VLAN ID	If you select VLAN ID for the Server Reject option, enter the VLAN ID to which the supplicant must be assigned.

Table 229: Fields on the Add Authentication Profile Page *(continued)*

Setting	Guideline
Guest	Select an action to be taken for a guest (corporate guest or supplicants that are not 802.1x enabled). • None—No action is taken and the supplicant is denied network access. • VLAN ID—Move the supplicants to a specified VLAN (guest VLAN) with limited network access (Internet only). If you select this option, enter the VLAN ID of the guest VLAN in the text box that appears below this field. **NOTE:** The guest VLAN should be already configured on the site containing the switch.
VLAN ID	Enter the VLAN ID of the guest VLAN.
Advanced Settings	
Transmit Period	Enter the number of seconds that the switch waits before retransmitting the initial authentication request to the supplicant. Range: 1 through 65,535 seconds. Default: 30 seconds.
Maximum Requests	Enter the maximum number of times that authentication request packets are retransmitted to a supplicant before the authentication session times out. Range: 1 through 10. Default: 2.
Retries	Enter the number of times that the switch attempts to contact an authentication server for authenticating a supplicant after an initial failure. Range: 1 through 10. Default: 3.
Quiet Period	Enter the number of seconds that the port remains in the wait state following a failed authentication exchange with the supplicant, before reattempting authentication. Range: 0 through 65,535 seconds. Default: 3 seconds.

Table 229: Fields on the Add Authentication Profile Page *(continued)*

Setting	Guideline
Reauthentication	Click to enable or disable (default) reauthentication of the supplicant after a specified interval. If you enable this option, you must provide the reauthentication interval.
Reauthentication Interval	If you enable reauthentication, enter the number of seconds after which a supplicant must be reauthenticated. Range: 1 through 65,535 seconds. Default: 3600 seconds.
Supplicant Timeout	Enter the number of seconds that the port must wait for a response from the supplicant, before considering a timing out and resending the request. Range: 1 through 60 seconds. Default: 30 seconds.
RADIUS Server Timeout	Enter the number of seconds that the port waits for a reply from the RADIUS server when authenticating a supplicant before timing out and invoking the server-fail action (action that the switch applies to supplicants when the authentication servers are not reachable). Range: 1 through 60 seconds. Default: 30 seconds.

RELATED DOCUMENTATION

Add Port Profiles | **674**

Add Access Profiles | **693**

Edit, Clone, and Delete an Authentication Profile

IN THIS SECTION

- Edit Authentication Profiles | 690
- Clone an Authentication Profile | 691
- Delete Authentication Profiles | 691

You can edit, clone, and delete authentication profiles from the Authentication Profiles page.

Edit Authentication Profiles

NOTE: If you edit an authentication profile that is used in a port profile and if the port profile is deployed on a port, you must redeploy the port profile on the port for the changes made in the authentication profile to take effect.

To edit the parameters of an authentication profile:

NOTE: You cannot edit the name of an authentication profile.

1. Select **Configuration > SD-LAN > Authentication Profiles** in Customer Portal.

   The Authentication Profiles page appears, displaying the existing authentication profiles.

2. Select the authentication profile that you want to edit and click the **Edit** icon (pencil).

   The Edit Authentication Profile page appears, displaying the same fields that were presented when you added the authentication profile.

3. Modify the authentication profile fields as needed. Refer to the "Add Authentication Profiles" on page 684 page.

4. Click **OK** to save your changes.

   You are taken to Authentication Profiles page. A confirmation message appears indicating the status of the edit operation.

Clone an Authentication Profile

To clone an authentication profile:

1. Select **Configuration > SD-LAN > Authentication Profiles** in Customer Portal.

   The Authentication Profiles page appears, displaying the configured authentication profiles.

2. Select the authentication profile that you want to edit and click **Clone**.

   The Clone Port Authentication Profile page appears. A default name appears for the cloned profile as CLONE-*Authentication Profile Name*

3. (Optional) Edit the name for the cloned profile.

   The name for the authentication profile should be unique and contain only alphanumeric characters and hyphen (-); 15-characters maximum.

4. Click **OK** to create a new profile.

   You are taken to the Authentication Profiles page. A confirmation message appears indicating the status of the clone job.

5. (Optional) After you clone the profile, select the profile and click **Edit** to modify the parameters as needed. See "Add Authentication Profiles" on page 684 to modify the parameters.

After you clone and modify the authentication profile, you can create a port profile by using it or modify an existing port profile to include the cloned authentication profile.

Delete Authentication Profiles

> **NOTE:** You cannot delete an authentication profile if the profile is assigned to one or more port profiles. To delete such an authentication profile, assign a different authentication profile to the port profile and then attempt deleting the authentication profile again.

To delete an authentication profile:

1. Select **Configuration > SD-LAN > Authentication Profiles** in Customer Portal.

   The Authentication Profiles page appears, displaying the configured authentication profiles.

2. Select an authentication profiles that you want to delete and click the **Delete** icon (dustbin).

   An alert message appears, asking you to confirm the delete operation.

3. Click **Yes** to delete the selected authentication profiles.

   A confirmation message appears, indicating the status of the delete operation.

RELATED DOCUMENTATION

NAC Overview | 671

Dissociate a Profile from a Port | 253

Add Port Profiles | 674

About the Access Profiles Page

To access this page, select **Configuration > SD-LAN > Access Profiles** in Customer Portal.

Use this page to view, edit, clone, and delete access profiles. Access profiles enable you to define the authentication and accounting servers and their priorities.

Tasks You Can Perform

You can perform the following tasks from this page:

- Add an access profile—See "Add Access Profiles" on page 693.
- Edit, clone, or delete an access profile—See "Edit, Clone, and Delete Access Profiles" on page 696.
- Clear the selected access profiles—Click **Clear All Selections** to clear any access profiles that you might have selected.
- Search for access profiles using keywords—Click the **Search** icon and enter the search term in the text box and press Enter. The search results are displayed on the same page.

Field Descriptions

Table 230 on page 692 describes the fields on the Access Profiles page.

Table 230: Fields on the Access Profiles Page

Field	Description
Profile Name	Name of the access profile.
Description	A description of the access profile.

Table 230: Fields on the Access Profiles Page *(continued)*

Field	Description
Authentication Servers	RADIUS authentication servers configured for providing authentication service. To provide the authentication service, the switch accesses the servers in the listed order and in a round-robin fashion until a valid response is received from one of the servers or until all the configured retry limits are reached.
Accounting Servers	RADIUS accounting servers configured for providing accounting service. RADIUS accounting servers store statistical data about users logging in to or out from a LAN through a switch. The statistical data is used for general network monitoring, analyzing and tracking usage patterns, or billing a user based upon the amount of time or type of services accessed. The RADIUS accounting server you specify can be the same server used for RADIUS authentication, or it can be a separate RADIUS server. You can specify more than one RADIUS accounting server. If the primary server (the first one configured) is unavailable, the switch tries to access each RADIUS server in the list in the order in which they are configured.
Revert Interval	The number of seconds that the switch waits after a RADIUS server has become unreachable before rechecking the connection. If the server is still unreachable, the server next in the priority list is used.

RELATED DOCUMENTATION

- Add RADIUS Server Profiles | 700
- Deploying an Access Profile on a Switch | 243
- Dissociating an Access Profile | 244

Add Access Profiles

An access profile defines a list of RADIUS authentication servers and RADIUS accounting servers, and their priorities. Use the Add Access Profile page in Customer Portal to add access profiles to CSO.

To add an access profile:

1. Select **Configuration > SD-LAN > Access Profiles** in Customer Portal.

   The Access Profiles page appears.

2. Click the **Add** icon (+) to add an access profile.

   The Add Access Profiles page appears.

3. Complete the configuration according to the guidelines provided in Table 231 on page 694.

   NOTE: Fields marked with * are mandatory.

4. Click **OK**.

   An access profile is added. You are returned to the Access Profiles page where a confirmation message is displayed.

After you add an access profile, you can deploy it on a switch; see "Deploying an Access Profile on a Switch" on page 243.

Table 231: Access Profile Settings

Setting	Guideline
General	
Profile Name	Enter a unique name for the access profile, which can contain only alphanumeric characters and hyphen (-); 64-character maximum.
Profile Description	Enter a description for the access profile.

Table 231: Access Profile Settings *(continued)*

Setting	Guideline
Authentication Servers	Add authentication servers to the access profile. To add authentication servers: 1. Click the **Add** icon (+). The Add Authentication Server(s) page appears. 2. Select a server from the Available list and click the **Right Arrow** icon to move it to the Selected list. Alternatively, if you want to add a new authentication server, click the **Add New RADIUS Server Profile** button and move it to the list of selected servers; see "Add RADIUS Server Profiles" on page 700 for details. 3. Click **OK** to save the list of authentication servers. The authentication servers are used according to their priority. The priority is assigned in the order in which the servers are listed, with the highest priority being assigned to the server on the top of the list. To change the priority, reorder the servers in the list by dragging and dropping them.
Accounting Servers	Add accounting servers to the access profile. To add accounting servers: 1. Click the **Add** icon (+). The Add Accounting Server(s) page appears. 2. Select a server from the Available list and click the **Right Arrow** icon to move it to the Selected list. Alternatively, if you want to add a new accounting server, click the **Add New RADIUS Server Profile** button and move it to the list of selected servers; see "Add RADIUS Server Profiles" on page 700 for details. 3. Click **OK** to save the list of accounting servers. The accounting servers are used according to their priority. The priority is assigned in the order in which the servers are listed, with the highest priority being assigned to the server on the top of the list. To change the priority, reorder the servers in the list by dragging and dropping them.

Table 231: Access Profile Settings *(continued)*

Setting	Guideline
Revert Interval	Enter a revert interval, which is the amount of time the switch waits after a RADIUS server has become unreachable before rechecking the connection. If the server is still not reachable, the server next in the priority list is used for authentication and accounting. Range: 0 through 604,800 seconds Default: 60 seconds

RELATED DOCUMENTATION

Dissociating an Access Profile | **244**

Deploying an Access Profile on a Switch | **243**

Edit, Clone, and Delete Access Profiles

IN THIS SECTION

- Edit Access Profiles | **696**
- Clone an Access Profile | **697**
- Delete Access Profiles | **698**

You can edit, clone, and delete access profiles from the Access Profiles page.

Edit Access Profiles

NOTE: If you edit an access profile, you must redeploy the access profile on a switch for the changes to take effect on the switch. See "Deploying an Access Profile on a Switch" on page 243 for details.

To edit an access profile:

> **NOTE:** You cannot modify the name of an access profile.

1. Select **Configuration > SD-LAN > Access Profiles** in Customer Portal.

   The Access Profiles page appears, displaying the configured access profiles.

2. Select the access profile that you want to edit and click the **Edit** icon (pencil).

   The Edit Access Profile page appears, displaying the same fields that were presented when you added the access profile.

3. Edit the access profile fields as needed; see "Add Access Profiles" on page 693 for details.

4. Click **OK** to save your changes.

   You are taken to the Access Profiles page. A confirmation message appears indicating the status of the edit operation.

Clone an Access Profile

To clone an access profile:

1. Select **Configuration > SD-LAN > Access Profiles** in Customer Portal.

   The Access Profiles page appears, displaying the configured access profiles.

2. Select the access profile that you want to clone and click **Clone**.

   The Clone Access Profile page appears. A default name is assigned for the access profile as CLONE-*Access Profile Name*.

3. (Optional) Edit the name for the cloned profile.

   The name for the access profile should be unique and contain only alphanumeric characters and hyphen (-); 64-characters maximum.

4. Click **OK** to create a new profile.

The Access Profiles page appears. A confirmation message appears indicating the status of the clone operation.

5. (Optional) After you clone the profile, select the profile and click **Edit** to modify the parameters as needed. See "Add Access Profiles" on page 693 to modify the parameters.

After you clone the access profile, you can deploy it on a switch. See "Deploying an Access Profile on a Switch" on page 243 for details.

Delete Access Profiles

> **NOTE:** You cannot delete an access profile that is deployed on a switch. To delete such an access profile, first dissociate the access profile from the switch and then attempt deleting the access profile.

To delete an access profile:

1. Select **Configuration > SD-LAN > Access Profiles** in Customer Portal.

   The Access Profiles page appears, displaying the existing access profiles.

2. Select an access profile that you want to delete and click the **Delete** icon (dustbin).

   An alert message appears, asking you to confirm the delete operation.

3. Click **Yes** to delete the selected access profiles.

   A confirmation message appears, indicating the status of the delete operation.

RELATED DOCUMENTATION

- About the Access Profiles Page | **692**
- Dissociating an Access Profile | **244**
- Deploying an Access Profile on a Switch | **243**

About the RADIUS Server Profiles Page

To access this page, select **Configuration > SD-LAN > RADIUS Server Profiles** in Customer Portal.

Use this page to view, clone, edit, and delete RADIUS server profiles. A RADIUS server profile defines the RADIUS server IP address, password, authorization and accounting ports, retry counts, and server timeout.

Tasks You Can Perform

You can perform the following tasks from this page:

- Add a RADIUS server profile—See "Add RADIUS Server Profiles" on page 700.
- Edit, clone, or delete a RADIUS server profile—See "Edit, Clone, and Delete RADIUS Server Profiles" on page 702.
- Clear the selected RADIUS server profiles—Click **Clear All Selections** to clear any RADIUS server profiles that you might have selected.
- Search for RADIUS server profiles using keywords—Click the search icon and enter the search term in the text box and press Enter. The search results are displayed on the same page.

Field Descriptions

Table 232 on page 699 describes the fields on the RADIUS Server Profiles page.

Table 232: Fields on the RADIUS Server Profiles Page

Field	Description
Server Name	Name of the RADIUS server profile.
Server Address	IPv4 address of the RADIUS server.
Authentication Port	The software port number used for authentication.
Authentication Retry	The number of times that the switch can contact an authentication server for authenticating a supplicant when no response is received from the server.
Authentication Timeout	The number of seconds that the switch can wait for a response from the RADIUS server for an authentication request before retrying authentication of a supplicant.
Accounting Port	The software port number used for accounting.
Accounting Retry	The number of times the switch can contact an accounting server when no response is received from the accounting server.

Table 232: Fields on the RADIUS Server Profiles Page *(continued)*

Field	Description
Accounting Timeout	The number of seconds the switch can wait for a response from the accounting server.

RELATED DOCUMENTATION

- Add Access Profiles | 693
- Dissociating an Access Profile | 244
- Deploying an Access Profile on a Switch | 243

Add RADIUS Server Profiles

Use the Add RADIUS Server Profile page in Customer Portal to configure RADIUS server profiles. In a RADIUS server profile, you can define the RADIUS server IP address, password, authorization and accounting ports, retry counts, and server timeout.

To add a RADIUS server profile:

1. Select **Configuration > SD-LAN > RADIUS Server Profiles** in Customer Portal.

   The RADIUS Server Profiles page appears.

2. Click the **Add** icon (**+**) to add a RADIUS server profile.

   The Add RADIUS Server Profile page appears.

3. Complete the configuration according to the guidelines provided in Table 233 on page 701.

   NOTE: Fields marked with * are mandatory.

4. Click **Finish**.

   A RADIUS server profile is added to CSO. You are returned to the RADIUS Server Profiles page where a confirmation message is displayed.

After you add a RADIUS server profile, you can add an access profile by using the RADIUS server profile. See “Add Access Profiles” on page 693 for details.

Table 233: RADIUS Server Profile Settings

Setting	Guideline
General	
Server Name	Enter a unique name for the RADIUS server, which can contain only alphanumeric characters and hyphen (-). 64-character maximum.
Server Address	Enter the IPv4 address of the RADIUS server.
Secret	Configure the password to use with the RADIUS server. The secret password used by the local switch must match the one configured on the server. The length must be between 6 and 21 characters. You can include spaces if the character string is enclosed in quotation marks.
Advanced Settings	
Authentication Port	Enter the software port number that you want to use for authentication. Range: 1 through 65,535. Default: 1812.
Authentication Retry Count	Enter the number of times that the switch can contact a RADIUS server for authenticating a supplicant when no response is received from the authentication server. Range:1 through 100. Default: 3.
Authentication Timeout	Enter the number of seconds that the switch can wait for a response from the RADIUS server for an authentication request before retrying authentication of a supplicant. Range: 1 through 1000 seconds. Default: 3 seconds.
Accounting Port	Enter the software port number that you want to use for accounting. Range: 1 through 65,535. Default: 1813.

Table 233: RADIUS Server Profile Settings *(continued)*

Setting	Guideline
Accounting Retry Count	Enter the number of times that the switch can contact an accounting server when no response is received from the accounting server. Range: 1 through 100. Default: 3.
Accounting Timeout	Enter the number of seconds that the switch can wait for a response from the accounting server before timing out. Range: 1 through 1000 seconds. Default: 3 seconds.

RELATED DOCUMENTATION

Add Authentication Profiles | 684
Add Access Profiles | 693

Edit, Clone, and Delete RADIUS Server Profiles

IN THIS SECTION

- Edit RADIUS Server Profiles | 703
- Clone a RADIUS Server Profile | 703
- Delete RADIUS Server Profiles | 704

You can edit, clone, and delete RADIUS server profiles from the RADIUS Server Profiles page.

Edit RADIUS Server Profiles

> **NOTE:** If you edit a RADIUS server profile that is used in an access profile deployed on a switch, you must redeploy the access profile on the switch for the changes to take effect.

To modify the parameters configured for a RADIUS server profile:

> **NOTE:** You cannot edit the name of the RADIUS server profile.

1. Select **Configuration > SD-LAN > RADIUS Server Profiles** in Customer Portal.

   The RADIUS Server Profiles page appears, displaying the configured profiles.

2. Select the RADIUS server profile that you want to edit and click the **Edit** icon (pencil).

   The Edit RADIUS Server Profiles page appears, displaying the same fields that were presented when you added a RADIUS server profile.

3. Edit the RADIUS server profile fields as needed. See “Add RADIUS Server Profiles” on page 700 for more information.

4. Click **OK** to save your changes.

   You are taken to the RADIUS Server Profiles page. A confirmation message appears indicating the status of the edit operation.

Clone a RADIUS Server Profile

To clone a RADIUS server profile:

1. Select **Configuration > SD-LAN > RADIUS Server Profiles** in Customer Portal.

   The RADIUS Server Profiles page appears, displaying the configured access profiles.

2. Select the profile that you want to clone and click **Clone**.

   The Clone RADIUS Server Profile page appears. A default name appears for the RADIUS server profile as CLONE-*RADIUS Server Profile Name*

3. (Optional) Edit the name for the cloned profile.

The name for the RADIUS Server profile should be unique and contain only alphanumeric characters and hyphen (-); 64-characters maximum.

4. Click **OK** to create a new profile.

   The RADIUS Server Profiles page appears. A confirmation message appears indicating the status of the clone operation.

5. (Optional) After you clone the profile, select the profile and click **Edit** to modify the parameters as needed. See "Add RADIUS Server Profiles" on page 700 topic to modify the parameters.

After you clone and modify the RADIUS server profile, you can create an access profile by using it or add it to an existing access profile.

Delete RADIUS Server Profiles

> **NOTE:** You cannot delete a RADIUS server profile that is used by one or more access profiles. To delete such a profile, first disassociate the RADIUS server profile from all the access profiles and then attempt to delete the RADIUS server profile again.

To delete a RADIUS Server profile:

1. Select **Configuration > SD-LAN > RADIUS Server Profiles** in Customer Portal.

   The RADIUS Server Profiles page appears, displaying the configured RADIUS Server profiles.

2. Select the RADIUS server profile that you want to delete and click the **Delete** icon (dustbin).

   An alert message appears, asking you to confirm the delete operation.

3. Click **Yes** to delete the selected RADIUS server profile.

   A confirmation message appears, indicating the status of the delete operation.

RELATED DOCUMENTATION

Add Access Profiles | 693

Deploying an Access Profile on a Switch | 243

Dissociating an Access Profile | 244

Firewall Filters Overview

Firewall filters provide rules that define whether to permit or deny packets that are transiting a port on a Juniper Networks EX Series Ethernet Switch from a source endpoint to a destination endpoint. You configure firewall filters to determine whether to permit or deny traffic before it enters or exits a port to which the firewall filter is applied. To apply a firewall filter, you must first configure the filter and then apply it to a port, either while manually configuring a port or through port profiles.

Firewall Filter

Each port or interface on the switch can have a maximum of only two filters:

- Ingress firewall filter—A filter that is applied to packets that are entering a network.
- Egress firewall filter—A filter that is applied to packets that are exiting a network.

You can configure firewall filters to subject packets to filtering, class-of-service (CoS) marking (grouping similar types of traffic together, and treating each type of traffic as a class with its own level of service priority), and traffic policing (controlling the maximum rate of traffic sent or received on an interface). You can create an ingress and an egress firewall filter and deploy the filter on a port.

> **NOTE:** If you apply ingress and egress filters to the same interface, the ingress filter is processed first.

Firewall Filter Components

In a firewall filter, you define one or more terms that specify the filtering criteria and the action to be taken if a match occurs. A firewall filter can have multiple terms.

Each term consists of the following components:

- Match conditions—Specify the values or fields that the packet must contain to be considered a match. You can define various match conditions, including the IP source address field, IP destination address field, MAC source address field, MAC destination address field, Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) source port field, and IP protocol field.
- Action—Specifies what to do if a packet matches the match conditions. Possible actions are to accept or discard the packet. In addition, packets can be counted to collect statistical information.
- Action modifier—Specifies one or more actions for the switch if a packet matches the match conditions. You can specify action modifiers such as count, and log.

Firewall Filter Processing

If there are multiple terms in a filter, the order of the terms is important. Packets are tested against each term in the order in which the terms are listed in the firewall filter configuration. If a packet matches the first term, the switch executes the action defined by that term, and no other terms are evaluated. If the switch does not find a match between the packet and the first term, it compares the packet to the next term. If no match occurs between the packet and the second term, the system continues to compare the packet to each successive term in the filter until a match is found. If the packet does not match any terms in the filter, the switch discards the packet by default.

Configure a Firewall Filter for an EX Series Switch

You configure firewall filters on EX Series switches to control traffic that enters or exits the switch ports. To configure a firewall filter, you must first configure the filter and then apply it to a port.

The following workflow describes the steps that are required to configure a firewall filter and assign it to a switch port.

1. Create a firewall filter. See "Add Firewall Filters" on page 707.
2. Create a term and associate the term to the firewall filter. See "Add Terms to Firewall Filters" on page 710.
3. Assign the firewall filter to a port profile. See "Add Port Profiles" on page 674.

   Or manually edit the port configurations. See "Edit Configuration of Ports" on page 251.
4. Deploy the port profile to a port. See "Deploy or Redeploy a Port Profile" on page 247.

About the EX Firewall Filters Page

To access this page, select **Configuration** > **SD-LAN** > **Firewall Filters**.

Use this page to view, add, and delete firewall filters on EX Series switches.

Tasks You Can Perform

You can perform the following tasks from this page:

- View details of the existing firewall filter. See Table 234 on page 707.
- Add a firewall filter—See "Add Firewall Filters" on page 707.

- Delete a firewall filter—See “Delete Firewall Filters” on page 708.
- Add terms to a new firewall filter—See “Add Terms to Firewall Filters” on page 710.

Field Descriptions

Table 234 on page 707 describes the fields on the EX Firewall Filters page.

Table 234: Fields on the EX Firewall Filters page

Field	Description
Name	Name of the firewall filter.
Description	Description of the firewall filter.
Number of Intents	Number of terms added to the firewall filter.

Add Firewall Filters

Use the Add Firewall Filter page to add a new ingress or egress firewall filter.

To add a firewall filter:

1. Select **Configuration > SD-LAN > Firewall Filters**.

   The EX Firewall Filters page appears.

2. Click the add icon (+) to add a new firewall filter.

   The Add Firewall Filter page appears.

3. Complete the configuration according to the guidelines provided in Table 235 on page 708.

   NOTE: Fields marked with * are mandatory.

4. Click **OK**.

   The firewall filter is created. You are returned to the EX Firewall Filters page where a confirmation message is displayed.

Table 235: Fields on the Add Firewall Filter Page

Field	Description
Name	Enter a unique name for the firewall filter. The name can contain only alphanumeric characters and hyphen (-); the maximum length allowed is 15 characters.
Description	Enter a description for the firewall filter.

WHAT'S NEXT

After creating a firewall filter, you must add terms to the firewall filter. See Add Terms to Firewall Filters | 710.

Delete Firewall Filters

Before deleting a firewall filter, check whether the firewall filter is associated with a port profile or port(s). If it is associated, perform the following steps before deleting a firewall filter:

1. Disassociate the port profile from the port. See "Dissociate a Profile from a Port" on page 253.
2. Deploy the port profile. See "Deploy or Redeploy a Port Profile" on page 247.
3. Delete the port profile. See "Edit, Clone, and Delete Port Profiles" on page 679.

To delete a firewall filter:

1. Select **Configuration > SD-LAN > Firewall Filters**.

   The EX Firewall Filters page appears, displaying the configured firewall filters.

2. Select the firewall filter that you want to delete and click the **Delete** icon.

   An alert message appears, asking you to confirm the delete operation.

3. Click **Yes** to delete the selected firewall filter.

   A confirmation message appears on top of the page indicating the status of the delete operation. If the delete operation is successful, the firewall filter is removed from the EX Firewall Filters page.

About the < Firewall-Filters-Name> / Terms Page

To access this page, select **Configuration** > **SD_LAN** > **Firewall Filters** > ***Firewall-Filters-Name***.

Use the *Firewall-Filters-Name* / Terms page to manage terms for the firewall filters.

Tasks You Can Perform

You can perform the following tasks from this page:

- Add a firewall filter term—See "Add Terms to Firewall Filters" on page 710.
- Edit, clone, or delete a firewall filter term—See "Edit, Clone, and Delete Terms" on page 713.
- Search for terms by using keywords—Click the search icon and enter the search term in the text box and press Enter. The search results are displayed on the same page.
- Filter the firewall filter terms—Click the filter icon (funnel) and select **Show Hide Column Filters**. Specify one or more filtering criteria. The filtered results are displayed on the same page.

  NOTE: Filtering is applicable only to some fields.

- Quickly filter the terms based on deployment status—Click the filter icon (funnel) and select **Undeployed Terms** to view the terms that are awaiting deployment.

Field Descriptions

Table 236 on page 709 describes the fields on the *Firewall-Filters-Name* / Terms page.

Table 236: Fields on the <Firewall-Filters- Name> / Terms Page

Field	Description
Name	Name of the firewall filter term.
Source	Displays the source endpoint such as IP address, MAC address, port, or protocol.
Destination	Displays the destination endpoint such as IP address, MAC address, port, or protocol.

RELATED DOCUMENTATION

About the EX Firewall Filters Page | 706

Add Terms to Firewall Filters

Use the *Firewall-Filter-Name* page to add a firewall term that controls the ingress and egress traffic. The traffic is classified by matching its source and destination IP addresses (for Layer 3), MAC addresses (for Layer 2), ports, or protocols.

To configure a firewall term:

1. Select **Configuration > SD-LAN > Firewall Filters**.

   The EX Firewall Filters page appears.

2. Click the firewall filter to which you want to add the term.

   The *Firewall-Filter-Term-Name* page appears.

3. Click the add icon (+).

   The option to create firewall term appears inline on the *Firewall-Filter-Term-Name* page.

4. Complete the configuration according to the guidelines provided in Table 237 on page 710.

5. Click **Save** to save the changes. If you want to discard your changes, click **Cancel** instead.

If you click **Save**, a new firewall term with the provided configuration is added and a confirmation message is displayed.

If a firewall filter contains multiple terms, then, by default, the new term is always added at the top of the list of terms in the *Firewall-Filter-Term-Name* page. The term that is at the top of the list has higher priority than the others in the list. You can re-order the term by dragging and dropping the term at a different level in the list.

Table 237: Fields on the <Firewall-Filter-Term-Name> Page

Field	Description
Name	Enter a unique string of alphanumeric characters, colons, periods, dashes, and underscores. No spaces are allowed and the maximum length is 255 characters. If you do not enter a name, the term is saved with a default name assigned by CSO.
Description	Enter a description for the firewall filter term; maximum length is 1024 characters.

Table 237: Fields on the <Firewall-Filter-Term-Name> Page *(continued)*

Field	Description
Counter	Click the toggle button to enable (default) or disable the counter. The counter counts the number of packets that pass this filter term. **NOTE:** If you have enabled counter for the firewall filter, you cannot add the firewall filter as an egress filter.
Logging	Click the toggle button to enable (default) or disable logging. By enabling logging, CSO logs the packet's header information in the Routing Engine. **NOTE:** If you have enabled logging for the firewall filter, you cannot add the firewall filter as an egress filter.
Source	Click the add icon (**+**) to select the source endpoints from the displayed list of IP addresses, MAC addresses, protocols, or ports to the firewall filter term. You can also select a source end point using the methods described in "Selecting Firewall Source" on page 389.
Destination	Click the add icon (**+**) to select the destination endpoints from the displayed list of IP addresses, MAC addresses, protocols, or ports to the firewall filter term. You can also select a destination end point using the methods described in "Selecting Firewall Destination" on page 393.
Select Action	Click the add icon (+) to choose whether you want to permit or deny the traffic between the source and destination endpoints. • **Allow**—Device permits the traffic. • **Deny**—Device silently drops all packets for the session.

Table 237: Fields on the <Firewall-Filter-Term-Name> Page *(continued)*

Field	Description
Endpoints	To add an endpoint to the source or destination: 1. Click **Select Source** or **Select Destination** text box and then click the lesser-than icon on the right side of the page to open the End Points panel. The End Points panel displays the endpoints from addresses, MAC, protocols, and ports relevant to the source or destination based on your selection. **NOTE:** You can also search for a specific end point using the search option. 2. Select the endpoint you want to add and click the check mark icon (√) to add it the source or destination. The selected endpoint is added to the source or destination. To add new source and destination end points: 1. Click the less-than icon (<) on the right side of the page to open the End Points panel. 2. Click the add icon (+) on the top right of the End Points panel. A list of endpoints that you can add is displayed. 3. Select the endpoint you want to add. You can add the following endpoints: • Address or address group. See "Creating Addresses or Address Groups" on page 737. • MAC address. See "Add a MAC Address Endpoint" on page 770. • Protocol. See "Add a Protocol Endpoint" on page 774. • Port. See "Add a Port Endpoint" on page 777. 4. Click **Save** to add the new endpoint. The endpoint that you created is listed in the **End Points** panel. 5. Select the endpoint that you want to add to the source or destination, and click on the check mark icon (√). The endpoint is added to the source or destination as specified.

WHAT'S NEXT

After adding terms to the firewall filter, assign the firewall filter as an ingress filter or egress filter in port profiles. See Add Port Profiles | 674.

Edit, Clone, and Delete Terms

IN THIS SECTION

- Edit a Firewall Filter Term | 713
- Clone a Firewall Filter Term | 714
- Delete a Firewall Filter Term | 715

You can edit, clone, and delete firewall filter terms from the *Firewall-Filter-Name* page.

Edit a Firewall Filter Term

Before editing a firewall filter term, check whether the respective firewall filter is associated with a port. If it is associated with a port, you must re-deploy the port profile after editing the firewall term. To deploy the port profile, see "Deploy or Redeploy a Port Profile" on page 247.

To modify the parameters configured for a firewall filter term:

1. Select **Configuration > SD-LAN > Firewall Filter**.

   The EX Firewall Filters page appears, displaying the list of firewall filters.

2. Click the firewall filter for which you want to edit the term.

3. Hover over the firewall filter term that you want to edit, and then click the ellipsis icon**...** that appears on the right side of the term.

4. Click **Edit**.

   The *Firewall-Filter-Term-Name* page appears with the values that you provided while creating the firewall filter term.

5. Modify the parameters by following the guidelines provided in "Add Terms to Firewall Filters" on page 710.

NOTE: You cannot modify the name of the firewall filter term.

6. Click **Save** to save the changes. If you want to discard your changes, click **Cancel** instead.

   A confirmation message appears, indicating the status of the edit operation.

Clone a Firewall Filter Term

Before cloning a firewall filter term, check whether the respective firewall filter is associated with a port. If it is associated with a port, you must re-deploy the port profile after cloning the firewall term. To deploy the port profile, see "Deploy or Redeploy a Port Profile" on page 247.

To clone a firewall filter term:

1. Select **Configuration > SD-LAN > Firewall Filter**.

   The EX Firewall Filters page appears, displaying the list of firewall filters.

2. Click the firewall filter for which you want to clone the term.

3. Hover over the firewall filter term that you want to clone, and then click the **...** icon that appears on the right side of the term.

4. Click **Clone**.

   The *Firewall-Filter-Term-Name* page appears with the values that you provided while creating the firewall filter term.

5. Enter a new name for the term and modify the parameters following the guidelines provided in "Add Terms to Firewall Filters" on page 710

6. Click **OK** to save your changes.

   A confirmation message appears indicating the status of the clone operation. If the operation is successful, the cloned term is listed in the *Firewall-Filters-Name* page.

Delete a Firewall Filter Term

Before deleting a firewall filter term, check whether the respective firewall filter is associated with a port. If it is associated with a port, you must re-deploy the port profile after deleting the firewall term. To deploy the port profile, see "Deploy or Redeploy a Port Profile" on page 247.

To delete a firewall filter term:

1. Select **Configuration > SD-LAN > Firewall Filter**.

   The EX Firewall Filters page appears, displaying the list of firewall filters.

2. Click the firewall filter for which you want to delete the term.

3. Hover over the firewall filter term that you want to delete and then click the **...** icon that appears on the right side of the term.

4. Select **Delete**.

   A message requesting confirmation for the deletion appears.

5. Click **Yes** to delete the selected firewall filter term. If you want to discard your changes, click **Cancel** instead.

If you click **Yes**, the selected firewall filter term is deleted from the Firewall Filter page.

A confirmation message appears, indicating the status of the delete operation.

Deploy or Redeploy a Port Profile

A port profile defines the authentication settings for the port and other port parameters such as flow control, link mode, storm control, MAC limit, and so on. The behavior of the port is defined by the values for parameters defined in the port profile. You can deploy a port profile on one or more ports at the same time.

You must redeploy a port profile when:

- the port profile that is assigned to a port is modified.
- the authentication profile or firewall filter that is assigned to the port profile is modified.

When a port profile or a profile associated with the port profile is modified, the deployment status of the port is changed to Pending Deployment.

The changes made to the port profile or the associated authentication profile or firewall filter are applied on the ports only when you redeploy the modified port profile.

To deploy or redeploy a port profile on one or more switch ports:

1. Select the port and click **More > Deploy**.

   The Deploy page appears.

2. For the **Type** field, do one of the following:
   - Click **Run now** to deploy the port profile immediately.
   - Click **Schedule at a later time** to deploy the port profile later.

     If you select this option, enter the date and time when you want to deploy in the Date and Time fields that appear.
3. Click **OK**.

   If you select the Run now option, a job is created to deploy the profile immediately; otherwise, the job to deploy is created on the date and at the time that you scheduled.

   When you deploy a port profile, the deployment status of the ports is set to Pending Deployment indicating that the profile is associated with the port. When the profile is in the process of being committed on the ports, the deployment status changes to In progress. If the deployment job completes successfully, the deployment status of the ports is set to Success and if the job fails, the deployment status is set to Failed.

Enable Ports

You enable a port to allow traffic through the port. When you enable a port, the Admin status of the port is changed to Up.

> **NOTE:** You can enable one or more ports at the same time.

To enable one or more ports:

1. Select the ports and click **More > Enable Port(s)**. Alternatively, right-click the ports and click **Disable Port(s)**.

   A job to enable the ports is initiated.

After a port is enabled, the Admin Status is changed to Up.

NOTE: The device is monitored once every five minutes. Therefore, it takes upto five minutes for the change in the Admin Status to reflect on the CSO GUI.

Disable Ports

You disable a port to block traffic through the port. When you disable a port, the Admin status and Link status of the port are changed to Down.

NOTE: You can disable one or more ports at the same time.

To disable one or more ports:

1. Select the ports and click **More > Disable Port(s)**. Alternatively, right-click the port and click **Disable Port(s)**

   A job to disable the ports is initiated.

After a port is disabled, the Admin and Link statuses for the port are set to Down.

NOTE: The device is monitored once every five minutes. Therefore, it takes upto five minutes for the change in the Admin Status and Link Status to reflect on the GUI.

Edit Configuration of Ports

You can edit the configuration of a port either by using a port profile or manually.

If a profile is already deployed, edit the profile, and then redeploy for the changes to take effect.

When using a port profile to edit ports, you can:

- assign a port profile or modify the profile assigned to the port
- modify the VLANs assigned to the port

To edit the configuration of one or more ports:

1. Select one or more ports and click **More > Edit Configuration**. Alternatively, right-click the ports and click **Edit Configuration**.

   The Edit Port(s) page appears.

2. Do one of the following:

   For **Options**, select one of the following:

   - To edit the port configuration by using a port profile:

     a. Click **Use Port Profile**.

        The Port Profile drop down list and an option to select VLAN appears.

     b. Select a port profile that you want to assign to the port from the **Port Profile** drop down list.

     c. In the VLAN field:

        - If the port is configured as a trunk port in the port profile, assign multiple VLANs by selecting the VLANs in the Available column and clicking the right-arrow to move them to the Selected column.
        - If the port is configured as an access port in the port profile, assign a single VLAN by selecting the VLAN in the Available column and clicking the right-arrow to move it to the Selected column.

   - To edit the port configuration manually:

     a. Click **Manually edit port configurations.**

        The port parameters such as port mode, link mode, flow control appear.

     b. Edit the parameters by referring to Table 227 on page 674.

3. Click **Next**.

   Deployment options appear.

4. (Optional) Select the **Do not deploy** option if you want to only save the edited configuration in CSO, but not push and commit the configuration to the switch.

5. For the **Type** field, do one of the following:

   - Click **Run now** to save the edited configuration in CSO and commit the edited configuration on the switch.
   - Click **Schedule at a later time** to schedule a time to commit the edited configuration on the switch.

If you select the Schedule at later time option, enter the date and time when you want to deploy, in the fields that appear when you select the option.

6. Click **Next**.

   A summary of all the port profile parameter appears.

7. (Optional) Click **Edit** to revisit the settings and make further changes.

8. Click **OK**.

   If you select the Run now option, a job is created to deploy the profile at once; otherwise, the job to deploy is created on the date and at the time that you scheduled.

   During the deployment (that is when the deploy is executing), the Deployment status of the port is set to Pending Deployment, on the Ports tab of the *Devices* page.

   If the deployment job completes successfully:

   - A message appears on the top of the Ports page indicating that the deployment is successful.
   - Deployment Status of the port is set to Deployed.
   - The Port Profile column displays:
     - the profile name when a port profile is used for editing.
     - Manually Configured when the configuration is manually edited.

5

PART

Managing Network Services and Shared Objects

Configuring Network Services in a Distributed Deployment | **721**

Managing Shared Objects | **734**

CHAPTER 23

Configuring Network Services in a Distributed Deployment

IN THIS CHAPTER

- Network Service Overview | 721
- About the Network Services Page | 722
- About the Service Overview Page | 724
- About the Service Instances Page | 726
- Configuring VNF Properties | 727
- vSRX VNF Configuration Settings | 728

Network Service Overview

A *network service* is a final product offered to end users with a full description of its functionality and specified performance.

Administrative users deploy network services between two locations in a virtual network, so that traffic traveling in a specific direction on that link is subject to action from that service. The term *network service* is defined in the ETSI Network Functions Virtualization (NFV) standard.

A network service consists of a *service chain* of one or more linked network functions, which are provided by specific virtualized network functions (VNFs), with a defined direction for traffic flow and defined ingress and egress points. The term service chain refers to the structure of a network service, and although not defined in the ETSI NFV standard, this term is regularly used in NFV and software-defined networking (SDN).

A network service designer creates network services in Network Service Designer. When the designer publishes the service to the network service catalog from Network Service Designer, administrators can see the network service in Administration Portal.

RELATED DOCUMENTATION

About the Network Services Page | 722
About the Service Overview Page | 724
About the Service Instances Page | 726

About the Network Services Page

To access this page, click **Configuration > Network Services**.

You can use the Network Services page to view the complete list of network services that service designers have published to the network service catalog from network service designer and to view information about the services. For an introduction to network services, see "Network Service Overview" on page 721.

Tasks You Can Perform

You can perform the following tasks from this page:

- Quickly view important data about network services and about instances of those services deployed at customers' sites in the widgets that appear at the top of the page. See Table 238 on page 722.
- View full information about a service and about instances of a service at customer sites. Click the name of a service in the list. See "About the Service Instances Page" on page 726.

Field Descriptions

Table 238 on page 722 shows the descriptions of the widgets that appear at the top of the Network Services page.

Table 238: Widgets on the Network Services Page

Widget	Description
Top Network Services Instantiated	View the numbers of instances of the three services that are most used by tenants in the network. This view helps you identify trends for network services, especially when you introduce a new service.
Services with Critical Alerts	View the top three network services receiving the maximum number of critical alerts.
Top Services by POP CPU Usage	View the top three network services using the largest percentage of CPU from the assigned CPU cores.

Table 239 on page 723 shows the descriptions of the fields on the Network Services page.

Table 239: Fields on the Network Services Page

Field	Description
Name	View the name of the service. Click the name to view full information about a service.
Tenants	View thenames of the tenants that have access to the network service.
Sites	View the total number of sites at which the service is deployed for the tenant. Example: 2
Instances	View the total number of occurrences of the service that administrative users have activated for the tenant. Example: 1
Last Update	View the date on which the network service designer last modified the service.

Table 240 on page 723 shows the descriptions of the fields on the Detail for *network service name* page.

Table 240: Fields on the Network Service Detail Page

Field	Description
General	
Configuration	View the settings that the network service designer or you have configured for this service.
Version	View the version number of the network service. Example: 1.1
State	View the status of the network service. Example: Published
Performance Goals	View performance parameters of the network service that include bandwidth, number of sessions, latency, and license cost.

RELATED DOCUMENTATION

Network Service Overview | 721
About the Service Overview Page | 724
About the Service Instances Page | 726

About the Service Overview Page

To access this page, click **Service > Service Name > Overview**.

You can use the Service Overview page to view information about a service that the service designer has published to the network service catalog from Network Service Designer.

Tasks You Can Perform

You can perform the following tasks from this page:

- View administrative details about the service. See *General Information* in Table 241 on page 724.
- View resources required for the service and its performance specification. See *Service Requirements* and *Service Performance* in Table 241 on page 724.
- View the service chain, with its constituent VNFs. See *Service Configuration* in Table 241 on page 724.
- Configure VNFs. Click a VNF in the service chain graphic. See "vSRX VNF Configuration Settings" on page 728.

Field Descriptions

Table 241 on page 724 provides guidelines on using the fields on the Service Overview page.

Table 241: Fields on the Service Overview Page

Field	Description
General Information	
Description	View a summary about the service's capabilities. The network service designer provides this summary.
State	View the state of the network service: • Discontinued—Service is no longer available for customers. • Published—Service designer has published service to network catalog, and it is available for customers.

Table 241: Fields on the Service Overview Page *(continued)*

Field	Description
Tenants	View the number of tenants using this service.
Service Requirements	
CPU	View the number of CPUs that the service needs (cores).
Memory	View the amount of RAM that the service needs in gigabytes (GB).
Service Performance	
Sessions	View the number of sessions concurrently supported by one instance of the service.
Bandwidth	View the data rate for the service in megabytes per second (Mbps) or gigabytes per second (Gbps).
Latency	View the time a packet takes to traverse the service in milliseconds (ms) or nanoseconds (ns).
License cost	Specify the license cost for the network service in USD.
Service Configuration (graphic of the service chain)	
I	View the ingress point—the point at which packets enter the service.
E	View the egress point—the point at which packets exit the service.
One or more VNFs	Click to view settings for the VNF. See "vSRX VNF Configuration Settings" on page 728. The service designer can configure the VNF settings in Network Service Designer and the administrative user can configure the VNF settings in Customer Portal. **BEST PRACTICE:** The network service designer configures settings for the virtual machine (VM) in which the virtualized network function (VNF) resides and the administrative user configures settings for the service, such as policies. The service designer can also configure a few example settings for the service. These example settings should be generic and not network-specific.

RELATED DOCUMENTATION

Network Service Overview | 721

About the Network Services Page | 722
About the Service Instances Page | 726

About the Service Instances Page

To access this page, click **Services** > *Service Name* > **Instances**

You can use the Service Instances page to view information about occurrences of the service at specific customer sites.

Tasks You Can Perform

You can perform the following tasks from this page:

- View details about a service instance. Click the details icon that appears when you hover over the name of a service. See Table 243 on page 727.
- Enable or disable a network service or virtualized network function (VNF) recovery. Select a service instance and click **Enable Auto Healing** to enable automatic recovery of a network service or VNF in a centralized deployment. By default, automatic recovery of a network service or VNFs is enabled. See "Configuring VNF Properties" on page 727.

Field Descriptions

Table 242 on page 726 shows the descriptions of the fields on the Service Instances page.

Table 242: Fields on the Service Instances Page

Field	Description
Name	View the name of the occurrence of a service at a specific tenant site.
Tenant	View the name of the tenant.
Status	View the state of the service at the customer site: • Created—Administrative user for the tenant has enabled this service instance, which is active. • Blank—Administrative user for the tenant has disabled this service instance.
Site	View the name of the site at which service occurrence is available.
POP	View the POP in which the site is located.

Table 242: Fields on the Service Instances Page *(continued)*

Field	Description
Functions	View network functions that the service offers; for example, Network Address Translation (NAT) or firewall.

Table 243 on page 727 shows the descriptions of the fields on the Detail for *Service-Instance-Name* page.

Table 243: Fields on the Service Instance Details Page

Field	Description
General	
Description	View information about this service instance. This information is generated from data in Customer Portal.

RELATED DOCUMENTATION

- Network Service Overview | **721**
- About the Network Services Page | **722**
- About the Service Overview Page | **724**

Configuring VNF Properties

You can specify whether to enable automatic recovery of a network service or virtualized network function (VNF) for a network service instance in a centralized deployment. Enabling automatic recovery of a network service or VNF improves reliability of the implementation.

Conversely, disabling automatic recovery of a network service or VNF allows you to quickly investigate a problem with a network service or VNF itself.

To enable or disable automatic recovery of a network service or VNF:

1. Select **Services** > *Services Name* > **Instances**.

   The Services Instances page appears.

2. Select a service instance for which you want to enable or disable automatic recovery.

3. Click **Enable Auto Healing**.

   The Service Properties page appears.

4. Select whether you want to enable or disable automatic recovery.

   NOTE: By default, automatic recovery of a network service or VNF is enabled.

5. Click **Save**.

RELATED DOCUMENTATION

About the Service Instances Page | 726

vSRX VNF Configuration Settings

You can configure the vSRX VNF from **Services > *Service Name* > Overview > Service Configuration**. Your service provider usually configures base settings for the virtual machine (VM) in which the virtualized network function (VNF) resides and you configure settings for the service, such as policies.

NOTE: A vSRX firewall virtualized network function (VNF) is always part of a service chain for a network service on a CPE device.

Use the information in the following tables to provide values for the available settings:

- Table 244 on page 729 shows the settings you can configure for the virtual machine (VM) that contains the VNF.

  NOTE: Your service provider usually configures the base settings and you should not need to change them.

- Table 245 on page 730 shows the firewall settings you can configure.

Table 244: Fields for the vSRX Base Settings

Field	Description
Host Name	For a cloud site, specify the hostname of the VM that contains the vSRX VNF. The field has no limit on the number of characters and accepts letters, numbers, and symbols. Example: vm-vsrx For an on-premise site, the vSRX application resides on the CPE device, and you cannot configure this setting.
Loopback Address	Specify an IPv4 loopback address for the management interface of the VM. Example: 192.0.2.25
DNS Servers	Specify the fully qualified domain names (FQDNs) or IP addresses of one or more DNS name servers. Example: 192.0.2.35
NTP Servers	Specify the FQDNs or IP addresses of one or more NTP servers. Example: 192.0.2.45
Syslog Servers	Specify the FQDNs or IP addresses of one or more system log servers. Example: 192.0.2.55
Enable Re-filter	Select **True** to enable a stateless firewall filter that protects the Routing Engine from denial-of-service (DoS) attacks or **False** to allow DoS attacks. Example: True
Enable Default Screens	For a cloud site, select **True** to enable the default screens security profile for the destination zone or **False** to disable default screening. Example: False You cannot configure this setting for an on-premise site.
Time Zone	Specify the time zone for the VM. Example: UTC

Table 244: Fields for the vSRX Base Settings *(continued)*

Field	Description
Right Interface	Specify the identifier of the VM interface that transmits data. Example: ge-0/0/1 For an on-premise site, the vSRX application resides on the CPE device, and you cannot configure this setting.
Left Interface	Specify the identifier of the VM interface that receives data. Example: ge-0/0/0 For an on-premise site, the vSRX application resides on the CPE device, and you cannot configure this setting.
SNMP Prefix List	If you set the Enable Re-filter field to **True**, specify the routes that the Junos Space Virtual Appliance uses for SNMP operations when it discovers the vSRX VNF. Example: 10.0.2.0/24
Ping Prefix List	If you set the Enable Re-filter field to **True**, specify the routes that the Junos Space Virtual Appliance uses for ping operations when it discovers the vSRX VNF. Example: 10.0.2.1/24
Space Servers	If you set the Enable Re-filter field to **True**, specify the IP addresses of the VMs that contain the Junos Space Virtual Appliances. Example: 10.0.2.50

Table 245: Fields for the vSRX Firewall Settings

Field	Description
Policy Name	Specify the name of the rule. The field has no limit on the number of characters and accepts letters, numbers, and symbols. Example: policy-1

Table 245: Fields for the vSRX Firewall Settings *(continued)*

Field	Description
Source Zone	Select the security zone from which packets originate. • **left**—Interface that transmits data to the host • **right**— Interface that receives data transmitted from the host Zone policies are applied to traffic traveling from one security zone (source zone) to another security zone (destination zone). This combination of a source zone and a destination zone is called a *context*. Example: left
Destination Zone	Select the security zone to which packets are delivered. • **left**—Interface that transmits data to the host • **right**—Interface that receives data transmitted from the host Zone policies are applied to traffic traveling from one security zone (source zone) to another security zone (destination zone). This combination of a source zone and a destination zone is called a *context*. Example: right
Source Address	Specify the source IP address prefixes that the network service uses as match criteria for incoming traffic. To add source addresses: 1. Click the **Source Address** column. The source-address page appears. 2. Select **any** to match any source IP address of packets or **ipp** to match a specific prefix in the source IP address for which the application enforces the policy. 3. If you select **ipp**, specify a prefix. 4. Click **OK**. Example: 10.0.2.30

Table 245: Fields for the vSRX Firewall Settings *(continued)*

Field	Description
Destination Address	Specify the destination IP address prefixes that the network service uses as match criteria for outgoing traffic. To add a destination address: 1. Click the **Destination Address** column. The destination-address page appears. 2. Select **any** to match any source IP address of packets or **ipp** to match a specific prefix in the source IP address for which the application enforces the policy. 3. If you select **ipp**, specify a prefix. 4. Click **OK**. Example: 192.0.2.0/24
Action	Select **permit** to transmit packets that match the rule or **deny** to drop packets that match the rule. Example: permit

Table 245: Fields for the vSRX Firewall Settings *(continued)*

Field	Description
Application	Specify the applications to which the policy applies. The applications are based on protocols and ports. To specify applications: 1. Click the **Application** column. The application page appears. 2. In the allowed_apps field, select **any** to match any application or **app** to choose specific applications. If you select **app**, press and hold the Ctrl key and click the required applications from the drop-down list. • junos-tcp-any • junos-udp-any • junos-ftp • junos-http • junos-https • junos-icmp-all • junos-icmp-ping • junos-telnet • junos-tftp 3. Click **OK**. Example: • junos-tcp-any • junos-udp-any

RELATED DOCUMENTATION

- About the Network Services Page | **722**
- About the Service Overview Page | **724**
- About the Service Instances Page | **726**
- Configuring VNF Properties | **727**

CHAPTER 24

Managing Shared Objects

IN THIS CHAPTER

- Addresses and Address Groups Overview | **735**
- About the Addresses Page | **736**
- Creating Addresses or Address Groups | **737**
- Editing, Cloning, and Deleting Addresses and Address Groups | **740**
- Services and Service Groups Overview | **742**
- About the Services Page | **743**
- Creating Services and Service Groups | **744**
- Creating Protocols | **746**
- Editing and Deleting Protocols | **749**
- Editing, Cloning, and Deleting Services and Service Groups | **751**
- Application Signatures Overview | **753**
- About the Application Signatures Page | **754**
- Understanding Custom Application Signatures | **755**
- Adding Application Signatures | **757**
- Editing, Cloning, and Deleting Application Signatures | **762**
- Adding Application Signature Groups | **764**
- Editing, Cloning, and Deleting Application Signature Groups | **765**
- About the Departments Page | **766**
- Adding a Department | **767**
- Deleting a Department | **768**
- About the MAC Addresses Page | **769**
- Add a MAC Address Endpoint | **770**
- Edit or Delete MAC Address Endpoint | **771**
- About the Protocols Page | **773**
- Add a Protocol Endpoint | **774**
- Edit or Delete Protocol Endpoint | **775**
- About the Ports Page | **776**

Add a Port Endpoint | 777
Edit or Delete Port Endpoint | 779

Addresses and Address Groups Overview

An address specifies an IP address or a hostname. You can create addresses that can be used across all policies. Addresses are used in firewall and NAT services and apply to the corresponding policies. If you know only the hostname, you enter it into the **Hostname** field and use the address resolution option to resolve it to an IP address. You can also resolve an IP address to the corresponding hostname.

After you create an address, you can combine it with other addresses to form an address group. Address groups are useful when you want to apply the same policy to multiple addresses.

Contrail Service Orchestration (CSO) manages its address book at the global level, assigning objects to devices that are required to create policies. An address book is a collection of addresses and address groups that are available in a security zone. If the device is capable of using a global address book, CSO pushes address objects used in the policies to the global address book of the device.

RELATED DOCUMENTATION

About the Addresses Page | 736
Creating Addresses or Address Groups | 737
Editing, Cloning, and Deleting Addresses and Address Groups | 740

About the Addresses Page

To access this page, select **Configuration** > **Shared Objects** > **Addresses**.

Use this page to create, edit, and delete addresses and address groups. Addresses and address groups are used in firewall and NAT services. After you create an address, you can combine it with other addresses to form an address group. Address groups are useful when you want to apply the same policy to multiple services.

Tasks You Can Perform

You can perform the following tasks from this page:

- Create an address or address group. See "Creating Addresses or Address Groups" on page 737.
- Modify, clone, or delete an address or address group. See "Editing, Cloning, and Deleting Addresses and Address Groups" on page 740.
- View the configured parameters of an address or address group. Click the details icon that appears when you hover over the name of an address or address group or select **More > Detailed View**.
- Show or hide columns about the address or address group. Click the **Show Hide columns** icon in the top right corner of the page and select columns that you want to view on the page.
- Search for an address or address group. Click the Search icon in the top right corner of the page to search for an address or address group.

  You can enter partial text or full text of the keyword in the text box and press Enter.

  The search results are displayed on the same page.

Field Descriptions

Table 246 on page 736 provides guidelines on using the fields on the Addresses page.

Table 246: Fields on the Addresses Page

Field	Description
Name	View the name of the address or address group.
Type	View the type of the address or address group.
Hostname	View the hostname of the address.
IP Address	View the IP address associated with the address.

Table 246: Fields on the Addresses Page *(continued)*

Field	Description
Description	View the description provided about the address or address group when it was created.

RELATED DOCUMENTATION

- Addresses and Address Groups Overview | 735
- Creating Addresses or Address Groups | 737
- Editing, Cloning, and Deleting Addresses and Address Groups | 740

Creating Addresses or Address Groups

Use the **Addresses** page to create addresses and address groups. Addresses and address groups are used in firewall and NAT services. After you create an address, you can combine it with other addresses to form an address group. Address groups are useful when you want to apply the same policy to multiple services.

To create an address or address group:

1. Select **Configure > Shared Objects > Addresses**.

   The **Addresses** page appears.

2. Click on the add icon (+).

   The **Create Addresses** page appears.

3. Complete the configuration according to the guidelines provided in Table 247 on page 738 and Table 248 on page 740.

4. Click **OK** to save the changes. If you want to discard your changes, click **Cancel** instead.

A new address or address group with your configurations is created. You can use this object in firewall or NAT policies.

Table 247: Fields on the Create Addresses Page

Field	Description
Object Type	Select Address or Address Group. If you select Address Group, then the screen changes so you can select the addresses you want to include in your address group. Table 248 on page 740 describes address group configuration parameters.
Name	Enter a unique name for the address. It must begin with an alphanumeric character and cannot exceed 63 characters. Dashes and underscores are allowed.
Description	Enter a description for your address; maximum length is 1,024 characters. You should make this description as useful as possible for all administrators.

Table 247: Fields on the Create Addresses Page *(continued)*

Field	Description

Type	Select a type of address and fill in the corresponding fields. Available types are:
	• Host
	  ◦ Host IP—Enter the IPv4 or IPv6 host IP address. For example: 192.0.2.0 or 2001:db8:4136:e378:8000:63bf:3fff:fdd2. If you do not know the IP address, you can enter the hostname and click **Look up hostname**.
	  ◦ Hostname—Enter the hostname. It must begin with an alphanumeric character and cannot exceed 63 characters. Dashes and underscores are allowed. If you do not know the host name, you can enter the IP address and click **Look up IP address**. For example, enter www.company.com and click **Look up IP address**. Hostname lookup is supported for IPv4 and IPv6 addresses.
	• Range
	  ◦ Start Address—Enter a starting IPv4 or IPv6 address for the address range. For example: 192.0.2.0 or 2001:db8:4136:e378:8000:63bf:3fff:fdd2.
	  ◦ End Address—Enter an ending IPv4 or IPv6 address for the address range. The range is validated after you enter the address.
	NOTE: An address range is configured on a managed device as an address set with one or more network address objects covering the specified address range.
	• Network
	  ◦ Network—Enter the network IP address. For example: 192.0.2.0. IPv6 is also supported. For example: 2001:db8:4136:e378:8000:63bf:3fff:fdd2.
	  ◦ Subnet Mask—Enter the subnet mask for the network range. For example, IPv4 netmask: 192.0.2.0/24. The subnet mask is validated as you enter it. You must enter the correct subnet mask in accordance with the network value. For example, IPv6 netmask: 2001:db8:4136:e378:8000:63bf:3fff:fdd2.
	• Wildcard
	  ◦ Network—Enter the network IPv4 or IPv6 address. For example: 192.0.2.0 or 2001:db8:4136:e378:8000:63bf:3fff:fdd2.
	  ◦ Wildcard Mask—Enter the wildcard mask for the network range. For example: 0.0.0.255.
	• DNS Host
	  ◦ DNS Name—Enter the DNS name. For example: company.com. Only alphanumeric characters, dashes, and periods are accepted. This name cannot exceed 69 characters in length, and must end with an alphanumeric character.

Table 248: Address Group Settings

Field	Description
Object Type	Select Address or Address Group. If you select Address Group, then the screen changes so you can select the addresses you want to include in your address group. Table 247 on page 738 describes address group configuration parameters.
Name	Enter a unique name for the address group. It must begin with an alphanumeric character and cannot exceed 63 characters. Dashes and underscores are allowed.
Description	Enter a description for your address group; maximum length is 1,024 characters. You should make this description as useful as possible for all administrators.
Addresses	Select the check box beside each address you want to include in the address group. Click the greater-than icon (>) to move the selected address or addresses from the **Available** column to the **Selected** column. Note that you can use the fields at the top of each column to search for addresses.

RELATED DOCUMENTATION

- Addresses and Address Groups Overview | 735
- About the Addresses Page | 736
- Editing, Cloning, and Deleting Addresses and Address Groups | 740

Editing, Cloning, and Deleting Addresses and Address Groups

IN THIS SECTION

- Editing Addresses and Address Groups | 741
- Cloning Addresses and Address Groups | 741
- Deleting Addresses and Address Groups | 742

You can edit, clone, and delete addresses and address groups from the **Addresses** page.

Editing Addresses and Address Groups

To modify the parameters configured for an address or address group:

1. Select **Configuration > Shared Objects > Addresses**.

   The **Addresses** page appears.

2. Select the address or address group that you want to edit, and then click **More > Edit**, or click the edit icon (pencil symbol) at the right top corner of the table, or right-click and select **Edit**.

   The **Edit** page appears, showing the same options as displayed when you create a new address or address group.

3. Modify the parameters according to the guidelines provided in "Creating Addresses or Address Groups" on page 737.

4. Click **OK** to save the changes. If you want to discard your changes, click **Cancel** instead.

When you click **OK**, the modified address or address group is displayed on the **Addresses** page.

> **NOTE:** When you edit an address that is a deployed as part of a policy, you will need to redeploy that policy in order for the changes to take effect. See "Deploying Policies" on page 668 for more information.

Cloning Addresses and Address Groups

To clone an address or address group:

1. Select **Configuration > Shared Objects > Addresses**.

   The **Addresses** page appears.

2. Right-click the address or address group that you want to clone and then click **Clone**, or select **More > Clone**.

   The **Clone** page appears with editable fields.

3. Modify the configured parameters of the address or address group, as required.

4. Click **OK** to save the changes. If you want to discard your changes, click **Cancel** instead.

If you select **OK**, the cloned address or address group is saved.

Deleting Addresses and Address Groups

> **NOTE:** Only addresses or address groups that have not been referenced in any policy can be deleted. If you try to delete such an address or address group, an error message will be displayed.

To delete an address or address group:

1. Select **Configuration > Shared Objects > Addresses**.

   The **Addresses** page appears.

2. Select the address or address group you want to delete and then click the delete icon **(X)** .

   An alert message appears verifying that you want to delete your selection.

3. Click **Yes** to delete the address or address group. If you do not want to delete, click **Cancel** instead.

If you select **Yes**, the selected address or address group is deleted, unless it is referenced in a policy.

RELATED DOCUMENTATION

- Addresses and Address Groups Overview | 735
- About the Addresses Page | 736
- Creating Addresses or Address Groups | 737

Services and Service Groups Overview

A service refers to an application on a device. For example, Domain Name Service (DNS). Services are based on protocols and ports used by an application, and when added to a policy, a configured service can be applied across all devices. Services are candidates for firewall policy end-points. The protocols used to create a service include: TCP, UDP, MS-RPC, SUN-RPC, ICMP, and ICMPv6. Contrail Service Orchestration (CSO) also includes predefined, commonly used services, and you cannot modify or delete them.

Once you create a service, you can combine it with other services to form a service group. Service groups are useful when you want to apply the same policy to multiple services, as this enables you create fewer policies.

RELATED DOCUMENTATION

About the Services Page | 743

Creating Services and Service Groups | 744

Editing, Cloning, and Deleting Services and Service Groups | 751

About the Services Page

To access this page, select **Configuration** > **Shared Objects** > **Services**.

Use the **Services** page to create, modify, clone and delete service or service groups. You can also create and manage protocols, that you use to create services.

A service refers to an application on a device, such as Domain Name Service (DNS). Services are based on protocols and ports used by an application. When added to a policy, a configured service can be applied across all devices. The protocols available to create a service include: TCP, UDP, SUN-RPC, MS-RPC, ICMP, ICMPv6, and so on.

Tasks You Can Perform

You can perform the following tasks from this page:

- Create a service or service group. See “Creating Services and Service Groups” on page 744.
- Modify, clone or delete a service or service group. See “Editing, Cloning, and Deleting Services and Service Groups” on page 751.
- View the configured parameters of a service or service group. Click the details icon that appears when you hover over the name of a service or service group, or click **More > Detailed View**.
- Show or hide columns about the services or service groups. Click the **Show Hide columns** icon in the top right corner of the page and select columns that you want to view on the page.
- Search a specific service or service group. Click the Search icon in the top right corner of the page to search for a service or service group.

  You can enter partial text or full text of the keyword in the text box and press Enter.

  The search results are displayed on the same page.

Field Descriptions

Table 249 on page 744 provides guidelines on using the fields on the **Services** page.

Table 249: Fields on the Service Page

Field	Description
Name	Name of the service or service group.
Type	Specifies whether the object is a service or service group.
Description	Description about the service or service group.
Predefined or Custom	List of predefined services and service groups, and a list of custom services or service groups that you created.

RELATED DOCUMENTATION

- Services and Service Groups Overview | 742
- Creating Services and Service Groups | 744
- Editing, Cloning, and Deleting Services and Service Groups | 751

Creating Services and Service Groups

Use the **Create Service** page to create a service. You can create services based on protocols and ports used by an application. The protocols used to create a service include: TCP, UDP, MS-RPC, SUN-RPC, ICMP, and ICMPv6. Once you create a service, you can combine it with other services to form a service group. Service groups are useful when you want to apply the same policy to multiple services.

You can also create or modify protocols that you base your services on, from the **Services** page.

To configure a service or service group:

1. Select **Configuration > Shared Objects > Services**.

   The **Services** page appears.

2. Click the add icon (+) to create service or service group.

   The **Create Services** page appears.

3. Complete the configuration of a service according to the guidelines provided in Table 250 on page 745.

If you want to configure a service group, see Table 251 on page 745.

4. Click **OK** to save the changes. If you want to discard your changes, click **Cancel** instead.

A new service or service group with the configuration you provided is created. You can use this service or service group as an endpoint in firewall policies.

Table 250 on page 745 provides guidelines on using the fields to create a service.

Table 250: Service Settings

Field	Description
Object Type	Select **Service** or **Service Group**. If you select **Service Group**, then the page changes so you can select the services you want to include in your service group.
Name	Enter a unique name for the service. It must begin with an alphanumeric character and cannot exceed 63 characters; dashes and underscores are allowed.
Description	Enter a description for your service. You should make this description as useful as possible for all administrators.
Protocols	Select the protocol you want to associate with the service. You can use existing protocols that are listed in the **Protocols** table. You can also create a new protocol, or edit existing protocols: • To create a new protocol, click on the add icon (+). See "Creating Protocols" on page 746. • To edit an existing protocol, click on the edit icon (pencil symbol). See "Editing and Deleting Protocols" on page 749.

Table 251 on page 745 provides guidelines on using the fields to create a service group.

Table 251: Service Group Settings

Field	Description
Object Type	Select **Service** or **Service Group**. If you select **Service Group**, then the screen changes so you can select the services you want to include in your service group.
Name	Enter a unique name for the service. It must begin with an alphanumeric character and cannot exceed 63 characters; dashes and underscores are allowed.
Description	Enter a description for your service group. You should make this description as useful as possible for all administrators.

Table 251: Service Group Settings *(continued)*

Field	Description
Services	Select the service you want to include in the service group and click the greater-than icon (>) to move the selected service or services from the **Available** column to the **Selected** column. You can use the search field at the top of each column to search for listed services.

RELATED DOCUMENTATION

- Services and Service Groups Overview | **742**
- About the Services Page | **743**
- Editing, Cloning, and Deleting Services and Service Groups | **751**
- Creating Protocols | **746**
- Editing and Deleting Protocols | **749**

Creating Protocols

Use the **Create Protocol** page to create TCP, UDP, MS-RPC, SUN-RPC, ICMP, and ICMPv6 protocols, that can be used in services. A service refers to an application on a device. Services are based on protocols and ports used by an application.

To create a protocol:

1. Select **Configuration > Shared Objects > Services**.

   The **Services** page appears.

2. Click the add icon (+) to create service or service group.

   The **Create Services** page appears.

3. Click the add icon (+) that appears about the **Protocols** table.

   The **Create Protocol** page appears.

4. Complete the configuration of the protocol according to the guidelines provided in Table 252 on page 747 and Table 253 on page 747.

5. Click **OK** to save the changes. If you want to discard your changes, click **Cancel** instead.

A new protocol with the configuration you provided is created. You can use this protocol to create services.

Table 252 on page 747 provides guidelines on using the fields to create a protocol.

Table 252: Fields on Create Protocol Page Settings

Field	Description
General Information	
Name	Enter a unique name for the protocol. It must begin with an alphanumeric character and cannot exceed 63 characters; dashes and underscores are allowed.
Description	Enter a description for your protocol. It cannot exceed 1,024 characters.
Type	Select the type of the protocol you want to create and fill in the corresponding fields. The available types of protocols are: TCP, UDP, ICMP, SUN-RPC, MS-RPC, ICMPv6, and so on. If you select TCP, continue with this table. See Table 253 on page 747 for the other protocol types.
Destination Port	Enter a destination port number for TCP. The range is from 0 to 65, 535.
Advanced Settings	
Enable Inactivity Timeout	Enabled by default. Enter a timeout value for this protocol in seconds or minutes. The maximum values are 129,600 seconds or 2,160 minutes.
ALG	Select an ALG (Application Layer Gateway) service option if applicable.
Source Ports and Port Ranges	Enter the source port or port range for the protocol.

Table 253 on page 747 includes the settings and guidelines for the various protocol types.

Table 253: Create Protocol Type Settings

Field	Description
UDP	
Destination Port	Enter a destination port number for UDP. This is a value or value range from 0 through 65,535.
Advanced Settings	
Enable Inactivity Timeout	Selected by default. Enter a timeout value for this protocol in seconds or minutes. The maximum values are 129,600 seconds and 2,160 minutes.

Table 253: Create Protocol Type Settings *(continued)*

Field	Description
ALG	Select an ALG (Application Layer Gateway) service option if applicable.
Source Ports and Port Ranges	Enter a source port or port range for UDP. This is a value or value range from 0 through 65,535.
ICMP	
Enable Inactivity Timeout	Enabled by default. Enter a timeout value for this protocol in seconds or minutes. The maximum values are 129,600 seconds and 2,160 minutes.
ICMP Type	Enter a value from 0 through 225 for the ICMP message type. For example, enter 1 for host unreachable. You can find these values in RFC 792.
ICMP Code	Enter a value from 0 through 225 for the ICMP code. For example, enter 0 for echo reply. You can find these values in RFC 792.
SUN-RPC	
Destination Port (available if Enable ALG is selected)	Enter a destination port for SUN-RPC. This is a value or value range from 0 through 65,535.
Enable Inactivity Timeout	Enabled by default. Enter a timeout value for this protocol in seconds or minutes. The maximum values are 129,600 seconds and 2,160 minutes.
Enable ALG	Not selected by default. If you enable ALG for this protocol, you must enter a destination port in the field that becomes available.
RPC Program Number	Enter a value or value range for the RPC (remote procedure call) service. For example, enter 100,017 for remote execution. You can find these values in RFC 5531.
Protocol Type	Select TCP or UDP for the protocol type.
MS-RPC	
Destination Port (available if Enable ALG is selected)	Enter a destination port for MS-RPC. This is a value or value range from 0 through 65,535.
Enable Inactivity Timeout	Enabled by default. Enter a timeout value for this protocol in seconds or minutes. The maximum values are 129,600 seconds and 2,160 minutes.
Enable ALG	Not selected by default. If you enable ALG for this protocol, you must enter a destination port number in the field that becomes available.

Table 253: Create Protocol Type Settings *(continued)*

Field	Description
UUID	Enter the corresponding UUID value for the MS-RPC service. For predefined values, refer to MS-RPC UUID Mappings.
Protocol Type	Select TCP or UDP for the protocol type.
ICMPv6	
Enable Inactivity Timeout	Selected by default. Enter a timeout value for this protocol in seconds or minutes. The maximum values are 129,600 seconds and 2,160 minutes.
ICMP Type	Enter a value from 0 through 225 for the ICMPv6 message type. You can find these values in RFC 4443.
ICMP Code	Enter a value from 0 through 225 for the ICMPv6 code. You can find these values in RFC 4443.
Destination Port	Use other to create protocols that do not match the provided type categories. Enter a destination port for the other protocol. This is a value or value range from 0 through 65,535.

RELATED DOCUMENTATION

- Editing and Deleting Protocols | 749
- About the Services Page | 743
- Creating Services and Service Groups | 744

Editing and Deleting Protocols

IN THIS SECTION

- Editing Protocols | 750
- Deleting Protocols | 750

You can edit and delete protocols through the **Services** page.

Editing Protocols

To modify the parameters configured for a protocol:

1. Select **Configuration > Shared Objects > Services** .

   The **Services** page appears.

2. Select the service to which the protocol you want to edit is associated, and click on the edit icon (pencil symbol) on the right top corner of the table, or right-click and select **Edit Service**.

   The **Edit Service** page appears, listing the protocols associated with the service in **Protocols** table.

3. Select the protocol that you want to edit, and then click on the edit icon (pencil symbol) on the right top corner of the **Protocols** table, or right-click and select **Edit Protocol**.

   The **Edit Protocol** page appears, showing the same fields as those seen when you create a new protocol.

4. Modify the parameters of the protocol according to the guidelines provided in "Creating Protocols" on page 746.

5. Click **OK** to save the changes. If you want to discard your changes, click **Cancel** instead.

If you click **OK**, the modified protocol appears in the **Protocols** table.

Deleting Protocols

To delete a protocol:

1. Select **Configuration > Shared Objects > Services**.

   The **Services** page appears.

2. Select the service to which the protocol you want to delete is associated, and click on the edit icon (pencil symbol) on the right top corner of the table, or right-click and select **Edit Service**.

   The **Edit Service** page appears, listing the protocols associated with the service in **Protocols** table.

3. Select the protocol you want to delete and then click the delete icon **(X)**.

   An alert message appears, verifying that you want to delete the protocol.

4. Click **Yes** to delete the protocol. If you do not want to delete, click **Cancel** instead.

If you click **Yes**, the selected protocol is deleted.

RELATED DOCUMENTATION

- Services and Service Groups Overview | 742
- About the Services Page | 743
- Creating Services and Service Groups | 744
- Editing, Cloning, and Deleting Services and Service Groups | 751
- Creating Protocols | 746

Editing, Cloning, and Deleting Services and Service Groups

IN THIS SECTION

- Editing Services and Service Groups | 751
- Cloning Services or Service Groups | 752
- Deleting Services and Service Groups | 752

You can edit, clone, and delete services and service groups from the **Services** page.

Editing Services and Service Groups

To modify the parameters configured for a service or service group:

1. Select **Configuration > Shared Objects > Services**.

   The **Services** page appears.

2. Select the service or service group that you want to edit, and click on the edit icon (pencil symbol) on the right top corner of the table, or right-click and select **Edit Service**.

   The **Edit Service** page appears, displaying the same options that are displayed when creating a new service or service group.

3. Modify the parameters according to the guidelines provided in "Creating Services and Service Groups" on page 744.

4. Click **OK** to save the changes. If you want to discard your changes, click **Cancel** instead.

If you click **OK**, you will see the modified service or service group in the **Services** page.

Cloning Services or Service Groups

To clone a service or service group:

1. Select **Configuration > Shared Objects > Services**.

   The **Services** page appears.

2. Right-click on the service or service group that you want to clone and then click **Clone**, or select **More > Clone**.

   The **Clone Service** page appears with editable fields.

3. Click **OK** to save the changes. If you want to discard your changes, click **Cancel** instead.

If you click **OK**, the cloned service or service group will appear beneath the selected service or service group.

Deleting Services and Service Groups

To delete a service or service group:

1. Select **Configuration > Shared Objects > Services**.

   The **Services** page appears.

2. Select the service or service group you want to delete and then click the delete icon **(X)**.

   An alert message appears, verifying that you want to delete the service or service group.

3. Click **Yes** to delete the service or service group. If you do not want to delete, click **Cancel** instead.

If you click **Yes**, the selected service or service group is deleted.

RELATED DOCUMENTATION

Services and Service Groups Overview | 742
About the Services Page | 743
Creating Services and Service Groups | 744

Application Signatures Overview

Juniper Networks regularly updates the predefined application signature database, making it available to subscribers on the Juniper Networks website. This database includes signature definitions of known application objects that can be used to identify applications for tracking, firewall policies, and quality-of-service prioritization.

Use the **Application Signatures** page to get an overall, high-level view of your application signature settings. You can filter and sort this information to get a better understanding of what you want to configure.

RELATED DOCUMENTATION

About the Application Signatures Page | 754
Adding Application Signatures | 757
Editing, Cloning, and Deleting Application Signatures | 762
Adding Application Signature Groups | 764
Editing, Cloning, and Deleting Application Signature Groups | 765
Signature Database Overview | 346

About the Application Signatures Page

To access this page, select **Configuration** > **Shared Objects** > **Application Signatures**.

Use the **Application Signatures** page to view application signatures that are already downloaded and to create, modify, clone, and delete custom application signature groups. The **Application Signatures** page displays the name, object type, category and subcategory, risk associated with, and characteristics of the signature. You can create custom application signature groups with a set of similar signatures for consistent reuse when defining policies.

Tasks You Can Perform

You can perform the following tasks from this page:

- Add an application signature. See "Adding Application Signatures" on page 757.
- Modify, clone, or delete an application signature. See "Editing, Cloning, and Deleting Application Signatures" on page 762.
- Add an application signature group. See "Adding Application Signature Groups" on page 764.
- Modify, clone, or delete an application signature group. See "Editing, Cloning, and Deleting Application Signature Groups" on page 765.
- View the configured parameters of an application signature or application signature group. Click the details icon that appears when you hover over the name of an image or click **More > Details**.
- Show or hide columns in the **Application Signatures**. Click the **Show Hide columns** icon in the top right corner of the page and select columns that you want to view on the page.
- Search for a specific application signature or application signature group. Click the Search icon in the top right corner of the page to search for an application signature or application signature group.
- Filter the application signature information based on select criteria. To do this, select the filter icon at the top right-hand corner of the table. The columns in the grid change to accept filter options. Select the filter options; the table displays only the data that fits the filtering criteria.

Field Descriptions

Table 254 on page 754 provides guidelines on using the fields on the **Application Signatures** page.

Table 254: Fields on the Application Signatures Page

Field	Description
Name	Name of the application signature or application signature group.

Table 254: Fields on the Application Signatures Page *(continued)*

Field	Description
Object Type	Signature type—either application signature or application signature group.
Category	UTM category of the application signature. For example, the value of **Category** can be Messaging, Web, Infrastructure, Remote-Access, Multimedia, and so on.
Subcategory	UTM subcategory of the application signature. For example, the value of **Subcategory** can be Wiki, File-Sharing, Multimedia, Social-Networking, News, and so on.
Risk	Level of risk associated with the application signature. For example, the value of **Risk** can be Low, High, unsafe, and so on.
Characteristic	One or more characteristics of the application signature.
Predefined or Custom	A list of predefined application signatures and application signature groups, and a list of custom application signature groups that you created.

RELATED DOCUMENTATION

- Application Signatures Overview | **753**
- Adding Application Signatures | **757**
- Editing, Cloning, and Deleting Application Signatures | **762**
- Adding Application Signature Groups | **764**
- Editing, Cloning, and Deleting Application Signature Groups | **765**
- Signature Database Overview | **346**
- About the Signature Database Page | **347**

Understanding Custom Application Signatures

Application identification supports user-defined custom application signatures to detect applications as they pass through the device. Custom application signatures are unique to your environment and are not included in the predefined application package. You use this custom application signature in SD-WAN policies and firewall policies to steer, and block traffic when a threat is detected.

Custom application signatures are required to:

- Control traffic particular to an environment.
- Bring visibility to unknown or unclassified applications.
- Identify Layer 7 applications or temporary applications, and to achieve further granularity of known applications.
- Perform QoS for your specific application.

CSO supports the following custom application signatures:

- ICMP-Based Mapping—The Internet Control Message Protocol (ICMP) mapping technique maps standard ICMP message types and optional codes to a unique application name. This mapping technique lets you differentiate between various types of ICMP messages.
- IP Address-Based Mapping—Layer 3 and Layer 4 address mapping defines an application by the IP address and optional port range of the traffic.

  To ensure adequate security, use address mapping when the configuration of your private network predicts application traffic to or from trusted servers. Address mapping provides efficiency and accuracy in handling traffic from a known application.

  With Layer 3 and Layer 4 address-based custom applications, you can match the IP address and port range to destination IP address and port range. When IP address and port range are configured, they must match the destination tuples (IP address and port range) of the packet.

  For example, consider a Session Initiation Protocol (SIP) server that initiates sessions from its known port 5060. Because all traffic from this IP address and port is generated by only the SIP application, the SIP application can be mapped to an IP address of the server and port 5060 for application identification. In this way, all traffic with this IP address and port is identified as SIP application traffic.
- IP Protocol-Based Mapping—Standard IP protocol numbers can map an application to IP traffic. As with address mapping, to ensure adequate security, use IP protocol mapping only in your private network for trusted servers.
- Layer 7-Based Signatures—Layer 7 custom signatures define an application running over TCP or UDP or Layer 7 applications. Layer 7-based custom application signatures are required for the identification of multiple applications running on the same Layer 7 protocols. For example, applications such as Facebook and Yahoo Messenger can run over HTTP, but there is a need to identify them as two different applications running on the same Layer 7 protocol. The custom signature is cacheable for Layer 7 signatures only. You can create multiple signatures and each signature can contain multiple members (maximum 15 members).

  Layer 7-based custom application signatures detect applications based on the patterns in HTTP contexts. However, some HTTP sessions are encrypted in SSL, also called Transport Layer Security (TLS). Application identification can extract the server name information or the server certification from the TLS or SSL sessions. It can also detect patterns in TCP or UDP payload in Layer 7 applications.

RELATED DOCUMENTATION

Adding Application Signatures | 757
Editing, Cloning, and Deleting Application Signatures | 762

Adding Application Signatures

You can add custom application signatures for applications that are not included in Juniper Networks predefined application database. When you add custom application signatures, make sure that your application signatures are unique, by providing a unique and relevant name.

You can add custom application signatures by specifying a name, protocol, port number where the application runs, and match criteria.

To create a custom application signature:

1. Select **Configuration > Shared Objects > Application Signatures**.
2. Click **Create > Signature**.
3. Complete the configuration according to the guidelines provided in Table 255 on page 757.
4. Click **OK** to save the changes. If you want to discard your changes, click **Cancel** instead.

A new application signature with your configurations is created. You use this application signature while creating SD-WAN policy intents.

Table 255 on page 757 provides guidelines on using the fields on the **Create Application Signature** page.

Table 255: Fields on the Create Application Signature Page

Field	Description
Name	Enter a unique name that is a string of alphanumeric characters, colons, periods, dashes, and underscores. No spaces are allowed and the maximum length is 63 characters.
Description	Enter a description for the application signature.
Signature Order and Priority	

Table 255: Fields on the Create Application Signature Page *(continued)*

Field	Description
Order	Enter the order for the custom application signature. A lower order value has higher priority. This option is used when multiple custom application signatures of the same type match the same traffic. However, you cannot use this option to prioritize among different type of applications such as TCP stream-based applications against TCP port-based applications or IP address-based applications against port-based applications. Range is 1-50000.
Priority	Specify the application signature priority (high or low) over other application signatures.
Signature Classification	
Category	Enter the category of the application signature. For example, Messaging, Web, Infrastructure, Remote-Access, Multimedia, and so on.
Sub Category	Enter the subcategory of the application signature. For example, Wiki, File-Sharing, Multimedia, Social-Networking, News, and so on.
Risk	Select the level of risk associated with the application signature. For example, low, moderate, high, critical, unsafe, and so on.
Characteristics	Enter one or more characteristics of the application signature. For example, supports file transfer, loss of productivity, and so on.
Application Criteria	Enable one or more application matching criteria: • ICMP Mapping • IP Protocol Mapping • Address Mapping • L7 Signature
ICMP Mapping	Click the toggle button to specify the Internet Control Message Protocol (ICMP) value for an application while configuring custom application signatures for application identification. The ICMP mapping technique maps standard ICMP message types and optional codes to a unique application name. The ICMP code and type provide additional specification, for packet matching in an application definition.
ICMP Type	Enter an ICMP value for the application. The ICMP mapping technique maps standard ICMP message types and optional codes to a unique application name. Range is 0-254.

Table 255: Fields on the Create Application Signature Page *(continued)*

Field	Description
ICMP Code	Enter an ICMP code for the application. The field provides further information (such as RFCs) about the ICMP type field. Range is 0-254.
IP Protocol Mapping	Click the toggle button to specify the IP protocol value for an application. This parameter is used to identify an application based on it's IP protocol value and is intended only for IP traffic. To ensure adequate security, use IP protocol mapping only in your private network for trusted servers.
IP Protocol	Enter an IP Protocol number for the application. Standard IP protocol numbers map an application to IP traffic. To ensure adequate security, use IP protocol mapping only in your private network for trusted servers. Range is 0-254. You can find a complete list of industry standard protocol numbers at the IANA website. **NOTE:** You cannot use IP protocol numbers 1(ICMP), 6(TCP) and 17(UDP) for custom application signature creation. Instead, we recommend you to use L7 signature policies for these protocols.
Address Mapping	Click the toggle button to specify address mapping information. Layer 3 and Layer 4 address mapping defines an application by matching the destination IP address or port range (optional) of the traffic. Use the address mapping option to configure custom applications signatures when the configuration of your private network predicts application traffic to or from trusted servers. Address mapping provides efficiency and accuracy while handling traffic from a known application. For more information, see Table 256 on page 760. **NOTE:** • You must specify either IP address or TCP/UDP port range for address mapping. • If both IP address and TCP/UDP ports are configured, both should match destination tuples (IP address and port range) of the packet.
L7 Signature	Click the toggle button to specify the Layer 7-based custom application signatures that are required to identify the multiple applications running on the same L7 protocols. Configure a custom signature based on L7 applications. You create Layer 7-based custom application signatures for the identification of multiple applications running on the same L7 protocols. For example, applications such as Facebook and Yahoo Messenger can both run over HTTP, but there is a need to identify them as two different applications running on the same Layer 7 protocol. For more information, see Table 257 on page 760.

Table 255: Fields on the Create Application Signature Page *(continued)*

Field	Description
Cacheable	Click the toggle button to enable caching of application identification results on the device. Enable this option to **True** only when L7 signatures are configured alone in a custom signature. This option is not supported for address-based, IP protocol-based, and ICMP-based custom application signatures.

Table 256: Fields on the Add IP Address Mapping Page

Field	Description
Name	Enter a unique string of alphanumeric characters, colons, periods, dashes, and underscores. No spaces are allowed; maximum length is 63 characters.
IP Address	Enter the destination IPv4 or IPv6 address of the application.
CIDR	Enter a CIDR value for the IP Address that you assign to the application. Range for IPv4 address is 1-32. Range for IPv6 address is 1-128.
TCP Port range	(Optional) Enter space-separated list of ports or port ranges to match a TCP destination port for Layer 3 and Layer 4 address-based custom applications. The range is 0-65535. Example: 80-82 443.
UDP port range	(Optional) Enter space-separated list of ports or port ranges ranges to match an UDP destination port for Layer 3 and Layer 4 address-based custom applications. The range is 0-65535. Example: 160-162 260.

Table 257: Fields on the Add Signature Page

Field	Description
Over Protocol	Displays the signature to match the application protocol. Example: HTTP.
Signature Name	Enter a unique name that is a string of alphanumeric characters, colons, periods, dashes, and underscores. No spaces are allowed and the maximum length is 63 characters.

Table 257: Fields on the Add Signature Page *(continued)*

Field	Description
Port Range	Enter the port range for the application. Range is 0-65535 Example: 80-82,443
Add Members	Click the plus icon (+) to add the member details.
Member No.	Enter the member name for a custom application signature. Custom signatures can contain multiple members that define attributes for an application. (The supported member name range is m01—m15.)

Context	Select the service-specific context.
	• For L7 Signatures over HTTP, select any of the following context:
	  ◦ http-get-url-parsed-param-parsed
	  ◦ http-header-content-type
	  ◦ http-header-cookie
	  ◦ http-header-host
	  ◦ http-header-user-agent
	  ◦ http-post-url-parsed-param-parsed
	  ◦ http-post-variable-parsed
	  ◦ http-url-parsed
	  ◦ http-url-parsed-param-parsed
	• For L7 Signatures over SSL, select the service-specific context as **ssl-server-name**.
	• For L7 Signatures over TCP, select the service-specific context as **stream**.
	• For L7 Signatures over UDP, select the service-specific context as **stream**.
	For possible combinations of context and direction for L7 application creation, refer context (Application Identification).

Direction	Select the direction of the packet flow to which the signature must be matched. • any—The direction of packet flow can either be from client-side to server-side or from server-side to client-side. • client-to-server—The direction of packet flow is from client-side to server-side. • server-to-client—The direction of packet flow is from server-side to client-side.
Pattern	Enter the deterministic finite automaton (DFA) pattern matched on the context. The DFA pattern specifies the pattern to be matched for the signature. Maximum length is 128.

RELATED DOCUMENTATION

- Understanding Custom Application Signatures | 755
- Editing, Cloning, and Deleting Application Signatures | 762
- Creating SD-WAN Policy Intents | 505
- Adding SLA-Based Steering Profiles | 517
- Adding Path-Based Steering Profiles | 528

Editing, Cloning, and Deleting Application Signatures

IN THIS SECTION

- Editing Application Signatures | 762
- Cloning Application Signatures | 763
- Deleting Application Signatures | 763

You can edit, clone, and delete application signatures from the **Application Signatures** page.

Editing Application Signatures

To modify the parameters configured for a cloned user-created (custom) application signature:

1. Select **Configuration > Shared Objects > Application Signatures**.

   The **Application Signatures** page appears.

2. Select the application signature that you want to edit, and then click on the edit icon (pencil), on the top right corner of the table, or right-click and select **Edit Application Signature**.

   The **Edit Application Signature** page appears, showing the same options as those displayed when you create a new application signature.

3. Modify the parameters according to the guidelines provided in *Adding Application Signatures*.

4. Click **Save** to save the changes. If you want to discard your changes, click **Cancel** instead.

The modified application signature appears on the **Application Signatures** page.

Cloning Application Signatures

You can clone a custom application signature when you want to reuse an existing application signature, but with a few minor changes. This way, you can save time recreating the application signature from scratch.

To clone a custom application signature:

1. Select **Configuration > Shared Objects > Application Signatures**.

   The **Application Signatures** page appears.

2. Select the application signature that you want to clone, and then select **More > Clone**, or right-click the application signature and then select **Clone**.

   The **Clone** page appears with editable fields.

3. Modify the fields as required. Refer to the guidelines provided in *Adding Application Signatures*.

4. Click **OK** to save the changes. If you want to discard your changes, click **Cancel** instead.

The cloned application signature is displayed on the **Application Signatures** page.

Deleting Application Signatures

To delete a cloned user-created (custom) application signature:

1. Select **Configuration > Shared Objects > Application Signatures**.

   The **Application Signatures** page appears.

2. Select the application signature you want to delete and then click the delete icon.

   An alert message appearsto verify that you want to delete the selected application signature.

3. Click **Yes** to delete the selected application signature. If you do not want to delete, click **Cancel** instead.

The deleted application signature is removed from the Application Signatures page.

RELATED DOCUMENTATION

Adding Application Signatures | 757

Editing, Cloning, and Deleting Application Signatures | 762

Adding Application Signature Groups

Application identification supports custom application signatures to detect applications as they pass through the device. When you add custom signature groups, make sure that your signature groups are unique, by providing a unique and relevant name.

To add an application signature group:

1. Select **Configure > Shared Objects > Application Signatures**.
2. Click the add icon (+).
3. Complete the configuration according to the guidelines provided in Table 258 on page 764.
4. Click **OK** to save the changes. If you want to discard your changes, click **Cancel** instead.

A new application signature group with your configurations is created. You can use this application signature group in firewall, NAT, and SD-WAN policies.

Table 258 on page 764 provides guidelines on using the fields on the **Create Application Signature Group** page.

Table 258: Fields on the Create Application Signature Group Page

Field	Description
Name	Enter a unique name that is a string of alphanumeric characters, colons, periods, dashes, and underscores. No spaces are allowed and the maximum length is 63 characters.
Group Members	Click the add icon (+) to add signatures to your application group. On the **Add Application Signatures** page, select the check boxes next to the signatures you want to add to the group.

RELATED DOCUMENTATION

- Application Signatures Overview | 753
- About the Application Signatures Page | 754
- Editing, Cloning, and Deleting Application Signature Groups | 765
- Signature Database Overview | 346
- About the Signature Database Page | 347

Editing, Cloning, and Deleting Application Signature Groups

IN THIS SECTION

- Editing Application Signature Groups | 765
- Cloning Application Signature Groups | 765
- Deleting Application Signature Groups | 766

You can edit, clone, and delete application signature groups from the **Application Signatures** page.

Editing Application Signature Groups

To modify the parameters configured for an application signature group:

1. Select **Configuration > Shared Objects > Application Signatures**.

   The **Application Signatures** page appears.

2. Select the application signature group that you want to edit, and then select **More > Edit**, or click on the edit icon (pencil symbol), on the top right corner of the table, or right-click and select **Edit**.

   The **Edit** page appears, showing the same options as those displayed when you create a new application signature group.

3. Modify the parameters according to the guidelines provided in "Adding Application Signature Groups" on page 764.

4. Click **Save** to save the changes. If you want to discard your changes, click **Cancel** instead.

The modified application signature group appears in the **Application Signatures** page.

Cloning Application Signature Groups

You can clone an application signature group when you want to reuse an existing application signature group, but with a few minor changes. This way, you can save time recreating the application signature group from the start.

To clone an application signature group:

1. Select **Configuration > Shared Objects > Application Signatures**.

The **Application Signatures** page appears.

2. Right-click the application signature group that you want to clone and then select **Clone**, or select **More > Clone**.

   The **Clone** page appears with editable fields.

3. Click **OK** to save the changes. If you want to discard your changes, click **Cancel** instead.

The cloned application signature group is displayed on the **Application Signatures** page.

Deleting Application Signature Groups

To delete an application signature group:

1. Select **Configuration > Shared Objects > Application Signatures**.

   The **Application Signatures** page appears.

2. Select the application signature group you want to delete and then click the delete icon **(X)**.

   An alert message appears, verifying that you want to delete the selected item.

3. Click **Yes** to delete the selected application signature group. If you do not want to delete, click **Cancel** instead.

RELATED DOCUMENTATION

- Application Signatures Overview | 753
- About the Application Signatures Page | 754
- Adding Application Signature Groups | 764
- Signature Database Overview | 346

About the Departments Page

To access this page, click **Configuration > Network Services > Shared Objects > Departments**.

You can use the Departments page to add, view, or delete departments. A department is a grouping of LAN segments within a site. You use departments to apply specific policies to LAN segments that are members of a department.

Tasks You Can Perform

You can perform the following tasks from this page:

- Add a Department. See "Adding a Department" on page 767.
- Delete a department. See "Deleting a Department" on page 768.

Field Descriptions

Table 259 on page 767 shows the descriptions of the fields on the **Departments** page.

Table 259: Fields on the Departments Page

Field	Description
Name	Displays the name of the department.
Site/LAN Segments	Displays the sites and LAN segments that are associated with the department.
VPN	Displays the VPN to which the department is assigned.
Data Center	Displays whether the department is a data center department or not.
Description	Displays the description of the department.
UUID	Displays the Universally Unique Identifier (UUID) of the department.
Network UUID	Displays an internal UUID used by Contrail Service Orchestration (CSO).

RELATED DOCUMENTATION

Adding a Department | 767
Deleting a Department | 768

Adding a Department

You can add departments from the **Configuration > Shared Objects > Departments** page.

To add a department:

1. Click the add icon (+) on the **Departments** page.

   The **Add Department** page appears.

2. Complete the configuration settings according to the guidelines provided inTable 260 on page 768.

3. Click **OK**.

   You are taken to the **Departments** page where the department that you added is displayed.

Table 260: Fields on the Add Department Page

Field	Description
Name	Enter a unique name for the department, which can contain alphanumeric characters and some special characters (. -). No spaces are allowed and the maximum length is 15 characters.
Description	Enter a description of the department.
VPN	The default VPN to which the department is assigned is displayed. **NOTE:** This field is displayed only if network segmentation is disabled for the tenant.
Data Center Department	Select whether the department is a data center department (True) or not (False). A data center department can be attached (by using LAN segments) only to enterprise hubs.

RELATED DOCUMENTATION

About the Departments Page | 766
Deleting a Department | 768

Deleting a Department

You can delete departments by clicking the delete icon (**X**) on the **Departments** page. You can delete only one department at a time. You cannot delete a department that is associated with one or more LAN segments. Before you delete the department, you must reassign the LAN segments assigned to that department.

To delete a department:

1. Click **Configuration > Shared Objects > Departments**.

   The Departments page appears.

2. Select the department that you want to delete.

3. Click the delete icon (**X**).

   The Delete Department page appears.

4. Click **OK** to confirm the deletion.

   The department is deleted and you are returned to the Departments page

RELATED DOCUMENTATION

About the Departments Page | **766**
Adding a Department | **767**

About the MAC Addresses Page

To access this page, click **Configuration > Network Services > Shared Objects > MAC**.

You can configure firewall filter terms with media access control (MAC) address as source and destination endpoints to permit or block packets to a port.

Use the MAC Addresses page to add, view, or delete MAC addresses.

Tasks You Can Perform

You can perform the following tasks from this page:

- Add a MAC Address. See “Add a MAC Address Endpoint” on page 770.
- Edit or delete a MAC Address. See “Edit or Delete MAC Address Endpoint” on page 771.

Field Descriptions

Table 261 on page 770 shows the descriptions of the fields on the **MAC Addresses** page.

Table 261: Fields on the MAC Addresses Page

Field	Description
MAC Address	Displays the MAC address of a Layer 2 packet.
Description	Displays the description of the MAC address.
UUID	Displays the Universally Unique Identifier (UUID) of the MAC address.

RELATED DOCUMENTATION

Add a MAC Address Endpoint | 770

Edit or Delete MAC Address Endpoint | 771

Add a MAC Address Endpoint

Use the MAC address page to add a MAC address, which you can specify as a source endpoint or destination endpoint in firewall filter terms.

To add a MAC address:

1. Select **Configuration > Shared Objects > MAC**.

   The MAC address page appears.

2. Click the add icon (+) to add a new MAC address.

   The Add MAC Endpoint page appears.

3. Complete the configuration according to the guidelines provided in Table 262 on page 771.

   NOTE: Fields marked with * are mandatory.

4. Click **OK** to save the changes. If you want to discard your changes, click **Cancel** instead.

   The MAC address is created. You are returned to the MAC Addresses page where a confirmation message is displayed.

Table 262 on page 771 provides guidelines on using the fields on the **Add MAC Endpoint** page.

Table 262: Fields on the Add MAC Endpoint Page

Field	Description
Name	Enter a valid MAC address. MAC address must be 12 digit hexadecimal number. Format: ([0-9a-fA-F][0-9a-fA-F][:|-|.]){5}([0-9a-fA-F][0-9a-fA-F]) Example: 00:11:22:33:44:55
Description	Enter a description for the MAC address.

WHAT'S NEXT

After adding MAC address, you can specify the MAC address as source endpoint or destination endpoint in firewall filter terms, see Add Terms to Firewall Filters | 710.

RELATED DOCUMENTATION

About the MAC Addresses Page | 769
Edit or Delete MAC Address Endpoint | 771

Edit or Delete MAC Address Endpoint

IN THIS SECTION

Edit MAC Address | 772
Delete MAC Address | 772

You can edit or delete MAC Addresses from the **MAC Addresses** page.

Edit MAC Address

> **NOTE:** You cannot modify the name of the MAC address.

To modify the parameters configured for a MAC address:

1. Select **Configuration > Shared Objects > MAC**.

   The **MAC Address** page appears.

2. Select the MAC address that you want to edit, and then click on the edit icon (pencil), on the top right corner of the table.

   The **Edit MAC Endpoint** page appears, showing the same options as those displayed when you create a new MAC address.

3. Modify the description for the MAC Endpoint.

4. Click **OK** to save the changes. If you want to discard your changes, click **Cancel** instead.

The modified MAC endpoint description appears on the **MAC Addresses** page.

Delete MAC Address

To delete a MAC address:

1. Select **Configuration > Shared Objects > MAC**.

   The **MAC Addresses** page appears.

2. Select the MAC address that you want to delete and then click the delete icon.

   An alert message appears to verify that you want to delete the selected MAC address.

3. Click **Yes** to delete the selected MAC address. If you do not want to delete, click **No** instead.

The deleted MAC address is removed from the MAC Addresses page.

RELATED DOCUMENTATION

About the MAC Addresses Page | 769

Add a MAC Address Endpoint | 770

About the Protocols Page

To access this page, click **Configuration > Network Services > Shared Objects > Protocols**.

You can configure firewall filter terms with protocols (name or value) as source and destination endpoints to permit or block traffic to a specific port.

Use the Protocols page to add, view, or delete protocols.

Tasks You Can Perform

You can perform the following tasks from this page:

- Add a protocol. See "Add a Protocol Endpoint" on page 774.
- Edit or delete a protocol. See "Edit or Delete Protocol Endpoint" on page 775.

Field Descriptions

Table 263 on page 773 shows the descriptions of the fields on the **Protocols** page.

Table 263: Fields on the Protocols Page

Field	Description
Protocol	Displays the IPv4 protocol (name or value) for the port.
Description	Displays the description of the protocol.
UUID	Displays the Universally Unique Identifier (UUID) of the protocol.

RELATED DOCUMENTATION

Add a Protocol Endpoint | 774
Edit or Delete Protocol Endpoint | 775

Add a Protocol Endpoint

Use the Protocols page to add a new protocol, which you can specify as a source endpoint or destination endpoint in firewall filter terms.

To add a protocol:

1. Select **Configuration > Shared Objects > Protocols**.

   The Protocols page appears.

2. Click the add icon (+) to add a new protocol.

   The Add Protocol Endpoint page appears.

3. Complete the configuration according to the guidelines provided in Table 264 on page 774.

   NOTE: Fields marked with * are mandatory.

4. Click **OK** to save the changes. If you want to discard your changes, click **Cancel** instead.

   The new protocol is created. You are returned to the Protocols page where a confirmation message is displayed.

Table 264 on page 774 provides guidelines on using the fields on the **Add Protocol Endpoint** page.

Table 264: Fields on the Add Protocol Endpoint Page

Field	Description
Protocol Number	Enter a valid protocol number. Range is 0-255. Example: egp (8), esp (50), gre (47), icmp (1), igmp (2), ipip (4), ospf (89), pim (103), rsvp (46), tcp (6), udp (17)
Description	Enter a description for the protocol.

WHAT'S NEXT

After adding protocol, you can specify the protocol as source endpoint or destination endpoint in firewall filter terms, see Add Terms to Firewall Filters | 710.

RELATED DOCUMENTATION

About the Protocols Page | 773

Edit or Delete Protocol Endpoint | 775

Edit or Delete Protocol Endpoint

IN THIS SECTION

Edit Protocols | 775

Delete Protocols | 776

You can edit or delete protocols from the **Protocols** page.

Edit Protocols

NOTE: You cannot modify the protocol number.

To modify the parameters configured for a protocol:

1. Select **Configuration > Shared Objects > Protocols**.

   The **Protocols** page appears.

2. Select the protocols that you want to edit, and then click on the edit icon (pencil), on the top right corner of the table.

   The **Edit Protocol Endpoint** page appears, showing the same options as those displayed when you create a new protocol.

3. Modify the description for the protocol.

4. Click **OK** to save the changes. If you want to discard your changes, click **Cancel** instead.

The modified protocol endpoint description appears on the **Protocols** page.

Delete Protocols

To delete a protocol:

1. Select **Configuration > Shared Objects > Protocols**.

   The **Protocols** page appears.

2. Select the protocols that you want to delete and then click the delete icon.

   An alert message appears to verify that you want to delete the selected protocols.

3. Click **Yes** to delete the selected protocols. If you do not want to delete, click **No** instead.

The deleted protocol is removed from the Protocols page.

RELATED DOCUMENTATION

About the Protocols Page | 773
Add a Protocol Endpoint | 774

About the Ports Page

To access this page, click **Configuration > Network Services > Shared Objects > Ports**.

You can configure firewall filter terms with ports (name or value) as source and destination endpoints to permit or block traffic to a specific port.

Use the Ports page to add, view, or delete ports.

Tasks You Can Perform

You can perform the following tasks from this page:

- Add a port. See “Add a Port Endpoint” on page 777.

- Edit or delete a port. See "Edit or Delete Port Endpoint" on page 779.

Field Descriptions

Table 265 on page 777 shows the descriptions of the fields on the **Ports** page.

Table 265: Fields on the Ports Page

Field	Description
Port	Displays the port name or value.
Description	Displays the description of the port.
UUID	Displays the Universally Unique Identifier (UUID) of the port.

RELATED DOCUMENTATION

Add a Port Endpoint | 777
Edit or Delete Port Endpoint | 779

Add a Port Endpoint

Use the Ports page to add a port, which you can specify as a source endpoint or destination endpoint in firewall filter terms.

To add a port:

1. Select **Configuration > Shared Objects > Ports**.

   The Ports page appears.

2. Click the add icon (+) to add a new port.

   The Add Port Endpoint page appears.

3. Complete the configuration according to the guidelines provided in Table 266 on page 778.

NOTE: Fields marked with * are mandatory.

4. Click **OK** to save the changes. If you want to discard your changes, click **Cancel** instead.

   The Port is created. You are returned to the Ports page where a confirmation message is displayed.

Table 266 on page 778 provides guidelines on using the fields on the **Add Port Endpoint** page.

Table 266: Fields on the Add Port Endpoint Page

Field	Description
Name	Enter a TCP or UDP port field. Range is 0-65535. Typically, you specify the port match condition in conjunction with the protocol match condition to determine which protocol is used on the port. For number, you can specify one of the following text synonyms (the port numbers are also listed): Example: afs (1483), bgp (179), biff (512), bootpc (68), bootps (67), cmd (514), cvspserver (2401), dhcp (67), domain (53), eklogin (2105), ekshell (2106), exec (512), finger (79), ftp (21), ftp-data (20), http (80), https (443), ident (113), imap (143), kerberos-sec (88), klogin (543), kpasswd (761), krb-prop (754), krbupdate (760), kshell (544), ldap (389), login (513), mobileip-agent (434), mobilip-mn (435), msdp (639), netbios-dgm (138), netbios-ns (137), netbios-ssn (139), nfsd (2049), nntp (119), ntalk (518), ntp (123), pop3 (110), pptp (1723), printer (515), radacct (1813),radius (1812), rip (520), rkinit (2108), smtp (25), snmp (161), snmptrap (162), snpp (444), socks (1080), ssh (22), sunrpc (111), syslog (514), tacacs-ds (65), talk (517), telnet (23), tftp (69), timed (525), who (513), xdmcp (177), zephyr-clt (2103), zephyr-hm (2104)
Description	Enter a description for the port.

WHAT'S NEXT

After adding port, you can specify the port as source endpoint or destination endpoint in firewall filter terms, see Add Terms to Firewall Filters | 710.

RELATED DOCUMENTATION

Add a Port Endpoint | 777

Edit or Delete Port Endpoint | 779

Edit or Delete Port Endpoint

IN THIS SECTION

Edit Ports | 779
Delete Ports | 780

You can edit or delete ports from the **Ports** page.

Edit Ports

> **NOTE:** You cannot modify the name of the port.

To modify the parameters configured for a port:

1. Select **Configuration > Shared Objects > Ports**.

   The **Ports** page appears.

2. Select the port that you want to edit, and then click on the edit icon (pencil), on the top right corner of the table.

   The **Edit Port Endpoint** page appears, showing the same options as those displayed when you create a new port.

3. Modify the description for the port.

4. Click **OK** to save the changes. If you want to discard your changes, click **Cancel** instead.

The modified port endpoint description appears on the **Ports** page.

Delete Ports

To delete a port:

1. Select **Configuration > Shared Objects > Ports**.

   The **Ports** page appears.

2. Select the port that you want to delete and then click the delete icon.

   An alert message appears to verify that you want to delete the selected port.

3. Click **Yes** to delete the selected port. If you do not want to delete, click **No** instead.

The deleted port is removed from the Ports page.

RELATED DOCUMENTATION

Add a Port Endpoint | 777

Edit or Delete Port Endpoint | 779

PART 6

Monitoring Jobs and Audit Logs

Managing Jobs | **782**

Managing Audit Logs | **788**

CHAPTER 25

Managing Jobs

IN THIS CHAPTER

- About the Jobs Page | 782
- Editing and Deleting Scheduled Jobs | 784
- Viewing Job Details | 786
- Retrying a Failed Job on Devices | 787

About the Jobs Page

To access this page, click **Monitor > Jobs**.

A job is an action that is performed on any object that is managed by CSO, such as a device, tenant, site, or user. You can monitor the status of jobs that have run or are scheduled to run in CSO. You can run the job immediately or schedule it for a later date and time. You can view the status of the job whether it is completed or failed. You can retry tssm.ztp type jobs that are failed.

Use this page to view the list of all jobs and the jobs that are scheduled to be executed. You can view general information about the jobs and the overall progress and status of the jobs. You can also edit and delete scheduled jobs.

Tasks You Can Perform

You can perform the following tasks from this page:

- View details about a job. See “Viewing Job Details” on page 786.
- Retry a job. See “Retrying a Failed Job on Devices” on page 787.
- Edit and delete scheduled jobs. See “Editing and Deleting Scheduled Jobs” on page 784.

Field Descriptions

Table 267 on page 783 provides guidelines on using the fields on the Jobs page.

Table 267: Fields on the Jobs Page

Field	Description
Job Name	View the name of the job. CSO automatically generates the job name. Example: MSEC_DOWNLOAD_IPS/APPLICATION_SIGNATURES_08_Jul_17_124229_024
Resource Name	View the resource name of the job. Example: Download IPS/Application Signatures
Status	View the status of the job to know whether the job succeeded, failed, or in progress. Example: Success
Owner	View the name of the owner who created the job. Example: cspadmin
Number of Tasks	View the number of tasks associated with the job. Example: 2 For example, the tasks **site.ucpe-32** and **customer.sdwan** are associated with the job.
Job ID	When a job is initiated from a object in CSO, CSO assigns a unique ID to that job, which serves to identify the job (along with the job type) on the Jobs page. The following is a list of some of the job types supported in CSO: • Configure Sites • Download Signature • Create Sites • Remove Site
Start Date	View the start date and time of a task associated with the job.
End State	View the end date and time of a task associated with the job.

Field Descriptions

Table 268 on page 784 provides guidelines on using the fields on the Scheduled Jobs page.

Table 268: Fields on the Scheduled Jobs Page

Field	Description
Schedule ID	View the unique ID of the scheduled job. The value is generated by the database when a new schedule record is inserted into the database. Example: 48
Name	View the unique name of the scheduled job. Example: Tenant Delete_csp.tssm_remove_site_e340354716ae43859fad5ba15669eee2
Status	View the status of the last triggered job. The default status is scheduled.
Record Type	View the job type. Example: tssm onboard tenant
Owner	View the name of the owner who scheduled the job. Example: cspadmin
Next Run Time	View the time when the job is scheduled to run next.

RELATED DOCUMENTATION

Editing and Deleting Scheduled Jobs | 784

Editing and Deleting Scheduled Jobs

IN THIS SECTION

Editing Scheduled Jobs | 785
Deleting Scheduled Jobs | 785

You can edit and delete scheduled jobs. This topic contains the following sections:

Editing Scheduled Jobs

You can modify the date and time of deployment of scheduled jobs.

To modify a scheduled job:

1. Select **Monitor > Jobs > Scheduled Jobs**.

   The Jobs page displays all scheduled jobs.

2. Select the job that you want to reschedule the deployment, and click the edit icon.

   The Edit Schedule page appears. This page displays the option that you have selected initially.

3. Modify the deployment type.

   To execute the job immediately, select the **Run now** option.

   To reschedule the job for a later date and time, select the **Schedule at a later time** option and select the date and time of deployment.

4. Click **Save** to save the changes.

   A success message is displayed indicating that the scheduled job is modified.

Deleting Scheduled Jobs

You can delete one or more scheduled jobs.

To delete a scheduled job:

1. Select **Monitor > Jobs> Scheduled Jobs**.

   The Jobs page displays all scheduled jobs.

2. Select the job that you want to delete and then click the delete icon (X). You can select one or more jobs

   The Confirm Delete page appears.

3. Click **Yes** to confirm.

   A success message is displayed indicating that the scheduled job is deleted.

RELATED DOCUMENTATION

About the Jobs Page | 782
Viewing Job Details | 786

Viewing Job Details

You can use the Detail for *Job-Name* page to view all the parameters of a job. This page has the following two tabs:

- **Details**—Displays the overall progress of the job and lists general information about the job (for example, the Job ID, Request ID, Created By, and so on). For more information about the field description on this page, see "About the Jobs Page" on page 782.
- **Tasks**—Displays the number of tasks associated with the job. A green check mark (success) or a red cross mark (failed) is displayed next to each task indicating the status of the task. You can click the Detailed View icon to view the summary of the task.

To view details of a job:

- Right-click the job name that you want to see the detailed view for and select **Detail View**.
- Select the job and click **More > Detail View**.
- Alternatively, hover over the job name and click the Detailed View icon that appears before it.

The Detail for *Job-Name* page appears, showing the details of the job and the number of tasks associated with the job. Click **View Logs** to view the status of the jobs. See "About the Jobs Page" on page 782 for a description of each fields on this page.

RELATED DOCUMENTATION

About the Jobs Page | 782

Retrying a Failed Job on Devices

As a tenant user with the Job Retry capability, you can retry a failed job instead of redoing the tasks involved in the job, to save time.

> **NOTE:** Before you retry a failed job, identify the reason for the failure and then fix it, before retrying the job.
>
> For example, if the bootstrap process failed because the device could not establish an outbound SSH connection, you must fix the problem and ensure that the outbound SSH connection is established before you retry the bootstrap job.

You can retry only the following jobs that did not complete successfully on your devices:

- ZTP jobs
- Bootstrap jobs

To retry a job that was not successful:

1. Select **Monitor > Jobs**.

   The Jobs page appears.

2. Select the failed job that you want to retry.

3. Click the **Retry Job** button on the top-right corner of the page.

   A retry job is created and executed.

   If the job is successful, a confirmation message appears and the job status changes to **Success** on the Jobs page.

RELATED DOCUMENTATION

About the Jobs Page | 782

Editing and Deleting Scheduled Jobs | 784

CHAPTER 26

Managing Audit Logs

IN THIS CHAPTER

- Audit Logs Overview | 788
- About the Audit Logs Page | 789
- Viewing the Details of an Audit Log | 790
- Exporting Audit Logs | 793
- Purging Audit Logs (After Archiving or Without Archiving) | 794

Audit Logs Overview

An audit log is a record of a sequence of activities that have affected a specific operation or procedure. Audit logs are useful for tracing events and for maintaining historical data.

Audit logs contain information about tasks initiated by using the Contrail Service Orchestration (CSO) GUI or APIs. In addition to providing information about the resources that were accessed, audit log entries usually include details about user-initiated tasks, such as the name, role, and IP address of the user who initiated a task, the status of the task, and date and time of execution.

NOTE: Device-driven tasks (that is, tasks not initiated by the user) are not recorded in audit logs.

Administrators can use audit logs to review events. For example, administrators can identify the user accounts associated with an event, determine the chronological sequence of events. For audit log entries that have an associated job, you can click the hyperlinked job ID to go to the Jobs page, where you can view the details of the job.

RELATED DOCUMENTATION

About the Audit Logs Page | 789

Exporting Audit Logs | 793

Purging Audit Logs (After Archiving or Without Archiving) | 794

About the Audit Logs Page

To access this page, select **Administration** > **Audit Logs**.

Use the Audit Logs page to view tasks that you have initiated either by using the Contrail Service Orchestration (CSO) GUI or APIs. You can also export audit logs as a comma-separated values (CSV) file and purge audit logs after archiving them or without archiving them.

Tasks You Can Perform

You can perform the following tasks from this page:

- View the details of various user-initiated tasks by selecting **More** > **Details**. You can also mouse over the audit log and click on the **Detailed View** icon. See "Viewing the Details of an Audit Log" on page 790.
- Export audit logs as a CSV file—See "Exporting Audit Logs" on page 793.
- Purge audit logs—See "Purging Audit Logs (After Archiving or Without Archiving)" on page 794.
- Search for audit logs by using keywords—Click the search icon and enter the search term in the text box and press Enter. The search results are displayed on the same page.
- Sort and filter audit logs:

  NOTE: Sorting and filtering is applicable only to some fields.

  - Click a column name to sort the audit logs based on the column name.
  - Click the filter icon and select whether you want to show or hide column filters or apply a quick filter. For example, you can use audit log filtering to track user accounts that were added on a specific date, track configuration changes across a particular type of device, view services that were provisioned on specific devices, monitor user login and logout activities over time, and so on.
- Show or hide columns—Click the **Show Hide Columns** icon at the top right corner of the page and select the columns that you displayed on the Audit Logs page.

Table 269 on page 789 provides description of the fields on the Audit Logs page.

Table 269: Fields on the Audit Logs Page

Field	Description
Username	Displays the username of the user who initiated the task.

Table 269: Fields on the Audit Logs Page *(continued)*

Field	Description
User IP	Displays the IP address of the client from which the user initiated the task. For tasks that do not have an associated user IP address, this field is blank.
Object Name	Displays the name of the object on which the task was initiated. An object can be a tenant, site, device, device image, template, and so on.
Task	Displays the name of the task that triggered the audit log. For example, tenant.create, device.create, site.configure, site.provision, tenant.update, and so on.
Description	Displays details about the task.
Status	Displays the status of the task that triggered the audit log: • Success—Job or task was completed successfully. • Failure—Job or task failed and was terminated. • Job Scheduled—Job is scheduled but has not yet started. • Recurring Job Scheduled—Recurring job is scheduled.
End Time	Displays the date and time at which the execution of the task was completed. This timestamp is stored in UTC time in the database, but is mapped to the local time zone of the client computer.
Job ID	For tasks that have associated jobs, displays the ID of the job associated with the task. You can click the job ID to go to the Jobs page, where you can view the status of the job.

RELATED DOCUMENTATION

About the Jobs Page | 782

Viewing the Details of an Audit Log

Use the Audit Log Details pane to view details of an audit log.

To view the details of an audit log:

1. Select **Administration** > **Audit Logs**.

The Audit Logs page appears displaying the audit logs.

2. Select the audit log for which you want to view details and click **More** > **Details**. Alternatively, you can mouse over the audit log, and click on the **Detailed View** icon.

   The Audit Log Details pane appears on the right side of the Audit Logs page. Table 270 on page 791 provides descriptions the fields on the Audit Log Details pane.

3. Click the close icon (X) to close the Audit Log Details pane.

   You are returned to the Audit Logs page.

Table 270 on page 791 provides descriptions the fields on the Audit Log Details pane.

Table 270: Fields on the Audit Log Details Pane

Field	Description
Details	
User	
Username	Displays the user who initiated the task.
User IP	Displays the IP address of the client from which the user initiated the task. For tasks that do not have an associated user IP address, this field is blank.
Task	
Task	Displays the name of the task that triggered the audit log. For example, tenant.create, device.create, site.configure, site.provision, tenant.update, and so on.
Status	Displays the status of the task that triggered the audit log: • Success—Job or task was completed successfully. • Failure—Job or task failed and was terminated. • Job Scheduled—Job is scheduled but has not yet started. • Recurring Job Scheduled—Recurring job is scheduled.
Description	Displays details about the task.
Affected Objects	

Table 270: Fields on the Audit Log Details Pane *(continued)*

Field	Description
Object Name	Displays the name of the affected object on which the task was initiated. An affected object can be a tenant, site, device, device image, template, and so on. Click the hyperlinked object name to view details of the object: **NOTE:** If the object is deleted or if you do not have permissions to view it, an error message is displayed.
Object UUID	Displays the Universally Unique Identifier (UUID) of the affected object.
Log Info	
Audit Log ID	Displays the automatically-generated unique ID of the audit log associated with the task.
Job ID	For tasks that have associated jobs, displays the ID of the job associated with the task. You can click the job ID to go to the Jobs page, where you can view the status of the job.
End Time	Displays the date and time at which the task completed execution. This timestamp is stored in UTC time in the database, but is mapped to the local time zone of the client computer.
Raw Audit Log	
Microservice	Displays the name of the microservice that initiated the task.
Raw Audit Log	Displays all the fields of the audit log that are stored in the database. The raw audit log typically contains additional details or parameters associated with the audit log.

RELATED DOCUMENTATION

Audit Logs Overview | 788

About the Audit Logs Page | 789

Exporting Audit Logs

You can export audit logs as a comma-separated values (CSV) file that can be opened or edited using an application such as Microsoft Excel. You can view and analyze the exported audit logs, as needed.

To export the audit logs:

1. Select **Administration** > **Audit Logs**.

   The Audit Logs page appears displaying the audit logs.

2. Click **Export**.

   The Export Audit Logs page appears.

3. Specify the time period for which you want to export the audit logs according to the guidelines provided in Table 271 on page 793.

   > **NOTE:** You can export audit logs for a maximum of 30 days prior to the current date and time. For example, if the current date is May 31, 2018, you can export the audit logs starting from May 1, 2018.

4. Click **OK** to export the audit logs.

   Depending on the settings of the browser that you are using, the CSV file containing the audit logs for the specified time period is either downloaded directly, or you are asked to open or save the file.

   You are returned to the Audit Logs page.

   After the file is downloaded, you can open the CSV file in an application such as Microsoft Excel and view and analyze the logs as required.

Table 271: Fields on the Export Audit Logs Page

Field	Description
Start Date and Time	Specify the date and time (in MM/DD/YYYY and HH:MM:SS 24-hour or AM/PM formats) from when the audit logs should be exported.
End Date and Time	Specify the date and time (in MM/DD/YYYY and HH:MM:SS 24-hour or AM/PM formats) up to when the audit logs should be exported.

RELATED DOCUMENTATION

Audit Logs Overview | 788

About the Audit Logs Page | 789

Viewing the Details of an Audit Log | 790

Purging Audit Logs (After Archiving or Without Archiving)

You can manage the volume of audit log data stored by purging log files from the CSO database without archiving them or by purging log files after archiving them. You can purge audit logs immediately or schedule the purging for a later date and schedule the purging on a recurring basis.

To purge audit logs after archiving or without archiving them:

1. Select **Administration** > **Audit Logs**.

   The Audit Logs page appears displaying the audit logs.

2. Click **Purge**.

   The Purge Audit Logs page appears.

3. Complete the configuration according to the guidelines provided in Table 272 on page 794.

   NOTE: Fields marked with an asterisk (*) are mandatory.

4. Click **OK**.

   You are returned to the Audit Logs page and one of the following operations occur:

   - If you triggered a purge of the audit logs without archiving, a job to purge the audit logs is created.
   - If you triggered a purge of the audit logs after archiving, a job is created to archive the audit logs and then purge the audit logs after archiving.

   After the audit logs are purged successfully, the Audit Logs page refreshes automatically and displays only the audit logs that were not purged.

Table 272: Purge Audit Logs Settings

Field	Description
Purge Options	

Table 272: Purge Audit Logs Settings *(continued)*

Field	Description
Purge Logs	Select one of the following options to purge audit logs: • Purge audit logs that were generated before a specified date and time—If you select this option, you must enter a date and time in the **Before** field. • Purge generated audit logs that are older than a specified number of days—If you select this option, you must specify the number of days in the **Older than** field.
Before	To purge audit logs before a specified date and time, enter the date (in MM/DD/YYYY format) and time (in HH:MM:SS 24-hour or AM/PM format) You specify the time in the local time zone of the client computer.
Older than	To purge generated audit logs older than a specified number of days, enter the number of days (from 1 through 90)
Archive Logs Before Purging	To archive audit logs *before* purging them, select this check box. By default, this check box is cleared, which means that audit logs are purged without archiving them. **CAUTION:** If you choose not to archive the audit logs before purging, the audit logs are deleted from the CSO database and cannot be recovered.
Archive Mode	Specify whether you want to archive the log files locally (**local**) or on a remote server (**remote**). If you archive the logs on a remote server, which is the default option, you must enter access and login details for the remote server. **NOTE:** • Archived log files are saved in a single file in compressed comma-separated values (CSV) format (extension .zip). • When you archive data locally, the archived log files are saved on the central microservices virtual machine (VM).
Username	Enter a valid username to access the remote server.
Password	Enter a valid password to access the remote server on which the audit logs will be archived.
Confirm Password	For confirmation, re-enter the password to access the remote server.
Remote Server IP Address	Enter the IPv4 address of the remote server; for example, 192.0.2.10.

Table 272: Purge Audit Logs Settings *(continued)*

Field	Description
Remote Server Path	Enter the directory path on the remote server on which to store the archived log files. The directory that you specify must already exist on the remote server.
Schedule Purge	
Type	Specify whether the audit logs should be purged immediately (**Run now**) or schedule the purge for later (**Schedule at a later time**). If you schedule the purge for later, enter the date (in MM/DD/YYYY format) and time (in HH:MM:SS 24-hour or AM/PM format) that you want the purge to occur. You specify the time in the local time zone of the client computer.
Recurrence	To specify whether the purge operation should occur on a recurring basis, select this check box. **NOTE:** This option is enabled only if you choose to archive and purge audit logs older than a specified number of days.
Repeat	Specify the periodicity of the recurrence. Currently, a weekly periodicity is the only option supported.
On	For purges that recur every week, specify one or more days on which you want the purge to recur.
Time	Enter the time (in HH:MM:SS 24-hour or AM/PM format) that you want the recurring purge to occur. By default, the purge recurs at 12.00 AM. You specify the time in the local time zone of the client computer.
Ends	Specify whether the recurring purge ends or not: • Select **Never** to continue (without an end date) the recurring purge operation at the specified recurrence interval. • Select **On** and enter the date (in MM/DD/YYYY format) and time (in HH:MM:SS 24-hour or AM/PM format) on which to stop the recurring purge operation. You specify the time in the local time zone of the client computer.

RELATED DOCUMENTATION

Audit Logs Overview | 788

About the Audit Logs Page | 789

PART 7

Monitoring Alarms, Events, and Threats

Monitoring Security Alerts and Alarms | **799**

Monitoring Security and Device Events | **813**

Monitoring SD-WAN Events | **845**

Monitoring Applications | **848**

Monitoring Threats | **866**

CHAPTER 27

Monitoring Security Alerts and Alarms

IN THIS CHAPTER

- About the Monitor Overview Page | 799
- Alerts Overview | 801
- About the Generated Alerts Page | 801
- About the Alert Definitions/Notifications Page | 803
- Managing Security Alerts Definitions | 804
- Creating Security Alert Definitions | 805
- Editing, Cloning, and Deleting Security Alert Definitions | 807
- About the Alarms Page | 809
- Enable E-mail Notifications for SD-LAN and SD-WAN Alarms | 810

About the Monitor Overview Page

To access this page, click **Monitor > Overview**.

You can use the Monitor Overview page to view information about the alarms and alerts for tenants, network services, connections, and sites on a geographical map. The network operator views the alarms and alerts, and then takes the necessary actions to resolve the issues.

You can also view the visual representation of the hub and link failure on this page.

- Hub Failure —The hub and the link connected to the hub appear in red color.
- Link Failure — The link connected to the hub appears in red color. However, the hub remains active and appears in green color.

Tasks You Can Perform

You can perform the following tasks from this page:

- View on-premise spoke site details.
- View on-premise hub site details.

- View cloud spoke sites.
- View provider hub sites.
- View multiple sites.

Field Descriptions

Table 273 on page 800 shows the descriptions of the fields on the Monitor Overview page.

Table 273: Fields on the Monitor Overview Page

Field	Description
Sites	View the sites at which the service is deployed. Click the **Sites** drop-down list and select **Show sites**
Connections	View the connections in the network. Click the **Connections** drop-down list and select **Show connections**.
Only the node with alerts	View the nodes with issues with the service. Click the drop-down list located next to the **Only the nodes with alerts** check box and select the type of alerts. • Critical—Issues that prevent the node from working and require action from the operator. The nodes with critical alerts are displayed in red. • Major—Issues that prevent the node from working at this time, but they do not require action from the operator. The nodes with major alerts are displayed in orange. • Minor—Issues that allow a node to continue working, but not optimally. The network operator may need to take action to resolve the issue. The nodes with minor alerts are displayed in yellow. **NOTE:** The nodes without any alerts are displayed in blue.

RELATED DOCUMENTATION

Managing Security Alerts Definitions | 804
Creating Security Alert Definitions | 805

Alerts Overview

Alerts and notifications are used to notify administrators about significant events within the system. Notifications can also be sent through e-mail. You will be notified when a predefined network traffic condition is met. The alert trigger threshold is the number of network traffic events crossing a predefined threshold within a period of time.

Alerts and notifications provide options for:

- Defining alert criteria based on a set of predefined filters. You can use the filters defined in the advanced search to create an alert. You can also save filters and add them to security alert definitions.
- Generating an alert message and notifying you when alert criteria are met.
- Searching for specific alerts on the Generated Alerts page based on alert ID, description, or alert type.
- Supporting event-based alerts.

For example, If you are an administrator, you can define a condition such that if the number of firewall-deny events crosses a predefined threshold in a given time range for a specific device, you will receive an e-mail alert.

> **NOTE:** If a threshold is crossed and remains so for a long duration, new alerts are not generated. Alerts are generated again when the number of logs matching the alert criteria drops below the threshold and crosses the threshold again.

RELATED DOCUMENTATION

About the Generated Alerts Page | **801**

About the Alert Definitions/Notifications Page | **803**

Managing Security Alerts Definitions | **804**

About the Generated Alerts Page

To access this page, click **Monitor > Alerts & Alarms > Alerts**.

Use this page to view the system event-based alerts in response to a configured alert definition. The generated alerts help you to identify problems that appear in your monitored network environment and

displays both security and CSO alerts. You can view statistics such as the number of critical and non-critical alerts.

Tasks You Can Perform

You can perform the following tasks from this page:

- Select the generated alert and then right-click or click **More > Jump to Events and Logs**. The corresponding events that triggered the alert are displayed.
- Select the generated alert and then right-click or click **More > Detail View**. The Alert Detail page appears displaying all the details of the alert.
- Select the generated alert and then right-click or click **More > Clear All Selections**.

Field Descriptions

Table 274 on page 802 provides guidelines on using the fields on the Generated Alerts page.

Table 274: Fields on the Generated Alerts Page

Field	Description
Severity	View the severity of the alert.
Time	View the date and time when the alert was generated.
Site	View the name of the tenant site.
Source	View the source of the alert. The source identifies whether an alert is a security alert or an SD-WAN alert.
Description	View the description of the alert.
Alert Type	View the type of alert.
ID	View the alert ID. Alert ID is an unique identification for each alert. For example, b4a3c027-7157-4861-8e3c-c872721cff2d.
Service Instance	View the service instance associated with the alert.
Object Type	View the object type.
Alert Name	View the name of the alert.
Tenant	View the name of the tenant.

RELATED DOCUMENTATION

Managing Security Alerts Definitions | 804

About the Alert Definitions/Notifications Page

To access this page, select **Monitor > Alarms & Alerts > Definitions/Notifications** in the Customer Portal.

Use the Alert Definitions page to view alert definitions for SD-WAN, SD-LAN, and manage alert definitions for security. An alert definition consists of data criterion for triggering alerts about issues in the SD-WAN environment. Alert definitions also define the necessary action required to resolve issues based on the severity of the alert. An alert is triggered when the event threshold exceeds the data criteria that is defined. You can create an alert definition to monitor your data in real time and identify issues and attacks before they impact your network.

You can also enable or disable SD-WAN/SD-LAN alarm notification.

Tasks You Can Perform

You can perform the following tasks from this page:

- View SD-WAN alert definitions. The SD-WAN alert definitions are loading by default when the Alert Definitions page is loaded. See Table 275 on page 804 for descriptions of the fields on the SD-WAN alert definitions pane.
- Manage Security alert definitions. See “Managing Security Alerts Definitions” on page 804.
- Enable or disable the e-mail notification for alarms. See “Enable E-mail Notifications for SD-LAN and SD-WAN Alarms” on page 810.
- Show or hide columns that contain information about security alert definitions. In the Security Alert Definitions tab, click the **Show Hide columns** icon in the top right corner of the page and select columns that you want to view on the page.
- Search for alert definitions using keywords. Click the search icon. Enter partial text or full text of the keyword in the search bar and press Enter. The search results are displayed.

Field Descriptions

Table 275 on page 804 describes the fields on the SD-WAN Alert Definitions pane.

Table 275: Fields on the SD-WAN Alert Definitions Pane

Field	Description
Rule Priority	View the priority of the alert definition. A value of one (1) indicates highest priority.
Alert Description	View the description of the alert.
Filter	View the matching alert criteria to trigger the alert.
Action	View the action to be performed to resolve issues.
Context	View the additional configuration parameters that you can pass on to the rule action function.

RELATED DOCUMENTATION

Managing Security Alerts Definitions | 804

About the Generated Alerts Page | 801

Managing Security Alerts Definitions

Use the Security pane to generate alerts that warn you of problems in your monitored environment. An alert definition consists of data criteria for triggering an alert. An alert is triggered when the event threshold exceeds the data criteria that is defined.

Tasks You Can Perform

You can perform the following tasks from this pane:

- Create security alert definition. See "Creating Security Alert Definitions" on page 805.
- Edit, clone, and delete security alert definition. See "Editing, Cloning, and Deleting Security Alert Definitions" on page 807.

Field Descriptions

Table 276 on page 805 provides guidelines on using the fields on the Security alert definitions pane.

Table 276: Fields on the Security Alert Definitions Pane

Field	Description
Alert Name	View the name of the alert.
Alert Description	View the description for the alert.
Filter	View filter values of the alert.
Recipients	View recipients' e-mail addresses where alert notifications are sent.
Status	View the status of the alert.
Alert Type	View the type of alert. Example: Event-based

RELATED DOCUMENTATION

Alerts Overview | 801

Creating Security Alert Definitions | 805

Creating Security Alert Definitions

You can create an alert definition to monitor your data in real time. You can identify issues and attacks before they impact your network.

For example, if you are an administrator, you can define a condition such that if the number of firewall deny events crosses a predefined threshold in a given time frame for a specific device, you receive an e-mail alert.

To create a security alert definition:

1. Select **Monitor > Alerts & Alarms > Definitions/Notifications > Security Alerts Definitions**.

   The Security alert definitions page appears.

2. Click the create icon (+) or add icon (+).

   The Create an Alert Definition page appears.

3. Complete the configuration according to the guidelines provided in Table 277 on page 806.

4. Click **OK**. If you want to discard the changes, click **Cancel** instead.

A new alert definition with the configured alert triggering condition is created. You can view the generated alerts from the alert definition to troubleshoot the issues with your system.

Table 277: Fields on the Security Alert Definitions Page

Field	Description
General	
Alert Name	Enter a unique string of alphanumeric characters, colons, periods, dashes, and underscores. No spaces are allowed and the maximum length is 63 characters.
Alert Description	Enter a description for the alerts; maximum length is 1024 characters.
Alert Type	Displays the type of alert that is system-based.
Status	Select the Active check box to view only the active alerts.
Severity	Select the severity level of the alert: info, minor, major, critical.
Trigger	
Use Data Criteria from Filters	Specifies the data criteria from the list of default and user-created filters that are saved from the Event Viewer. To add saved filters: • Click the **Use data criteria** from filters link. The Add Saved Filters page appears. • Select the filters to be added. • Click **OK**.
Add Data Criteria	Specifies the data criteria based on the Time Span period, Group By, and Filter By option. Filtered data only displays the subset of data that meets the criteria that you specify.
Recipient(s)	
E-mail Address(es)	Specify the e-mail addresses for the recipients of the alert notification.
Custom Message	Enter a custom string for identifying the type of alert in the alert notification e-mail.

RELATED DOCUMENTATION

Managing Security Alerts Definitions | 804
Editing, Cloning, and Deleting Security Alert Definitions | 807

Editing, Cloning, and Deleting Security Alert Definitions

IN THIS SECTION

- Editing Security Alert Definitions | 807
- Cloning Security Alert Definitions | 807
- Deleting Security Alert Definitions | 808

You can edit, clone, and delete security alert definitions.

Editing Security Alert Definitions

To edit the security alert definition:

1. Select **Monitor > Alerts & Alarms > Definitions/Notifications > Security Alerts Definitions**.

   The Security Alerts Definition page appears.

2. Select the check box of the security alert definition that you want to modify, and click the edit icon.

   The Edit Alert Definition page appears. The options available on the Create Alert Definition page are available for editing.

3. Update the configuration as needed.

4. Click **OK** to save the changes. If you want to discard your changes, click **Cancel** instead.

Cloning Security Alert Definitions

You can clone an alert definition when you want to quickly create a copy of an alert definition and modify its parameters including the name of the alert.

To clone an alert definition:

1. Select **Monitor > Alerts & Alarms > Definitions/Notifications > Security Alerts Definitions**.

   The Security Alert Definitions page appears.

2. Select the alert definition that you want to clone, and click **More > Clone** at the top right corner of the page.

   The Clone Alert Definition page appears. The options available on the Create Alert Definition page are available for editing.

3. Click **OK** to save the configuration.

   A new alert definition is created.

Deleting Security Alert Definitions

You can click the delete icon (X) to delete one or more alert definitions.

To delete the alert definition:

1. Select **Monitor > Alerts & Alarms > Definitions/Notifications > Security Alerts Definitions**.

   The Security Alerts Definition page appears.

2. Select the alert definition that you want to delete and click the delete icon (X icon).

   The Confirm Delete page appears.

3. Click **Yes** to delete the alert definition or **No** to cancel the deletion.

   If you click **Yes**, then the alert definition is deleted from the main page.

RELATED DOCUMENTATION

Managing Security Alerts Definitions | **804**

Creating Security Alert Definitions | **805**

About the Alarms Page

To access this page, select **Monitor > Alerts & Alarms > Alarms**.

Use the Alarms page to view system-generated alarms. Alarms notify you of conditions that might prevent the device from operating normally. Alarm conditions for a system are preset and are based on the fault monitoring and performance monitoring (FMPM) being performed on a device. For example, conditions such as hardware issues, drop in throughput and latency of data, temperature variations, and capacity optimization issues automatically trigger an alarm.

The difference between alerts and alarms lies in the type of events that are being monitored. An alert is used to notify administrators about significant events within the system. For example, when a predefined network traffic condition is met. For more information about alerts, see "Alerts Overview" on page 801.

Tasks You Can Perform

You can perform the following tasks from this page:

- View alarm activity within a specific time range:
  - Select either 2 hours (2h), 4 hours (4h), 8 hours (8h), 16 hours (16h), 24 hours (24h), or 1 week (1w), or Custom as the time range to view alarm activity. By default, alarm activity is displayed for a time range of 1 week.

    If you click Custom, the Custom Time Range Selection page appears.

    You must specify the **From** date and time, and **To** date and time (in MM/DD/YYYY and HH:MM:SS formats).
- View details of an alarm—Select a generated alarm on the page and right-click to select **Detail View** or click **More > Detail View** to view more details (such as alarm type, severity, and so on) of the alarm.
- Delete an alarm—Select an alarm that you want to delete and click the **Delete** icon. The selected alarm is deleted from the page.
- Apply a filter to view specific alarms—Click the **Filter** icon and select the filter criteria from the list of available options, to view only specific alarms. You can filter the alarms based on severity (critical, major, minor, normal), tenant name, site name, and source of the alarm.
- Show or hide columns on the page—Click the **Show or Hide Columns** icon to select or clear columns that you want to display or hide on the page.
- Select the number of alarms that you want to view on the page—From the **Details** list, select either **20**, **50** or **100** as the number of alarms that you want to view on the page.

Field Descriptions

Table 278 on page 810 describes the fields on the Alarms page.

Table 278: Fields on the Alarms Page

Field	Description
Severity	Severity of the alarm.
Time	Date and time when the alarm was generated.
Tenant	Name of the tenant.
Site	Name of the tenant site for which the alarm was generated.
Source	Source from where the alarm originated.
Description	Description of the alarm.
UUID	Universally Unique Identifier (UUID) of the alarm. You can use the UUID to identify an alarm on the **Monitor > Logs** page.
Link Name	Name of the link that generated the alarm.
Service Instance	Service instance associated with the alarm.
Object Type	Type of device from which the alarm originated. Example: Hub

Enable E-mail Notifications for SD-LAN and SD-WAN Alarms

Starting from CSO Release 5.1.1, you now notify the user (tenant administrators and tenant operators) about SD-WAN and SD-LAN alarms. You can also specify the minimum severity level of alarms that must be notified to the users. Alarm notifications enable users to take action to ensure that the network runs smoothly.

NOTE: You can enable or disable the e-mail notification for SD-WAN and SD-LAN alarms if you are an SP administrator, or OpCo administrator, or tenant administrator.

To enable e-mails notification for SD-WAN and SD-LAN alarms:

1. Select **Monitor > Alerts & Alarms > Definitions/Notifications**.

   The Definitions/Notifications page appears.

2. Select the **SD-WAN/SD-LAN Alarm Notifications** tab.

   The SD-WAN/SD-LAN Alarm Notifications page appears.

3. Complete the configuration according to the guidelines provided in Table 279 on page 811.

   NOTE: Fields marked with an asterisk (*) are mandatory.

4. Click **Save** to save the changes.

If you have enabled e-mail notifications, an e-mail will be sent to the user based on the severity level that you specified for an alarm.

If you have disabled e-mail notifications, the users will not receive e-mail notifications in case of alarms.

Table 279: SD-WAN/SD-LAN Alarm Notifications Settings

Field	Description
Send Email Notifications	Click the toggle button to enable or disable the e-mail notifications of alarms to users. By default, e-mail notifications are disabled. After enabling this field, you must specify the minimum severity level of the alarm and select the e-mail addresses of the users.
Minimum Severity to Report	Select the minimum severity level (critical, major, minor) of the alarms to users through an e-mail. • Critical—If you select this option, e-mail notifications are sent to users only for alarms with the severity level critical. • Major—If you select this option, e-mail notifications are sent to users only for alarms with the severity levels major or critical. • Minor—If you select this option, e-mail notifications are sent to users only for alarms with the severity levels minor, major, or critical.

Table 279: SD-WAN/SD-LAN Alarm Notifications Settings *(continued)*

Field	Description
Recipients	Select one or more e-mail addresses of the users from the list. Only users with tenant administrator or tenant operator roles are listed. The e-mail addresses listed are based on the users that are listed in the Administration > Users page.

Release History Table

Release	Description
5.1.1	Starting from CSO Release 5.1.1, you now notify the user (tenant administrators and tenant operators) about SD-WAN and SD-LAN alarms. You can also specify the minimum severity level of alarms that must be notified to the users. Alarm notifications enable users to take action to ensure that the network runs smoothly.

RELATED DOCUMENTATION

About the Alert Definitions/Notifications Page | 803

CHAPTER 28

Monitoring Security and Device Events

IN THIS CHAPTER

- About the All Security Events Page | 813
- About the Firewall Events Page | 818
- About the Web Filtering Events Page | 821
- About the IPsec VPNs Events Page | 824
- About the Content Filtering Events Page | 826
- About the Antispam Events Page | 828
- About the Antivirus Events Page | 830
- About the IPS Events Page | 833
- About the Device Events Page | 836
- About the Screen Events Page | 840

About the All Security Events Page

To access this page, click **Monitoring** > **Security Events** > **All Events**.

Use this page to get an overall, high-level view of your network environment. You can view abnormal events, attacks, viruses, or worms when log data is correlated and analyzed.

This page provides administrators with an advanced filtering mechanism and provides visibility into actual events collected by the Log Collector. Using the time-range slider, you can instantly focus on areas of unusual activity by dragging the time slider to the area of interest to you. The slider and the Custom button under Time Range remain at the top of each tab. Users select the time range, and then they can decide how to view the data, using the summary view or detail view tabs.

Tasks You Can Perform

You can perform the following tasks from this page:

- View a brief summary of all events in your network. See “Summary View” on page 814.

- View the comprehensive details of events in a tabular format that includes sortable columns. See “Detail View” on page 814.

Summary View

You can view a brief summary of all the events in your network. At the center of the page is critical information, including total number of events, viruses found, total number of interfaces that are down, number of attacks, CPU spikes, and system reboots. This data is refreshed automatically based on the selected time range. At the bottom of the page is a swim lane view of different events that are happening at a specific time. The events include firewall, web filtering, VPN, content filtering, antispam, antivirus, and IPS. Each event is color-coded, with darker shades representing a higher level of activity. Each tab provides deep information like type, and number of events occurring at that specific time.

Table 280 on page 814 describes the widgets on the All Events Summary View page.

Table 280: Widgets on the All Events Summary View Page

Field	Description
Total Events	View the total number of all the events that includes firewall, web filtering, IPS, IPSec VPNs, content filtering, antispam, and antivirus events.
Virus Instances	View the total number of virtual instances running in the system.
Attacks	View the total number of attacks on the firewall.
Interface Down	View the total number of interfaces that are down.
CPU Spikes	View the total number of times a CPU utilization spike has occurred.
Reboots	View the total number of system reboots.
Sessions	View the total number of sessions established through firewall.

Detail View

Click **Detail View** for comprehensive details of events in a tabular format that includes sortable columns. You can sort the events using the Group By option. For example, you can sort the events based on severity. The table includes information such as the rule that caused the event, severity for the event, event ID, traffic information, and how and when the event was detected.

Advanced Search

You can perform advanced search of all events using the text field present above the tabular column. It includes the logical operators as part of the filter string. Enter the search string in the text field and based on your input, a list of items from the filter context menu is displayed. . You can select a value from the list and then select a valid logical operator to perform the advanced search operation Press Enter to display the search result in the tabular column below.

To delete the search string in the text field, click the delete icon (X icon).

Examples of event log filters are shown in the following list:

- Specific events originating from or landing within United States

  Source Country = United States OR Destination Country = United States AND Event Name = IDP_ATTACK_LOG_EVENT, IDP_ATTACK_LOG_EVENT_LS, IDP_APPDDOS_APP_ATTACK_EVENT_LS, IDP_APPDDOS_APP_STATE_EVENT, IDP_APPDDOS_APP_STATE_EVENT_LS, AV_VIRUS_DETECTED_MT, AV_VIRUS_DETECTED, ANTISPAM_SPAM_DETECTED_MT, ANTISPAM_SPAM_DETECTED_MT_LS, FWAUTH_FTP_USER_AUTH_FAIL, FWAUTH_FTP_USER_AUTH_FAIL_LS, FWAUTH_HTTP_USER_AUTH_FAIL, FWAUTH_HTTP_USER_AUTH_FAIL_LS, FWAUTH_TELNET_USER_AUTH_FAIL, FWAUTH_TELNET_USER_AUTH_FAIL_LS, FWAUTH_WEBAUTH_FAIL,FWAUTH_WEBAUTH_FAIL_LS

- User wants to filter all RT flow sessions originating from IP addresses in specific countries and landing on IPs in specific countries

  Event Name = RT_FLOW_SESSION_CREATE,RT_FLOW_SESSION_CLOSE AND Source IP = 177.1.1.1,220.194.0.150,14.1.1.2,196.194.56.4 AND Destination IP = 255.255.255.255,10.207.99.75,10.207.99.72,223.165.27.13 AND Source Country = Brazil, United States, China, Russia, Algeria AND Destination Country = Germany, India, United States

- Traffic between zone pairs for policy – IDP2

  Source Zone = trust AND Destination Zone = untrust, internal AND Policy Name = IDP2

- UTM logs coming from specific source country, destination country, source IP addresses with or without specific destination IP addresses.

  Event Category = antispam, antivirus, contentfilter, webfilter AND Source Country = Australia AND Destination Country = Turkey, United States, Australia AND Source IP = 1.0.0.0,1.1.1.3 OR Destination IP = 74.125.224.47,5.56.17.61

- Events with specific sources IPs or events hitting HTP, FTP, HTTP, and unknown applications coming from host DC-SRX1400-1 or VSRX-75.

  Application = tftp, ftp, http, unknown OR Source IP = 192.168.34.10,192.168.1.26 AND Hostname = dc-srx1400-1,vsrx-75

Table 281 on page 816 describes the fields on the All Events Detail View Page.

Table 281: Fields on the All Events Detail View Page

Field	Description
Time	View the time when the log was received.
Event Name	View the event name of the log.
Site	View the name of the tenant site.
Source Country	View the source country name.
Source IP	View the source IP address from where the event occurred.
Destination Country	View the destination country name from where the event occurred.
Destination IP	View the destination IP address of the event.
Source Port	View the source port of the event.
Destination Port	View the destination port of the event.
Description	View the description of the log.
Attack Name	View the attack name of the log: Trojan, worm, virus, and so on.
Threat Severity	View the severity level of the threat.
Policy Name	View the policy name in the log.
UTM Category or Virus Name	View the UTM category of the log.
URL	View the accessed URL name that triggered the event.
Event Category	View the event category of the log.
User Name	View the username of the log.
Action	View the action taken for the event: warning, allow, and block.
Log Source	View the IP address of the log source.
Application	View the application name from which the events or logs are generated
Hostname	View the hostname in the log.

Table 281: Fields on the All Events Detail View Page *(continued)*

Field	Description
Service Name	The name of the application service. For example, FTP, HTTP, SSH, and so on.
Nested Application	View the nested application in the log.
Source Zone	View the source zone of the log.
Destination Zone	View the destination zone of the log.
Protocol ID	View the protocol ID in the log.
Roles	View the role name associated with the log.
Reason	View the reason for the log generation. For example, a connection tear down may have an associated reason such as “authentication failed”.
NAT Source Port	View the translated source port.
NAT Destination Port	View the translated destination port.
NAT Source Rule Name	View the NAT source rule name.
NAT Destination Rule Name	View the NAT destination rule name.
NAT Source IP	View the translated (or natted) source IP address. It can contain IPv4 or IPv6 addresses.
NAT Destination IP	View the translated (also called natted) destination IP address.
Traffic Session ID	View the traffic session ID of the log.
Path Name	View the path name of the log.
Logical system Name	View the name of the logical system.
Rule Name	View the name of the rule.
Profile Name	View the name of the All events profile that triggered the event.

RELATED DOCUMENTATION

- About the Firewall Events Page | 818
- About the Web Filtering Events Page | 821
- About the IPsec VPNs Events Page | 824
- About the Content Filtering Events Page | 826
- About the Antispam Events Page | 828
- About the Antivirus Events Page | 830
- About the IPS Events Page | 833

About the Firewall Events Page

To access this page, click **Monitor > Security Events > Firewall**.

Use the Firewall Events page to view information about security events based on firewall policies. Analyzing firewall logs yields useful security management information, such as attempts to breach your network and observing the inherent characteristics of your traffic in real-time. Using the time-range slider, you can quickly focus on the area of activity that you are most interested in. Once the time range is selected, all of the data presented in your view is refreshed automatically. You can also use the Custom button to set a custom time range.

There are two ways to view your data. You can select either the **Summary View** tab or the **Detail View** tab.

Tasks You Can Perform

You can perform the following tasks from this page:

- View a brief summary of all the firewall events in your network. See "Summary View" on page 818
- View the comprehensive details of events in a tabular format that includes sortable columns. See "Detail View" on page 819.

Summary View

The data presented in the line graph (also known as swim lanes) is refreshed automatically based on the selected time range. The line graph shows light blue lanes that represent all firewall events and dark blue lanes represent blocked firewall events.

Below the swim lanes are widgets displaying critical information such as top sources, top destinations, top users, and top reporting devices.

Table 282 on page 819 describes the widgets on the Summary View page.

Table 282: Widgets on the Summary View Page

Widget	Description
Top Sources	View the top source IP addresses of the network traffic; sorted by event count.
Top Destinations	View the top destination IP addresses of the network traffic; sorted by event count.
Top Users	View then top users of the network traffic; sorted by event count.
Top Reporting Devices	View the top reporting devices in the network; sorted by event count.

Detail View

Detail view includes information such as the rule that caused the event, severity for the event, event ID, traffic information, and how and when the event was detected

Table 283 on page 819 provides guidelines on using the fields on the Detail View page.

Table 283: Fields on the Detail View Page

Field	Description
Time	View the time when the log was received.
Event Name	View the event name of the log.
Source Country	View the source country name from where the event originated.
Source IP	View the source IP address from where the event occurred.
Destination Country	View the destination country name from where the event occurred.
Destination IP	View the destination IP address of the event.
Source Port	View the source port of the event.
Destination Port	View the destination port of the event.
Description	View the description of the log.
Policy Name	View the policy name in the log.

Table 283: Fields on the Detail View Page *(continued)*

Field	Description
User Name	View the username of the log.
Action	View the action taken for the event: warning, allow, and block.
Log Source	View the IP address of the log source (IPv4 or IPv6).
Application	View the application name from which the events or logs are generated.
Hostname	View the hostname in the log.
Service Name	The name of the application service. For example, FTP, HTTP, SSH, and so on.
Nested Application	View the nested application in the log.
Source Zone	View the user traffic received from the zone.
Destination Zone	View the destination zone of the log.
Protocol ID	View the protocol ID in the log.
Roles	View the role names associated with the event.
NAT Source Port	View the translated source port.
NAT Destination Port	View the translated destination port.
NAT Source Rule Name	View the NAT source rule name.
NAT Destination Rule Name	View the NAT destination rule name.
NAT Source IP	View the translated (or natted) source IP address. It can contain IPv4 or IPv6 addresses.
NAT Destination IP	View the translated (also called natted) destination IP address.
Traffic Session ID	View the traffic session ID of the log.
Rule Name	View the rule name of the log.

RELATED DOCUMENTATION

- About the All Security Events Page | 813
- About the Web Filtering Events Page | 821
- About the IPsec VPNs Events Page | 824
- About the Content Filtering Events Page | 826
- About the Antispam Events Page | 828
- About the Antivirus Events Page | 830
- About the IPS Events Page | 833

About the Web Filtering Events Page

To access this page, click **Monitor > Security Events > Web Filtering**.

Use the Web Filtering page to view information about security events based on Web filtering policies. Web filtering allows you to permit or block access to specific websites by URL or by URL category using cloud-based lookups, a local database, or an external Websense server. Analyzing Web filtering logs yields useful security management information such as users detected accessing restricted URLs and actions taken by the system. Using the time-range slider, you can quickly focus on the area of activity that you are most interested in. Once the time range is selected, all of the data presented in your view is refreshed automatically. You can also use the Custom button to set a custom time range.

There are two ways to view your data. You can select either the **Summary View** tab or the **Detail View** tab.

Tasks You Can Perform

You can perform the following tasks from this page:

- View a brief summary of all the Web filtering events in your network. See "Summary View" on page 821.
- View the comprehensive details of events in a tabular format that includes sortable columns. See "Detail View" on page 822.

Summary View

The top of the page has a swim lane graph of all the Web filtering events against the blocked events.

Below the swim lanes are widgets displaying critical information such as top sources, top destinations, top users, and top reporting devices.

You can use the widgets at the bottom of the page to view critical information such as top URLs blocked, top matched profiles, top sources, and top destinations.

Table 284 on page 822 describes the widgets on the Summary View page.

Table 284: Widgets on the Summary View Page

Widget	Description
Top URLs blocked	View the URL names that are blocked; sorted by event count.
Top Matched Profiles	View the web filtering profile names; sorted by event count.
Top Sources	View the top source IP addresses of the network traffic; sorted by event count.
Top Destinations	View the top destination IP addresses of the network traffic; sorted by event count.

Detail View

You can aggregate the events using the Group By option. For example, you can group the events based on source country. The table includes information such as the event name, UTM category, source IP address, source country, and so on.

Table 285 on page 822 provides guidelines on using the fields on the Detail View page.

Table 285: Fields on the Detail View Page

Fields	Description
Time	View the time when the event occurred.
Event Name	View the event name of the log.
Source Country	View the source country name from where the event originated.
Source IP	View the source IP address from where the event occurred (IPv4 or IPv6).
Destination Country	View the destination country name from where the event occurred.
Destination IP	View the destination IP address of the event (IPv4 or IPv6).
Source Port	View the source port of the event.
Destination Port	View the destination port of the event.

Table 285: Fields on the Detail View Page *(continued)*

Fields	Description
Description	View the description of the log.
UTM category or Virus Name	View the UTM category of the log: enhanced, local, and redirect.
URL	View the accessed URL name that triggered the event.
Action	View the action taken for the event: warning, allow, and block.
Log Source	View the IP address of the log source (IPv4 or IPv6).
Host Name	View the hostname in the log.
Source Zone	View the user traffic received from the zone.
Roles	View the role names associated with the event.
Reason	View the reason for the log generation. For example, unrestricted access.
Path Name	View the path name of the log.
Profile Name	View the name of the Web filtering profile that triggered the event.

RELATED DOCUMENTATION

- About the All Security Events Page | **813**
- About the Firewall Events Page | **818**
- About the IPsec VPNs Events Page | **824**
- About the Content Filtering Events Page | **826**
- About the Antispam Events Page | **828**
- About the Antivirus Events Page | **830**
- About the IPS Events Page | **833**

About the IPsec VPNs Events Page

To access this page, click **Monitor > Security Events > IPsec VPNs**.

Use this page to view information about security events based on IPSec VPN policies. The event viewer provides a view of all IPsec VPN events.

Using the time-range slider, you can quickly focus on the area of activity that you are most interested in. Once the time range is selected, all of the data presented in your view is refreshed automatically. You can also use the custom button to set a custom time range.

There are two ways to view your data. You can select either the **Summary View** tab or the **Detail View** tab.

Tasks You Can Perform

You can perform the following tasks from this page:

- View a brief summary of all the IPsec VPN events in your network. See "Summary View" on page 824.
- View the comprehensive details of events in a tabular format that includes sortable columns. See "Detail View" on page 825.

Summary View

The top of the page has a swim lane graph of all the VPN events. You can use the widgets at the bottom of the page to view critical information such as top sources, top destinations, and top reporting devices.

Table 286 on page 824 describes the widgets on the Summary View page.

Table 286: Widgets on the Summary View Page

Widget	Description
Top Sources	View the top source IP addresses of the network traffic; sorted by event count.
Top Destinations	View the top destination IP addresses of the network traffic; sorted by event count.
Top Reporting Devices	View the top reporting device IP addresses; sorted by event count.

Detail View

You can aggregate the events using the Group By option. For example, you can group the events based on source country. The table includes information such as the event name, log source, host name, source country, and so on.

Table 287 on page 825 provides guidelines on using the fields on the Detail View page.

Table 287: Fields on the Detail View Page

Fields	Description
Time	View the time when the event occurred.
Event Name	View the event name of the log.
Source Country	View the source country name from where the event originated.
Destination Country	View the destination country name from where the event occurred.
Destination Port	View the destination port of the event.
Description	View the description of the log.
Log Source	View the IP address of the log source (IPv4 or IPv6).
Host Name	View the hostname in the log.
Rule Name	View the name of the antivirus profile that triggered the event.

RELATED DOCUMENTATION

- About the All Security Events Page | 813
- About the Firewall Events Page | 818
- About the Web Filtering Events Page | 821
- About the Content Filtering Events Page | 826
- About the Antispam Events Page | 828
- About the Antivirus Events Page | 830
- About the IPS Events Page | 833

About the Content Filtering Events Page

To access this page, click **Monitor > Security Events > Content Filtering**.

Use this page to view information about security events based on content filtering policies. The event viewer provides a view of all content filtering events and how the events are handled by content filter. This page can be used to view traffic on the network in real time or as a debugging tool to view how content filtering is operating.

Content filtering provides basic data loss prevention functionality. Content filtering screens traffic based on MIME type, file extension, protocol commands, and embedded object type. It either permits or blocks specific commands or extensions on a protocol-by-protocol basis.

Using the time-range slider, you can quickly focus on the area of activity that you are most interested in. Once the time range is selected, all of the data presented in your view is refreshed automatically. You can also use the Custom button to set a custom time range.

There are two ways to view your data. You can select either the **Summary View** tab or the **Detail View** tab.

Tasks You Can Perform

You can perform the following tasks from this page:

- View a brief summary of all the content filtering events in your network. See "Summary View" on page 826.
- View the comprehensive details of events in a tabular format that includes sortable columns. See "Detail View" on page 827.

Summary View

The top of the page has a swim lane graph of all the content filtering events against the blocked events. You can use the widgets at the bottom of the page to view critical information such as top blocked protocol commands, top reasons, and top sources.

Table 288 on page 826 describes the widgets on the Summary View page.

Table 288: Widgets on the Summary View Page

Widget	Description
Top Blocked Protocol commands	View the top command names or file extensions blocked on a protocol-by-protocol basis.
Top Reasons	View the top reasons for blocking the content. For example: Inappropriate or harmful communication.

Table 288: Widgets on the Summary View Page *(continued)*

Widget	Description
Top Sources	View the top source IP addresses of the network traffic; sorted by event count.

Detail View

You can aggregate the events using the Group By option. For example, you can group the events based on source country. The table includes information such as the event name, UTM category, source IP address, source country, and so on.

Table 289 on page 827 provides guidelines on using the fields on the Detail View page.

Table 289: Fields on the Detail View Page

Fields	Description
Time	View the time when the event occurred.
Event Name	View the event name of the log.
Source Country	View the source country name from where the event originated.
Source IP	View the source IP address from where the event occurred (IPv4 or IPv6).
Description	View the description of the log.
UTM Category or Virus Name	View the UTM category of the log: enhanced, local, and redirect.
URL	View the accessed URL name that triggered the event.
Argument	View the type of traffic. For example, FTP and HTTP.
Action	View the action taken for the event: warning, allow, and block.
Log Source	View the IP address of the log source (IPv4 or IPv6).
Host Name	View the hostname in the log.
Source Zone	View the user traffic received from the zone.
Roles	View the role names associated with the event.
Reason	View the reason for the log generation. For example, unrestricted access

Table 289: Fields on the Detail View Page *(continued)*

Fields	Description
Profile Name	View the name of the content filtering profile that triggered the event.

RELATED DOCUMENTATION

- About the All Security Events Page | **813**
- About the Firewall Events Page | **818**
- About the Web Filtering Events Page | **821**
- About the IPsec VPNs Events Page | **824**
- About the Antispam Events Page | **828**
- About the Antivirus Events Page | **830**
- About the IPS Events Page | **833**

About the Antispam Events Page

To access this page, click **Monitor > Security Events > Antispam**.

Use this page to view information about security events based on antispam policies. The event viewer provides a view of all antispam events and the action taken by the antispam scanner.

The antispam scanner inspects and block spam by scanning inbound and outbound SMTP e-mail traffic. The filtering can be server-based using an external spam block list server or local-based using local lists (blocklists and allowlists) for matching.

Using the time-range slider, you can quickly focus on the area of activity that you are most interested in. Once the time range is selected, all of the data presented in your view is refreshed automatically. You can also use the Custom button to set a custom time range.

There are two ways to view your data. You can select either the **Summary View** tab or the **Detail View** tab.

Tasks You Can Perform

You can perform the following tasks from this page:

- View a brief summary of all the antispam events in your network. See "Summary View" on page 829.

- View the comprehensive details of events in a tabular format that includes sortable columns. See "Detail View" on page 829.

Summary View

The top of the page has a swim lane graph of all antispam events. You can use the widget at the bottom of the page to view source IP addresses of the network traffic, sorted by event count.

Detail View

You can aggregate the events using the Group by option. For example, you can group the events based on source country. The table includes information such as the event name, UTM category, source IP address, source country, and so on.

Table 290 on page 829 provides guidelines on using the fields on the Detail View page.

Table 290: Fields on the Detail View Page

Fields	Description
Time	View the time when the event occurred.
Event Name	View the event name of the log.
Source Country	View the source country name from where the event originated.
Source IP	View the source IP address from where the event occurred (IPv4 or IPv6).
Description	View the description of the log.
UTM Category or Virus Name	View the UTM category of the log: enhanced, local, and redirect.
URL	View the accessed URL name that triggered the event.
Argument	View the type of traffic. For example, FTP and HTTP.
Action	View the action taken for the event: warning, allow, and block.
Log Source	View the IP address of the log source (IPv4 or IPv6).
Host Name	View the hostname in the log.
Source Zone	View the user traffic received from the zone.
Roles	View the role names associated with the event.

Table 290: Fields on the Detail View Page *(continued)*

Fields	Description
Reason	View the reason for the log generation. For example, unrestricted access
Profile Name	View the name of the content filtering profile that triggered the event.

RELATED DOCUMENTATION

- About the All Security Events Page | 813
- About the Firewall Events Page | 818
- About the Web Filtering Events Page | 821
- About the IPsec VPNs Events Page | 824
- About the Content Filtering Events Page | 826
- About the Antivirus Events Page | 830
- About the IPS Events Page | 833

About the Antivirus Events Page

To access this page, click **Monitor > Security Events > Antivirus**.

Use this page to view information about security events based on antivirus policies. The event viewer provides a view of all antivirus events and the action taken by the virus scanner.

The antivirus scanner inspects files transmitted over several protocols to determine if the files exchanged are malicious (for example, viruses, Trojans, rootkits, and worms).

Using the time-range slider, you can quickly focus on the area of activity that you are most interested in. Once the time range is selected, all of the data presented in your view is refreshed automatically. You can also use the Custom button to set a custom time range.

There are two ways to view your data. You can select either the **Summary View** tab or the **Detail View** tab.

Tasks You Can Perform

You can perform the following tasks from this page:

- View a brief summary of all the antivirus events in your network. See "Summary View" on page 831.

- View the comprehensive details of events in a tabular format that includes sortable columns. See "Detail View" on page 831.

Summary View

The top of the page has a swim lane graph of all the antivirus events against the blocked events. You can use the widgets at the bottom of the page to view critical information such as top blocked protocol commands, top reasons, and top sources.

Table 291 on page 831 provides guidelines on using the widgets on the Detail View page.

Table 291: Widgets on the Summary Page

Field	Description
Top Sources	View the top source IP addresses of the network traffic; sorted by event count.
Top Destinations	View the top destination IP addresses of the network traffic; sorted by event count.
Top Reporting/Attacked Devices	View the top reporting/attacked device IP addresses; sorted by event count.
Top Viruses	View the top virus names detected; sorted by event count.
Top Source Countries	View the top source country names where the events originated; sorted by event count.
Top Destination Countries	View the top destination country names where the events occurred; sorted by event count.

Detail View

You can aggregate the events using the Group By option. For example, you can group the events based on source country. The table includes information such as the event name, UTM category, source IP address, source country, and so on.

Table 292 on page 831 provides guidelines on using the fields on the Detail View page.

Table 292: Fields on the Detail View Page

Fields	Description
Time	View the time when the event occurred.

Table 292: Fields on the Detail View Page *(continued)*

Fields	Description
Event Name	View the event name of the log.
Source Country	View the source country name from where the event originated.
Source IP	View the source IP address from where the event occurred (IPv4 or IPv6).
Destination Country	View the destination country name from where the event occurred.
Destination IP	View the destination IP address of the event (IPv4 or IPv6).
Source Port	View the source port of the event.
Destination Port	View the destination port of the event.
Description	View the description of the log.
UTM Category or Virus Name	View the UTM category of the log: enhanced, local, and redirect.
URL	View the accessed URL name that triggered the event.
Action	View the action taken for the event: warning, allow, and block.
Log Source	View the IP address of the log source (IPv4 or IPv6).
Host Name	View the hostname in the log.
Source Zone	View the user traffic received from the zone.
Roles	View the role names associated with the event.
Reason	View the reason for the log generation. For example, unrestricted access.
Profile Name	View the name of the antivirus profile that triggered the event.

RELATED DOCUMENTATION

About the All Security Events Page | **813**

About the Firewall Events Page | **818**

About the Web Filtering Events Page | 821
About the IPsec VPNs Events Page | 824
About the Content Filtering Events Page | 826
About the Antispam Events Page | 828
About the IPS Events Page | 833

About the IPS Events Page

To access this page, click **Monitor > Security Events > IPS**.

Use the IPS Events page to view information about security events based on IPS policies. Analyzing IPS logs yields useful security management information, such as abnormal events, attacks, viruses, or worms.

Using the time-range slider, you can quickly focus on the area of activity that you are most interested in. Once the time range is selected, all of the data presented in your view is refreshed automatically. You can also use the custom button to set a custom time range.

There are two ways to view your data. You can select either the **Summary View** tab or the **Detail View** tab.

Tasks You Can Perform

You can perform the following tasks from this page:

- View a brief summary of all the all the IPS events in your network. See "Summary View" on page 833.
- View the comprehensive details of events in a tabular format that includes sortable columns. See "Detail View" on page 834.

Summary View

The data presented in the area graph is refreshed automatically based on the selected time range. You can use widgets to view critical information such as IPS severities, top sources, top destinations, top reporting devices, top IPS attacks, top source countries, and top destination countries.

Table 293 on page 833 provides guidelines on using the widgets on the Detail View page.

Table 293: Widgets on the Summary Page

Field	Description
IPS Severities	View the top IPS severities of the events based on the severity level: high, medium, low.

Table 293: Widgets on the Summary Page *(continued)*

Field	Description
Top Sources	View the top source IP addresses of the network traffic; sorted by the number of event occurrences.
Top Destinations	View the top destination IP addresses of the network traffic; sorted by the number of event occurrences.
Top Reporting/Attacked Devices	View the top devices that are attacked by IPS events; sorted by the number of times users are active on the network.
Top IPS attacks	View the top IPS attacks in the network traffic; sorted by the times devices are attacked.
Top Source Countries	View the top source countries from where the event source originated; sorted by the number of IP addresses.
Top Destination Countries	View the top source countries from where the event source originated; sorted by the number of IP addresses.

Detail View

You can sort the events using the Group By option. For example, you can sort the events based on severity. The table includes information such as the rule that caused the event, severity for the event, event ID, traffic information, and how and when the event was detected.

Table 294 on page 834 provides guidelines on using the fields on the Detail View page.

Table 294: Fields on the Detail View Page

Column	Description
Time	View the time when the log was received.
Event Name	View the event name of the log.
Source Country	View the source country name from where the event originated.
Source IP	View the source IP address from where the event occurred.
Destination Country	View the destination country name from where the event occurred.
Destination IP	View the destination IP address of the event.

Table 294: Fields on the Detail View Page *(continued)*

Column	Description
Source Port	View the source port of the event.
Destination Port	View the destination port of the event.
Description	View the description of the log.
Attack name	View the attack name of the log: Trojan, worm, virus, and so on.
Threat Severity	View the threat severity of the event.
Policy Name	View the policy name in the log.
Action	View the action taken for the event: warning, allow, and block.
Log Source	View the IP address of the log source.
Application	View the application name from which the events or logs are generated.
Hostname	View the host name in the log.
Service Name	View the name of the application service. For example, FTP, HTTP, SSH, and so on.
Nested Application	View the nested application name in the log.
Source Zone	View the source zone of the log.
Destination Zone	View the destination zone of the log.
Protocol ID	View the protocol ID in the log.
NAT Source Port	View the translated source port.
NAT Destination Port	View the translated destination port
NAT Source IP	View the NAT source IP address of the log.
NAT Destination IP	View the NAT destination IP address of the log.
Rule Name	View the name of the rule.

RELATED DOCUMENTATION

- About the All Security Events Page | 813
- About the Firewall Events Page | 818
- About the Web Filtering Events Page | 821
- About the IPsec VPNs Events Page | 824
- About the Content Filtering Events Page | 826
- About the Antispam Events Page | 828
- About the Antivirus Events Page | 830

About the Device Events Page

To access this page, click **Monitor** > **Device Events**.

Use the Device Events page to view information about device events such as routine operations, failure and error conditions, and emergency or critical conditions.

You can view comprehensive details of device events in a tabular format that includes sortable columns and a line graph (also known as swim lanes). The data presented in the line graph is refreshed automatically based on the selected time range. The line graph shows light blue areas that represent all device events and dark blue areas represent blocked device events

Tasks You Can Perform

You can perform the following tasks from this page:

- Click **Custom** button to select the date and time range to generate the device event.
- Show or hide time range in the carousel by clicking **show** or **hide** buttons at the top of the page.

Advanced Search

You can perform advanced search of all events using the text field present above the tabular column. It includes the logical operators as part of the filter string. Enter the search string in the text field and based on your input, a list of items from the filter context menu is displayed. . You can select a value from the list and then select a valid logical operator to perform the advanced search operation Press Enter to display the search result in the tabular column below.

To delete the search string in the text field, click the delete icon (X icon).

Examples of event log filters are shown in the following list:

- Specific events originating from or landing within United States

  Source Country = United States OR Destination Country = United States AND Event Name = IDP_ATTACK_LOG_EVENT, IDP_ATTACK_LOG_EVENT_LS, IDP_APPDDOS_APP_ATTACK_EVENT_LS, IDP_APPDDOS_APP_STATE_EVENT, IDP_APPDDOS_APP_STATE_EVENT_LS, AV_VIRUS_DETECTED_MT, AV_VIRUS_DETECTED, ANTISPAM_SPAM_DETECTED_MT, ANTISPAM_SPAM_DETECTED_MT_LS, FWAUTH_FTP_USER_AUTH_FAIL, FWAUTH_FTP_USER_AUTH_FAIL_LS, FWAUTH_HTTP_USER_AUTH_FAIL, FWAUTH_HTTP_USER_AUTH_FAIL_LS, FWAUTH_TELNET_USER_AUTH_FAIL, FWAUTH_TELNET_USER_AUTH_FAIL_LS, FWAUTH_WEBAUTH_FAIL,FWAUTH_WEBAUTH_FAIL_LS

- User wants to filter all RT flow sessions originating from IPs in specific countries and landing on IPs in specific countries

  Event Name = RT_FLOW_SESSION_CREATE,RT_FLOW_SESSION_CLOSE AND Source IP = 177.1.1.1,220.194.0.150,14.1.1.2,196.194.56.4 AND Destination IP = 255.255.255.255,10.207.99.75,10.207.99.72,223.165.27.13 AND Source Country = Brazil,United States,China,Russia,Algeria AND Destination Country = Germany,India,United States

- Traffic between zone pairs for policy – IDP2

  Source Zone = trust AND Destination Zone = untrust, internal AND Policy Name = IDP2

- UTM logs coming from specific source country, destination country, source IPs with or without specific destination IPs

  Event Category = antispam, antivirus, contentfilter, webfilter AND Source Country = Australia AND Destination Country = Turkey, United States, Australia AND Source IP = 1.0.0.0,1.1.1.3 OR Destination IP = 74.125.224.47,5.56.17.61

- Events with specific sources IPs or events hitting HTP, FTP, HTTP, and unknown applications coming from host DC-SRX1400-1 or VSRX-75.

  Application = tftp, ftp, http, unknown OR Source IP = 192.168.34.10,192.168.1.26 AND Hostname = dc-srx1400-1,vsrx-75

Field Descriptions

Table 295 on page 837 provides guidelines on using the fields on the Device Events page.

Table 295: Fields on the Device Events Detailed View Page

Field	Description
Time	View the time when the log was received.
Event Name	View the event name of the log.
Site	View the name of the tenant site.

Table 295: Fields on the Device Events Detailed View Page *(continued)*

Field	Description
Source Country	View the name of source country from where the event originated.
Source IP	View the source IP address from where the event occurred.
Destination Country	View the name of destination country from where the event occurred.
Destination IP	View the destination IP address of the event.
Source Port	View the source port of the device event.
Destination Port	View the destination port of the device event.
Description	View the description of the log.
Attack Name	View the attack name of the log. For example, Trojan, worm, virus, and so on.
Threat Severity	View the severity level of the threat.
Policy Name	View the policy name in the log.
UTM Category or Virus Name	View the UTM category of the log.
URL	View the accessed URL name that triggered the event.
Event Category	View the event category of the log.
User Name	View the username of the log.
Argument	View the type of traffic. For example, ftp and http.
Action	View the action taken for the event. For example, warning, allow, or block.
Log Source	View the IP address of the log source.
Application	View the application name from which the events or logs are generated.
Hostname	View the hostname in the log.
Service Name	View the name of the application service. For example, FTP, HTTP, SSH, and so on.

Table 295: Fields on the Device Events Detailed View Page *(continued)*

Field	Description
Nested Application	View the nested application in the log.
Source Zone	View the source zone of the log.
Destination Zone	View the destination zone of the log.
Protocol ID	View the protocol ID in the log.
Roles	View the role name associated with the log.
Reason	View the reason for the log generation. For example, a connection tear down may have an associated reason such as authentication failed.
NAT Source Port	View the translated source port.
NAT Destination Port	View the translated destination port.
NAT Source Rule Name	View the NAT source rule name.
NAT Destination Rule Name	View the NAT destination rule name.
NAT Source IP	View the translated (or natted) source IP address. It can contain IPv4 or IPv6 addresses.
NAT Destination IP	View the translated (also called natted) destination IP address.
Traffic Session ID	View the traffic session ID of the log.
Path Name	View the path name of the log.
Logical System Name	View the name of the logical system.
Rule Name	View the name of the rule.
Profile Name	The name of the profile that triggered the event.
Event Count	View the number of events occurred.
Tenant	View the name of the tenant from which the event originated.

RELATED DOCUMENTATION

About the All Security Events Page | 813

About the Screen Events Page

To access this page, click **Monitor > Security Events > Screen**.

Use this page to view information about screen events that occur as a result of the screen options configured on SRX Series or vSRX security devices. Screen options are a detection and defense mechanism configured to filter the connection attempts bound towards a security zone. Screen options are used to prevent attacks, such as IP address sweeps, port scans, denial of service (DOS) attacks, Internet Control Message Protocol (ICMP), UDP, and SYN (Synchronize) floods.

You can view information related to screen events, including ICMP screening, IP screening, TCP screening, and UDP screening.

Using the time-range slider, you can quickly focus on the time and area of activity that you are most interested in. Once the time range is selected, all of the data presented in your view is refreshed automatically. You can also use the **Custom** button to set a custom time range.

There are two ways to view your data. You can select either the **Summary View** tab or the **Detail View** tab.

Tasks You Can Perform

You can perform the following tasks from this page:

- View a brief summary of all the screen events in your network. See “Summary View” on page 840.
- View the comprehensive details of events in a tabular format that includes sortable columns. See “Detail View” on page 841.

Summary View

The top of the page has a swim lane graph of all the screen events. You can use the widgets at the bottom of the page to view critical information such as, top sources, top source countries, top destinations, and top destination countries.

Table 296 on page 841 describes the widgets on the Detail View page.

Table 296: Widgets on the Summary Page

Field	Description
Top Sources	Top five source IP addresses with highest network traffic.
Top Destinations	Top five destination IP addresses with highest network traffic.
Top Source Countries	Top five countries from which the traffic that triggered the highest number of events originated and the number of events per country.
Top Destination Countries	Top five countries to which the traffic that triggered the highest number events was sent and the number of events per country.

Detail View

You can group the events using the **Group By** option. For example, you can group the events based on source country. The table includes information such as the event name, UTM category, source IP address, source country, and so on.

Table 297 on page 841 describes the fields on the Detail View page.

Table 297: Fields on the Detail View Page

Fields	Description
Log Generated Time	Time when the event occurred.
Log Received Time	Time the log was received at the log collector.
Site	Name of the tenant site from which the event originated.
Event Name	Name of the device event in the log.
Source Country	Country from which the traffic that triggered the event originated.
Source IP	Source IP address for the traffic that triggered the event (IPv4 or IPv6).
Destination Country	Country to which the traffic that triggered the event was sent.
Destination IP	Destination IP address for the traffic that triggered the event (IPv4 or IPv6).
Source Port	Source TCP/UDP port number of the traffic that triggered the event.
Destination Port	Destination TCP/UDP port number of the traffic that triggered the event.

Table 297: Fields on the Detail View Page *(continued)*

Fields	Description
Attack Name	Name of the attack in the log for threat event. For example, trojan, worm, virus, and so on.
Description	Brief description of the event.
Threat Severity	Level of severity of the threat. For example, minor, major, critical, and so on.
Policy Name	Name of the policy which generates the log. The policy is configured on the SRX Series or vSRX device.
Virus Name	This field is not applicable for screen events.
URL	Accessed URL that triggered the event.
Event Category	Event category in the log. For example, screen.
User Name	User name identified by the SRX Series or vSRX device, if user identity is enabled on the device.
Argument	Type of traffic. For example, FTP and HTTP.
Action	Action taken for the event. For example, warning, allow, and block.
Log Source	IP address of the device where the log is received (IPv4 or IPv6).
Application	Name of the application associated with the traffic that triggered the event.
Host Name	Hostname of the device where the log was generated.
Service Name	Name of the application service used for the traffic that triggered the event. For example, FTP, HTTP, SSH, and so on.
Nested Application	Nested application associated with the traffic that triggered the event.
Source Zone	Source security zone of the traffic that triggered the event.
Destination Zone	Destination security zone of the traffic that triggered the event.
Protocol ID	Protocol ID of the traffic that triggered the event.
Roles	Roles of the user as defined in the Active Directory, if available.

Table 297: Fields on the Detail View Page *(continued)*

Fields	Description
Reason	Reason for the log generation. For example, unrestricted access.
NAT Source Port	Translated source port.
NAT Destination Port	Translated destination port.
NAT Source Rule Name	NAT source rule name configured on the SRX Series or vSRX device.
NAT Destination Rule Name	NAT destination rule name configured on the SRX Series or vSRX device.
NAT Source IP	Translated source IP address for the traffic that triggered the event (IPv4 or IPv6).
NAT Destination IP	Translated destination IP address for the traffic that triggered the event (IPv4 or IPv6).
Traffic Session ID	Traffic session ID of the log.
Path Name	This field is not applicable for screen events.
Logical System Name	Name of the logical system which received the log.
Rule Name	Name of the rule which generates the log. This rule is configured on the SRX Series or vSRX device.
Profile Name	Name of the profile which filters the traffic that triggered the event.
Client Host Name	Hostname of the client associated with the traffic that triggered the event. For example, if a specific computer is infected, the name of that computer is displayed.
Malware info	Information about the malware causing the event.

RELATED DOCUMENTATION

- About the All Security Events Page | **813**
- About the Firewall Events Page | **818**
- About the Web Filtering Events Page | **821**

About the IPsec VPNs Events Page | 824
About the Content Filtering Events Page | 826
About the Antispam Events Page | 828
About the Antivirus Events Page | 830
About the IPS Events Page | 833

CHAPTER 29

Monitoring SD-WAN Events

IN THIS CHAPTER

- SD-WAN Events Overview | 845
- About the SD-WAN Events Page | 846

SD-WAN Events Overview

Service-level agreements (SLAs) define the expected class of service (CoS) for all applications and application groups in a site. The network operator needs tools to measure and monitor the performance metrics for all applications to determine the quality of the network and adherence to an assured CoS. To ensure compliance with SLAs, the network operator also needs tools to take remedial action when network performance deteriorates and SLAs are not being met. SD-WAN link-switch events enable the network to switch WAN links to meet the site's SLA requirements when the network-designated WAN link is unable to meet the site's SLA requirements.

Because SLA parameters override the path preference, in dynamic SD-WAN policies, the SD-WAN network chooses the best possible WAN link for traffic management. The WAN link is chosen is based on the SLA parameters defined in the SLA profile. If multiple links match the SLA profile, the least loaded link is chosen. When a policy intent is deployed on a site, if the WAN link chosen by the SD-WAN network is unable to meet the SLA requirements in runtime, then the site switches WAN links to meet the SLA requirements. This link switching is called an SD-WAN event. Link switching also takes into account the priority defined in the SLA profile and SLA profiles with higher priority are given precedence while finding alternate WAN links. The ability of a site to switch WAN links ensures that SLA requirements are met and instances of not meeting the SLA requirements are minimized.

In static policies, link switching cannot occur even if the designated WAN link is unable to meet the SLA requirements, because path preference is defined.

RELATED DOCUMENTATION

- About the SD-WAN Events Page | 846
- SLA Profiles and SD-WAN Policies Overview | 499

About the SD-WAN Events Page

To access this page, click **Monitor > Link Switch Events** in the Customer Portal.

You can use the SD-WAN Events page to view information about SD-WAN events. An SD-WAN event is triggered when the SLA requirements for a site are not met on its network-designated WAN link and the site switches WAN links to meet the SLA requirements.

Tasks You Can Perform

You can perform the following tasks from this page:

- View a graphical representation of SD-WAN events in a specified time range (Time Range widget)

  The x-axis represents the defined time and the y-axis represents number of SD-WAN events.

  Use the slider to decrease or increase the time range within which you want to view SD-WAN events. You can also select from pre-defined time ranges such as 2h, 4h, 8h, 16h, 24h, or Custom.

  If you select Custom, you must specify the dates and times (in MM/DD/YYYY and HH:MM:SS 24-hour or AM/PM formats) from when and up to when you want the SD-WAN events displayed.
- View the SD-WAN events that occurred and information related to the events; see Table 298 on page 847.
- Show or hide the columns displayed on the page—Click the Show Hide Columns icon at the top right corner of the page and select the columns that you want displayed in the grid.
- Sort and filter SD-WAN events:

  NOTE: Sorting and filtering is applicable only to some fields.

  - Click a column name to sort the SD-WAN events based on the column name.
  - Click the filter icon (funnel) to toggle the filtering. You can enter the filter parameters in one or more fields and press Enter to display the filtered results.
- Search for SD-WAN events using keywords. Click the search icon. Enter partial text or full text of the keyword in the search bar and press Enter. The search results are displayed.

Field Descriptions

Table 298 on page 847 describes the fields on the SD-WAN Events page.

Table 298: Fields on the SD-WAN Events Page

Field	Description
Time Range	View a graphical representation of SD-WAN events against a defined time range. The x-axis represents the defined time and the y-axis represents SD-WAN events. Use the slider to decrease or increase the time range within which you want to view SD-WAN events. You can also choose from pre-defined time ranges such as 2h, 4h, 8h, 16h, 24h, or Custom. For custom time, you must enter from and to dates in MM/DD/YYYY format and the time in HH:MM:SS format. By default, Previous 1 day is selected.
SLA Violation Time	Date and time at which the SLA violation occurred.
Link Switch Time	Date and time at which the link was switched.
Site	Name of the site for which the link was switched.
Connected To	Displays the name of the destination spoke or hub site to which the traffic is being sent.
SLA Profile	Name of the SLA-based steering profile associated with the site.
Reason	Indicates the reason for the link switch. Mouse over the reason to view details of the SLA metrics violated.
Apps	Name of the applications for which the SLA violation occurred.
Department	Name of the department for which the SLA violation occurred.
Source Tunnel	Overlay tunnel of the device *from* which the link switch took place.
Destination Tunnel	Overlay tunnel of the device *to* which the link switch took place.
Duration (Sec)	Duration (in seconds) for which the SLA requirement for a site was not met before the site switched WAN links. A duration of 0 indicates that the site switched WAN links before it failed to meet the SLA requirements, and the SLA requirements were met immediately on the new WAN link with no loss in meeting SLA requirements.

RELATED DOCUMENTATION

SD-WAN Events Overview | 845

CHAPTER 30

Monitoring Applications

IN THIS CHAPTER

- About the SLA Performance of a Single Tenant Page | 848
- Viewing the SLA Performance of a Site | 851
- Viewing the SLA Performance of an Application or Application Group | 855
- Application Visibility Overview | 857
- About the Application Visibility Page | 857
- About the User Visibility Page | 861
- Viewing Application or User Visibility Data for Specific Sites | 864

About the SLA Performance of a Single Tenant Page

To access this page, select **Monitor > Application SLA Performance > *Tenant-Name* SLA Performance** in the Customer Portal.

You can use the *Tenant-Name* SLA Performance page to view performance reports for all sites in a tenant. You can view the SLA performance of all sites that have met and all the sites that have not met the defined SLA target values for the specified time range. You can customize your view and also the time range for which you want to view the SLA performance.

Tasks You Can Perform

You can perform the following tasks from this page:

- View the SLA performance for all sites in the tenant that have met the defined SLA target values, without switching WAN links, for the specified time range.
- View the SLA performance for all sites in the tenant that have met the defined SLA target values, after switching WAN links, for the specified time range.
- View the SLA performance for all sites in a tenant that have not met the defined SLA target values for the specified time range.
- View the SLA performance for all sites in a tenant in grid or card views.

Select card view or grid view at the top right of the page. By default, card view is selected.

- Customize the time range to view the SLA performance for all sites in a tenant.
- View the SLA performance for multiple departments within a single tenant.

  Select the specific department for which you want to view the SLA performance from the drop-down list at the top right of the page.

Field Descriptions

Table 299 on page 849 describes the fields on the *Tenant-Name* SLA Performance page.

Table 299: Fields on the SLA Performance of a Single Tenant Page

Field	Description
Time range	The time range for which you want to view the SLA performance. You can choose from Previous 1 hour, Previous 1 day, Previous 1 week, Previous 1 month, and Custom. For custom time, you must enter from and to dates in MM/DD/YYYY format and the time in HH:MM:SS format. By default, **Previous 1 day** is selected.
View	The view in which you want to display the SLA performance for all sites in the tenant. You can choose between card and grid views. By default, card view is selected.
Sites Not Meeting SLAs	The sites that did not meet the defined SLA target values in the selected time range. Click each site to view more information about the SLA performance of the applications and application groups in the site. See “Viewing the SLA Performance of a Site” on page 851.
Sites Meeting SLAs With Switch	The sites that switched WAN links to meet the defined SLA target values in the selected time range. Click each site to view more information about the SLA performance of the applications and application groups in the site. See “Viewing the SLA Performance of a Site” on page 851.
Sites Meeting SLAs Without Switch	The sites that met the defined SLA target values in the selected time range without switching WAN links. Click each site to view more information about the SLA performance of the applications and application groups in the site. See “Viewing the SLA Performance of a Site” on page 851.

Table 300 on page 850 describes the fields in the card and grid views.

Table 300: Fields on the SLA Performance of a Single Tenant Page in Card and Grid Views

Field	View	Description
Name	Card and Grid	View the name of the site.
SLA not met (Time)	Card and Grid	View the average time (in %) during which all the sites in a tenant did not meet the defined SLA target values.
Profiles	Card	View the time (in %) during which defined SLA target values were not met for each SLA profile. The top two profiles with highest priority and the percentage of time during which SLA target values were not met are listed. The remaining profiles and their combined sum of time (in %) for which SLA target values were not met are listed under **Others**. The SLA profile priority is indicated inside a circle circle. You can define priority of the SLA profile when you create an SLA profile. Hover over the profile priority to view the SLA profile name.
Profile SLA Not Met	Grid	
App - Groups	Card and Grid	View the total number of applications and application groups in the site.
Switch Events	Card and Grid	View the number of times the site switched WAN links over the number of designated WAN links. A switch event, also called SD-WAN event, occurs when a site switches WAN links to meet the SLA requirements.
Switch Events Per Profile	Card and Grid	View the number of times the site switched WAN links for each profile. You can view the switch events for the top two SLA profiles in the decreasing order of switch events for each profile.

RELATED DOCUMENTATION

- Viewing the SLA Performance of a Site | 851
- Viewing the SLA Performance of an Application or Application Group | 855
- SD-WAN Events Overview | 845
- Adding SLA-Based Steering Profiles | 517
- Adding Path-Based Steering Profiles | 528

Viewing the SLA Performance of a Site

IN THIS SECTION

- SLA Not Met by SLA Profiles | 851
- Applications SLA Performance by Throughput | 852
- SLA Performance for ALL | 854

You can use the **Monitor > Applications > *Tenant_name* SLA Performance > *Site_name* SLA Performance** page in the Customer Portal to view the SLA performance for all applications and application groups in a site. You can view the SLA performance for all applications and application groups in a site for a specified time range and in graph or grid views.

The ***Site_name* SLA Performance** page is divided into the following sections:

SLA Not Met by SLA Profiles

You can use the **SLA Not Met by SLA Profiles** section on the ***Site_name* SLA Performance** page to view the SLA profiles for which SLA requirements were not met and the time at which they were not met. The y-axis represents the SLA profiles and the x-axis represents the specified time range. The **SLA Not Met by SLA Profiles** section can be viewed and remains the same in both graph and grid views.

To view a graphical representation of SLA profiles for which SLA target values were not met:

1. Select the time range for which you want to view the SLA profiles for which SLA target values were not met. You can choose from Previous 1 hour, Previous 1 day, Previous 1 week, Previous 1 month, and Custom. For custom time, you must enter from and to dates in MM/DD/YYYY format and the time in HH:MM:SS format. By default, Previous 1 day is selected.

The graphical representation of SLA profiles for which SLA target values were not met is displayed for the selected time range.

2. (Optional) You can use the sliders at the sides of the graph to further customize the time range.

   The graphical representation of SLA profiles for which SLA target values were not met is refreshed and displayed for the customized time range. The graphical representation of SLA performance data in the subsequent sections on the page is also refreshed and displayed for the customized time range.

Applications SLA Performance by Throughput

You can use the **Applications SLA Performance by Throughput** section on the ***Site_name*** **SLA Performance** page to view average throughput performance of all applications and application groups in a site. You can also customize your view by selecting graph or grid views. In the graph view, you can further select scatter plot or tree map.

To view a graphical representation of average throughput performance of all applications and application groups in a site:

1. Select **Graph View** at the top right of the page. By default, Graph View is selected.

   A graphical representation of average throughput performance of all applications and application groups in a site against the target throughput is displayed in the **Scatter Plot** view. The y-axis represents the average throughput. 0% on the x-axis represents the target throughput (in %) defined in the SLA profiles, while the regions on the left and right of the target represent percentages below and above the target throughput, respectively.

   A carousel at the bottom of the section also displays the list of all applications and application groups with their SLA profiles, target throughput, and average throughput values.

2. Click **Legend** at the bottom right of the section to view the plotting legend.

   The items described in the **Legend** are:

   - A single application is represented by a blue circle.
   - An application group is represented by a blue square.
   - An application or application group whose target throughput value in the SLA profile was modified during runtime is represented by an uncolored circle and uncolored square, respectively.
   - The SLA profiles are represented by their priority numbers within the colored or uncolored circles and squares.

3. (Optional) You can use the sliders at the sides of the graph further to customize the time range.

   The carousel is refreshed for the customized time range.

4. Click the circles or squares to view more information about the application or application groups. See "Viewing the SLA Performance of an Application or Application Group" on page 855.

5. Select **Tree Map** at the top right of the section to view a list of all applications and application groups in a site and their average throughput values.

   A list of all applications and application groups in a site along with their associated SLA profiles and the average throughput values is displayed.

To view a tabular representation of average throughput performance of all applications and application groups in a site:

1. Select **Grid View** at the top right of the page.

   A list of all applications and application groups along with their SLA profiles, average throughput, and target throughput values is displayed in a tabular format.

   Table 301 on page 853 describes the fields on the Applications SLA Performance by Throughput grid view.

   Table 301: Fields on the Applications SLA Performance by Throughput Grid View

Field	Description
Name	View name of the application or application group.
SLA Profile	View the SLA profile associated with the application or application group.
Type	View the type—application or application group
Category	View the category of the application or application group. The value of Category can be Messaging, Web, Infrastructure, Remote-Access, Multimedia, Video, and so on.
Sessions	View number of sessions consumed by the application or application group.
Throughput Avg. Performance	View the average throughput performance value (in %) of the application or application group. The upward triangle on the left of the average throughput performance value indicates that the average throughput is higher than the target throughput configured in the SLA profile of the application or application group. The value (in %) denotes the percentage above the target throughput value. Similarly, the downward triangle on the left of the average throughput performance value indicates that the average throughput is lower than the target throughput configured in the SLA profile of the application or application group. The value (in %) denotes the percentage below the target throughput value.

2. (Optional) Click the details icon to the left of the application or application group name to view more information about the application or application group. See "Viewing the SLA Performance of an Application or Application Group" on page 855.

SLA Performance for ALL

View a graphical representation of the performance of the SLA parameters such as round-trip time (RTT), latency, packet loss, and jitter for the specified time range for MPLS and Internet WAN links for all SLA profiles. The y-axis represents the SLA parameters and the x-axis represents the specified time range. You can also view the respective target SLA parameters in the graphs.

> **NOTE:** The graphical representation of the performance of all SLA parameters for the WAN links is available only in the graph view.

To view a graphical representation of the performance of all SLA parameters for the WAN links:

- Select **All** at the top right of the section. By default, All is selected.

  A graphical representation of the performance of the SLA parameters such as RTT, latency, packet loss, and jitter for the specified time range for all WAN links is displayed.

- Select **wan_0**, **wan_1**, and so on at the top right of the section to view the performance of the SLA parameters for the MPLS and Internet WAN links. You can enable and configure **wan_0**, **wan_1**, and so on and map them to MPLS or Internet links when you create a site.

  The graphical representation of the performance of the SLA parameters such as RTT, latency, packet loss, and jitter for the specified time range is refreshed and only the performance for the selected WAN link is displayed.

- (Optional) Click **Legend** at the bottom right of the section to view the plotting legend for the horizontal dotted lines parallel to the x-axis in the graphs. The horizontal dotted lines represent the respective target SLA parameters of the SLA profiles.

RELATED DOCUMENTATION

About the SLA Performance of a Single Tenant Page | 848

Viewing the SLA Performance of an Application or Application Group | 855

Viewing the SLA Performance of an Application or Application Group

You can use the **Monitor > Applications > *Tenant-Name* SLA Performance > *Site-Name* SLA Performance** page in the Customer Portal to view the SLA performance for individual applications and application groups in a site. You can also view the SLA performance of the associated SLA profile for all SLA parameters.

To view SLA performance of an application or application groups:

- Click one of the circles or squares in the **Applications SLA Performance by Throughput** section on the ***Site-Name* SLA Performance** page.

  The page that appears displays SLA performance details of the application or application group.

Table 302 on page 855 describes the fields on the application or application group SLA Performance details page.

Table 302: Fields on the Application or Application Group Details Page

Field	Description
Category and Description	View the category of the application or application group. The category can be Messaging, Web, Infrastructure, Remote-Access, Multimedia, Video, and so on. You can also view a description of the application or application group.
SLA	View the name of the SLA profile associated with the application or application group.
Target	View the current target throughput defined in the SLA profile associated with the application or application group. If the target throughput was modified during runtime, the date and time when the throughput was modified and the previously defined throughput value are also displayed.
Avg. Performance	View the average throughout performance (in %) above or below the configured target throughput. The average throughput (in Mbps) is displayed within parentheses.
SLA Metrics by Throughput	View a graphical representation of the SLA metrics by throughput during the specified time range for that application or application group. The y-axis represents the throughput (in Mbps). The x-axis represents the specified time range. Hover over the graph to view the throughput value and time at any specified point. You can also view the sessions consumed by the WAN links for the application or application group time range.

Table 302: Fields on the Application or Application Group Details Page *(continued)*

Field	Description
Global SLA Profile Performance	View the performance for all the SLA parameters of the SLA profile associated with the application or application group. The SLA performance is represented by a color-coded donut chart. The section in blue in the donut chart indicates the percentage of time during which SLA requirements for the SLA profile were met. The section in red in the donut chart indicates the percentage of time during which SLA requirements for the SLA profile were not met. Click the red colored section of the donut chart to view more information about when SLA requirements for the SLA profile were not met. The **SLA Profile Performance** page appears. The SLA Profile Performance page displays the following fields: • SLA Profile—SLA profile associated with the application or application group • Target—Target throughput configured in the SLA profile • SLAs Not Met—Percentage of time SLA requirements were not met for the SLA profile • Sessions—Number of sessions consumed by the application or application group • Start Time—Time at which the WAN links associated with the application or application groups started to fail meeting the SLA requirements • End Time—Time at which SLA profile requirements started to be met again • Avg Val—Average throughput (in Mbps) when the SLA requirements started to fail • Duration—Total duration (in seconds) during which SLA requirements were not met • From—Source WAN link • To—Destination WAN link

RELATED DOCUMENTATION

About the SLA Performance of a Single Tenant Page | 848

Viewing the SLA Performance of a Site | 851

Application Visibility Overview

You can use the **Application Visibility** page to view information about bandwidth consumption, session establishment, and the risks associated with your applications.

Analyzing your network applications yields useful security management information, such as abnormal applications that can lead to data loss, heavy bandwidth usage, time-consuming applications, and personal applications that can elevate business risks.

RELATED DOCUMENTATION

About the Application Visibility Page | **857**

Viewing Application or User Visibility Data for Specific Sites | **864**

About the Application Visibility Page

To access this page, select **Monitor** > **Applications** > **Visibility**.

There are two ways in which you can view your application visibility data—**Chart View** or **Grid View**. By default, the data is displayed in **Chart View**.

Tasks You Can Perform

You can perform the following tasks from this page:

- View application visibility data in **Chart View**. See "Chart View" on page 857.
- View application visibility data in **Grid View**. See "Grid View" on page 859.
- Select a device to which the application visibility settings are applicable. See "Viewing Application or User Visibility Data for Specific Sites" on page 864.

Chart View

Click the **Chart View** link for a brief summary of the top 50 applications consuming the maximum bandwidth in your network. The data can be presented graphically as a bubble graph, heat map, or a zoomable bubble graph. The data is refreshed automatically based on the selected time range. You can also use the **Custom** button to set a custom time range.

You can hover over your applications to view critical information such as total number of sessions, total number of blocks, category, bandwidth consumed, risk levels, and characteristics. You can also view the top five users accessing your application.

Table 303 on page 858 provides guidelines on using the fields on the **Chart View** of the **Application Visibility** page.

Table 303: Fields on the Chart View

Field	Description
All Devices	Displays application visibility data for all the sites managed by CSO. Click **Edit** to select individual devices for which you want to view the data.
Show By	Select from the following options to view a user's data: • Bandwidth—Shows data based on the amount of bandwidth the application has consumed for a particular time range. • Number of Sessions—Shows data based on the number of sessions consumed by the application.
Time Span	Select the required time range to view a user's data. Use the custom option to choose the time range if you want to view data for more than one day. The time range is from 00:00 through 23:59.
Select graph	Select from the following graphical representations to view an application's data: • Bubble Graph • Heat Map • Zoomable Bubble Graph By default, data is shown in the Bubble Graph format.
Group By	Select from the following options to view the application's data: • Risk—Grouped by critical, high, unsafe, and so on. • Category—Grouped by categories such as web, infrastructure, and so on.
Number of Sessions	Displays the total number of application sessions.
Number of Blocks	Displays the total number of times the application was blocked.
Bandwidth	Displays the bandwidth usage of the application.
Risk Level	Displays the risk associated with the application. For example, critical, high, unsafe, and so on.
Category	Displays the category of the application. For example, web, infrastructure, and so on.

Table 303: Fields on the Chart View *(continued)*

Field	Description
Characteristics	Displays the characteristics of the application. For example, prone to misuse, bandwidth consumer, capable of tunneling, and so on.

Grid View

Click the **Grid View** link to obtain comprehensive details about applications. You can view top users by volume, top applications by volume, top category by volume, top characteristics by volume, and sessions by risk. You can also view the data in a tabular format that includes sortable columns. You can sort the applications in ascending or descending order based on application name, risk level, and so on. Table 304 on page 859 describes the widgets in this view. Use these widgets to get an overall, high-level view of your applications, users, and the content traversing your network.

Table 304 on page 859 provides guidelines on using the fields on the **Grid View** of the **Application Visibility** page.

Table 304: Widgets on the Grid View

Field	Description
Top Users By Volume	Top users of the application; sorted by bandwidth consumption.
Top Apps By Volume	Top applications using the network traffic, such as Amazon, Facebook, and so on, sorted by bandwidth consumption.
Top Category By Volume	The top category of the application, such as Web, infrastructure, and so on; sorted by bandwidth consumption.
Top Characteristics By Volume	Top behavioral characteristics of the application, such as whether it is highly prone to misuse, the top bandwidth consumer, and so on.
Sessions By Risk	Number of events or sessions received; grouped by risk.

Table 305 on page 860 describes the fields in the table below the widgets. Users are displayed by usernames or IP addresses. When you click a link, the **User Visibility** page appears in a grid view, with the correct filter applied. Sessions are also displayed as links and when you click a link, the **All Events** page appears with all security events.

Table 305: Detailed View of Applications

Field	Description
Application Name	Name of the application, such as Amazon, Facebook, and so on.
Risk Level	Risk associated with the application: critical, high, unsafe, moderate, low, and unknown.
Users	Total number of users accessing the application.
Volume	Bandwidth used by the application.
Total Sessions	Total number of application sessions.
No of Rejects	Total number of sessions blocked.
Category	Category of the application, such as Web, infrastructure, and so on.
Sub Category	Subcategory of the application. For example, social networking, news, and advertisements.
Characteristics	Characteristics of the application. For example, prone to misuse, bandwidth consumer, capable of tunneling.

RELATED DOCUMENTATION

- Application Visibility Overview | **857**
- Viewing Application or User Visibility Data for Specific Sites | **864**
- About the SLA Performance of a Single Tenant Page | **848**

About the User Visibility Page

To access this page, select **Monitor** > **User Visibility**.

Use the User Visibility page to view information about devices (such as top 50 devices accessing high bandwidth-consuming applications and establishing higher number of sessions) on your network. Based on this information, network administrators can choose to rate-limit a device that is accessing applications which consume large bandwidth or create maximum traffic.

Tasks You Can Perform

You can perform the following tasks from this page:

- View user visibility data in **Chart View**. See "Chart View" on page 861.
- View user visibility data in **Grid View**. See "Grid View" on page 863.
- Select one or more sites for which you want to view the user visibility data. See "Viewing Application or User Visibility Data for Specific Sites" on page 864.

Chart View

Click the **Chart View** tab to view the data graphically as a bubble graph, heat map, or a zoomable bubble graph. The data is refreshed automatically based on the selected time span.

You can hover over the chart to view critical information (such as the total number of sessions established and bandwidth consumed) about each user.

Users are represented by the IP address or usernames of their devices on the network.

You can also view the top five applications of each user, based on either their bandwidth consumption or number of sessions established.

Table 306 on page 861 provides guidelines on using the fields on the **Chart View** tab of the **User Visibility** page.

Table 306: Fields on the Chart View

Field	Description
All Sites	By default, the chart displays user visibility data for all the sites managed by CSO. Click **Edit** to select one or more sites for which you want to view the user visibility data. See "Viewing Application or User Visibility Data for Specific Sites" on page 864 for more information.

Table 306: Fields on the Chart View *(continued)*

Field	Description
Show By	Select the criterion to display information regarding the bandwidth consumed and number of sessions established by applications in the selected time span: • Bandwidth—Displays users based on their bandwidth consumption. Users running applications that consume larger bandwidth are represented by larger bubbles or matrices. • Number of Sessions—Displays users based on the number of sessions established. Users running applications that have higher number of sessions established are represented by larger bubbles or matrices.
Time Span	Select the duration (last 15 minutes, last 30 minutes, last 45 minutes, last 1 hour, last 4 hours, last 8 hours, last 12 hours, last 1 day, or custom) for which you want to view the user visibility data. Select **Custom** to view data for more than one day. The **Custom Time** page appears. Specify the **From date** and **To date** (in MM/DD/YYYY format).The time span is from 00:00 through 23:59.
Select Graph	Select one of the following options to view data graphically: • Bubble Graph (default) • Heat Map • Zoomable Bubble Graph

Table 307 on page 862 describes the parameters that are displayed when you hover your cursor over the chart.

Table 307: Parameters on the Chart

Parameter	Description
Number of Sessions	Total number of application sessions established by the user (device).
Bandwidth	Total Bandwidth consumed by the user (device).
View All Applications	Click the *View All Applications* link to view details (such as risk level and category) of all the applications on the network. The Application Visibility page in grid view appears. See "About the Application Visibility Page" on page 857 for more information.

Grid View

Click the **Grid View** tab to view high-level details of the users on your network. You can view widgets that provide information about top users by volume and top applications that create network traffic by volume. The data is also displayed in a tabular format with sortable columns.

Table 304 on page 859 describes the widgets on the **Grid View** of the **User Visibility** page.

Table 308: Widgets on the Grid View

Field	Description
Top Users by Volume	Top users of applications, based on bandwidth consumption, for the selected time span.
Top Apps by Volume	Top applications accessed by users on the network, based on bandwidth consumption, for the selected time span. For example: Amazon

Table 309 on page 863 describes the fields in the table below the widgets.

The table includes sortable columns, with the users (devices) represented by usernames or IP addresses.

Click the Search icon to enter the search text that can include a specific application or user name, or IP address of a device on the network.

The search results are displayed. Click **Clear All** to clear the search results.

Table 309: Detailed View of Users

Field	Description
Applications	Name of the application accessed by a specific user (device). For example: Google **NOTE:** By default, this column lists only one application per user. If a user accesses more than one application, a +<integer>icon (for example: +2) appears to the right of the application name. The integer indicates the number of additional applications accessed by the user. Click the integer to view all applications accessed by a user.
User Name	IP address or username of the user (device) accessing the applications.
Volume	Bandwidth consumed by a user (who is represented by a user name or IP address).
Total Sessions	Total number of application sessions established by a specific user (device).

Viewing Application or User Visibility Data for Specific Sites

IN THIS SECTION

- Viewing Application Visibility Data for Specific Sites | 864
- Viewing User Visibility Data for Specific Sites | 865

You can select one or more sites for which you want to view application visibility or user visibility data (such as bandwidth consumption and number of sessions). By default, the application visibility and user visibility data is displayed for all sites in a tenant.

Viewing Application Visibility Data for Specific Sites

To select the sites for which you want to view the application visibility data:

1. Select **Monitor > Application Visibility**.

   The **Application Visibility** page appears.

2. Click the **Edit** link (next to the Show By field).

   The **Select Sites** page appears.

3. From the Sites field:

   - Click **Selective** to select the sites for which you want to view the application visibility data:

     The available sites are displayed in the **Available** column.

   - Click **All** to view application visibility data for all sites in the tenant.

     If you click All, proceed to step 5.

4. Select the sites from the **Available** column and click the right arrow to move them to the **Selected** column.

5. Click **OK** to save your changes.

   You are returned to the Application Visibility page and the application visibility data is displayed for the sites that you selected.

   To view application visibility data for other sites, repeat step 2.

Viewing User Visibility Data for Specific Sites

To select the sites for which you want to view the user visibility data:

1. Select **Monitor > User Visibility**.

   The **User Visibility** page appears.

2. Click the **Edit** link (next to the Show By field).

   The **Select Sites** page appears.

3. From the Sites field:

   - Click **Selective** to select the sites for which you want to view the user visibility data.

     The available sites are displayed in the **Available** column.

   - Click **All** to view user visibility data for all sites in the tenant.

     If you click All, proceed to step 5.

4. Select the sites from the **Available** column and click the right arrow to move them to the **Selected** column.

5. Click **OK** to save your changes.

   You are returned to the User Visibility page and the user visibility data is displayed for the sites that you selected.

   To view user visibility data for other sites, repeat step 2.

CHAPTER 31

Monitoring Threats

IN THIS CHAPTER

- About the Threats Map (Live) Page | 866

About the Threats Map (Live) Page

IN THIS SECTION

- Tasks You Can Perform | 867
- Field Descriptions | 868
- Threat Types | 869

To access this page, select **Monitor > Threats Map (Live)** in Customer Portal.

Use this page to visualize incoming and outgoing threats between geographic regions. You can view blocked and allowed threat events based on feeds from intrusion prevention systems (IPS), antivirus, and antispam engines, unsuccessful login attempts, and screen options. You can also click a specific geographical location to view the event count and the top five inbound and outbound IP addresses.

The threat data is displayed starting from 12:00 AM (midnight) up to the current time (in your time zone) on that day and is updated every 30 seconds. The current date and time is displayed at the top right and a legend is displayed at the bottom left of the page.

If a threat occurs when you are viewing the page, an animation shows the country from which the threat originated (source) and the country in which the threat occurred (destination).

NOTE: For threats with unknown geographical IP addresses (private IP addresses), the animation shows the threat originating from the bottom center of the geographical map.

Tasks You Can Perform

You can perform the following tasks from this page:

- Toggle between updating the data and allowing live updates—Click the **Pause** icon to stop the page from updating the threat map data and to stop animations. Click the **Play** icon to update the page data and resume animations.
- Zoom in and out of the page—Click the zoom in (**+**) and zoom out (**–**) icons to zoom in and out of the page.
- Pan the page—Click and drag the mouse to pan the page.
- View country-specific details:
  - Click a country on the threat map to view threat information specific to that country. A *Country-Name* pop-up appears displaying country-specific information.
  - Click the **View Details** link in the *Country-Name* pop-up to view additional details. The *Country-Name* (Details) panel appears.

  For more information, see Table 310 on page 867.

Table 310: Country-Specific Threat Information

Field	Description	Displayed In
Number-of-threat-events **Threat Events since 12:00 am**	Displays the total number of threat events (inbound and outbound) since midnight for that country. Click the hyperlinked number to go to the All Events page, where you can view more information about the events.	*Country-Name* pop-up
Inbound (*Number-of-threat-events*)	Displays the total number of inbound threats for the country and the IP address and the number of events for that IP address for the top five inbound events.	*Country-Name* pop-up
Outbound (*Number-of-threat-events*)	Displays the total number of outbound threats for the country and the IP address and the number of events for that IP address for the top five outbound events.	*Country-Name* pop-up
Number-of-threat-events **Events since 12:00 am**	Displays the total number of threat events (inbound and outbound) since midnight for that country. Click the hyperlinked number to go to the All Events page, where you can view more information about the events.	*Country-Name* (Details) panel

Table 310: Country-Specific Threat Information *(continued)*

Field	Description	Displayed In
***Number-of* Inbound Events**	Displays the total number of inbound threats for the country and the number of inbound threat events for each of the following categories: • IPS Threats • Virus • Spam • Device Authentication • Screen Click the hyperlinked number for a category to go to the page for that category, where you can view more information about that category. For example, clicking the hyperlinked number for IPS threats takes you to the IPS Events page. Click the **Top 5 IP Addresses (Inbound)** to view the IP address and the number of events for that IP address for the top five inbound events.	*Country-Name* (Details) panel
***Number-of* Outbound Events**	Displays the total number of outbound threats for the country and the number of outbound threat events for each of the following categories: • IPS Threats • Virus • Spam • Device Authentication • Screen Click the hyperlinked number for a category to go to the page for that category, where you can view more information about that category. For example, clicking the hyperlinked number for screens takes you to the Screen Events page. Click the **Top 5 IP Addresses (Outbound)** to view the IP address and the number of events for that IP address for the top five outbound events.	*Country-Name* (Details) panel

Field Descriptions

Table 311 on page 869 displays the fields the Threats Map (Live) page.

Table 311: Fields on the Threats Map (Live) Page

Field	Description
Total Threats Blocked & Allowed	Displays the total number of threats blocked and allowed. Click the hyperlinked number to go to the All Events page (filtered view of the Detail View tab), where you can view more information about the IPS, virus, spam, device authentication, and screen events.
Threats Blocked & Allowed	Displays the total number of threats blocked and allowed by the following categories: • IPS Threats • Virus • Spam • Device Authentication • Screen Click the hyperlinked number for a category to go to the page for that category, where you can view more information about that category. For example, clicking the hyperlinked number for IPS threats takes you to the IPS Events page (filtered view of the Detail View tab).
Top Target Devices	Displays the top five targeted devices and the number of threats per device. Click the hyperlink for a device to go to the All Events page (filtered view of the Detail View tab), where you can view more information about the IPS, virus, spam, device authentication, and screen events for that device.
Top Destination Countries	Displays the top five destination countries and the number of threats per country. Click the hyperlink for a country to go to the All Events page (filtered view of the Detail View tab), where you can view more information about the IPS, virus, spam, device authentication, and screen events for that country.
Top Source Countries	Displays the top five source countries and the number of threats per country. Click the hyperlink for a country to go to the All Events page (filtered view of the Detail View tab), where you can view more information about the IPS, virus, spam, device authentication, and screen events for that country. **NOTE:** For threats with unknown geographical IP addresses (private IP addresses), the country name is displayed as *Undefined*. So, when you click the hyperlinked threat count and go to the All Events page, the filter query uses **Undefined** as the source country.

Threat Types

The Threats Map (Live) page displays blocked and allowed threat events based on feeds from IPS, antivirus, and antispam engines, unsuccessful login attempts, and screen options. Table 312 on page 870 describes different types of threats blocked and allowed.

Table 312: Types of Threats

Attack	Description
IPS threat events	Intrusion detection and prevention (IDP) attacks detected by the IDP module. The information reported about the attack (displayed on the IPS Events page) includes information about: • Source of attack • Destination of attack • Type of attack • Session information • Severity • Policy information that permitted the traffic. • Action: traffic permitted or dropped.
Virus events	Virus attacks detected by the antivirus engine. The information reported about the attack (displayed on the Antivirus Events page) includes information about: • Source of the infected file • Destination • Filename • URL used for accessing the file
Spam events	E-mail spam that is detected based on the blocklist spam e-mails. The information reported about the attack (displayed on the Antispam Events page) includes information about: • Source • Action: E-mail is rejected or allowed. • Reason for identifying as e-mail spam.
Device authentications	The firewall authentication messages generated due to unauthorized attempts to access the network. The reported information (displayed on the All Events page) contains the reason for authentication failure and the source of the request.

Table 312: Types of Threats *(continued)*

Attack	Description
Screen events	Events that are detected based on screen options. The information reported about the attack (displayed on the Screen Events page) includes information about: • Internet Control Message Protocol (ICMP) screening • IP screening • TCP screening • UDP screening

RELATED DOCUMENTATION

About the All Security Events Page | 813

8

PART

Managing Reports

Security Reports | **873**

SD-WAN Reports | **897**

CHAPTER 32

Security Reports

IN THIS CHAPTER

- Reports Overview | 873
- About the Security Report Definitions Page | 874
- Scheduling, Generating, Previewing, and Sharing Security Reports | 877
- About the Security Generated Reports Page | 880
- Creating Log Report Definition | 881
- Creating Bandwidth Report Definition | 885
- Creating ANR Report Definition | 887
- Editing, Deleting, and Cloning Log Report Definitions | 890
- Editing, Deleting, and Cloning Bandwidth Report Definitions | 892
- Editing, Deleting, and Cloning ANR Report Definitions | 894

Reports Overview

Reports are generated based on the summary of network activity (such as top web applications or viruses detected) and overall network status. To generate reports, you can use the predefined report definitions as is, or you can create custom report definitions that meet your needs for specific data.

The generated report contains a table of contents (TOC) with links to each section of the report. The designated recipients, whose e-mail addresses are included in the report definition, receive the report in PDF format.

You can generate two categories of reports:

Security reports—Provide information about network activity and network status.

SD-WAN reports—Provide information about SLA performance of all sites or specific sites in a tenant.

The following are the types of security reports:

- **Log Reports**—Enable you to analyze event history based on the data criteria (such as filters, aggregation criteria, time span, etc.) that you select.

- **Bandwidth Reports**—Enable you to analyze the bandwidth usage of an application or a user.
- **ANR Reports**—Enable you to analyze business risks in the network, based on application usage and resource usage.

The following are the types of SD-WAN reports:

- **SD-WAN Tenant Performance Reports**—Enable you to analyze tenant performance based on the parameters (top applications by bandwidth, top sites not meeting the SLA, top sites meeting the SLA with switching, sites meeting the SLA without switching, top sites by highest packet loss, and top sites by highest latency, top sites by highest jitter, and current active tunnels) that measure the SLA performance across all sites in a tenant.
- **SD-WAN Site Performance Reports**—Enable you to analyze site performance based on the parameters (top 10 applications and link utilization, top profiles not meeting the SLA, top SLA profiles switching links, top applications by highest packet loss, top applications by highest latency, top applications by highest jitter, SLA performance between two sites, and tunnels created and deleted) that measure the SLA performance of specific sites in a tenant. You can select a maximum of five sites for which you want to generate the report.

RELATED DOCUMENTATION

About the Security Report Definitions Page | 874
About the SD-WAN Report Definitions Page | 897

About the Security Report Definitions Page

To access this page, click **Customer Portal > Reports > Report Definitions > Security**.

The **Security Report Definitions** page appears.

The Security Report Definitions page displays a list of predefined and custom report definitions. To generate reports, you can use the predefined report definitions as is, or you can create custom report definitions.

> **NOTE:** From CSO Release 4.1.0 onward, an Application and Network Risk (ANR) report is the only predefined report definition available on the Security Report Definitions page.
>
> The ANR report provides information about data usage and network risks. The information is consolidated from various predefined report definitions available in the release prior to CSO Release 4.1.0.

Tasks You Can Perform

You can perform the following tasks from this page:

- Create a report definition:
  - To create a log report definition, see "Creating Log Report Definition" on page 881.
  - To create a bandwidth report definition, see "Creating Bandwidth Report Definition" on page 885.
  - To create an application and network risk report definition, see "Creating ANR Report Definition" on page 887
- Edit, delete, or clone report definitions:
  - To edit, delete, or clone a log report definition, see "Editing, Deleting, and Cloning Log Report Definitions" on page 890.
  - To edit, delete, or clone a bandwidth report definition, see "Editing, Deleting, and Cloning Bandwidth Report Definitions" on page 892.
  - To edit, delete, or clone an application and network risk report definition, see "Editing, Deleting, and Cloning ANR Report Definitions" on page 894.
- To schedule, generate reports, preview reports as PDF, and send the reports through e-mail, see "Scheduling, Generating, Previewing, and Sharing Security Reports" on page 877.
- To view all the parameters of a report definition, right click the report definition that you want to see the detailed view for and select **Detailed View**, or select the report definition and click **More** > **Detailed View**. Alternatively, hover over the report definition name and click the Detailed View icon that appears before it.

  The Report Definition Details page appears, displaying the same values that you specified for each parameter in the selected report definition.
- To search for a report definition from the list of available report definitions, click the **Search** icon in the top right corner of the page.

  Enter the name of the report definition in the search bar and click the Search icon.

  The search results are displayed.

  Click **Clear All** to clear the search results.

Field Descriptions

Table 313 on page 876 describes the fields on the Security Report Definitions page.

Table 313: Fields on the Security Report Definitions Page

Field	Description
Name	Name of the report definition.
Description	Description of the report definition. For example: Report of data usage by application and network risk.
Type	Type of report definition—ANR, Bandwidth, or Log.
Definition Type	Indicates whether the report definition is predefined (system-generated) or custom (user-created).
Report Content	Details of the sections in the report. For example: Count, Time Duration.
Schedule	Indicates whether the report generation is scheduled at the current time (Now) or for a later date and time (Once).
Recipients	E-mail addresses of recipients to whom the generated report is sent.
Last Generated	Date and time when the report was last generated.
Job ID	Use the job ID to view the status of the task on the Jobs (Monitor > Jobs) page.

RELATED DOCUMENTATION

- Reports Overview | 873
- Creating Log Report Definition | 881
- Creating Bandwidth Report Definition | 885
- Creating ANR Report Definition | 887

Scheduling, Generating, Previewing, and Sharing Security Reports

IN THIS SECTION

- Editing Report Generation Schedule | 877
- Generating Reports | 878
- Previewing Reports in PDF | 879
- Sharing Reports through E-mail | 879

You can schedule report generation, generate reports, preview reports as PDF, and share the reports through e-mail.

To perform these actions on a report definition:

1. Select **Reports** > **Report Definitions > Security**.

   The Security Report Definitions page appears.

2. Select or right-click the report definition on which you want to perform an action and click **More**.

   A list of actions that you can perform on the report definition is displayed.

3. Select the appropriate action from the list:

Editing Report Generation Schedule

You can edit the report generation schedule of the selected report definition from the Security Report Definitions page:

1. Select the report definition for which you want to edit the report generation schedule.

2. Click **More** > **Edit Schedule**. Alternatively, right-click on the selected report definition and select **Edit Schedule**.

   The Edit Report Schedule page appears.

3. Specify whether you want to generate the report at the current time or schedule it for a later date and time:

   - Run now—Select this option to schedule the report generation at the current time.

- Schedule at a later time—Select this option to schedule the report generation for a later date and time in MM/DD/YYYY and HH:MM:SS formats.

4. Click **OK** to save your changes.

   You are returned to the Security Report Definitions page on which a confirmation message, indicating that the report generation schedule is updated successfully, appears.

Generating Reports

You can generate a report at the current time, with either saved settings or custom settings. You can select Saved Settings to generate a report based on the values specified in the report definition for the selected report or select Custom Settings to modify the values for the Number of Top Logs and the Time Span settings, and generate a report based on the modified values.

> **NOTE:** The modified values are applicable only for the report that is being generated and are not saved in the report definition.

To generate a report:

1. From the Security Report Definitions page, select the report definition based on which you want to generate the report.

2. Click **Run Now** on the Security Report Definitions page. Alternatively, click **More** > **Run Now** or right-click on the report definition and click **Run Now**.

   The Run Report page appears.

3. Do one of the following:

   - Select **Saved Settings** to generate a report based on the values already specified in the report definition.
   - Select **Custom Settings** to modify the values for the Number of Top Logs and Time Span settings, and generate a report based on the modified values.

     > **NOTE:** The values that you modify are applicable only for the report that is being generated and are not saved in the report definition.

4. Click **OK** to save your changes.

The Run Report page appears indicating the progress of the report generation. After the report is generated, the **Download PDF Report** link appears on the Run Report page.

5. Click **Download PDF Report**.

   Follow your browser instructions to view or save the report in PDF.

Previewing Reports in PDF

You can preview and download the selected report in PDF:

1. Select the report definition based on which you want to generate the PDF of the report.

2. Click **More** > **Preview as PDF**. Alternatively, right-click on the report definition and select **Preview as PDF**.

   The Preview as PDF page appears.

3. Click **Download PDF Report** to view the report in PDF. or click **Cancel** to cancel previewing the report.

   The Security Report Definitions page appears.

Sharing Reports through E-mail

You can share the generated report through e-mail:

1. Select the report definition based on which the report is to be generated and shared through e-mail.

2. Click **More** > **Send Report**. Alternatively, right-click on the report definition and select **Send Report**.

   The Edit Recipients page appears:

   - Recipients—Enter or select one or more e-mail addresses of users to whom you want to send the report.

     By default, you can search by first name and select registered users. You can also enter external e-mail addresses.

   - Subject—Enter the subject line for the e-mail. The maximum length is 2048 characters.
   - Comments—Enter the text to be included in the body of the e-mail.

     The maximum length allowed is 2048 characters.

3. Click **OK** to save your changes.

   The Security Report Definitions page appears.

RELATED DOCUMENTATION

About the Security Report Definitions Page | 874

About the Security Generated Reports Page

To access this page, click **Customer Portal** > **Reports** > **Generated Reports** > **Security**.

Use this page to view the list of reports that are generated from the Security Report Definitions page. You must click on the report to view the report in PDF.

You can also delete one or more generated reports.

Field Descriptions

Table 314 on page 880 describes the fields on the Generated Reports page.

Table 314: Fields on the Generated Reports Page

Field	Description
Report PDF Name	Type of the report (user created or predefined).
Generated Time	Date and time when the report was generated.
Description	Description of the report.
Definition Name	Name of the report definition.
Generated By	Name of the user who generated the report.
Recipients	Recipients of the generated report.

RELATED DOCUMENTATION

Reports Overview | 873
About the Security Report Definitions Page | 874

Creating Log Report Definition

Use the Create Log Report Definition page to create log report definitions and generate the corresponding log reports.

Log reports are generated based on the data criteria, which are derived from one or more filters that you select. These reports help you to analyze business risks based on logs from services such as unified threat management (UTM) and firewalls.

To create a log report definition:

1. Select **Reports > Report Definitions > Security**.

   The Security Report Definitions page appears.

2. Click **Add > Log Report Definitions**.

   The Create Log Report Definition page appears.

3. Complete the configuration according to the guidelines provided in Table 315 on page 881.

   NOTE: Fields marked with * are mandatory.

4. Click **OK** to save the log report definition.

   The report definition is saved and the Security Report Definitions page appears.

   A confirmation message appears on this page, indicating that the log report definition was successfully created.

   You can perform various actions on the report definition. See "Scheduling, Generating, Previewing, and Sharing Security Reports" on page 877.

Table 315: Fields on the Create Log Report Definition Page

Field	Description
General	
Report Name	Enter a unique name for the report definition. The name can contain a string of alphanumeric characters and some special characters (colons, periods, dashes, and underscores); no spaces are allowed and the maximum length allowed is 63 characters.

Table 315: Fields on the Create Log Report Definition Page *(continued)*

Field	Description
Description	Enter a description for the report definition; the maximum length (including spaces) allowed is 1024 characters.
Content	
Data Criteria	Click **Filters** to select one or more filters. The **Use Data Criteria From Filter** page appears. The list of default and custom filters, which are saved from the Security Events page, is displayed in a tabular format. The table displays the Filter Name, Filter Description, Time Span, and Grouping and Filtering criteria for each filter. Select one or more filters from the list as per your requirement, and click **OK**. The Create Log Report Definition page appears. When you select one or more filters, new fields appear on the Create Log Report Definition page. The fields are populated with values from the filters. You can either retain the values or change the values if needed. See Table 316 on page 883 for an explanation of the fields.
Schedule	
Schedule Report	Click **Add Schedule** to schedule the report generation. The **Add Report Schedule** page appears. Specify whether you want to generate the report immediately or schedule it for a later date and time: • Run now—Select this option to schedule the report generation at the current time, and click **OK**. • Schedule at a later time—Select this option to schedule the report generation for a later date and time (in MM/DD/YYYY and HH:MM:SS formats) and click **OK**. The Create Log Report Definition page appears with details of the report generation schedule.
E-Mail	

Table 315: Fields on the Create Log Report Definition Page *(continued)*

Field	Description
E-Mail Recipients	Click **Add Email Recipients** to add e-mail addresses of recipients to whom you want to send the log report. The Add Recipients page appears. • Recipients—Enter or select one or more e-mail addresses of users to whom you want to send the report. By default, you can search by first name and select registered users. You can also enter external e-mail addresses (e-mail addresses that are not registered with CSO). • Subject—Enter the subject line for the e-mail that is sent with the generated report. The maximum length allowed is 2048 characters. • Comment—Enter the text to be included in the body of the e-mail that is sent with the generated report. The maximum length allowed is 2048 characters.

Table 316 on page 883 displays the additional fields that appear on the Create Log Report Definition page when you select one or more filters.

Table 316: Additional Fields on the Create Log Report Definition Page

Section	Section number in the log report for a selected filter. Click **Delete Section** to remove the section and the corresponding filter.
Section Title	Name of the section in the log report. The section title is based on the selected filter.
Section Description	Description for the section in the log report.
Group By	Criteria, such as **Nested Application**, based on which logs are aggregated. You can select a maximum of two data criteria from the **Group By** drop-down list.
Time Span (Last)	Duration for which the report is to be generated. The default time span is 3 hours. You can specify the duration in minutes, hours, days, weeks, months, or specify a custom duration.

Table 316: Additional Fields on the Create Log Report Definition Page *(continued)*

	If you select **Custom**, the **Custom Time Range Selection** page appears. You must specify the **From** date and time, and **To** date and time (in MM/DD/YYYY and HH:MM:SS formats).
Filter By	Filter criteria (such as filtering applications based on http and https protocols) based on which the log report is to be generated. You can use AND, OR, Equal to (=), and Not Equal to (!=) logical operators as values to generate the report. For example: If you want to generate a report with the event category as antivirus and event name as AV_VIRUS_Detected_MT, then the value must be: Event Category = antivirus AND Event Name = AV_VIRUS_DETECTED_MT
Chart	Type of chart to graphically present data on the report. The available options are Bar (default), Comparison Bar, Timeline, Grid, Grouped Grid, Donut, and Bubble chart.
Number of Top Logs	Specify the number of records that you want to retrieve and display for each section in the report. Range: 1 through 20. Default: 10.

RELATED DOCUMENTATION

About the Security Report Definitions Page | **874**

Creating Bandwidth Report Definition | **885**

Creating ANR Report Definition | **887**

Creating Bandwidth Report Definition

You can use the Create Bandwidth Report Definition page to create bandwidth report definitions and generate the corresponding bandwidth reports. Bandwidth reports are used to analyze the bandwidth usage of an application or a user.

To create a bandwidth report definition:

1. Select **Reports > Report Definitions > Security**.

   The Security Report Definitions page appears.

2. Click **Add > Bandwidth Report Definitions**.

   The Create Bandwidth Report Definition page appears.

3. Complete the configuration according to the guidelines provided in Table 317 on page 885.

   NOTE: Fields marked with * are mandatory.

4. Click **OK** to save the bandwidth report definition.

   The report definition is saved and the Security Report Definitions page appears.

   A confirmation message appears on this page, indicating that the bandwidth report definition was successfully created.

   You can perform various actions on the report definition. See "Scheduling, Generating, Previewing, and Sharing Security Reports" on page 877.

Table 317: Fields on the Create Bandwidth Report Definition Page

Field	Description
General	
Report Name	Enter a unique name for the report definition. The name can contain a string of alphanumeric characters and some special characters (colons, periods, dashes, and underscores); no spaces are allowed and the maximum length allowed is 63 characters.
Description	Enter a description for the report definition; the maximum length (including spaces) allowed is 1024 characters.

Table 317: Fields on the Create Bandwidth Report Definition Page *(continued)*

Field	Description
Content	
Number of Top Logs	Specify the number of records that you want to retrieve and display for each section in the report. Range: 1 through 20. Default: 10.
Time Span (Last)	Specify the duration (Custom, last 3 hours, last 6 hours, last 12 hours, or last 24 hours) for which you want the report to be generated. If you select **Custom**, the **Custom Time Range Selection** page appears. You must specify the **From** date and time, and **To** date and time (in MM/DD/YYYY and HH:MM:SS formats).
Schedule	
Add Schedule	Click **Add Schedule** to schedule the report generation. The **Add Report Schedule** page appears. Specify whether you want to generate the report immediately or schedule it for a later date and time: • Run now—Select this option to schedule the report generation at the current time, and click **OK**. • Schedule at a later time—Select this option to schedule the report generation for a later date and time (in MM/DD/YYYY and HH:MM:SS formats), and click **OK**. The Create Bandwidth Report Definition page appears with details of the report generation schedule.
E-Mail	

Table 317: Fields on the Create Bandwidth Report Definition Page *(continued)*

Field	Description
Add E-Mail Recipients	Click **Add Email Recipients** to add e-mail addresses of recipients to whom you want to send the Bandwidth report. The Add Recipients page appears. • Recipients—Enter or select one or more e-mail addresses of users to whom you want to send the report. By default, you can search by first name and select registered users. You can also enter external e-mail addresses. • Subject—Enter the subject line for the e-mail that is sent with the generated report. The maximum length is 2048 characters. • Comment—Enter the text to be included in the body of the e-mail that is sent with the generated report. The maximum length allowed is 2048 characters.

RELATED DOCUMENTATION

About the Security Report Definitions Page | 874

Creating Log Report Definition | 881

Creating ANR Report Definition | 887

Creating ANR Report Definition

You can use the Create ANR Report Definition page to create Application and Network Risk (ANR) report definitions and generate the corresponding ANR reports. ANR reports help you to analyze business risks in a network, based on application usage and resource usage.

To create an ANR report definition:

1. Select **Reports > Report Definitions > Security**.

   The Security Report Definitions page appears.

2. Click **Add > ANR Report Definition**.

   The Create ANR Report Definition page appears.

3. Complete the configuration according to the guidelines provided in Table 318 on page 888.

NOTE: Fields marked with * are mandatory.

4. Click **OK** to save the ANR report definition.

   The report definition is saved and the Security Report Definitions page appears.

   A confirmation message appears on this page, indicating that the ANR report definition is successfully created.

   You can perform various actions on the report definition. See "Scheduling, Generating, Previewing, and Sharing Security Reports" on page 877.

Table 318: Fields on the Create ANR Report Definition Page

Field	Description
General	
Report Name	Enter a unique name for the report definition. The name can contain a string of alphanumeric characters and some special characters (colons, periods, dashes, and underscores); no spaces are allowed and the maximum length allowed is 63 characters.
Description	Enter a description for the report definition; the maximum length (including spaces) allowed is 1024 characters.
Content	
Number of Top Logs	Specify the number of records that you want to retrieve and display for each section in the report. Range: 1 through 20. Default: 10.
Time Span (Last)	Specify the duration (Custom, last 3 hours, last 6 hours, last 12 hours, or last 24 hours) for which you want the report to be generated. If you select **Custom**, the **Custom Time Range Selection** page appears. You must specify the **From** date and time, and **To** date and time (in MM/DD/YYYY and HH:MM:SS formats).
Schedule	

Table 318: Fields on the Create ANR Report Definition Page *(continued)*

Field	Description
Schedule Report	Click **Add Schedule** to schedule the report generation. The **Add Report Schedule** page appears. Specify whether you want to generate the report immediately or schedule it for a later date and time: • Run now—Select this option to schedule the report generation at the current time, and click **OK**. • Schedule at a later time—Select this option to schedule the report generation for a later date and time (in MM/DD/YYYY and HH:MM:SS formats), and click **OK**. The Create ANR Report Definition page appears with details of the report generation schedule.
E-Mail	
E-Mail Recipients	Click **Add Email Recipients** to add e-mail addresses of recipients to whom you want to send the ANR report. The Add Recipients page appears. • Recipients—Enter or select one or more e-mail addresses of users to whom you want to send the report. By default, you can search by first name and select registered users. You can also enter external e-mail addresses. • Subject—Enter the subject line for the e-mail that is sent with the generated report. The maximum length allowed is 2048 characters. • Comment—Enter the text to be included in the body of the e-mail that is sent with the generated report. The maximum length allowed is 2048 characters.

RELATED DOCUMENTATION

About the Security Report Definitions Page | 874

Creating Bandwidth Report Definition | 885

Creating Log Report Definition | 881

Editing, Deleting, and Cloning Log Report Definitions

IN THIS SECTION

- Editing the Log Report Definition | 890
- Deleting Log Report Definitions | 890
- Cloning Log Report Definitions | 891

You can edit, delete, and clone log report definitions.

Editing the Log Report Definition

You can edit custom log report definitions from the Security Report Definitions page.

To edit a custom log report definition:

1. Select **Reports > Report Definitions > Security**.

   The Security Report Definitions page appears.

2. Select the log report definition that you want to edit, and click the edit icon (pencil).

   The **Edit Log Report Definition** page appears, displaying the same fields that are presented when you create a log report definition.

3. Modify the log report definition fields as needed.

4. Click **OK** to save the changes or click **Cancel** to discard the changes.

   The Security Report Definitions page appears.

   If you click OK, a confirmation message appears on top of this page.

Deleting Log Report Definitions

You can use the Security Report Definitions page to delete one or more custom log report definitions.

To delete one or more log report definitions:

1. Select **Reports > Report Definitions > Security**.

   The **Security Report Definitions** page appears.

2. Select the log report definitions that you want to delete, and click the delete icon. Alternatively, right click on the report definitions that you want to delete and select **Delete Report**.

   The **Delete Report Definition** page appears.

3. Click **Yes** to delete the selected log report definitions or click **No** to cancel the deletion.

   The Security Report Definitions page appears.

   If you click Yes, the selected log report definitions are deleted and a confirmation message appears on top of this page.

Cloning Log Report Definitions

Cloning enables you to create a new log report definition based on an existing one.

> **NOTE:** You can clone predefined and custom log report definitions.

To clone a log report definition:

1. Select **Reports > Report Definitions > Security**.

   The **Security Report Definitions** page appears.

2. Right-click the log report definition that you want to clone and select **Clone**. Alternatively, select the log report definition and then select **More > Clone**.

   The **Clone log Report Definition** page appears, displaying the same fields that are presented when you create a log report definition.

3. Modify the log Report Definition fields as needed.

4. Click **OK** to save your changes or click **Cancel** to discard the changes.

   The Security Report Definitions page appears.

   If you click OK, a confirmation message appears on top of this page.

   You can perform various actions on the cloned report definition. See "Scheduling, Generating, Previewing, and Sharing Security Reports" on page 877.

RELATED DOCUMENTATION

Creating Log Report Definition | 881

About the Security Report Definitions Page | 874

Editing, Deleting, and Cloning Bandwidth Report Definitions

IN THIS SECTION

- Editing Bandwidth Report Definitions | 892
- Deleting Bandwidth Report Definitions | 892
- Cloning Bandwidth Report Definitions | 893

You can edit, delete, and clone bandwidth report definitions.

Editing Bandwidth Report Definitions

You can edit custom bandwidth report definitions from the Security Report Definitions page.

To edit a custom bandwidth report definition:

1. Select **Reports > Report Definitions > Security**.

   The Security Report Definitions page appears.

2. Select the bandwidth report definition that you want to edit, and click the edit icon (pencil).

   The **Edit Bandwidth Report Definition** page appears, displaying the same fields that are presented when you create a bandwidth report definition.

3. Modify the bandwidth report definition fields as needed.

4. Click **OK** to save the changes or click **Cancel** to discard the changes.

   The Security Report Definitions page appears.

   If you click OK, a confirmation message appears on top of this page.

Deleting Bandwidth Report Definitions

You can use the Security Report Definitions page to delete one or more custom bandwidth report definitions.

To delete one or more bandwidth report definitions:

1. Select **Reports > Report Definitions > Security**.

   The **Security Report Definitions** page appears.

2. Select the bandwidth report definitions that you want to delete, and click the delete icon. Alternatively, right click on the report definitions that you want to delete and select **Delete Report**.

   The **Delete Report Definition** page appears.

3. Click **Yes** to delete the selected bandwidth report definitions or click **No** to cancel the deletion.

   If you click Yes, the selected log report definitions are deleted and the Security Report Definitions page appears.

   A confirmation message appears on top of this page.

Cloning Bandwidth Report Definitions

Cloning enables you to create a new bandwidth report definition based on an existing one.

> **NOTE:** You can clone predefined and custom bandwidth report definitions.

To clone a bandwidth report definition:

1. Select **Reports > Report Definitions > Security**.

   The **Security Report Definitions** page appears.

2. Right-click on the bandwidth report definition that you want to clone and select **Clone**. Alternatively, select the bandwidth report definition and then select **More > Clone**.

   The **Clone Bandwidth Report Definition** page appears, displaying the same fields that are presented when you create a bandwidth report definition.

3. Modify the bandwidth report definition fields as needed.

4. Click **OK** to save your changes or click **Cancel** to discard the changes.

   The Security Report Definitions page appears.

   If you click OK, a confirmation message appears on top of this page.

   You can perform various actions on the cloned report definition. See "Scheduling, Generating, Previewing, and Sharing Security Reports" on page 877.

RELATED DOCUMENTATION

About the Security Report Definitions Page | 874
Creating Bandwidth Report Definition | 885

Editing, Deleting, and Cloning ANR Report Definitions

IN THIS SECTION

- Editing ANR Report Definitions | 894
- Deleting ANR Report Definitions | 895
- Cloning ANR Report Definitions | 895

You can edit, delete, and clone ANR report definitions.

Editing ANR Report Definitions

You can edit custom ANR report definitions from the Security Report Definitions page.

To edit the custom ANR report definition:

1. Select **Reports > Report Definitions > Security**.

   The Security Report Definitions page appears.

2. Select the ANR report definition that you want to edit, and click the edit icon (pencil).

   The **Edit ANR Report Definition** page appears, displaying the same fields that are presented when you create an ANR report definition.

3. Modify the ANR Report Definition fields as needed.

4. Click **OK** to save the changes or click **Cancel** to discard the changes.

   The Security Report Definitions page appears.

   If you click OK, a confirmation message appears on top of this page.

Deleting ANR Report Definitions

You can use the Security Report Definitions page to delete one or more custom ANR report definitions.

To delete one or more ANR report definitions:

1. Select **Reports > Report Definitions > Security**.

   The **Security Report Definitions** page appears.

2. Select the ANR report definitions that you want to delete, and click the delete icon. Alternatively, right click on the report definitions that you want to delete and select **Delete Report**.

   The **Delete Report Definition** page appears.

3. Click **Yes** to delete the selected ANR report definitions or click **No** to cancel the deletion.

   If you click Yes, the selected ANR report definitions are deleted and the Security Report Definitions page appears.

   A confirmation message appears on top of this page.

Cloning ANR Report Definitions

Cloning enables you to create a new ANR report definition based on an existing one.

> **NOTE:** You can clone predefined and custom ANR report definitions.

To clone an ANR report definition:

1. Select **Reports > Report Definitions > Security**.

   The **Security Report Definitions** page appears.

2. Right-click on the ANR report definition that you want to clone and select **Clone**. Alternatively, select the ANR report definition and then select **More > Clone**.

   The **Clone ANR Report Definition** page appears, displaying the same fields that are presented when you create an ANR report definition.

3. Modify the ANR Report Definition fields as needed.

4. Click **OK** to save your changes or click **Cancel** to discard the changes.

   The Security Report Definitions page appears.

If you click OK, a confirmation message appears on top of this page.

You can perform various actions on the cloned report definition. See “Scheduling, Generating, Previewing, and Sharing Security Reports” on page 877.

RELATED DOCUMENTATION

About the Security Report Definitions Page | **874**

Creating ANR Report Definition | **887**

CHAPTER 33

SD-WAN Reports

IN THIS CHAPTER

- About the SD-WAN Report Definitions Page | 897
- Editing, Deleting, and Cloning SD-WAN Report Definitions | 899
- Creating SD-WAN Tenant Performance Report Definitions | 901
- Creating SD-WAN Site Performance Report Definitions | 905
- About the SD-WAN Generated Reports Page | 908

About the SD-WAN Report Definitions Page

To access this page, click **Customer Portal > Reports > Report Definitions > SD-WAN.**

The **SD-WAN Report Definitions** page appears.

The SD-WAN Report Definitions page displays a list of predefined and custom report definitions. To generate reports, you can use the predefined report definitions as is, or you can create custom report definitions.

> **NOTE:** The SD-WAN Performance Report Definition is the only predefined report definition available on the SD-WAN Report Definitions page.

Tasks You Can Perform

You can perform the following tasks from this page:

- Create SD-WAN tenant performance report definitions. See “Creating SD-WAN Tenant Performance Report Definitions” on page 901
- Create SD-WAN site performance report definitions. See “Creating SD-WAN Site Performance Report Definitions” on page 905

- Run a report immediately, edit a schedule, edit e-mail recipients, preview a report in PDF, send reports, and clone reports. See "Scheduling, Generating, Previewing, and Sharing Security Reports" on page 877
- View details about an SD-WAN report definition—Right-click a report definition and then select **Detailed View** or select the report definition and click **More > Detailed View**. Alternatively, hover over the report definition name and click the Detailed View icon that appears before it.

  The Report Definition Details page appears, displaying the same values that you specified for each parameter in the selected report definition.
- To search for a report definition from the list of available report definitions, click the **Search** icon in the top right corner of the page.

  Enter the name of the report definition in the search bar and click the search icon.

  The search results are displayed.

  Click **Clear All** to clear the search results.

Field Descriptions

Table 319 on page 898 describes the fields on the SD-WAN Report Definitions page.

Table 319: Fields on the SD-WAN Report Definitions Page

Field	Description
Name	Name of the SD-WAN report definition.
Description	Description of the SD-WAN report definition.
Type	Type of SD-WAN report definition—Tenant Performance or Site Performance.
Definition Type	Indicates whether the report definition is predefined (system-generated) or custom (user-created).
Schedule	Indicates whether the report generation is scheduled at the current time (Now) or for a later date and time (Once).
Recipients	E-mail addresses of recipients to whom the generated report is sent.
Job ID	Use the job ID to view the status of the task on the Jobs (Monitor > Jobs) page.

RELATED DOCUMENTATION

Editing, Deleting, and Cloning SD-WAN Report Definitions | 899

Creating SD-WAN Tenant Performance Report Definitions | 901
Creating SD-WAN Site Performance Report Definitions | 905

Editing, Deleting, and Cloning SD-WAN Report Definitions

IN THIS SECTION

- Editing the SD-WAN Report Definition | 899
- Deleting SD-WAN Report Definitions | 899
- Cloning SD-WAN Report Definitions | 900

You can edit, delete, and clone SD-WAN report definitions from the SD-WAN Report Definitions page.

Editing the SD-WAN Report Definition

You can edit custom SD-WAN report definitions from the SD-WAN Report Definitions page.

To edit an SD-WAN report definition:

1. Select **Reports > Report Definitions > SD-WAN**.

   The SD-WAN Report Definitions page appears.

2. Select the SD-WAN report definition that you want to modify, and click the edit icon (pencil).

   The Update SD-WAN Performance Report Definition page appears, displaying the same fields that are presented when you create an SD-WAN report definition.

3. Modify the report definition fields as needed.

4. Click **OK** to save the changes or click **Cancel** to discard the changes

   The SD-WAN Report Definitions page appears.

   If you click OK, a confirmation message appears on top of this page.

Deleting SD-WAN Report Definitions

You can use the SD-WAN Report Definitions page to delete one or more custom SD-WAN report definitions.

To delete one or more SD-WAN report definitions:

1. Select **Reports > Report Definitions > SD-WAN**.

   The SD-WAN Report Definitions page appears.

2. Select the SD-WAN report definitions that you want to delete and click the Delete icon. Alternatively, right click the report definitions that you want to delete and select **Delete Report**.

   The Confirm Delete page appears.

3. Click **Yes** to delete the selected SD-WAN report definitions or **No** to cancel the deletion.

   The SD-WAN report definitions page appears.

   If you click Yes, the selected SD-WAN report definitions are deleted and a confirmation message **Successfully deleted report template** appears on top of this page.

Cloning SD-WAN Report Definitions

Cloning enables you to create a new SD-WAN report definition based on an existing one.

> **NOTE:** You can clone predefined and custom SD-WAN report definitions.

To clone an SD-WAN report definition:

1. Select **Reports > Report Definitions > SD-WAN**.

   The SD-WAN Report Definitions page appears.

2. Right-click on the SD-WAN report definition that you want to clone and select **Clone**. Alternatively, select the SD-WAN report definition and then select **More > Clone**.

   The **Clone SD-WAN Performance Report Definition** page appears, displaying the same fields that are presented when you create an SD-WAN report definition.

3. Modify the SD-WAN Report Definition fields as needed.

4. Click **OK** to save your changes or click **Cancel** to discard the changes.

   You are returned to the SD-WAN Report Definitions page.

   If you click OK, a confirmation message appears on top of this page.

RELATED DOCUMENTATION

About the SD-WAN Report Definitions Page | 897

Creating SD-WAN Tenant Performance Report Definitions

Use the SD-WAN Report Definitions page to create SD-WAN tenant performance report definitions for all sites in a tenant and generate reports based on the definitions. SD-WAN tenant performance reports enable you to analyze tenant performance based on the following parameters that measure the SLA performance across all sites in a tenant:

- Top applications by bandwidth
- Top sites not meeting SLA
- Top sites meeting SLA with switching
- Sites meeting SLA without switching
- Top sites by current active tunnels
- Top sites by highest packet loss
- Top sites by highest latency
- Top sites by highest jitter

> **NOTE:** Only users with the Tenant Administrator role can create SD-WAN tenant performance report definitions.

To create an SD-WAN tenant performance report definition:

1. Select **Reports > Report Definitions > SD-WAN**.

   The SD-WAN Report Definitions page appears.

2. Click **Add > Tenant Performance** .

   The Add SD-WAN Tenant Performance Report Definition page appears.

3. Complete the configuration according to the guidelines provided in Table 320 on page 902.

NOTE: Fields marked with * are mandatory.

4. Click **OK** to save the report definition.

   The report definition is saved and the SD-WAN Report Definitions page appears.

   A confirmation message appears on top of this page, indicating that the report definition is successfully created.

Table 320: Fields on the Create Tenant Performance Report Definition

Field	Description
General	
Report Name	Enter a unique name for the report definition. The name can contain a string of alphanumeric characters and some special characters (colons, periods, dashes, and underscores); no spaces are allowed and the maximum length allowed is 63 characters.
Description	Enter a description for the report definition; maximum length allowed is 1024 characters.
Content	
Time Span	Specify the duration (last 24 hours, last 7 days, last 30 days, or custom) for which you want the report to be generated. If you select **Custom**, the **From** and **To** fields appear: • From—Specify the start date and time from which the report should be generated. • To—Specify the end date and time up to which the report should be generated.
Number of Top Logs	Enter the number top of SLA records (1 through 20) that you want to retrieve and display for each section in the report.

Table 320: Fields on the Create Tenant Performance Report Definition *(continued)*

Field	Description
Report Content	Select the content that you want to view in the report. • **Top Applications By Bandwidth**—Displays a report on top applications by bandwidth. • **Top Sites Not Meeting SLA**—Displays a report on top sites not meeting the SLA performance. • **Top Sites Meeting SLA with Switching**—Displays a report on top sites meeting SLA performance with link switching. • **Sites Meeting SLA without Switching**—Displays report on sites meeting SLA performance without switching. • **Current Active Tunnels**—Displays a report on top 10 sites that have the maximum number of active dynamic mesh tunnels. • **Top Sites by Highest Packet Loss**—Displays a report on top 10 sites based on the highest cumulative average packet loss across all the links. • **Top Sites by Highest Latency**—Displays report on top 10 sites based on the highest cumulative average latency across all the links. • **Total Sites by Highest Jitter**—Displays report on top 10 sites based on the highest cumulative average jitter across all the links. For more information about SLA parameters and dynamic mesh tunnels, see "SLA Profiles and SD-WAN Policies Overview" on page 499 and "Dynamic Mesh Tunnels Overview" on page 221.
Schedule Report	
Add Schedule	Click **Add Schedule** to schedule the report generation. The Add Report Schedule page appears. Specify whether you want to generate the report immediately or schedule it for a later date and time. • **Run now**—Select this option to generate the report immediately. • **Schedule at a later time**— Select this option to generate the report at a later date and time (in MM/DD/YYYY and HH:MM:SS format).
Email Recipients	

Table 320: Fields on the Create Tenant Performance Report Definition *(continued)*

Field	Description
Add Email Recipients	Click **Add Email Recipients** to add e-mail addresses of recipients to whom you want to send the SD-WAN reports. The Add Recipients page appears. • **Recipients**—Select e-mail addresses of users to whom you want to send the report. You can select more than one e-mail address. • **Subject**—Enter the subject line for the e-mail that is sent with the generated report. The maximum length allowed is 2048 characters. • **Comment**—Enter the text to be included in the body of the e-mail that is sent with the generated report. The maximum length allowed is 2048 characters.

RELATED DOCUMENTATION

Creating SD-WAN Site Performance Report Definitions | **905**

About the SD-WAN Report Definitions Page | **897**

Editing, Deleting, and Cloning SD-WAN Report Definitions | **899**

Creating SD-WAN Site Performance Report Definitions

Use the SD-WAN Report Definitions page to create SD-WAN site performance report definitions for specific sites of a tenant and generate the report based on the definitions.

SD-WAN site performance reports enable you to analyze site performance based on the following parameters that measure the SLA performance of specific sites in a tenant:

- Top 10 applications and link utilization for site
- Top profiles not meeting SLA
- Top profiles switching links
- Top applications by highest packet loss
- Top applications by highest latency
- Top applications by highest jitter
- SLA performance between two sites

> **NOTE:** Only users with the Tenant Administrator role can create SD-WAN site performance report definitions.

To create an SD-WAN site performance report definition:

1. Select **Reports > Report Definitions > SD-WAN**.

   The SD-WAN Report Definitions page appears.

2. Click **Add > Site Performance**.

   The Add SD-WAN Site Performance Report Definition page appears.

3. Complete the configuration according to the guidelines provided in Table 321 on page 906.

   > **NOTE:** Fields marked with * are mandatory.

4. Click **OK** to save the report definition.

   The report definition is saved and the SD-WAN Report Definition page appears.

   A confirmation message appears on top of this page.

Table 321: Fields on the Site Performance Report Definition Page

Field	Description
General	
Report Name	Enter a unique string of alphanumeric characters and some special characters (: . -). No spaces are allowed and the maximum length allowed is 63 characters.
Description	Enter a description for the report definition; maximum length allowed is 1024 characters.
Content	
Time Span	Specify the duration (last 24 hours, last 7 days, last 30 days, or custom) for which you want the report to be generated. If you select **Custom**, the **From** and **To** fields appear: • From—Specify the start date and time from which the report should be generated. • To—Specify the end date and time up to which the report should be generated.
Number of Top Logs	Enter the number of top SLA events (1 through 20) that you want to retrieve and display for each section in the report.
Sites	Select one or more sites for which you want to generate the report. You can select up to five sites.
Report Content	Select the content that you want to view in the report. • **Top 10 Applications and Link Utilization**—Displays a report on top 10 applications and link utilization for the selected sites. • **Top Profiles Not Meeting SLA**—Displays a report on top SLA profiles not meeting SLA for the selected sites. • **Top Profiles Switching Links**—Displays a report on top SLA profiles switching links for the selected sites. • **Top Applications by Highest Packet Loss**—Displays report on top 10 applications based on the selected site that has the highest average packet loss across SLA profiles. • **Top Applications by Highest Latency**—Displays report on top 10 applications based on the selected site that has the highest average latency across SLA profiles. • **Total Applications by Highest Jitter**—Displays report on top 10 applications based on the selected site that has the highest average jitter across SLA profiles. • **SLA Performance Between Two Sites**—Displays report on top 20 applications based on the performance of SLA parameters (latency, jitter, and packet loss) between the source and destination site that you have selected. For more information about SLA parameters and dynamic mesh tunnels, see "SLA Profiles and SD-WAN Policies Overview" on page 499 and "Dynamic Mesh Tunnels Overview" on page 221.

Table 321: Fields on the Site Performance Report Definition Page *(continued)*

Field	Description
Sites	Select one or more sites, from the list of available sites, for which you want to generate the report.
Schedule	
Schedule Report Generation	Click **Add Schedule** to schedule the report generation. The **Add Report Schedule** page appears. Specify whether you want to generate the report immediately or schedule it for a later date and time: • Run now—Select this option to schedule the report generation at the current time, and click **OK**. • Schedule at a later time—Select this option to schedule the report generation for a later date and time (in MM/DD/YYYY and HH:MM:SS formats), and click **OK**. The Add SD-WAN Tenant Performance Report Definition page appears with details of the report generation schedule.
E-Mail	
E-Mail Recipients	Click **Add Email Recipients** to add e-mail addresses of recipients to whom you want to send the SD-WAN report. The Add Recipients page appears. • Recipients—Enter or select one or more e-mail addresses of users to whom you want to send the report. By default, you can search by first name and select registered users. You can also enter external e-mail addresses. • Subject—Enter the subject line for the e-mail that is sent with the generated report. The maximum length allowed is 2048 characters. • Comments—Enter the text to be included in the body of the e-mail that is sent with the generated report. The maximum length allowed is 2048 characters.

RELATED DOCUMENTATION

Creating SD-WAN Tenant Performance Report Definitions | 901

About the SD-WAN Report Definitions Page | 897

Editing, Deleting, and Cloning SD-WAN Report Definitions | 899

About the SD-WAN Generated Reports Page

To access this page, click **Customer Portal** > **Reports** > **Generated Reports** > **SD-WAN**.

Use this page to view the list of tenant and site performance reports that are generated from the SD-WAN Report Definitions page.

You must click on the report to view the report in PDF. You can view the generated report for up to 30 days and the report will be deleted after 30 days.

You can also delete one or more generated reports.

Field Descriptions

Table 322 on page 908 describes the fields on the SD-WAN Generated Reports page.

Table 322: Fields on the SD-WAN Generated Reports Page

Field	Description
Name	Name of the SD-WAN report.
Description	Description of the report.
Generated Time	Date and time when the report was generated.
Definition Name	Name of the report definition.
Generated By	Name of the tenant administrator who generated the report.
Recipients	Recipients of the generated report.

RELATED DOCUMENTATION

Reports Overview | 873

About the SD-WAN Report Definitions Page | 897