

Junos[®] Space Network Director

Network Director User Guide

Published
2020-12-10

Release
4.1

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.

Junos[®] Space Network Director Network Director User Guide

4.1

Copyright © 2020 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 | xxix

Documentation and Release Notes | xxix

Documentation Conventions | xxix

Documentation Feedback | xxxii

Requesting Technical Support | xxxii

Self-Help Online Tools and Resources | xxxiii

Creating a Service Request with JTAC | xxxiii

1

Working With Network Director

About Network Director | 2

Understanding Network Director and the Management Life-Cycle Modes | 2

Benefits of Network Director | 3

Understanding the Network Director User Interface | 3

Network Director Banner | 4

View Pane | 6

Displaying Devices in Various Network Views | 7

Filtering the Network Tree | 7

Expanding or Collapsing Nodes in the Network Tree | 8

Searching the Network Tree | 9

Tasks Pane | 9

Alarms | 11

Main Window or Workspace | 11

Tables in Network Director | 11

Moving and Resizing Columns | 11

Displaying the Column Drop-Down Menu | 12

Sorting on a Column | 12

Hiding and Exposing Columns | 13

Searching Table Contents | 13

Filtering Table Contents | 15

Accessing Network Director | 16

Logging In to Network Director | 16

Logging Out of Network Director | 17

Changing Your Password | 18

Understanding Network Director System Administration and Preferences | 19

Understanding Network Director User Administration | 19

Understanding the System Tasks Pane | 20

Audit Logs Overview | 21

Viewing Audit Logs From Network Director | 22

Managing Jobs | 23

Collecting Logs for Troubleshooting | 25

Setting Up User and System Preferences | 26

Accessing the Preferences Page | 27

Choosing Server Time or Local Time | 27

Specifying Search Preferences | 28

Enabling Import of Configuration Group Data from Ethernet Design | 28

Selecting the Approval Mode | 28

Setting up Auto-resynchronization Preferences | 29

Retaining Network Director Reports | 30

Changing Monitor Mode Settings | 30

Disabling Data Collection for Monitors | 31

Changing the Polling Interval | 32

Enabling and Disabling Collection for Managed Devices | 33

Specifying Database History Retention | 34

Specifying the Data Learning Engine (DLE) Settings | 34

Changing Alarm Settings | 36

Configuring Global Alarm Notifications | 36

Retaining Alarm History | 36

Specifying Event History | 37

Enabling Alarms | 37

Changing the Severity of Individual Alarms | 46

Configuring Threshold Alarms | 46

Configuring Individual Alarm Notifications | 46

Modifying Data Center Synchronization Interval Using the Virtualization Tab | 47

Getting Started with Network Director | 49

Getting Started with Junos Space Network Director | 49

Building Your Network | 49

Creating Profiles in Network Director | 50

Managing Software Images using Network Director | 51

Configuring Approval Modes for Device Configurations | 51

Resynchronizing Device Configuration | 51

Creating the Baseline Configuration | 52

Monitoring Your Network | 52

Setting up Network Traffic Analysis and Analyzing the Traffic | 52

Managing Network Faults and Notifications | 53

Generating Network Reports | 53

2

Working with the Dashboard

About the Dashboard | 56

Understanding the Dashboard | 56

Using the Dashboard | 57

Using Dashboard Widgets | 57

Dashboard Widget Reference | 59

Alarms Widget | 59

Alarms Widget Summary | 59

Alarms Widget Details | 60

Config Deployment Jobs Status Widget | 61

Config Deployment Jobs Status Widget Summary | 61

Config Deployment Jobs Status Widget Details | 62

Device & Port Latency Widget | 62

Device & Port Utilization Widget | 63

Using the Global Controls | 64

Interacting with the Heat Maps | 64

Viewing Active Flows on a Port | 65

Flow Analysis Details Window | 65

Viewing Traffic on a Device | 67

Equipment By Type Widget | 69

Equipment By Type Widget Summary | 69

Equipment By Type Widget Details | 70

Port Status - Physical Widget | 70

Port Status - Physical Widget Summary | 71

Port Status - Physical Widget Details | 71

Top Talker - Wired Devices Widget | 71

Top Talker - Wired Devices Widget Summary | 72

Top Talker - Wired Devices Widget Details | 72

Top Overlay Networks Widget | 73

Top Overlay Networks Widget Summary | 73

Top Overlay Networks Widget Details | 74

3

Working in Build Mode

About Build Mode | 77

Understanding Build Mode in Network Director | 77

Discovering Devices | 77

Building the Logical, Location, and Custom Views | 78

Configuring Devices | 79

Deploying Device Configurations | 80

Importing Device Configurations | 80

Out-of-Band Configuration Changes | 81

Managing Devices | 81

Understanding the Build Mode Tasks Pane | 82

Understanding Network Configuration Profiles | 88

Assigning Profiles to an Interface, Device, or a Group of Devices | 91

Discovering Devices | 94

Discovering Devices in a Physical Network | 94

Specifying Target Devices | 95

Specifying Discovery Options | 97

- Specifying Schedule Options | 99
- Reviewing Device Discovery Options | 99
- Viewing the Discovery Status | 99

Understanding the Device Discovery Process | 101

- Benefits of the Device Discovery Process | 102

Troubleshooting Device Discovery Error Messages | 102

Setting Up Sites and Locations Using the Location View | 105

Understanding the Location View | 105

Setting Up the Location View | 107

Creating a Site | 111

- How to Add or Edit a Location Site | 111
- Creating or Editing a Site | 111

Configuring Buildings | 112

- How to Add or Edit a Building | 112
- Adding or Editing a Building for a Location | 113

Configuring Floors | 114

- How to Add or Edit a Floor | 114
- Adding or Editing a Building Floor for a Location | 115

Setting Up Closets | 116

- How to Add or Edit a Closet | 116
- Adding or Editing a Wiring Closet | 117

Assigning and Unassigning Devices to a Location | 117

- How to Assign or Unassign Devices | 118
- Assigning Devices | 118

Changing the Location of a Device | 119

- How to Move a Device to a New Location | 120
- Changing the Location of a Device | 120

Deleting Sites, Buildings, Floors, Wiring Closets, and Devices | 121

- How to Delete a Location Object | 121
- Deleting Sites | 121
- Deleting Buildings | 122
- Deleting Floors | 122
- Deleting Closets | 122

Deleting Devices | 122

Configuring Outdoor Areas | 123

How to Configure an Outdoor Area | 123

Configuring an Outdoor Area | 124

Building a Topology View of the Network | 125

Understanding the Network Topology in Network Director | 126

Understanding the Topology View Tasks pane | 130

Setting Up the Topology View | 133

Managing the Topology View | 134

Viewing the Network Topology | 134

Refreshing the Topology | 137

Viewing Topology | 138

Viewing Topology Discovery Job | 139

Setting Up Locations | 139

Viewing the Alarm Details | 140

Discovering the Linux Hosts | 140

Displaying Device Connectivity | 140

Displaying Virtual Chassis and Virtual Chassis Fabric Connectivity | 145

Uploading Floor Plans | 148

Uploading Topology Map | 149

Adding and Managing OUI Data in Network Director | 150

Creating Custom Device Groups | 152

Understanding Custom Device Groups | 152

Where Is the Custom Group Function Located in Network Director? | 153

How Do Custom Group Rules Work? | 153

What Happens When I Edit a Custom Group Rule? | 154

When Are Rules Executed? | 154

Creating Custom Device Groups | 155

Creating Custom Groups | 155

Creating a Custom Group | 156

Configuring Quick Templates | 161

Understanding Quick Templates | 161

Benefits of Quick Templates | 162

Configuring and Managing Quick Templates | 163

Creating a Quick Template | 164

Applying Templates to Devices | 165

Editing a Quick Template | 166

Deleting a Quick Template | 166

Cloning a Quick Template | 166

Using the Quick Template Details Window | 167

Viewing Deployed Quick Templates | 167

Configuring Device Settings | 169

Understanding Device Common Settings Profiles | 169

Creating and Managing Device Common Settings | 170

Managing Device Common Settings | 170

Creating a Device Common Settings Profile | 172

Specifying Basic Settings for Device Common Settings | 174

Specifying Management Settings for EX Switching Device Common Settings | 177

Specifying Management Settings for Campus Switching ELS Device Common Settings | 181

Specifying Management Settings for Data Center ELS Device Common Settings | 185

Specifying Protocol Settings for EX Switching Device Common Settings | 188

Specifying DHCP Relay/DNS Settings for Campus Switching ELS Device Common Settings | 192

Specifying Protocol Settings for Data Center Switching ELS Device Common Settings | 194

Reviewing and Saving a Device Common Settings Configuration | 197

What to Do Next | 197

Assigning Device Common Settings to Devices | 198

Assigning Device Common Settings | 199

Editing the Assignments of the Device Common Setting | 201

Configuring Authentication, Authorization, and Access for Your Network | 202

Understanding Central Network Access Using RADIUS and TACACS+ | 202

Why Do I Want Remote Authentication ? | 203

Why Not Just Rely on Firewalls and Filters for Access Control? | 203

What About Using LDAP For Authentication? | 203

Where Is RADIUS Installed on the Network? | 204

How Is TACACS+ Installed on the Network? | 204

A Comparison of RADIUS and TACACS+ | 205

Creating and Managing RADIUS Profiles | 206

Managing RADIUS Profiles | 206

Creating RADIUS Profiles | 208

Specifying Settings for a RADIUS Profile | 208

What to Do Next | 211

Creating and Managing LDAP Profiles | 212

Managing LDAP Profiles | 212

Creating LDAP Profiles | 214

Specifying Settings for an LDAP Profile | 214

What to Do Next | 217

Understanding Access Profiles | 217

Creating and Managing Access Profiles | 218

Managing Access Profiles | 219

Creating an Access Profile | 220

Specifying Basic Settings for an EX Series Switching Access Profile | 222

Specifying RADIUS Accounting Settings for an EX Switching Access Profile | 224

Specifying Basic Settings for a Campus Switching ELS Access Profile | 227

Specifying RADIUS and LDAP Settings for Campus Switching ELS | 227

Reviewing and Modifying the Access Profile Settings | 234

What To Do Next | 235

Understanding Authentication Profiles | 235

802.1X Authentication | 236

MAC RADIUS Authentication | 236

Captive Portal Authentication | 237

Creating and Managing Authentication Profiles | 237

- Managing Authentication Profiles | 238**

- Creating an Authentication Profile | 239**

- Specifying Authentication Settings for Switches | 240**

- What To Do Next | 245**

Configuring Interfaces and VLANs | 246

Understanding Port Profiles | 246

- Interface Settings Configured in the Port Profile | 247**

- Interface Settings Configured by Referencing Other Profiles | 248**

- Default Port Profiles | 248**

Creating and Managing Port Profiles | 251

- Managing Port Profiles | 252**

- Creating Port Profiles | 254**

- Specifying Settings for an EX Switching Port Profile | 255**

- Specifying Settings for a Campus Switching ELS Port Profile | 269**

- Specifying Settings for a Data Center Switching ELS Port Profile | 285**

- What to Do Next | 303**

Assigning and Unassigning Port Profiles from Interfaces | 304

- Selecting Devices for Assignment | 305**

- Selecting Interfaces for Assignment | 306**

- Reviewing and Accepting the Assignments | 309**

- Editing Profile Assignments | 310**

- Unassigning a Port Profile from an Interface | 311**

Managing Auto Assignment Policies | 312

Creating Auto Assignments | 314

- Adding Port Profiles using the Select Port Profiles Page | 314**

- Adding Devices and Ports for Auto Assignment | 315**

- Viewing the Auto Assignment Policy Summary | 316**

Configuring Easy Config Setup | 317

- Configuring Interface Settings | 317**

Understanding Port Groups | 323

Creating and Managing Port Groups | 323**Managing Port Groups | 324****Creating Port Groups | 325****Specifying Settings for a Port Group | 326****What to Do Next | 326****Understanding VLAN Profiles | 327****Creating and Managing VLAN Profiles | 329****Managing VLAN Profiles | 330****Creating a VLAN Profile | 331****Specifying Basic EX Switching VLAN Settings | 333****Specifying Basic Campus Switching ELS VLAN Settings | 334****Specifying Basic VLAN Settings for Data Center Switching ELS | 336****Specifying Advanced VLAN Profile Settings for EX Series Switches | 337****Specifying Advanced VLAN Settings for Campus Switching ELS | 339****Specifying Advanced VLAN Settings for Data Center Switching ELS | 341****Reviewing and Saving the VLAN Profile Configuration | 344****What to Do Next | 344****Assigning a VLAN Profile to Devices or Ports | 345****Assigning a VLAN Profile | 345****Editing Profile Assignments | 347****Configuring Firewall Filters (ACLs) | 349****Understanding Filter Profiles | 349****Creating and Managing Wired Filter Profiles | 350****Managing Wired Filter Profiles | 351****Creating a Wired Filter Profile | 352****Specifying Settings for an EX Series Switch Filter Profile | 353****Specifying Settings for a Campus Switching ELS Switch Filter Profile | 365****Specifying Settings for a Data Center Switching ELS Filter Profile | 379****What to Do Next | 394**

Configuring Class of Service (CoS) | 396

Understanding Class of Service (CoS) Profiles | 396

How Would I Use CoS (also known as QoS)? | 397

How Do I Create CoS Groups? | 397

How Is CoS Different From QoS? | 397

What CoS Parameters Can I Control? | 397

What Are the Default CoS Traffic Types? | 398

Data Center Switching CoS Configuration | 399

How Do I Implement Class of Service? | 399

Editing Discovered CoS Profiles | 399

Creating and Managing Wired CoS Profiles | 400

Managing Wired CoS Profiles | 400

Using the Default CoS Profiles for Switches | 401

Using the Default CoS Profiles for Data Center Switching | 402

Creating a Wired CoS Profile | 402

Specifying Settings for a Switching and Campus Switching ELS CoS Profile | 403

Specifying Settings for a Data Center Switching CoS Profile | 407

What to Do Next | 416

Configuring Media Access Control Security (MACsec) | 417

Media Access Control Security Overview | 417

Configuring and Managing MACsec Profiles | 418

Creating a MACsec Profile | 419

Specifying Settings for a MACsec Profile | 420

What to Do Next | 423

Assigning the MACsec Profiles | 423

Assigning a MACsec Profile to a Device | 423

Editing the MACsec Profile Assignments | 424

Configuring Link Aggregation Groups (LAGs) | 425

Understanding Link Aggregation | 425

Managing and Creating a Link Aggregation Group | 426

Link Aggregation Group Options | 427

Creating a Link Aggregation Group | 428

Managing ICCP Settings | 430

What To Do Next | 431

Understanding Multichassis Link Aggregation | 431

Creating and Managing Multichassis Link Aggregation Groups (MC-LAGs) | 432

Accessing the MC-LAG Page | 433

Creating an MC-LAG | 434

Selecting Peer Devices and Configuring Peer-to-Peer Link Settings | 434

Selecting Client Devices and Configuring Client-to-Peer Link Settings | 436

Saving MC-LAG Settings | 439

Deploying MC-LAG Configuration | 439

MC-LAG Automation Parameters | 440

Editing an MC-LAG | 441

Managing Peer Devices and Peer-to-Peer Link Settings | 442

Managing Client Devices and Client-to-Peer Link Settings | 443

Deleting an MC-LAG | 447

Managing an MC-LAG Created Through CLI Mode | 447

MC-LAG Peer Pairing | 447

Mapping Client Devices to Peer Devices | 448

Ports Mapping Between Peer-to-Peer and Client-to-Peer Devices | 448

Creating and Managing ESI Link Aggregation Groups (ESI-LAGs) | 448

Accessing the ESI-LAG Page | 449

Creating an ESI-LAG | 449

Selecting Peer Devices and Configuring Peer-to-Peer Link Settings | 450

Selecting Client Devices and Configuring Client-to-Peer Link Settings | 452

Saving ESI-LAG Settings | 453

Deploying ESI-LAG Configuration | 454

Editing an ESI-LAG | 454

Managing Peer Devices and Peer-to-Peer Link Settings | 455

Managing Client Devices and Client-to-Peer Link Settings | 456

Deleting an ESI-LAG | 457

ESI-LAG Automation Parameters | 458

Creating and Managing Fabrics | 461

Understanding Junos Fusion | 461

Understanding Junos Fusion Enterprise | 463

Software Requirements for Junos Fusion | 465

Creating and Managing Fusion Configuration Templates | 466

Create a Configuration Template for Junos Fusion Enterprise | 467

Clone a Configuration Template | 473

Apply Configuration Template to Devices | 474

View Details about a Configuration Template | 479

Delete a Configuration Template | 480

Managing Fusion Fabrics | 480

Modify the Fusion Fabric | 481

Edit Aggregation Device Details | 482

Edit Satellite Device Details | 483

Enable Uplink Failure Detection | 483

View the Cabling Plan | 484

View Fabric Connectivity | 484

Replace Aggregation Device or Satellite Device in Junos Fusion | 485

Creating and Managing Satellite Software Upgrade Groups | 486

Create a Software Upgrade Group | 487

Edit a Software Upgrade Group | 488

View Details of a Software Upgrade Group | 488

Delete a Software Upgrade Group | 488

Understanding Layer 3 Fabrics | 489

User Privileges Required for the DHCP and File Server While Using Zero Touch Provisioning | 490

Managing Layer 3 Fabrics | 491

Creating Layer 3 Fabrics | 493

Specifying the Fabric Requirements | 494

Specifying the Device Details | 499

Specifying Configuration Details | 500

Viewing the Cabling Plan | 501

Specifying Zero Touch Provisioning Details	504
Reviewing the Layer 3 Fabric Settings	507
Editing Layer 3 Fabrics	508
Viewing Layer 3 Fabric Connectivity	511
Performing Layer 3 Fabric Connectivity Checks	512
Setting Up Virtual Chassis Fabrics	513
Selecting a Provisioning Method	514
Adding Devices	515
Assigning Devices to a Virtual Chassis Fabric by Using a Graph	515
Assigning Devices to a Virtual Chassis Fabric by Using a Grid	518
Specifying Software Packages	519
Reviewing the Virtual Chassis Fabric Summary and Deploying Changes	520
Managing Virtual Chassis Fabrics	521
Adding Spine Device(s) to the Virtual Chassis Fabric	521
Adding Leaf Node(s) to the Virtual Chassis Fabric	522
Replacing Spine or Leaf Devices	523
Removing Spine or Leaf Devices	524
Configuring Software Package for Upgrade	525
Deploying VCF Configuration Changes	525
Configuring Overlay Networks and Tenants 	527
VXLAN—EVPN Overlay Overview	527
Create a Layer 3 Fabric based Underlay Network	528
Creating and Managing Overlay Fabrics	531
Creating an Overlay Fabric	532
Modifying an Overlay Fabric	533
Viewing Overlay Fabric Configuration Details	534
Setting Up a VXLAN—EVPN-Based Data Center	535
Creating and Managing Tenants	536
Creating a Tenant	537
Modifying a Tenant	540
Viewing Tenant Configuration Details	541

Configuring VRRP Profiles | 544

Understanding VRRP Profiles | 544

Creating and Managing VRRP Profiles | 545

Managing VRRP Profiles | 546

Creating VRRP Profiles | 547

Specifying VRRP Settings for an EX Switching or Campus Switching ELS or Data Center Switching ELS | 547

Managing Network Devices | 550

Viewing the Device Inventory Page | 551

Physical Topology | 553

Viewing Profiles Assigned to a Device | 558

Viewing the Physical Inventory of Devices | 560

Viewing Licenses With Network Director | 561

Viewing a Device's Current Configuration from Network Director | 563

Assigning Devices to Logical Category | 564

Accessing a Device's CLI from Network Director | 565

Accessing a Device's Web-Based Interface from Network Director | 566

Deleting Devices | 568

Rebooting Devices | 569

Viewing Virtual Machines | 569

4

Working in Deploy Mode

About Deploy Mode | 572

Understanding Deploy Mode in Network Director | 572

Deploying Configuration Changes | 573

Managing Software Images | 574

Zero Touch Provisioning | 575

Managing Devices | 575

Managing Device Configuration Files | 575

Managing Baseline Configuration | 575

Understanding the Deploy Mode Tasks Pane | 577

Deploying and Managing Device Configurations | 580

Deploying Configuration to Devices | 580

Selecting Configuration Deployment Options | 581

Using the Change Request Details Page | 585

Creating a Change Request | 586

Validating Configuration | 586

Discarding the Pending Configurations | 587

Viewing Pending Configuration Changes | 587

Using the Pending Changes Window | 587

Using the Configuration or Pending Configuration Window | 588

Using the Deploy Configuration Errors/Warnings Window | 589

Using the Configuration Validation Window | 589

Deploying Configuration Changes to Devices Immediately | 589

Scheduling Configuration Deployment | 590

Specifying Configuration Deployment Scheduling Options | 590

Editing Change Requests | 591

Deleting Change Request | 592

Resubmitting a Change Request | 592

Performing a Rollback | 593

Managing Configuration Deployment Jobs | 594

Selecting Configuration Deployment Job Options | 594

Viewing Configuration Deployment Job Details | 595

Canceling Configuration Deployment Jobs | 596

Deploy Configuration Window | 596

Importing Configuration Data from Junos OS Configuration Groups | 598

Enabling Import of Configuration Group Data | 598

Viewing Configuration Group Data | 599

Using the Configuration or Pending Configuration Window | 600

Deploying Configuration Group Changes to Devices Immediately | 601

Scheduling Configuration Group Change Deployment | 601

Specifying Configuration Deployment Scheduling Options | 601

Using the Deploy Configuration Errors/Warnings Window	602
Enabling High-Frequency Traffic Statistics Monitoring on Devices	602
Configuring Network Traffic Analysis	604
Approving Change Requests	606
Enabling SNMP Categories and Setting Trap Destinations	608
Viewing Eligible Devices for Trap Forwarding	608
Enabling Trap Forwarding	609
Deploying SNMP Trap Configurations	610
Understanding Resynchronization of Device Configuration	611
The Resynchronize Device Configuration Task	612
How Resynchronization Works in NSOR Mode	613
How Resynchronization Works in SSOR Mode	614
How Network Director Resynchronizes the Build Mode Configuration	616
Resynchronizing Device Configuration	616
The Resynchronize Device Configuration List of Devices	618
Resynchronizing Devices When Junos Space Is in NSOR Mode	619
Resynchronizing Devices When Junos Space Is in SSOR Mode	619
Resynchronizing Devices in Manual Approval Mode	620
Viewing the Network Changes	620
Viewing Resynchronization Job Status	621
Managing Device Configuration Files	622
Selecting Device Configuration File Management Options	622
Backing Up Device Configuration Files	623
Restoring Device Configuration Files	624
Viewing Device Configuration Files	624
Comparing Device Configuration Files	625
Deleting Device Configuration Files	625
Managing Device Configuration File Management Jobs	625
Creating and Managing Baseline of Device Configuration Files	626
Selecting Baseline Management Options	627
Baselining Device Configuration Files	628
Restoring Baseline Device Configuration Files	629
Viewing Baseline Configuration Files	629
Comparing Baseline Configuration with Current Configuration	629

- Deleting Baseline | 630
- Managing Baseline Management Jobs | 630

Deploying and Managing Software Images | 632

Managing Software Images | 632

- Selecting Software Image Management Options | 633
- Adding Software Images to the Repository | 633
- Using the Device Image Upload Window | 634
- Viewing Software Image Details | 634
- Using the Device Image Summary Window | 634
- Deleting Software Images | 635

Deploying Software Images | 635

- Specifying Software Deployment Job Options | 636
- Selecting Software Images To Deploy | 637
- Selecting Options for Software Deployment | 638
- Summary of Software Deployment | 640

Managing Software Image Deployment Jobs | 640

- Selecting Software Image Management Options | 641
- Viewing Software Image Job Details | 642
- Using the Device Image Staging Window | 642
- Canceling Software Image Jobs | 643

Managing Devices | 645

- Enabling or Disabling Network Ports on Switches | 645
- Converting the QSFP+ Ports on QFX Series Devices | 646
 - Selecting Devices | 646
 - Converting Ports | 648
 - Reviewing and Deploying Port Conversions | 648

Setting Up Zero Touch Provisioning for Devices | 649

- Understanding Zero Touch Provisioning in Network Director | 649
- Configuring and Monitoring Zero Touch Provisioning | 650
 - Configuring Zero Touch Provisioning | 651
 - Specifying the Server Details | 652
 - Specifying the Software Image and Configuration Details | 654

- Reviewing and Modifying Zero Touch Provisioning Settings | 655
- What To Do Next | 655
- Configuration Statements for Custom Configuration of DHCP Server | 655
- Monitoring Zero Touch Provisioning Profiles | 656

Monitoring Devices and Traffic

About Monitor Mode | 658

Understanding Monitor Mode in Network Director | 658

- Scope and Monitor Tab Availability | 659
- Monitors and Tasks | 660
- Scope and Data Aggregation | 661
- How Network Director Collects and Displays Monitoring Data | 662
- How Network Director Displays and Stores Trend Data | 662
- More About the Monitor Tabs | 663

- The Summary Tab | 663
- The Traffic Tab | 663
- The Client Tab | 664
- The Equipment Tab | 664
- The Fabric Analysis Tab | 664
- VC Equipment Summary Tab | 664

Understanding the Monitor Mode Tasks Pane | 665

Monitoring Traffic | 670

Monitoring Traffic on Devices | 670

Monitoring Port Traffic Statistics | 671

- Procedure for Monitoring Port Traffic Statistics | 671
- Port on Device Window | 672
- Port Traffic Stats Window | 673

Monitoring Traffic on Layer 3 VLANs | 674

- Procedure for Monitoring Layer 3 VLAN Traffic Statistics | 674
- L3 VLAN Traffic Stats Window | 675

Monitoring Routing Instances | 676

- Procedure for Monitoring Routing Instances | 677

- Show Routing Instances Window | 677

- Show Interfaces Window | 678

- Show Bridge Domains Window | 679

- Show Connections | 680

- Show Routing Tables | 683

- Show MAC Table | 686

Monitoring Port Utilization | 687

- How to Access the Port Utilization Task | 687

- Port Utilization Details Window | 688

- Utilization for Device Window | 688

- Device View | 689

- Port View | 689

- Utilization for Fabric Devices Window | 690

- Device View | 690

- Port View | 691

Monitoring Tenant Details | 691

- Viewing the List of Tenants | 693

- View Port Details of Tenants | 694

- View Endpoints | 694

- View the Port Utilization Trend for a VXLAN Port | 696

Monitoring Virtual Chassis Protocol Statistics | 697

- Procedure for Monitoring Virtual Chassis Protocol Statistics | 697

- Virtual Chassis Protocol Statistics Window | 697

Viewing Congestion Events | 699**Monitoring Client Sessions | 701****Finding User Sessions | 701**

- Procedure for Finding User Sessions | 701

- Search User Session Window | 702

Finding End Points | 704

- Procedure for Finding End Points | 704

- Find End Point Window | 705

Refreshing End Point Information | 705

Monitoring Client Sessions | 706

Monitoring Devices | 708

Comparing Device Statistics | 708

Procedure for Comparing Device Statistics | 708

Compare Interfaces Window | 709

Showing ARP Table Information | 710

Procedure for Showing ARP Table Information | 710

Show ARP Table Information Window | 710

Viewing PoE Information | 711

Procedure for Viewing PoE Information | 711

Show PoE Information Window | 712

Monitoring the Status of Logical Interfaces | 713

Locating Information about Logical Interfaces | 713

Show Logical Interface Information Table | 713

Monitoring the Status of a Virtual Chassis | 715

Monitoring the Status of Virtual Chassis Members | 716

Monitoring and Analyzing Fabrics | 717

Monitoring Junos Fusion Fabric Systems and Components | 717

Analyzing Virtual Chassis Fabrics | 718

Procedure for Analyzing a Virtual Chassis Fabric | 719

Using the Fabric Health Check Tab | 719

Using the Topology Tab | 719

Monitoring Virtual Networks | 723

Using Monitor Mode for Virtual Devices | 723

Current Active Alarms Monitor | 724

Viewing vMotion History in Network Director | 726

General Monitoring | 728

Selecting Monitors To Display on the Summary Tab | 728

Changing Monitor Polling Interval and Data Collection | 729

Pinging Host Devices | 729

Troubleshooting Network Connections Using Traceroute | 731

Monitor Reference | 733

802.11 Packet Errors Monitor | 734

Access vs. Uplink Port Utilization Trend Monitor | 734

Current Sessions Monitor | 735

Current Sessions by Type Monitor | 735

Current Sessions by Type | 736

Current Session Details | 736

Error Trend Monitor | 736

Error Trend | 737

Error Trend Details | 737

Equipment Summary By Type Monitor | 739

Equipment Summary By Type | 739

Equipment Summary By Type Details | 739

Node Device Summary Monitor | 740

Port Status Monitor | 740

Port Status Summary | 741

Port Status Details | 741

Port Status for IP Fabric Monitor | 743

Port Utilization Monitor | 743

Power Supply and Fan Status Monitor | 744

Power Supply and Fan Status | 744

Power Supply and Fan Status Details | 745

Resource Utilization Monitor for Switches, Routers, Virtual Chassis, and Virtual Chassis Fabrics | 745

Resource Utilization Summary | 746

Resource Utilization Details | 746

Status Monitor for Junos Fusion Systems | 747

Status Monitor for Layer 3 Fabrics | 748

Status Monitor for Switches and Routers | 748

Status Monitor for Virtual Chassis | 749

Status Monitor for Virtual Chassis Members | 751

Top Talker - Wired Devices Monitor | 752

Top Talker - Wired Devices Summary | 752

Top Talker - Wired Devices Details | 752

Traffic Trend Monitor | 753

Unicast vs Broadcast/Multicast Monitor | 753

Unicast vs Broadcast/Multicast Trend Monitor | 754

User Session Details Window | 755

Virtual Chassis Topology Monitor | 756

VC Equipment Summary By Type Monitor | 758

VC Equipment Summary By Type | 758

VC Equipment Member Status By Type Details | 758

Using Fault Mode

About Fault Mode | 761

Understanding Fault Mode in Network Director | 761

What Are Events and Alarms? | 761

Alarm Severity | 762

Alarm Classification | 762

Alarm State | 764

Alarm Notifications | 764

Threshold Alarms | 764

Understanding the Fault Mode Tasks Pane | 765

Using Fault Mode | 766

Customizing Alarms | 766

Searching Alarms | 766

Changing Alarm State | 769

Fault Reference | 771

Alarm Detail Monitor | 771

Finding Specific Alarms | 772

Sorting Alarms | 773

Reading Events | 774

Investigating Event Attributes | 775

Changing the Alarm State | 775

Current Active Alarms Monitor | 776

Alarms by Category Monitor | 778

Alarms by Severity Monitor | 778

Alarms by State Monitor | 779

Alarm Trend Monitor | 780

Working in Report Mode

About Report Mode | 782

Understanding Report Mode in Network Director | 782

Understanding the Report Mode Tasks Pane | 784

Understanding the Types of Reports You Can Create | 786

Creating and Managing Reports | 787

Managing Reports in Network Director | 787

How to Locate and Manage Reports | 788

Managing Report Definitions | 788

Creating Reports | 790

How to Create a Report Definition | 790

Creating a Report Definition | 792

Setting Report Options | 794

Reviewing the Report Definition | 795

Changing a Report Definition | 796

Scheduling Reports | 796

How to Create or Manage Schedules | 797

Managing Schedules | 797

Creating New Schedules | 798

Editing Schedules | 800

Deleting Schedules | 800

Managing Generated Reports | 801

Reviewing Generated Reports | 801

Viewing Report Details | 802

Exporting Reports | 803

Deleting Generated Reports | 803

Retaining Reports | 803

Managing Reports on SCP Servers | 804

How to Configure SCP Servers | 804

Managing SCP Servers | 805

Mailing Reports | 806

How to Configure SMTP Servers | 807

Managing SMTP Servers | 807

Adding or Editing SMTP Server Settings | 809

Report Reference | 810

Active User Sessions Report | 810

Alarm History Report | 812

Alarm History Header | 812

Alarm History Tables | 813

Alarm Summary Report | 815

Alarm Summary Header | 815

Alarm Summary Charts | 816

Audit Trail Report | 817

Client Devices Report | 819

Device Inventory Report | 820

Fabric Analyzer Report | 822

Network Device Traffic Report | 824

Network Device Traffic Report Header | 824

Network Device Traffic Charts | 825

Port Bandwidth Utilization Report | 826

Top Users by Data Usage Report | 827

Top Users by Data Usage Header | 827

Top Users of Data Table | 828

Traffic and Congestion Summary Report | 828

Working with Network Director Mobile

About Network Director Mobile | 831

Overview of Network Director Mobile | 831

Getting Started with Network Director Mobile | 832

Network Director Mobile System Requirements | 832

Logging Into Network Director Mobile | 833

Understanding the Network Director Mobile User Interface | 833

- Dashboard Mode | 833

- Devices Mode | 833

Configuring Network Director Mobile Settings | 834

Working in the Network Director Mobile Dashboard Mode | 835

Monitoring Network-Wide Activity Using Network Director Mobile | 835

Network Director Mobile Dashboard Reference | 835

- Network Summary Monitor | 836

- Alarms Monitor | 836

- Top Sessions Monitor | 837

- Ports Monitor | 838

- Session Count Monitor | 838

- Session Trend Monitor | 838

Working in the Network Director Mobile Devices Mode | 839

Locating a Device and Viewing Device Properties Using Network Director Mobile | 839

Monitoring Sessions on a Device Using Network Director Mobile | 840

About the Documentation

IN THIS SECTION

- Documentation and Release Notes | xxix
- Documentation Conventions | xxix
- Documentation Feedback | xxxii
- Requesting Technical Support | xxxii

Use this guide to familiarize with the Junos Space Network Director GUI and how to use Network Director to provision and manage services on EX Series and QFX Series switches.

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 xxx](#) 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 xxx 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
<i>Italic text like this</i>	<ul style="list-style-type: none"> Introduces or emphasizes important new terms. Identifies guide names. Identifies RFC and Internet draft titles. 	<ul style="list-style-type: none"> A policy <i>term</i> is a named structure that defines match conditions and actions. <i>Junos OS CLI User Guide</i> RFC 1997, <i>BGP Communities Attribute</i>

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

Convention	Description	Examples
<i>Italic text like this</i>	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 <i>domain-name</i>
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	<ul style="list-style-type: none">• 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 <i>metric</i> >;
(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 (<i>string1</i> <i>string2</i> <i>string3</i>)
# (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 [<i>community-ids</i>]
Indentation and braces ({ })	Identifies a level in the configuration hierarchy.	[edit] routing-options { static { route default { nexthop <i>address</i> ; 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.	<ul style="list-style-type: none"> 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

Working With Network Director

About Network Director | **2**

Accessing Network Director | **16**

Understanding Network Director System Administration and Preferences | **19**

Getting Started with Network Director | **49**

About Network Director

IN THIS CHAPTER

- [Understanding Network Director and the Management Life-Cycle Modes | 2](#)
- [Understanding the Network Director User Interface | 3](#)

Understanding Network Director and the Management Life-Cycle Modes

Junos Space Network Director enables unified management of EX Series Ethernet Switches, MX Series routers, QFX Series switches, Virtual Chassis fabrics, Layer 3 fabrics, and Junos Fusion Enterprise in your network. It provides for full network life cycle management by simplifying the discovery, configuration, visualization, monitoring, and administration of large networks containing physical and virtual devices. You can quickly deploy a network by using it, configure it optimally to improve network uptime and maximize resources, and respond agilely to the needs of applications and users.

The Network Director user interface is based on the network management life-cycle. The interface provides five main working modes that are aligned to the network management life-cycle, and a sixth mode for working with Network Director itself. Each mode provides access to different tasks:

- **Build mode**—Use Build mode to build your network in Network Director. You use Build mode to discover the devices in your network, to create and manage device configurations, and to manage devices. You can also organize your devices into hierarchical groupings based on logical relationships or by physical locations.
- **Deploy mode**—Use Deploy mode to deploy and manage changes to devices. In Deploy mode, you deploy the configurations you built in Build mode, install new software images on your devices, and manage device configuration files.
- **Monitor mode**—Use Monitor mode to gain visibility into your network performance and health. Monitor mode provides a host of information about your network such as the operational status of devices, traffic patterns and trends, client session statistics, port capacity, and interference patterns. You can also search for a user and view a history of the user sessions.

- **Fault mode**—Use Fault mode to gain visibility into unexpected network events and to manage faults or notifications.
- **Report mode**—Use Report mode to generate reports from the data that Network Director stores on network performance, status, and activity.

In addition to these modes, Network Director enables you to perform system-level tasks from the System button and the Preferences button. System-level tasks include viewing the Network Director user and system audit trail, managing jobs, and gathering logs for troubleshooting.

Benefits of Network Director

- Enables unified management of wired, and virtual infrastructures for the campus.
- Improves operational efficiency by automating routine management tasks such as device and port provisioning.
- Supports flexible, large-scale deployment of devices. For example, Build mode enables you to apply configurations across multiple devices grouped by logical relationships, physical locations, or type.

RELATED DOCUMENTATION

Understanding the Network Director User Interface 3
Understanding Build Mode in Network Director 77
Understanding Deploy Mode in Network Director 572
Understanding Monitor Mode in Network Director 658
Understanding Fault Mode in Network Director 761
Understanding Report Mode in Network Director 782
Understanding Network Director User Administration 19

Understanding the Network Director User Interface

IN THIS SECTION


- [Network Director Banner | 4](#)
- [View Pane | 6](#)
- [Tasks Pane | 9](#)
- [Alarms | 11](#)

- Main Window or Workspace | 11
- Tables in Network Director | 11

Junos Space Network Director provides a simple to use, HTML5-based, Web 2.0 user interface that you can access through standard Web browsers. The user interface is task-oriented, using task-based workflows to help you accomplish administrative tasks quickly and efficiently. It provides you the flexibility to work with single devices or with multiple devices grouped by logical relationship, location, or device type. You can filter, sort, and select columns in tables, making looking for specific information easy.

Figure 1 on page 4 illustrates the main components of the interface.

Figure 1: The Network Director User Interface Components



This topic describes:

Network Director Banner

Use the Network Director banner, shown in Figure 2 on page 5, to select the working mode. You can also use the Network Director banner to perform other global tasks, such as setting up your preferences or accessing Junos Space. Table 3 on page 5 describes the functions available to you on the banner.


Figure 2: Network Director Banner



Table 3: Network Director Banner Functions

Item	Function
Accessing Junos Space Platform	Click to exit Network Director and open the Junos Space Network Application Platform. You can switch back and forth between Network Director and Junos Space without logging in again.
Network View Selector	<p>Select the network view that you want to work in. You can choose from one of the following views:</p> <ul style="list-style-type: none"> • Dashboard View • Logical View • Location View • Device View • Custom Group View • Datacenter View • Topology View <p>For more details, see “Displaying Devices in Various Network Views” on page 7.</p>
Mode Icons	<p>Select the mode you want to work in.</p> <p>NOTE: You might not have access to all the Network Director modes. What modes you have access to depends on your assigned user role.</p>
User Log out	<p>Displays the username using which you logged in to Network Director.</p> <p>Click the Down arrow next to the username and select Logout to log out of Network Director and Junos Space.</p>

Table 3: Network Director Banner Functions (*continued*)

Item	Function
System Tasks	<p>Access the system tasks such as viewing audit logs, jobs, and to collect troubleshooting logs.</p> <p>Click the Down arrow next to System and select Preferences to set your Network Director user and system preferences.</p>
System Preferences, Product Information, and Online Help 	<p>Click this button and select an appropriate option:</p> <ul style="list-style-type: none"> • Preferences—Enables you to set your Network Director user and system preferences. • Help—Open searchable help. This help icon is not context-sensitive—it always opens help to the first page. From here, you can browse or search the help. Context-sensitive help is available from the help icon provided on each pane or page. • About—Displays information about Network Director, such as the currently running version.

In addition to this, Network Director displays the date and time in the local time zone in the bottom right corner.

View Pane

IN THIS SECTION

- [Displaying Devices in Various Network Views | 7](#)
- [Filtering the Network Tree | 7](#)
- [Expanding or Collapsing Nodes in the Network Tree | 8](#)
- [Searching the Network Tree | 9](#)

In the View pane, Network Director provides you a unified, hierarchal view of your wired networks in the form of a expand tree that is expandable and collapsible. By selecting both a view and a node in the tree, you indicate the *scope* over which you want an operation or task to occur. For example:

- By selecting the Access node in Logical View, you indicate that the scope for a task is all access switches under the Access node.
- By selecting a floor node in Location View, you indicate that the scope for a task is all devices belonging to that floor.

- By selecting the EX2300 node in Device View, you indicate that the scope for a task is all EX2300 switches in your network.

You can perform the following actions in the View pane:

Displaying Devices in Various Network Views

Use the selection box in the Network Director banner to choose one of the following network views:


- **Dashboard View**—This is a customizeable view that provides information about your network, and is the default view that opens when you log in. You can select and add monitoring widgets to the Dashboard View based on your requirements. This is the default view that opens when you log in to Network Director.
- **Logical View**—Devices are organized by their logical relationships in the network. All switches appear in the Switching Network and are categorized by their role in the network: access, aggregation, or core.
Network Director builds most of this view for you as you discover devices. However, you need to manually assign switches to the access, aggregation, or core categories.
- **Location View**—Devices are organized by their physical locations. You build this view by creating sites, building, floors, aisles, racks, outdoor areas, and then assigning your switches to these locations.
- **Device View**—Devices are organized by device type: switches. Within device type, devices are organized by device model. For example, all models of EX2300 switches are grouped together under one node in the tree.
- **View**—Displays devices that are part of your network
- **Custom Group View**—If you have defined one or more custom groups, Network Director displays these custom groups in this view. You can manually add devices to a custom group or define a rule to automatically add devices to the custom group once they are discovered in Network Director. The devices are grouped under each custom group.
- **Topology View**—Topology enables you to view all the discovered devices in your network, overlaid on a map where the devices are located across sites, buildings, floors, closets, aisles, and racks along with their physical interconnection with other devices in your network. Topology also provides visualization around physical and logical connectivity between various discovered interconnected devices.
- **Datacenter View**—Datacenter view enables you to view all the data centers in Network Director. Expand each data center to view servers that are part of the data center.

Filtering the Network Tree

To make it easier for you to focus on selected aspects of your network, you can apply predefined filters to your network tree so that only nodes and devices that meet the filter criteria are shown. For example, you can apply a filter so that only devices in a specific building are shown in the network tree in all views.

To apply filters:

1. Click the filter icon:



2. In the Filters dialog box, click **Show available filters**.

The Available Filters section of the dialog box appears.

3. Under Available Filters, click the tab for the view you want to use to define your filter. For example, if you want to filter on devices—that is, show only certain types of devices—click the **Device** tab.


The filters that you can apply are listed below the tab.

4. To select a filter, click its associated plus icon.

The filter appears in the Selected Filters section of the dialog box. You can repeat Steps 3 and 4 until you have selected all the filters you want apply.

5. Click **Apply**.

The Filters dialog box closes and the filters are applied. The filter icon changes appearance to indicate that filters have been applied:



To remove a filter, click the filter icon, click the trash can next to the filter in the Selected Filters list, and click **Apply**.

Expanding or Collapsing Nodes in the Network Tree

To expand a node in the network tree, select the node and then click the **Expand All** icon:



The node you selected and any child nodes under the selected node are expanded to show their contents.

Similarly, to collapse a node in the network tree, select the node and then click the **Collapse All** icon (next to the Expand All icon). The node you selected is collapsed and no nodes under it are shown.

Searching the Network Tree

To quickly find and select a device or device group, use the search function.

To perform a search, type three or more characters into the Search box and click the **Search** icon, as shown in [Figure 3 on page 9](#).

Figure 3: Performing Search in the View pane



Network Director finds the first instance of a node whose name contains the characters. To find the next instance, click the right arrow.

Searches are not case-insensitive: a search on *wla115* and one on *WLA115* return the same results.


Tasks Pane

The Tasks pane is available in every mode and lists tasks specific to that mode. In addition to changing according to the mode selected, tasks listed in the Tasks pane can change as you select different scopes in the View pane. For example, some tasks are appropriate only at the device level and thus appear only when you have selected an individual device.

Clicking a task brings up task-specific content in the main window.

In general, to perform a task in Network Director, you navigate to the task as shown in [Figure 4 on page 10](#). You select your mode, your scope, and then your task.

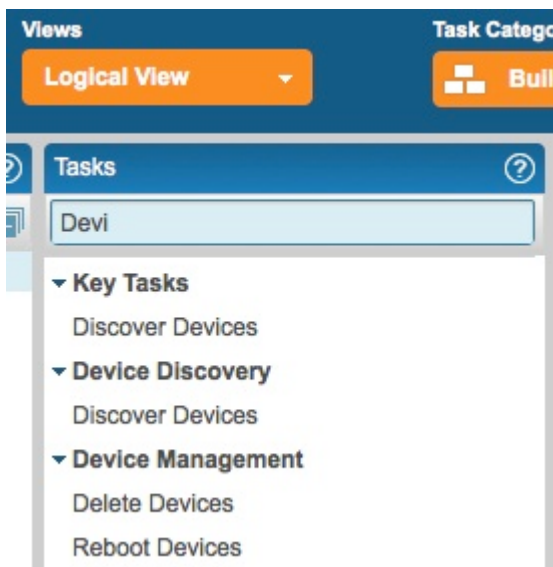
Figure 4: Navigating to a Task in a Tasks Pane



TIP: The location of the Tasks pane changes with mode. In Build and Deploy mode, it is adjacent to the View pane. In Monitor, Fault, and Report mode, it is located to the right of the main window.

Use the Search box in the Tasks pane to easily locate a task, as shown in [Figure 5 on page 10](#). To perform a search, type three or more characters into the Search box and press Enter.

Figure 5: Performing Search in the Tasks pane



Alarms

The Alarms bar that is displayed at the bottom of your browser window provides a quick summary of how many critical, major, minor, and info alarms are currently active in the network and is visible in every mode. To display more information about alarms, click the alarm count or the Alarms banner. You are automatically placed in Fault mode and the Fault mode monitors are displayed.

Main Window or Workspace

The main window or workspace displays the content relevant to the mode, scope, and task you have selected. When you log in to Network Director, this pane displays the Device Inventory page. The Device Inventory page is the default landing page for Build and Deploy modes. It contains a list of the devices for your current scope. It includes pie charts that permit you to see at a glance the connection states, configuration synchronization states, and device-type distribution for your devices.

Tables in Network Director

IN THIS SECTION

- [Moving and Resizing Columns | 11](#)
- [Displaying the Column Drop-Down Menu | 12](#)
- [Sorting on a Column | 12](#)
- [Hiding and Exposing Columns | 13](#)
- [Searching Table Contents | 13](#)
- [Filtering Table Contents | 15](#)

Tables are used throughout Network Director to display data. These tables share common features. By becoming familiar with these features, you can navigate and manipulate tabular data quickly and efficiently. The following sections describe:

Moving and Resizing Columns

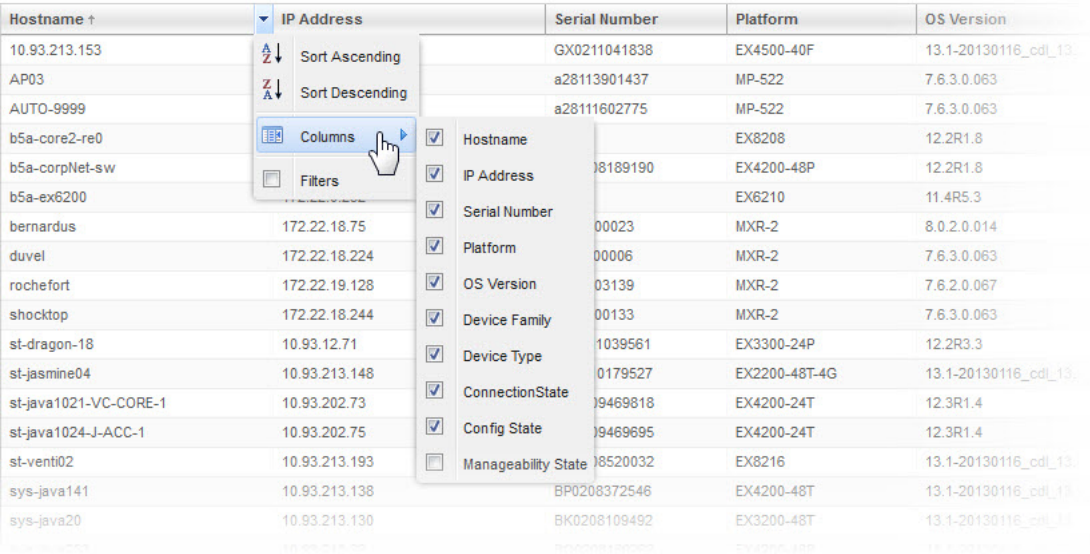
You can reposition and resize columns in a table. To move a column, drag and drop the column head to the new location. Network Director displays a green check mark when you mouse over a valid column location.

To resize a column, mouse over the edge of a column until the cursor becomes two vertical lines with outward arrows. Drag the column width to the new size.

Displaying the Column Drop-Down Menu

A drop-down menu is available from each column head, allowing you to perform additional operations on columns. To display the column drop-down menu, mouse over the column head. A downward arrow appears. By clicking the arrow, you display the drop-down menu, as shown in [Figure 6 on page 12](#).

Figure 6: Column Drop-Down Menu



Hostname	IP Address	Serial Number	Platform	OS Version
10.93.213.153		GX0211041838	EX4500-40F	13.1-20130116_cdl_13.1
AP03		a28113901437	MP-522	7.6.3.0.063
AUTO-9999		a28111602775	MP-522	7.6.3.0.063
b5a-core2-re0			EX8208	12.2R1.8
b5a-corpNet-sw		8189190	EX4200-48P	12.2R1.8
b5a-ex6200			EX6210	11.4R5.3
bernardus	172.22.18.75	00023	MXR-2	8.0.2.0.014
duvel	172.22.18.224	00006	MXR-2	7.6.3.0.063
rochefort	172.22.19.128	03139	MXR-2	7.6.2.0.067
shocktop	172.22.18.244	00133	MXR-2	7.6.3.0.063
st-dragon-18	10.93.12.71	1039561	EX3300-24P	12.2R3.3
st-jasmine04	10.93.213.148	0179527	EX2200-48T-4G	13.1-20130116_cdl_13.1
st-java1021-VC-CORE-1	10.93.202.73	09469818	EX4200-24T	12.3R1.4
st-java1024-J-ACC-1	10.93.202.75	09469695	EX4200-24T	12.3R1.4
st-venti02	10.93.213.193	08520032	EX8216	13.1-20130116_cdl_13.1
sys-java141	10.93.213.138	BP0208372546	EX4200-48T	13.1-20130116_cdl_13.1
sys-java20	10.93.213.130	BK0208109492	EX3200-48T	13.1-20130116_cdl_13.1

Sorting on a Column

You can sort the table on a column by clicking the column head—each click changes the direction of the sort. In addition, you can use the Sort Ascending and Sort Descending options in the drop-down menu.

When you sort on a column, a small arrow appears next to the column name to indicate that the table is being sorted by the column and the direction of the sort.

Network Director uses a lexical sort for tabular data that is not strict numeric data, which means that data such as IP addresses do not sort in numerical sequence, as shown in [Table 4 on page 13](#).

Table 4: Numerical Sorts and Lexical Sorts

Numerical Sort	Lexical Sort
10.93.200.65	10.93.200.129
10.93.200.129	10.93.200.199
10.93.200.199	10.93.200.65

Hiding and Exposing Columns

You can customize your tables by hiding or exposing columns. This way, you can choose to see only relevant information.

To hide or expose columns, display the drop-down menu for any column head and mouse over the Columns option, as shown in [Figure 6 on page 12](#). Select a column to expose it—clear a column to hide it.

As a general rule, Network Director displays all columns in a table by default. However, some tables have more columns than can fit easily within the page. In these tables, some columns are hidden by default.

Searching Table Contents

You can search for specific data in large tables by using search criteria.

To search for an item in a table, enter the search term in the text box. Select ANY for Network Director to search for the term in all columns in the table. Every table has a predefined default column that the system searches first; before it proceeds to search other columns.

You can also choose to search a particular column for a term. Network Director displays a list of all the columns in a table. To search a particular column for a term, select that column for the list.

NOTE: When you enter a search expression, note the following:

- You must add a back slash “\” if you want to use the following special characters in the search text:

+ ~ && || ! () { } [] ^ “ ~ * ? : \

- Field names are case-sensitive.

For example, if you have a few systems running on Junos OS 19.1 Release 1, then `os: 19.1R1` returns search results, whereas `OS: 19.1R1` does not return search results. This is because the field name that is indexed is `os` and not `OS`.

- If you want to search for a term that includes a space, enclose the term within double quotation marks.

For example, to search for all devices that are synchronized (that is, In Sync), enter “In Sync” in the Search field.

- You must append “*” if you want to search using partial keywords. Otherwise, the search returns 0 (zero) matches or hits.

You can filter search results by specifying one or more search terms. Network Director uses the AND operator for each search term that you enter. Network Director lists the search results in the table, depending on the search criteria that you specified.

For example, perform the following steps to search for an EX2300 switch that is running Junos OS Release 19.1:

1. Enter **EX2300** as the search term in the text box.
2. From the list that appears, select to search the Platform column.
Network Director lists all the EX2300 switches in your network.
3. Enter **19.1** as the search term after the comma separator in the text box.
4. From the list, select to search from the OS Version column.

Network Director lists all the EX2300 switches in your network that are running Junos OS Release 19.1.

Filtering Table Contents

For large tables, it is helpful to be able to sort data to show only relevant entries. When you mouse over the Filters option in the column drop-down menu, a fill-in box appears where you can type filter criteria. If you type a text string and click **Go**, entries that do not contain the text string (filter criterion) are removed from the table. A red asterisk appears on the column head to indicate that the column has been filtered. To restore all entries to the table, clear the Filters option.

For example, to filter the Device Inventory page so that only devices in the **192.168.1.0** subnet are displayed:

1. Mouse over the right side of the IP Address column head to expose the downward arrow.
2. Click the arrow to display the column drop-down menu.
3. Mouse over **Filters** to display the Filter field.
4. Type **192.168.1.** in the field and click **Go**.

Only the devices in the **192.168.1.0** subnet are shown.

RELATED DOCUMENTATION

[Understanding Network Director and the Management Life-Cycle Modes | 2](#)

[Network Director Documentation home page](#)

Accessing Network Director

IN THIS CHAPTER

- [Logging In to Network Director | 16](#)
- [Logging Out of Network Director | 17](#)
- [Changing Your Password | 18](#)

Logging In to Network Director

You connect to Network Director using your Web browser. The following Web browsers are supported: Mozilla Firefox version 72.0.2 (64-bit) or later, and Google Chrome version 86 and later. The minimum screen resolution is 1280 x 1024.

You can connect to Network Director one of two ways:

- Log in to Network Director directly by using the following URL:

```
https://<n.n.n.n>/networkdirector
```

where *n.n.n.n* is the IP address of the Junos Space Web interface. You can bookmark the login page for future use.

- Log in to Junos Space first by using the following URL:

```
https://<n.n.n.n>/mainui
```

where *n.n.n.n* is the IP address of the Junos Space Web interface.

You can then switch to the Network Director interface by selecting Network Director from the Applications list in the left pane of the Junos Space user interface.

The default username and password is the same for both Junos Space and Network Director:

- username—super
- password—juniper123

RELATED DOCUMENTATION

[Logging Out of Network Director | 17](#)

[Changing Your Password | 18](#)

[Understanding the Network Director User Interface | 3](#)

[Network Director Documentation home page](#)

Logging Out of Network Director

When you are finished using Network Director, log out to prevent unauthorized access. To log out of Network Director, click the username in the Network Director banner and select Logout from the list.

Network Director automatically logs you out if you have not performed any action, such as keystrokes or mouse clicks, for a set period of time. This automatic logout conserves server resources and protects the system from unauthorized access. By default, automatic logout occurs if a session has been idle for 60 minutes.

Network Director uses the same automatic logout period as Junos Space. To change the automatic logout period:

1. Click the System Platform icon.
2. Navigate to **Administration > Applications**.
3. Right-click **Network Management Platform** and select **Modify Application Settings..**
4. In the Modify Network Management Settings page, select **User**.

RELATED DOCUMENTATION

[Logging In to Network Director | 16](#)

[Changing Your Password | 18](#)

[Understanding the Network Director User Interface | 3](#)

[Network Director Documentation home page](#)

Changing Your Password

Any user, regardless of user role, can change his or her password.

Your username and password are the same in Junos Space and Network Director. To change your password, change it in Junos Space:

1. From Network Director, click the Junos Space icon in the Network Director banner.
2. Click the User Password icon in the Junos Space banner.
3. Follow the instructions to change your password.

RELATED DOCUMENTATION

[Logging In to Network Director | 16](#)

[Logging Out of Network Director | 17](#)

[Understanding the Network Director User Interface | 3](#)

[Network Director Documentation home page](#)

Understanding Network Director System Administration and Preferences

IN THIS CHAPTER

- [Understanding Network Director User Administration | 19](#)
- [Understanding the System Tasks Pane | 20](#)
- [Audit Logs Overview | 21](#)
- [Viewing Audit Logs From Network Director | 22](#)
- [Managing Jobs | 23](#)
- [Collecting Logs for Troubleshooting | 25](#)
- [Setting Up User and System Preferences | 26](#)

Understanding Network Director User Administration

Network Director uses the user administration features of the Junos Space platform on which it runs. Use Junos Space for tasks such as adding, deleting, and editing user accounts and roles, and changing user passwords. Refer to the Junos Space documentation for information about user administration.

When Network Director is installed, some additional user administration options are available in Junos Space, which are specific to Network Director:

- In addition to the Super Administrator role, the following predefined roles are available for Network Director users:

Network Director - Admin—Has access to all the Network Director tasks. This role is the system administrator role and has full privileges.

Network Director - Engineer—Has access to either all the device management tasks or only those device management sub-tasks to which the Engineer role is mapped. These users can also view the device monitoring and fault management tasks.

Network Director - Monitor—Has access to monitor the network status and performance or view the faults to determine the health of your network and take appropriate action.

Network Director - Configuration Approver—Has access to provide additional privileges to approve the configuration changes in addition to all the tasks that a Config Admin can perform.

Network Director - Image Admin—Has access to all the image management tasks.

Network Director - Config Admin—Has access to create, edit, delete, assign, deploy profiles, and manage fabrics (VCF and IP Fabric).

- You can create custom roles to grant users different access rights to the Network Director modes, group, dashboard widgets, and tasks. Users can access only those portions of the navigation hierarchy to which they are explicitly granted access through access privileges.

If you try to log in to Network Director using an account that does not have access rights to any Network Director modes, you will be redirected to Junos Space instead.

Access to Network Director system preferences is controlled by user access rights. For more information, see [“Setting Up User and System Preferences” on page 26](#).

RELATED DOCUMENTATION

[Understanding the Network Director User Interface | 3](#)

[Setting Up User and System Preferences | 26](#)

[Network Director Documentation home page](#)

Understanding the System Tasks Pane

The System Tasks pane provides tasks for viewing audit logs of Network Director user activities, for managing jobs, and for collecting troubleshooting logs.

To access the System Tasks pane, click **System** in the Network Director banner. The tasks are described in [Table 5 on page 20](#).

Table 5: System Tasks

Task	Description
View Audit Logs	View a history of user activities on Network Director, including log in, log out, and task initiation and completion.
Manage Jobs	View all jobs that are scheduled to run or have been run by Network Director. You can cancel jobs that are in progress or scheduled to run in the future.

Table 5: System Tasks *(continued)*

Task	Description
Collect Jobs for Troubleshooting	Download a zip file containing logs and troubleshooting data from both Network Director and Junos Space.

RELATED DOCUMENTATION

Viewing Audit Logs From Network Director 22
Managing Jobs 23
Collecting Logs for Troubleshooting 25
Network Director Documentation home page

Audit Logs Overview

Audit logs provide a record of login history and user-initiated tasks that are performed from the user interface. From the Audit Logs page, you can monitor user login-logout activity over time, track device management tasks, view services that were provisioned on devices, and so forth. Audit logging does not record non-user initiated activities, such as device-driven activities, and is not designed for debugging purposes.

Administrators can sort and filter on audit logs to determine which users performed what actions on what objects at what time. For example, an administrator can use audit log filtering to track the 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, or monitor user login-logout activity over time.

Over time, Network Director will archive a large volume of log entries. Such log entries might or might not be reviewed, but they must be retained for a period of time.

The audit logs can be saved to a local server (the server that functions as the active node for Network Director) or a remote network host or media.

RELATED DOCUMENTATION

Viewing Audit Logs From Network Director 22
Network Director Documentation home page

Viewing Audit Logs From Network Director

Audit logs are generated for login activity and tasks that are initiated from the Network Director application. The Audit Logs page displays the logs for all user-initiated activities.

You can do the following on the Audit Logs page:

- Sort, filter, and search the log entries using the standard table manipulation features in Network Director.
- Obtain more information about a log entry by double-clicking the entry or by selecting the entry and clicking **Show Details**. The Audit Log Details window is displayed.
- For a user-initiated task that runs as a job, you can obtain more information about the job by clicking the job ID in the Job ID column.

To display the Audit Logs page:

1. Click **System** in the Network Director banner.
2. Select **View Audit Logs** from the Tasks pane.

The Audit Logs page is displayed with the fields listed in [Table 6 on page 22](#).

Table 6: Audit Logs Page Fields

Field	Description
User Name	The login ID of the user that initiated the task
User IP	The IP address of the client computer from which the user initiated the task
Task	The name of the task that triggered the audit log
Time	The data and time when the user initiated the task
Result	<p>The execution result of the task that triggered the audit log:</p> <ul style="list-style-type: none"> • Success—Job completed successfully • Failure—Job failed and was terminated • Job Scheduled—Job is scheduled but has not yet started
Description	A description of the audit log
Job ID	The job ID for any task that runs as a job

RELATED DOCUMENTATION

[Audit Logs Overview | 21](#)

[Managing Jobs | 23](#)

[Network Director Documentation home page](#)

Managing Jobs

Network Director enables you to view and manage jobs. You can view the status of completed jobs and cancel the jobs that are scheduled to execute at a later time or jobs that are in progress.

The Job Management page, accessible as a System task, enables you to view and manage all jobs. In addition, Network Director enables you to view special pre-filtered versions of this page from various other tasks, such as View Discovery Status, View Image Deployment Jobs, or View Baseline Mgmt Jobs. These pages contain the same fields (although some fields might be hidden) and have the same functionality as the Job Management page, but they list only those jobs relevant to particular tasks.

To display the Job Management page:

1. Click **System** on the Network Director banner.
2. Select **Manage Jobs** from the Tasks pane. The Job Management page appears.
3. To view the details of a job, select a row and click **Show Details** or double-click a row.
4. To cancel a scheduled job, select a job that is scheduled for a later time or a job that is in progress and click **Cancel**.

The fields in the Job Management page are described in [Table 7 on page 24](#). To view any hidden column, keep the mouse on any column heading and select the down arrow and then click Columns. Select the check box to display the hidden columns.

NOTE: Details of jobs initiated from Network Director will be available only from Network Director. These jobs will not be listed in the Job Management pane in Junos Space platform and vice-versa.

Table 7: Job Management Page Fields

Field	Description
Job ID	The unique ID assigned to the job
Name	The name of the job
Percent	The percentage of completion of the job
State	<p>The status of the job:</p> <ul style="list-style-type: none"> • Success—Job completed successfully • Failure—Job failed and was terminated • Job Scheduled—Job is scheduled but has not yet started • In progress—Job is has started, but not completed • Cancelled—Job is cancelled
Job Type	The type of the job
Summary	Summary of the job scheduled and executed with status
Scheduled Start Time	The time when the job is scheduled to start
Actual Start Time	The actual time when the job started
End Time	The time when the job was completed
User	The login ID of the user that initiated the task
Recurrence	The recurrent time when the job will be restarted.

RELATED DOCUMENTATION

[Audit Logs Overview | 21](#)
[Viewing Audit Logs From Network Director | 22](#)
[Network Director Documentation home page](#)

Collecting Logs for Troubleshooting

Network Director enables you to collect logs and other data from both Network Director and Junos Space that can assist in managing and monitoring Network Director servers.

Network Director collects the logs and troubleshooting data into a compressed file that you can download. This file is named **troubleshoot_yyyy-mm-dd_hh-mm-ss.zip**—for example, **troubleshoot_2012-12-21_11-25-12.zip**. The date and time in the file name is the server Coordinated Universal Time (UTC) date and time.

To retrieve troubleshooting data and log files, follow these steps:

1. Click **System** on the Network Director banner.
2. From the Tasks pane, click **Collect Logs for Troubleshooting**. The Collect Logs for Troubleshooting page appears.
3. Click the **Download troubleshooting data and logs from Network Director and Junos Space** link.
Network Director begins collecting the logs and data. It can take a few minutes for Network Director to collect the information and create the zip file.
4. When the standard file download window for your browser opens, save the **troubleshoot_yyyy-mm-dd_hh-mm-ss.zip** file.
5. When you contact the Juniper Technical Assistance Center, describe the problem you encountered and provide the JTAC representative with the **troubleshoot.zip** file.

[Table 8 on page 25](#) lists the files included in the **troubleshoot_yyyy-mm-dd_hh-mm-ss.zip** file.

Table 8: Log Files in the troubleshooting.zip File

Description	Location
Jboss log files	<code>/var/log/jboss/servers/server1</code>
MSS OS adapter log files	<code>/home/jmp/mssosadpater/var/errorLog/</code>
Daemon log files	<code>/opt/opennms/logs/daemon/</code>
Platform log files	<code>/var/log/platform</code>
Access Log Files	<code>/var/log/httpd</code>
Log files for Apache, NMA, Webproxy	<code>/var/log/httpd/</code>

Table 8: Log Files in the troubleshooting.zip File (*continued*)

Description	Location
Watchdog log file	<code>/var/log/</code>

RELATED DOCUMENTATION

[Managing Jobs | 23](#)

[Audit Logs Overview | 21](#)

[Viewing Audit Logs From Network Director | 22](#)

[Network Director Documentation home page](#)

Setting Up User and System Preferences

IN THIS SECTION

- [Accessing the Preferences Page | 27](#)
- [Choosing Server Time or Local Time | 27](#)
- [Specifying Search Preferences | 28](#)
- [Enabling Import of Configuration Group Data from Ethernet Design | 28](#)
- [Selecting the Approval Mode | 28](#)
- [Setting up Auto-resynchronization Preferences | 29](#)
- [Retaining Network Director Reports | 30](#)
- [Changing Monitor Mode Settings | 30](#)
- [Changing Alarm Settings | 36](#)
- [Modifying Data Center Synchronization Interval Using the Virtualization Tab | 47](#)

Depending on your privileges, the Preferences page displays either user settings or a combination of user settings and system settings. One or more of the following preference tabs appear when you open the Preferences page:

- User—All users can choose whether monitors and reports display the local time or the server time.
- Search—Network administrators can configure options for search indexing.
- Config & Deploy—Network administrators can:
 - choose to enable or disable import of configuration group data into Network Director.
 - specify the Auto Approval or Manual Approval mode for device configuration deployments.
- Monitoring—As a network administrator you can change the polling interval for data collection for Monitor mode monitors and enable or disable the internal processes used for data collection. You can also specify the IP address of the Data Learning Engine server, if installed, and the database record retention periods.
- Fault—As a network administrator you can enable or disable alarms. They can also set the retention period for alarms and the number of events per alarm.
- Report—Network administrators can specify the period of time for which Network Director reports are retained.
- Topology—Network administrators can specify a retention period for the deleted links in Topology view.
- Virtualization—Network administrators can modify the synchronization time interval between Network Director and the cloud infrastructure.

This topic describes:

Accessing the Preferences Page


To open the Preferences page, click  in the Network Director banner and select **Preferences** as shown in [Figure 7 on page 27](#).

Figure 7: Accessing the Preferences Page



The Preferences page opens with User Preferences as the default tab.

Choosing Server Time or Local Time

All users can specify whether Network Director displays local time or the server's time in monitors and reports on the User Preferences tab. The default setting is to display local time. To change the setting to display the server's time:

1. In the Preferences page, select **Use Server Time** from the list.
2. Click **OK** to save your changes or click **Cancel** to close Preferences.

Specifying Search Preferences

Network Director indexes the device inventory data periodically to enable users to perform efficient searches. You can specify a time interval after which Network Director initiates the next indexing on the Search tab. You can also specify to stop indexing while devices are imported into Network Director. If you are running short of system memory, selecting this option helps save some memory and speed up the discovery and import of new devices. By default, this option is selected and the search index update interval is set to 900 seconds.

Enabling Import of Configuration Group Data from Ethernet Design

For Network Director to be able to import configuration group data.

To enable the import of configuration group data:

1. In the Preferences window, select the **Config & Deploy** tab.
2. Select the **Enable migration from Ethernet Design** check box to enable import of configuration group data. By default this check box is not selected.
3. Click **Save** to save and close the preferences.

For detailed steps on importing configuration group data from Ethernet Design, see [“Importing Configuration Data from Junos OS Configuration Groups” on page 598](#).

Selecting the Approval Mode

Use the Config & Deploy tab of System Preferences to configure the approval mode:

1. Select the **Manual Approval** mode if you want an approver to review and approve the changes before they are deployed.

By default, **Auto Approval** mode is selected. Use this mode if you want to deploy the configuration changes without a prior approval.
2. If you select the Manual Approval mode, add one or more approvers' e-mail addresses to notify the approvers every time a change request is submitted.
3. Specify the rollback limit, which is the number of change requests that can be rolled back.

The default value is 50. You can roll back a maximum of 1000 change requests.

4. Specify the time after which a change request elapses after the time it was created.

The minimum and maximum number of days that you can specify after which a change request elapses is 1 day and 365 days respectively. The change requests are highlighted in the following colors that indicate their overdue status.

- Red color—Indicates that the change request is in overdue status.
- Orange color—Indicates that the change request is due in less than 2 days.
- Green color—Indicates that the change request is not yet due.

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

BEST PRACTICE: Configuring the approval mode must be a one-time operation. Do not change the approval mode frequently.

To change the approval mode from Auto Approval to Manual Approval, you must either deploy or discard the device configuration changes. You are unable to change the approval mode to Manual Approval, or from Manual Approval to Auto Approval if local changes are in pending deployment state. The message: **Do you want to retain the Change Request history?** is displayed when you change the approval mode. If you choose to retain the change request history, all the existing change requests are retained by the system. Hence, even if you switch to the Auto Approval mode, you can view the change requests that were created in Manual Approval mode.

NOTE: While configuring the Manual Approval mode, you can specify any number of approvers. If you specify more than one approver while configuring the Manual Approval mode, after any approver accepts or rejects a proposed change, the change request is not listed for the other approvers and they cannot approve or reject the same change request.

Setting up Auto-resynchronization Preferences

If you enable auto-resynchronization in Network Director, any configuration changes made on the physical device, including out-of-band CLI commits and change-request updates, automatically trigger resynchronization on the device.

To set up auto-resynchronization:

1. Select the **Config & Deploy** tab in the Preferences window.
2. Select the option **Purge unassigned system profiles after resynchronizing configuration**, which removes unassigned profiles generated by Network Director after resynchronization or deletion of a device.

NOTE: While upgrading Network Director, the profiles that are in unassigned state are not removed even if you select this option.

3. Specify the time interval in **Auto Resync TriggerWait Interval(sec)**. Network Director waits for this time interval before triggering auto-resynchronization.

The default time interval is 120 seconds.

4. Click **OK**.

Retaining Network Director Reports

By default, Network Director retains reports for 30 days. However, Network Administrators can change the retention period within the range 0 through 365 days. To change the setting, move the slider right or left on the Report tab of Preferences to the new setting. Click **OK** to save the setting.

Changing Monitor Mode Settings

IN THIS SECTION

- [Disabling Data Collection for Monitors | 31](#)
- [Changing the Polling Interval | 32](#)
- [Enabling and Disabling Collection for Managed Devices | 33](#)
- [Specifying Database History Retention | 34](#)
- [Specifying the Data Learning Engine \(DLE\) Settings | 34](#)

The Monitoring tab of Preferences has three tabs under it. These are:

- **Monitoring Settings**—Enables you to change the default polling interval for data collection for Monitor mode monitors. You can also disable or reenable the internal processes used for data collection on this sub-tab.
- **Client Session History**—Enables you to set the retention period for history records and the frequency that these records are checked for deletion.
- **Data Learning Engine Settings**—Enables you to specify the IP address of the Data Learning Engine (DLE) server or servers, which is a component of Cloud Analytics Engine that supports the flow path analysis and high-frequency statistics features in Network Director.
- **Device Settings**—Allows you to enable or disable data collection for one or more devices.

This section describes:

Disabling Data Collection for Monitors

Network Director internally gathers data for monitors by using a set of data collection processes. You can disable these data collectors if they do not pertain to your installation. For example, if you do not use Virtual Chassis, you can disable the data collection processes used for Virtual Chassis.

The data collection processes are divided into the following categories:

- Client
- Equipment
- FM
- Traffic
- Virtual

One data collector can be used by multiple monitors. Likewise, some monitors can be supported by multiple data collectors. These data collectors are enabled by default. To ensure proper data collection, if you enable the equipment data collectors, you must also enable the traffic collectors..

To disable or reenable a data collector:

1. Determine which monitors are used by the data collectors. Use [Table 9 on page 31](#) to determine the relationship between the data collectors and the monitors.

Table 9: Monitor Mapping for Data Collectors

Monitor	Data Collector	Category
Current Sessions	Client Monitor Collector and SessionCountCollector	Client
Error Trend	PortTrafficMonitorCollector	Traffic
Logical Interfaces	EquipmentMonitorDeviceStatusCollector	Equipment

Table 9: Monitor Mapping for Data Collectors (*continued*)

Monitor	Data Collector	Category
Find End Point	EquipmentMonitorEndPointCollector	Equipment
Port Status (physical)	EquipmentMonitorDeviceStatusCollector	Equipment
Resource Utilization	EquipmentMonitorDeviceStatusCollector	Equipment
Session Trend	ClientMonitorCollector and SessionCountCollector	Client
Switch Status	EquipmentMonitorDeviceStatusCollector	Equipment
Traffic Trend	PortTrafficMonitorCollector	Traffic
Top Sessions by MAC Address	ClientMonitorCollector	Client
Top Users	ClientMonitorCollector	Client
Unicast vs Broadcast/Multicast	PortTrafficMonitorCollector	Traffic
Unicast vs Broadcast/Multicast Trend	PortTrafficMonitorCollector	Traffic
Virtual Chassis Topology	EquipmentMonitorVCStatsCollector and EquipmentMonitorDeviceStatusCollector	Equipment
Virtual Chassis Protocol	EquipmentMonitorVCStatsCollector and EquipmentMonitorDeviceStatusCollector	Equipment
Virtual Chassis Statistics	EquipmentMonitorVCStatsCollector and EquipmentMonitorDeviceStatusCollector	Equipment

2. Clear the check box to disable the collector or select to enable the collector.
3. Click **Save** and **Close** to save the configuration and to close the window.

Changing the Polling Interval

The frequency at which data is collected is determined by the polling interval. [Table 10 on page 33](#) shows the default polling intervals used by each data collector.

Table 10: Default Polling Intervals

Collector	Polling Interval
ClientMonitorCollector	10 minutes
SessionCountCollector	10 minutes
EquipmentMonitorVCStatsCollector	30 minutes
EquipmentMonitorAPCollector	10 minutes
EquipmentMonitorEndPointCollector	1440 minutes
EquipmentMonitorDeviceStatusCollector	10 minutes
FMAAlarmCountCollector	10 minutes
PortTrafficMonitorCollector	10 minutes
VirtualHostPMCollector	10 minutes
VirtualNICStatsCollector	10 minutes
VirtualMachineStatsCollector	10 minutes
VirtualMachineWeeklyStatsCollector	30 minutes

To change the polling interval:

1. Select the polling interval for a data collector in the Monitor Settings table.
2. Type the new interval level in whole minutes. For example, do not specify 1.5 minutes. Recommended intervals are 5, 10, or 20 minutes.
3. Click **OK** and then **Yes** to verify the change to the configuration.

Enabling and Disabling Collection for Managed Devices

By default all the devices that are discovered and managed by Network Director are enabled for data collection. You can disable or re-enable data collectors across all categories for devices that are managed by Network Director from this tab.

To enable or disable data collectors for devices

1. Open the **Device Settings** sub-tab in the Monitor tab.

All the devices that are managed by Network Director is displayed in the Device Settings section. The last column of the device table indicates the status of data collection as Enabled or Disabled.

2. Select the devices for which you want to enable or disable data collection and do one of the following:
 - Click **Enable** to enable data collection for the selected devices.
 - Click **Disable** to disable data collection for the selected devices.

Specifying Database History Retention

To keep the database manageable, the system periodically checks the age of the records and retires those that have past an expiration date. By default, Network Director ages database records off at 90 days and runs a database cleanup every 6 hours.

Use the Client Session History sub-tab to change the default values:

1. Select from the lists new values.
 - Age of history records (in days) from 1 to 365 days.
 - Cleanup job frequency (in hours) from 1 through 24 hours.
2. Click **OK** to save the changes.

Specifying the Data Learning Engine (DLE) Settings

The Data Learning Engine (DLE) is the component of Cloud Analytics Engine that enables Network Director to collect and analyze high-frequency statistics data from devices and to perform flow path analysis. For Network Director to use Cloud Analytics Engine, you must specify on what server or servers the Data Learning Engine (DLE) is running.

Each DLE supports up to a specific number of Compute Agents (CAs) running on network devices. If you have more CAs in a network than a single DLE can support, you might require multiple DLEs.

Use the Data Learning Engine Settings sub-tab under the Monitoring tab to specify which Data Learning Engine (DLE) server or servers Network Director uses. You can also change the default ports used by the DLE.

To configure DLE in Network Director:

1. Log in to Network Director.
2. Select **Preferences** from the list next to the **System** button in the Network Director banner.
The Preferences page is displayed.
3. Select the **Monitoring** tab and then select **Data Learning Engine Settings**.
4. In the **DLE IP Address** field, enter the IP address of the DLE server.

NOTE: Before you configure DLE in Network Director, make sure that there are no errors in the monitor.log file. The log file is stored in the `/var/log/jboss/server/server1` directory.

- If you want to change the ports used by the DLE, click **View/Edit DLE Ports** to edit the ports and then click **OK**.

NOTE: If you change the default DLE ports (8282, 8081, and 50006), you must ensure that the new ports are open between DLE and Junos Space Network Management Platform or Network Director.

You can use the `netstat -anp | grep port-number` command to verify that the new ports are open (in listening mode) between DLE and Junos Space Network Management Platform or Network Director.

Table 11 on page 35 describes the default DLE ports.

Table 11: Default DLE Port Descriptions

Port	Description
Flow Analysis API Port	Used by the flow path analysis feature and network traffic analysis feature to communicate with the DLE. Default value is 8282.
HFS API Port	Used by the high-frequency statistics feature to communicate with the DLE. Default value is 8081.
HFS Control Channel Port	Used by the high-frequency statistic feature for communication about threshold-related events with the DLE. Default value is 50006.

- Click **OK** to save the DLE settings.

The message **Preferences saved successfully** is displayed.

NOTE: After you configure the DLE settings, check whether the DLE connection state is **UP** in the DLE settings page.

- Click **Add Another** to add a new DLE server.

RELATED DOCUMENTATION

| [Understanding Monitor Mode in Network Director](#) | 658

Changing Alarm Settings

IN THIS SECTION

- [Configuring Global Alarm Notifications](#) | 36
- [Retaining Alarm History](#) | 36
- [Specifying Event History](#) | 37
- [Enabling Alarms](#) | 37
- [Changing the Severity of Individual Alarms](#) | 46
- [Configuring Threshold Alarms](#) | 46
- [Configuring Individual Alarm Notifications](#) | 46

Use the Fault tab to enable individual alarms, set the retention period for alarms, configure alarm notifications, configure threshold alarms, and to specify the number of events to keep for each alarm. The Fault tab has multiple sections, which you can expand and collapse by clicking the arrow next to the section title:

- Global Settings, for configuring Faults settings such as global alarm notifications and alarm data retention.
- Individual Alarms and Threshold Settings, for configuring settings for individual alarms and threshold alarms.

This section describes the following tasks that you can perform by using the Fault tab:

Configuring Global Alarm Notifications

You can configure global e-mail notifications to be sent when any alarm with notifications enabled is generated. To configure global e-mail notifications, enter the e-mail addresses to receive global alarm notifications in the Alarm Notifications Destinations field in the Global Settings section. Separate addresses with a comma (,). For information about enabling notification for an alarm, see [“Configuring Individual Alarm Notifications” on page 46](#).

Retaining Alarm History

Use the **No. of days to keep Alarm** field in the Global Settings section to specify the number of days to keep alarm history. The default retention time is 120 days; but you can specify a period of 7 through 1000 days. Specifying a longer retention time consumes more database resources. To change the alarm retention duration, type a new value and click **OK** and **Yes** to confirm the change.

Specifying Event History

Use the **Events/Alarm** field in the Global Settings section to specify the number of event entries that are kept in the alarm history. The default setting for events is 20. To change the setting, type a new value and click **OK** and **Yes** to confirm the change.

Enabling Alarms

Ensure all devices are configured to send traps to Network Director. This task is performed for the devices in Deploy mode through Set SNMP Trap Configuration.

Use the Individual Alarms and Threshold Settings section to disable and re-enable individual alarms or all alarms. Alarms appear on both tabs in the section: Alarm Settings and Threshold Settings. Fault alarms are preconfigured and initially enabled. To enable or disable alarms:

1. (Optional) Sort the alarms. By default, the list of alarms is sorted alphabetically within each category. You can also sort by description or alarm severity within a category by clicking a column heading.
2. Review the alarms and either select the check box in the heading to select all of the alarms or select the check box for the individual alarms you want to enable. For a full description of each of the alarms, see [Table 12 on page 37](#).
3. Click **OK** and **Yes** to confirm the alarm change.

Table 12: Alarm Descriptions

Alarm Name	Description	Device Type
<i>BFD</i>		
BfdSessionDetectionTimeAlarm	Generated when the threshold value for detection time is set and the BFD session detection-time adapts to a value greater than the threshold.	EX Series Switch
BfdSessionTxAlarm	Generated when the threshold value for transmit interval (in microseconds) is exceeded.	EX Series Switch
<i>BGP</i>		
BgpM2BackwardTransitionAlarm	Generated when the BGP FSM moves from a higher-numbered state to a lower-numbered state.	EX Series Switch
BgpM2EstablishedAlarm	Generated when the BGP Finite State Machine (FSM) enters the ESTABLISHED state.	EX Series Switch

Table 12: Alarm Descriptions (*continued*)

Alarm Name	Description	Device Type
<i>Chassis</i>		
FanFailureAlarm	Generated when the specified cooling fan or impeller has failed (is not spinning).	EX Series Switch
FEBSwitchoverAlarm	Generated when the Forwarding Engine Board (FEB) has switched over.	EX Series Switch
FRUCheckAlarm	Generated when the device has detected that a field-replaceable unit (FRU) has some operational errors and has gone into check state.	EX Series Switch
FRUFailedAlarm	Generated when a FRU has failed.	EX Series Switch
FRUInsertionAlarm	Generated when the system detects that the specified FRU is inserted into the chassis.	EX Series Switch
FRUOfflineAlarm	Generated when the specified FRU goes offline.	EX Series Switch
FRUOnlineAlarm	Generated when the specified FRU goes online.	EX Series Switch
FRUPowerOffAlarm	Generated when the specified FRU is powered off.	EX Series Switch
FRUPowerOnAlarm	Generated when the specified FRU is powered on.	EX Series Switch
FRURemovalAlarm	Generated when the system detects that the specified FRU was removed from the chassis.	EX Series Switch
HardDiskFailedAlarm	Generated when the hard disk for the specified routing engine has failed.	EX Series Switch
HardDiskMissingAlarm	Generated when the hard disk in the specified routing engine is missing from the boot device list.	EX Series Switch

Table 12: Alarm Descriptions (*continued*)

Alarm Name	Description	Device Type
PowerSupplyFailureAlarm	Generated when the specified power supply has failed (bad DC output).	EX Series Switch
RedundancySwitchOverAlarm	Generated when a graceful Routing Engine switchover (GRES) occurs on a switch with dual Routing Engines or on a Virtual Chassis.	EX Series Switch
TemperatureAlarm	Generated when the device has over heated.	EX Series Switch
<i>Configuration (Configuration)</i>		
CmCfgChangeAlarm	Generated when the jnxCMCfgChgEventTable records a configuration management event.	EX Series Switch
CMRescueChangeAlarm	Generated when a change is made to the rescue configuration.	EX Series Switch
<i>Core</i>		
Device alarm	Generated when the device status changes (up to down or down to up).	EX Series Switch
<i>CoS</i>		
CoSAlmostOutOfDedicatedQueuesAlarm	Generated when only 10% of CoS queues are available.	EX Series Switch
CoSOutOfDedicatedQueuesAlarm	Generated when there are no more available dedicated CoS queues.	EX Series Switch
<i>DHCP</i>		
JdhcpLocalServerDupClientAlarm	Generated when a DHCP client is detected changing interfaces.	EX Series Switch
JdhcpLocalServerIfLimitExceededAlarm	Generated when the client limit is reached on an interface.	EX Series Switch
Jdhcpv6LocalServerLimitExceededAlarm	Generated when the client limit is reached on an interface for DHCPv6.	EX Series Switch

Table 12: Alarm Descriptions (*continued*)

Alarm Name	Description	Device Type
<i>DOM</i>		
DomAlertSetAlarm	Generated when an interface detects Digital Optical Monitor (DOM) alarm conditions.	EX Series Switch
<i>Flow Collection (FlowCollection)</i>		
CollFlowOverloadAlarm	Generated when a collector PIC detects a hard or soft flow overload.	EX Series Switch
CollFtpSwitchOverAlarm	Generated when an FTP server switchover occurs.	EX Series Switch
CollMemoryUnavailableAlarm	Generated when a PIC is out of memory or the memory is unavailable.	EX Series Switch
CollUnavailableDestAlarm	Generated when a file transfer destination is unavailable.	EX Series Switch
CollUnsuccessfulTransferAlarm	Generated when a collector file is unable to transfer because the destination is unavailable.	EX Series Switch
<i>General</i>		
Authentication Failure Alarm	Generated when a protocol message is received that is not properly authenticated.	EX Series Switch
Cold Start Alarm	Generated when a device is re-initializing and its configuration might have changed.	EX Series Switch
Link Down Alarm	Generated when a link is down. The trap is generated when the ifOperStatus object for a communication link is about to enter the down state from another state other than notPresent. This other state is indicated by the included value of ifOperStatus.	EX Series Switch

Table 12: Alarm Descriptions (*continued*)

Alarm Name	Description	Device Type
Link Up Alarm	Generated when a link comes up that was previously in the down state. The trap is generated when the ifOperStatus object for a communication link left the down state and transitioned into another state other than notPresent state. This other state is indicated by the included value of ifOperStatus.	EX Series Switch
Warm Start Alarm	Generated when a device is re-initializing and its configuration has not changed.	EX Series Switch
<i>Generic (GenericEvent)</i>		
GenericEventTrapAlarm	Generated by an Op script or event policies. This notification can include one or more attribute-value pairs. The pairs are identified by the jnxEventAvAttribute and jnxEventAvValue objects.	EX Series Switch
<i>L2ALD</i>		
L2aldGlobalMacLimitAlarm	Generated when the MAC limit is reached for the entire system. This trap is sent only once, when the limit is reached.	EX Series Switch
L2aldInterfaceMacLimitAlarm	Generated when the given interface reaches the MAC limit (jnxl2aldInterfaceMacLimit).	EX Series Switch
L2aldRoutingInstMacLimitAlarm	Generated when the MAC limit is reached for a given routing instance (jnxl2aldRoutingInst).	EX Series Switch
<i>L2CP</i>		
LacpTimeoutAlarm	Generated when LACP has timed out.	EX Series Switch
PortBpduErrorStatusChangeTrapAlarm	Generated when the port's BPDU error state (no-error or detected) changes.	EX Series Switch
PortLoopProtectStateChangeTrapAlarm	Generated when the port's loop-protect state (no-error or loop-prevented) changes.	EX Series Switch

Table 12: Alarm Descriptions (*continued*)

Alarm Name	Description	Device Type
PortRootProtectStateChangeTrapAlarm	Generated when the port's root-protect state (no-error or root-prevented) changes.	EX Series Switch
<i>MAC Forwarding Database (MACFDB)</i>		
MacChangedNotificationAlarm	Generated when MAC addresses of the monitored devices are learned or removed from the forwarding database (FDB).	EX Series Switch
<i>PoE (Power over Ethernet)</i>		
PoE Port ON-OFF Alarm	Generated when the PoE power is turned on or off.	EX Series Switch
PoE Power Usage High	Generated when Power over Ethernet (PoE) used is below or above the defined threshold.	EX Series Switch
<i>Passive Monitoring (PassiveMonitoring)</i>		
PMonOverloadSetAlarm	Generated when an overload condition is detected on a Passive Monitoring Interface.	EX Series Switch
<i>Ping</i>		
PingEgressJitterThresholdExceededAlarm	Generated when egress time jitter (jnxPingMaxEgressUs minus jnxPingResultsMinEgressUs) exceeds the configured threshold (jnxPingCtlEgressJitterThreshold) causing the egressJitterThreshold bit to be set.	EX Series Switch
PingEgressStdDevThresholdExceededAlarm	Generated when the standard deviation of the egress time (jnxPingResultsStddevEgressUs) exceeds the configured threshold (jnxPingCtlEgressTimeThreshold) and causes the egress bit to be set.	EX Series Switch

Table 12: Alarm Descriptions (*continued*)

Alarm Name	Description	Device Type
PingEgressThresholdExceededAlarm	Generated when the egress time (jnxPingResultsStddevEgressUs) exceeds the configured threshold (jnxPingCtlEgressTimeThreshold) and the egress threshold bit is set in jnxPingCtlTrapGeneration.	EX Series Switch
PingIngressJitterThresholdExceededAlarm	Generated when ingress time jitter (jnxPingResultsMaxIngressUs minus jnxPingResultsMinIngressUs) exceeds the configured threshold (jnxPingCtlIngressJitterThreshold) and the ingressJitterThreshold bit is set in jnxPingCtlTrapGeneration.	EX Series Switch
PingIngressStddevThresholdExceededAlarm	Generated when the standard deviation of the ingress time (jnxPingResultsStdDevIngressUs) exceeds the configured threshold (jnxPingCtlIngressStddevThreshold) and the ingress StdDevThreshold bit is set in jnxPingCtlTrapGeneration.	EX Series Switch
PingIngressThresholdExceededAlarm	Generated when the ingress time jitter (jnxPingResultsIngressUs) exceeds the configured threshold (jnxPingCtlIngressTimeThreshold) and the ingress threshold bit (jnxPingIngressThresholdExceeded) is set in jnxPingCtlTrapGeneration.	EX Series Switch
PingRttJitterThresholdExceededAlarm	Generated when the round trip time jitter (jnxPingResultsMaxRttUs minus jnxPingResultsMinRttUs) exceeds the configured threshold (jnxPingCtlRttJitterThreshold) and the rttJitterThreshold bit is set in jnxPingCtlTrapGeneration.	EX Series Switch

Table 12: Alarm Descriptions (*continued*)

Alarm Name	Description	Device Type
PingRttStdDevThresholdExceededAlarm	Generated when the standard deviation of the round trip time (jnxPingResultsStdDevRttUs) exceeds the configured threshold (jnxPingCtlRTTStdDev) and the rttStdDevThreshold bit is set in jnxPingCtlTrapGeneration.	EX Series Switch
PingRttThresholdExceededAlarm	Generated when the round trip time (jnxPingCtlRttThreshold) exceeds the configured threshold (jnxPingCtlRttThreshold) and the rttThreshold bit is set in jnxPingCtlTrapGeneration.	EX Series Switch
<i>RMon</i>		
RmonAlarmGetFailureAlarm	Generated when a GET request for an alarm variable returns an error. The specific error is identified by a varbind in jnxRmonAlarmGetFailReason.	EX Series Switch
<i>SONET</i>		
SonetAlarmSetAlarm	Generated when there is a notification of a recently set SONET or SDH alarm on an interface.	EX Series Switch
<i>SONET APS (SONETAPS)</i>		
APSEventChannelMismatchAlarm	Generated when the value of an instance of apsStatusChannelMismatches increments.	EX Series Switch
APSEventFEPLFAlarm	Generated when the value of an instance of apsEventFEPLFs increments.	EX Series Switch
APSEventModeMismatchAlarm	Generated when the value of an instance of apsEventModeMismatch increments.	EX Series Switch
APSEventPSBFAlarm	Generated when the value of an instance of apsStatusPSBFs increments.	EX Series Switch
APSEventSwitchoverAlarm	Generated when the value of an instance of apChanStatusSwitchover increments.	EX Series Switch

Table 12: Alarm Descriptions (*continued*)

Alarm Name	Description	Device Type
<i>Virtual Chassis (VirtualChassis)</i>		
VccpMemberAlarm	Generated when a member has completed transition from the down state to another state.	EX Series Switch
VccpPortAlarm	Generated when one of the member's communication links has completed transition from the down state to another state.	EX Series Switch
<i>VNetwork</i>		
HostConnectivityLostAlarm	Generated when all the uplink ports of a virtual switch residing in a host loses network connectivity.	Host
HostNetworkRedundancyLostAlarm	Generated when some uplink ports of a virtual switch residing in a host loses network connectivity. It indicates that there are one or more ports that still has network connectivity.	Host
VNetworkConnectivityLostAlarm	Generated when Network Director loses network connectivity with the vCenter server.	Virtual Network

Changing the Severity of Individual Alarms

You can change the severity of the alarms to match your corporate procedures and guidelines. For example, at your company a DoS attack might be considered a critical alarm, while Network Director has a default severity for DoS attacks as a major alarm. Alarms appear on both tabs in the Individual Alarms and Threshold Settings section: Alarm Settings and Threshold Settings.

To change the severity of an alarm:

1. Select the current severity in the **Severity** column. A list of the severity levels appear.
2. Select the new severity level for the alarm.
3. Click **OK** and **Yes** to confirm the change to the severity setting.

To configure alarm notifications, see [“Configuring Individual Alarm Notifications” on page 46](#).

Configuring Threshold Alarms

Threshold alarms are alarms that are generated when a monitored value crosses the configured threshold. They provide enhanced visibility into potential issues on the network. You configure and manage threshold alarms the same way as other alarms. You also have the option of setting the threshold level of individual threshold alarms.

To edit the threshold of threshold alarms:

1. Select the **Threshold Settings** tab in the Individual Alarms and Threshold settings section of the Fault tab.
2. Click **Edit Settings** in the Threshold Settings column of the alarm threshold you want to edit.
3. Set the threshold in the window that opens.
4. Click **Save** to save the new threshold.

To configure alarm notifications, see [“Configuring Individual Alarm Notifications” on page 46](#).

Configuring Individual Alarm Notifications

You can configure e-mail notifications to be sent when an individual alarm is generated. When you enable notification for an alarm, the notifications are sent to the e-mail addresses configured for the alarm and the addresses configured for global alarm notifications. Alarms appear on both tabs in the Individual Alarms and Threshold Settings section: Alarm Settings and Threshold Settings.

To configure e-mail notification for an alarm name:

1. Select the check box in the alarm's Notification column.
If you later want to disable notification for the alarm, clear the check box.
2. Click **Edit Notification** in the Notification column. The Alarm Notification Details window opens.
3. Type one or more e-mail addresses in the **Notification Email Addresses** box. Separate addresses with a comma (,).
You can later edit the addresses to send notifications to different addresses.
4. (Optional) Type a comment in the Comments box. This comment is included in the e-mail notification message.
5. Click **Save**.

RELATED DOCUMENTATION

[Understanding the Fault Mode Tasks Pane | 765](#)

[Current Active Alarms Monitor | 776](#)

[Alarms by State Monitor | 779](#)

[Alarms by Severity Monitor | 778](#)

[Alarms by Category Monitor | 778](#)

[Network Director Documentation home page](#)

Modifying Data Center Synchronization Interval Using the Virtualization Tab

You can configure and manage data centers by using the Data Center View in Network Director. Network Director synchronizes data from the cloud infrastructure that is part of your data center every 24 hours. However, depending on the size of your data center network, you can modify this interval from the Virtualization tab in the Preferences page. Modify the interval by changing the **Periodic Cloud Infrastructure Synchronization Interval (hours)** field.

For data centers that use OpenStack as the cloud infrastructure provider, Network Director displays the links between hypervisors and switches only in the Data Center View. If you want to view the link details using the View Virtual Network Connectivity task, then you must select the **Synchronize links from Topology to Virtualization application for OpenStack** check box.

NOTE: Selecting the **Synchronize links from Topology to Virtualization application for OpenStack** check box might impact the performance of Network Director, while the synchronization is in process.

Getting Started with Network Director

IN THIS CHAPTER

- [Getting Started with Junos Space Network Director | 49](#)

Getting Started with Junos Space Network Director

IN THIS SECTION

- [Building Your Network | 49](#)
- [Creating Profiles in Network Director | 50](#)
- [Managing Software Images using Network Director | 51](#)
- [Configuring Approval Modes for Device Configurations | 51](#)
- [Resynchronizing Device Configuration | 51](#)
- [Creating the Baseline Configuration | 52](#)
- [Monitoring Your Network | 52](#)
- [Setting up Network Traffic Analysis and Analyzing the Traffic | 52](#)
- [Managing Network Faults and Notifications | 53](#)
- [Generating Network Reports | 53](#)

This section describes a series of steps that you must perform after installing Network Director to manage and troubleshoot your network.

Building Your Network

The first step after you install and log in to Network Director is to build your network. Even with large networks, Network Director has made this step relatively easy and straightforward. The steps that you need to perform depend on whether your network contains legacy devices, or new devices, or a combination of both.

You add legacy devices, which already have some configurations, to Network Director by using a process called *device discovery*. Once such a device is successfully discovered, Network Director reads the device configurations and replicates these configurations in the form of profiles in Network Director. You can use device discovery to add Juniper Networks switches to Network Director. For more information on device discovery, see [“Discovering Devices in a Physical Network” on page 94](#) and [“Understanding the Device Discovery Process” on page 101](#).

With new devices or devices that are set to factory-default configuration, you can use the *zero touch provisioning (ZTP)* feature to provision the device. ZTP enables you to auto-discover, auto-upgrade, and load the requisite default configuration on Juniper Networks switches in your network automatically—without manual intervention. When you physically connect a switch that has the factory-default configuration to a network and boot the switch, the switch attempts to upgrade Junos OS automatically and autoinstall a configuration file from the network. For more information, see [“Configuring and Monitoring Zero Touch Provisioning” on page 650](#).

Creating Profiles in Network Director

Profiles in Network Director are a group of feature-specific configurations that you can assign to devices. For example, you can create a CoS profile that combines all the supported class-of-service configurations for a particular device family, and assign it to a port on a device.

- You can create a new profile for an interface or device by defining the custom configuration. You can use the Tasks pane in Build mode to manually create profiles.
- Network Director automatically creates profiles when a supported configuration of a device that is already discovered and managed by Network Director is modified outside Network Director (also known as out-of-band configuration changes). For more information, see [“Understanding Resynchronization of Device Configuration” on page 611](#).

Following are some advantages of using profiles:

- Bulk provisioning—You can combine a group of configurations as a profile and apply it to one or more ports or devices in one go, thereby saving a lot of time and effort. Profiles ensure that the configurations are error-free as most configuration value ranges are set in the profile workflow. Network Director prompts the user if there are any errors. You must fix the errors before you can create a profile.
- Editing—For profiles that are already deployed on devices, if you want to make changes to the configuration values, you can modify the configuration values in the profile and redeploy the profile. Network Director updates the new configuration value on each device where the profile is deployed.
- Cloning—If you already have a set of profiles defined for your network and want to apply a different configuration for a set of devices or ports in your network, you use the clone feature. The clone feature enables you to make a copy of any profile and make the necessary modifications. You can then apply these to devices and ports that require the different set of configuration.

For more information on profiles, see [“Understanding Network Configuration Profiles” on page 88](#).

Managing Software Images using Network Director

As a Network Administrator, you can store different versions of Junos OS software images in the Network Director image repository. You can then deploy these images on one or more managed devices manually or have the system deploy the images by using zero touch provisioning (ZTP).

For more information on managing and deploying software images, see [“Managing Software Images” on page 632](#) and [“Deploying Software Images” on page 635](#).

Configuring Approval Modes for Device Configurations

When you make configuration changes in Build mode, the changes are not deployed to devices automatically. You must manually deploy the changes to devices in Deploy mode. When you deploy configuration changes to a device, all pending configuration changes for that device are deployed. You can deploy the device configurations in the following two ways:

- **Auto Approval**—In this mode, the device configuration changes are approved automatically by the system and do not require explicit (manual) approval by a configuration approver before they can be deployed. This is the default approval mode.
- **Manual Approval**—In this mode, the device configuration changes must be explicitly approved by a configuration approver before the changes can be deployed to the device. An operator performs device configurations and creates a change request for that configuration and submits it for approval to one or more approvers. The approvers are notified by e-mail whenever a change request is created. If a configuration or a change to it is approved by an approver, then the operator is able to deploy it. If a configuration is rejected, the operator must make the necessary changes, resubmit the change request, and procure an approval before the configuration can be deployed.

NOTE: For manual approval, the **Network Director - Configuration Approver** role is available in Junos Space, which is specific to Network Director. A user with this role reviews device configurations and proposed changes to device configurations and can either approve or reject them.

For more information about deploying configuration to devices, see [“Deploying Configuration to Devices” on page 580](#).

Resynchronizing Device Configuration

A network managed by Network Director has three repositories of information about the configuration of a network device—the configuration stored on the device itself, the device configuration record maintained by Junos Space, and the Build mode configuration maintained by Network Director.

When the configuration contained in all three repositories match, the device configuration state is shown as In Sync in Network Director. When the repositories do not match, the configuration state is shown as Out of Sync. A common cause for this state is out-of-band configuration changes—that is, configuration changes made to a device outside of Network Director.

When a device state is Out of Sync, you cannot deploy configuration changes on the device in Deploy mode. Use the Resynchronize Device Configuration task to resynchronize the three configuration repositories and change the device configuration state back to In Sync.

For more information about device resynchronization, see [“Understanding Resynchronization of Device Configuration” on page 611](#).

Creating the Baseline Configuration

You can create a baseline of configuration and the Junos OS version of the devices on the Network Director server. By creating a baseline configuration file for a device you define a reference point to save the device configuration and its Junos OS version to a particular known state and later restore the configuration to that known state.

For more information about device resynchronization, see [“Creating and Managing Baseline of Device Configuration Files” on page 626](#).

Monitoring Your Network

Network Director provides the visibility into your network status and performance by using the Monitor Mode.

Network Director monitors the devices it manages and maintains the information it collects from the devices in a database. You can view this data as easy-to-understand graphs and tables—known as monitoring widgets—to quickly visualize the state of your network, spot trends developing over time, and view important details.

For more information about the monitor mode, see [“Understanding Monitor Mode in Network Director” on page 658](#).

You can also use the Dashboard widgets to monitor your network performance. For more information about the Dashboard widgets, see [“Understanding the Dashboard” on page 56](#).

Setting up Network Traffic Analysis and Analyzing the Traffic

The Network Traffic Analysis feature of Network Director monitors high-speed switched or routed networks. Once enabled, Network Director randomly samples network packets and sends the samples to a data learning engine (DLE) for analysis. Network traffic analysis uses packet-based sampling. Network Director samples one packet out of a specified number of packets from an interface enabled for network traffic analysis and sends the packet to the DLE. DLE uses this sampling information to create a picture

of the network traffic, which includes the applications that contribute to the traffic, traffic statistics, and the top applications. You can enable network traffic analysis on all devices that are managed by Network Director.

For more information about installing and configuring DLE, see [Installing and Configuring Data Learning Engine for Network Director](#).

Managing Network Faults and Notifications

In Fault mode, Network Director informs you of unexpected, significant events happening in your network. Examples of such events include link up or link down, power supply failure, client authentication failure, and so on.

Network Director receives information about events from its managed devices in the form of SNMP notifications. A single event can often generate multiple SNMP notifications. To simplify management of events, Network Director correlates these notifications, creating high-level alarms of different severity levels for the events. For example, a power supply failure might generate a number of notifications. Network Director correlates these notifications and raises a single power supply failure alarm for the device. Network Director also automatically clears an alarm if it receives notification from the device that the error condition has been resolved.

To tailor Network Director fault management to your organization's requirements, you can enable or disable the receipt of specific alarms and change the default severity level of alarms.

For more information about the fault mode in Network Director, see [“Understanding Fault Mode in Network Director” on page 761](#).

Generating Network Reports

Use the Report mode in Network Director to create standardized reports from the monitoring and fault data collected by Network Director. An essential part of the network management life cycle, reporting provides administrators and management insight into the network for maintenance, troubleshooting, trend and capacity analysis, and provides records that can be archived for compliance requirements.

Network Director provides reports in PDF and HTML formats that use graphs and tables to clearly convey data. Reports are also available in CSV format for importing into spreadsheets.

For more information about managing reports in Network Director, see [“Managing Reports in Network Director” on page 787](#).

RELATED DOCUMENTATION

[Understanding Build Mode in Network Director](#) | 77

[Understanding Deploy Mode in Network Director | 572](#)

[Understanding Monitor Mode in Network Director | 658](#)

[Understanding Fault Mode in Network Director | 761](#)

[Understanding Report Mode in Network Director | 782](#)

2

PART

Working with the Dashboard

[About the Dashboard](#) | 56

[Using the Dashboard](#) | 57

[Dashboard Widget Reference](#) | 59

About the Dashboard

IN THIS CHAPTER

- [Understanding the Dashboard | 56](#)

Understanding the Dashboard

The Dashboard is a customizable page to view information about the network, and is the default page that opens when you log in. You select monitoring widgets to display on the Dashboard that show various information about the network. The Dashboard is a view. To open a different view, select a view from the Views list in the Network Director banner.

RELATED DOCUMENTATION

[Using Dashboard Widgets | 57](#)

[Network Director Documentation home page](#)

Using the Dashboard

IN THIS CHAPTER

- [Using Dashboard Widgets | 57](#)

Using Dashboard Widgets

The Dashboard is a customizable page for viewing information about the network. You select monitoring widgets to display on the Dashboard that show various information about the network. The Dashboard is the default view that opens when you log in. When a different view is selected, select **Dashboard View** from the Select View list in the Network Director banner to open the Dashboard.

To select what appears on the Dashboard:

1. To add a monitor to the Dashboard:
 - a. Select **Add Widgets**. Thumbnails of the available widgets appear.
 - b. To add a widget to the Dashboard, mouse over the widget's thumbnail, then click the **Add** button that appears on the widgets.
 - c. When you are finished adding widgets, click **Done**. The new widgets appear on the Home page.
2. To refresh a widget's data, click the **Refresh** button in its title bar.
3. To see additional information for a widget, click the **Maximize** button in the widget's title bar.
4. To remove a widget from the Dashboard, click the Close button (X) in its title bar.
5. To open online help for a widget, click the Help button (?) in its title bar.
6. To move a widget, click its title bar and drag it to the new location.

RELATED DOCUMENTATION

[Understanding the Dashboard | 56](#)

[Network Director Documentation home page](#)

Dashboard Widget Reference

IN THIS CHAPTER

- [Alarms Widget | 59](#)
- [Config Deployment Jobs Status Widget | 61](#)
- [Device & Port Latency Widget | 62](#)
- [Device & Port Utilization Widget | 63](#)
- [Equipment By Type Widget | 69](#)
- [Port Status - Physical Widget | 70](#)
- [Top Talker - Wired Devices Widget | 71](#)
- [Top Overlay Networks Widget | 73](#)

Alarms Widget

IN THIS SECTION

- [Alarms Widget Summary | 59](#)
- [Alarms Widget Details | 60](#)

The Alarms widget provides summary and detailed information about network alarms.

This topic describes:

Alarms Widget Summary

The summary view of the Alarms widget displays summary information about network alarms and their location. The number of active alarms of each severity is shown in colored circles on the left side of the widget. The distribution of alarms by site is shown on a map. The alarms count for each site is shown as

a pie chart. The color of each pie chart segment indicates severity level. The colored circles to the left of the map also serve as the legend for the color coding.

Mouse over a pie chart to see more information about the alarms for that site.

Alarms Widget Details

To open the Alarms widget details page, click the **Maximize** button in the widget's title bar. The Alarms widget details window displays detailed information active alarms. The top of the page contains a larger view of the widget. The bottom of the page contains a table of detailed information about active alarms. [Table 13 on page 60](#) describes the columns in this table. Click an alarm severity level circle to filter the table to show only alarms of that severity. To close the details page, click the **Minimize** button in the title bar.

Table 13: Alarm Widget Details Table

Column	Description
Name	The alarm name.
ID	A system and sequentially-generated identification number.
Assigned To	If assigned to an individual, it shows the name of the person assigned; otherwise, it shows System to mark that the alarm is still unassigned.
Severity	<p>The severity of the alarm. Severity levels are:</p> <ul style="list-style-type: none"> ● Critical—A critical condition exists; immediate action is necessary. ● Major—A major error has occurred; escalate or notify as necessary. ● Minor—A minor error has occurred; notify or monitor the condition. ● Info—An informational message; no action is necessary.
Entity ID	The identification of the entity responsible for causing this alarm. The Entity ID is the key for correlation of events into an alarm. The Entity ID could be an IP address of the device.
Reporting Device IP	The IP address of the device or network element that generated the alarm. The SNMP agent is located at the source IP. In most cases, the source IP is the IP address of the switch.
Reporting Device	The hostname of the reporting device.
Creation Date	The date and time the alarm was first reported.
Last Updated	The date and time that the information for the alarm was last modified.

Table 13: Alarm Widget Details Table (continued)

Column	Description
Updated By	Either the system or the last user who modified the alarm.
Acknowledged	Indicates if the alarm has been acknowledged.

RELATED DOCUMENTATION

Understanding the Dashboard 56
Using Dashboard Widgets 57
Network Director Documentation home page

Config Deployment Jobs Status Widget

IN THIS SECTION

- [Config Deployment Jobs Status Widget Summary | 61](#)
- [Config Deployment Jobs Status Widget Details | 62](#)

The Config Deployment Jobs Status widget provides summary and detailed information about the status of configuration deployment jobs.

This topic describes:

Config Deployment Jobs Status Widget Summary

The Config Deployment Jobs Status widget displays summary information about the status of configuration deployment jobs. The information appears in a table. The vertical axis lists the job statuses. The horizontal axis shows the times when job status data was collected. You can do the following tasks:

- Select a time period to view from the **Deployment Trend** list.
- Click the **Refresh** button to refresh the information displayed.

Config Deployment Jobs Status Widget Details

To open the Config Deployment Jobs Status widget details page, click the **Maximize** button in the widget's title bar. The Config Deployment Jobs Status widget details window displays detailed information about the status of configuration deployment jobs. The page shows the same summary information table as the widget. It also shows a table of detailed configuration job status information. To close the details page, click the **Minimize** button in the title bar.

RELATED DOCUMENTATION

[Understanding the Dashboard | 56](#)

[Using Dashboard Widgets | 57](#)

[Network Director Documentation home page](#)

Device & Port Latency Widget

The Device & Port Latency widget provides a graphical view of latency on devices. The heat map represents each device as a color-coded box. The color coding indicates the level of latency on a device. Cooler colors (for example, green) indicate lower latency, while hotter colors (for example, red) indicate higher latency.

The Device & Port Latency widget can show information only for devices that support Cloud Analytics Engine and that have the high-frequency traffic statistics feature enabled in Network Director.

To use the Device & Port Latency Heat Map widget, you must do the following first:

- Configure the Data Learning Engine (DLE) settings under **Preferences > Monitoring > Data Learning Engine Settings**. The DLE is a component of Cloud Analytics Engine. For information on configuring the DLE settings, see [“Specifying the Data Learning Engine \(DLE\) Settings” on page 34](#).
- Enable high-frequency traffic statistics on the devices you want to monitor. For information, see [“Enabling High-Frequency Traffic Statistics Monitoring on Devices” on page 602](#).

You can do the following with the heat map:

- Select the time period to view from the list in the title bar.
- Select how to organize the heat map by clicking the Settings icon and then selecting an option from the **Group Devices By** list. Each option opens a different view of the heat map, with device boxes grouped according to your selection.
- Select a filter for which devices to show by clicking the Settings icon and then selecting an option from the **Show** list.

- Drill down into the heat map's hierarchy by clicking one of the device container names (for example, a site or building). To move back up the hierarchy, click the navigation arrows above the heat map.
- Mouse over a device box to view detailed latency and congestion information about the device in a pop-up window.
- Click a device box to view detailed latency and congestion information about the device's ports. In this view, each port is represented by a box that is color-coded to show its level of latency and congestion. Mouse over a port box to see more information about the port. To move back up the hierarchy, click the navigation arrows above the heat map.
- Slide the circular controls along the bar under the heat map to filter the devices or ports shown in the heat map by degree of latency.

RELATED DOCUMENTATION

[Understanding the Dashboard | 56](#)

[Using Dashboard Widgets | 57](#)

Device & Port Utilization Widget

IN THIS SECTION




- [Using the Global Controls | 64](#)
- [Interacting with the Heat Maps | 64](#)
- [Viewing Active Flows on a Port | 65](#)
- [Flow Analysis Details Window | 65](#)
- [Viewing Traffic on a Device | 67](#)

The Device & Port Utilization Heatmap widget provides a graphical view of device port utilization percentage. The heat map represents each device as a color-coded box. The color coding indicates the overall level of port utilization on a device. Cooler colors (for example, green) indicate lower port utilization, while hotter colors (for example, red or dark red) indicate higher port utilization.

You can view the utilization level for each port on a device by clicking the box representing the device. A heat map is displayed that represents each port on the device as a color-coded box, with the color coding representing the level of port utilization.

Using the Global Controls

Use the controls in the upper right corner to make global changes to how the device and port heat maps are displayed. You can:

- Select the time period over which device utilization and port utilization are shown.
- Display information about the devices or the ports in either graphical heat map or tabular format by clicking either  (graphical) or  (tabular).
- Select how to organize the heat map by clicking the Settings icon (), and then selecting an option from the **Group Devices By** list. Each option creates a different view of the heat map, with device boxes grouped according to your selection.

Interacting with the Heat Maps

You can interact with the device and port heat maps as follows:

- If you have grouped the devices by location, you can drill down into the heat map's hierarchy by clicking one of the device container names (for example, a site or building). To move back up the hierarchy, click the navigation arrows above the heat map.
- Mouse over a device box to see detailed device-level port utilization information in a pop-up window. In the pop-up window, click the **View top 5 ports** link to view the top five ports that use the most bandwidth on the device.
- Click on a device box to display a heat map of the ports on the device. In this port-level heat map, each port is represented by a box that is color-coded to show its level of utilization. To return to the device view, click the navigation arrows above the heat map.
- Slide the bandwidth utilization control to filter and view devices based on utilization.
- Mouse over a port box to display information about the port—such as port name, status, speed, and percent utilization—in a pop-up window.
 - For ports on devices that support Cloud Analytics Engine, you can view any existing flow analysis results on flows through the port by clicking **View active flows through this link**. See [“Viewing Active Flows on a Port” on page 65](#) for more information.
 - For ports on devices that are configured for traffic analysis, you can view the traffic analysis data by clicking **Analyze Traffic on the Port**.
- Slide the circular controls along the bar under the heat map to Filter the devices or ports shown in the heat map by degree of port utilization.

Viewing Active Flows on a Port

For devices that support Cloud Analytics Engine, you can view the results of the most recent flow analysis traces on application flows on the port by mousing over the port and clicking **View active flows through this link**. The Current Active Flows window is displayed.

The Current Active Flows window lists only application flows for which flow analysis traces exist—there might be other active application flows on the port that are not shown. Each flow is uniquely defined by source IP address and TCP/UDP port, destination IP address and TCP/UDP port, and transport protocol. [Table 14 on page 65](#) describes the fields in this window.

Table 14: Fields in the Current Active Flows Window

Field	Description
Source IP Source Port	Source IP address and source TCP/UDP port for the flow. In the case of a flow between two VMS, the IP address is the source VTEP address. If the port is associated with a well-known service, the service name is also shown.
Destination IP Destination Port	Destination IP address and destination TCP/UDP port for the flow. In the case of a flow between two VMS, the IP address is the destination VTEP address. If the port is associated with a well-known service, the service name is also shown.
Protocol	Either TCP or UDP.
Bandwidth	Bandwidth used by the flow. This is a count of the number of packets through the port for the flow up to this point in time. For a value to be displayed in this field, flow analysis must have been performed on flow with the Capture Bandwidth option enabled.
Flow Analysis	Click View Results to see the results of the most recent flow analysis trace. The Flow Analysis Details window opens. NOTE: The View Results link is not available for VM to VM flows.

Flow Analysis Details Window

The Flow Analysis Details window provides detailed information about a flow trace.


The Flow Analysis Details window is divided into three sections:

- The flow path diagram—This diagram shows the path taken by a probe through the network. By default, the path shown is the path taken by the probe that experienced the highest per-hop latency in the trace.

You can change this diagram to reflect the path taken by a different probe by selecting the probe from the top Latency Trend chart.

- **Latency Trend charts**—These charts show the change in latency experienced by the probes during the trace. The bars in the top chart are grouped by completed probes, with each bar in a probe group representing the latency experienced by the probe at a hop. By clicking a probe group, you can change the flow path diagram and the Analysis Results section to reflect the results of that probe. For traces of long duration, the bar chart shows only a portion of the trace results.

The bottom area chart graphs the highest latency experienced by each probe over the entire duration of the trace. You can use the provided controls to focus on a portion of the trace—the portion you choose is reflected in the top bar chart. By default, the focus is on the portion of the trace that had the highest latency. If the trace is ongoing, a rotating circle appears at the end of the plotted area and the chart is periodically refreshed to show new results.

Both charts display a path change icon () when the path a probe takes through the network differs from the path taken by the previous probe.

- **Analysis Results**—This section provides details about the overall trace results and about the selected probe:
 - The Latency table provides overall latency information for the trace: the highest and lowest latency experienced at a single hop and the average latency of all hops.
 - The Latency for Selected Path table shows the latency experienced by the selected probe at each hop.

You can perform the following actions in the Flow Analysis Details window.

General actions:

- For bidirectional traces, you can select the direction for which you want results by clicking one of the arrows at the top of the window (these arrows do not appear for unidirectional traces).
- To stop an active flow analysis, click **Stop Flow Analysis** at the bottom of the window. When you stop an active flow analysis, the results up to the time you stopped the flow analysis are retained and the previously active trace is marked as complete.


On the flow path diagram, you can:

- Reposition the topology diagram by dragging it or reposition devices by dragging them.
- Zoom in or out by clicking the plus or minus signs on the left.
- Mouse over the link connecting two devices to get the connecting port names. The names are displayed in green if the link is up and in red if the link is down.
- Mouse over a device to view details about the device, such as name, connection state, and IP address. The details shown depends on the device type.

If a device in the flow path does not support Cloud Analytics Engine, it is shown in the diagram in light grey color and minimal details, such as IP address, are available.

- Display the traffic statistics for switches by mousing over the device to display the device details and clicking the **Show Traffic Data** link. If you selected the Capture Bandwidth option when you started the flow analysis, the flow bandwidth is also displayed along with the traffic statistics.
- Display the active flows associated with a VM, BMS, or virtualized host by mousing over the device and clicking **Show Active Flows** in the details box.

On the Latency Trend charts, you can:

- Mouse over a bar group in the top bar chart. A pop-up box displays the latency figures for each hop taken by the probe.
- Click a bar group in the top bar chart. The flow path diagram and the Analysis Results change to reflect the information for the probe.
- Mouse over a path change icon in the top bar chart. Information about the old and new paths is displayed.
- Change the span and position of the focus indicator on the bottom area chart:
 - To increase or decrease the time span of the focus—in other words, to zoom in or zoom out on a portion of the trace—click one of the handle controls () and move it in either direction.
 - To change the focus to another time period, click on the arrows at either end of the slider bar.

Viewing Traffic on a Device

The Traffic on Device window displays the details of traffic that flows through the selected port on a device, such as the applications that are running on the client system, IP address of the client system and the destination, protocol used by the application, data usage, and the data usage percentage. You can choose to view the real-time traffic analysis data on an interface or to view data over a specified period of time.

The Traffic on Device window displays traffic details in two modes—top applications and top conversations. Network Director displays this data in graphical and tabular format, for each mode.

To view details of traffic that passes through a port on a device:

1. Log in to Network Director.

The Dashboard View is selected by default. All the devices that are managed by Network Director in a particular network are represented as cells in the Device & Port utilization widget.

2. Click a device cell to view the ports associated with that device.

All the ports in the selected device are represented as cells. Mouse over the cells to open up a pop-up displaying the port information.

3. Mouse over a cell (port) to view the port information in a pop-up.

- Click **Analyse Traffic on this Port** in the pop-up.

The Traffic on Device :<port name> page is displayed.

- Select **Top Applications** (default) or **Top Conversations** to view traffic details sorted based on applications or conversations respectively.

- Select real-time or a time period for which you want to view the traffic analysis data.

Network Director displays the traffic details on the selected port for the time period you specified. If you chose to view the real-time data, the data in the graph refreshes after each sampling interval.

The graphical view displays traffic from each application or conversation plotted against time (x-axis) and data usage (y-axis). In the Top Applications mode, Network Director displays the names of well-known applications such as *http*, *ftp*, and *ssh*.

[Table 15 on page 68](#) describes the fields that are displayed in the traffic details table.

- If you are viewing traffic data in the Top Applications mode and if you know the application that uses a particular protocol port, then select that corresponding port number from the list and click **Associate Application with Port**.

- Enter the name of the application and the protocol that the application uses. Click **Add**.

The name you entered replaces the name of the application in the list.

Table 15: Traffic on Device—Port Traffic Details Table Fields

Name	Mode	Description
Application	Top Applications	Name of the application.
Protocol	Top Applications	Protocol that ths application uses.
Ingress Bytes	Top Applications Top Conversations	Number of bytes that enter the device through the ingress interface for the given application or conversation for the selected duration.
Egress Bytes	Top Applications Top Conversations	Number of bytes that leave the device through the egress interface for the given application or conversation for the selected duration.

Table 15: Traffic on Device—Port Traffic Details Table Fields *(continued)*

Name	Mode	Description
Total Bytes	Top Applications Top Conversations	Total number of bytes that traversed through the port for the given application or conversation for the selected duration.
Percentage of Total Traffic	Top Applications Top Conversations	Percentage of traffic that the application or conversation uses with respect to the total traffic that traverses the port for the selected duration.

RELATED DOCUMENTATION

Understanding the Dashboard 56
Using Dashboard Widgets 57

Equipment By Type Widget

IN THIS SECTION

- [Equipment By Type Widget Summary | 69](#)
- [Equipment By Type Widget Details | 70](#)

The Equipment By Type widget provides summary and detailed information about the types of devices Network Director is managing.

This topic describes:

Equipment By Type Widget Summary

The Equipment By Type widget displays summary information about the types of devices Network Director is managing. The diagram represents the managed devices as a set of nested rings. The circle in the center of the diagram shows information about the ring segments when you mouse over them. The inner ring

divides the devices into segments that represent wired device types. The outer ring divides each of those types into more specific device type segments. Mouse over any diagram segment to see the device type and number of those devices that it represents in the center circle.

Equipment By Type Widget Details

To open the Equipment By Type widget details page, click the **Maximize** button in the widget's title bar. The Equipment By Type widget details window has a table containing detailed information about the devices Network Director is managing. [Table 16 on page 70](#) describes the columns in the table. To close the details page, click the **Minimize** button in the title bar.

Table 16: Equipment By Type Widget Details Table

Column	Description
Equipment Type	Device equipment type.
Platform	Device platform (model name).
Device Type	Device type.
Software Version	Software version running on the device.
Count	Number of devices of that platform in the inventory.

RELATED DOCUMENTATION

Understanding the Dashboard 56
Using Dashboard Widgets 57
Network Director Documentation home page

Port Status - Physical Widget

IN THIS SECTION

- [Port Status - Physical Widget Summary | 71](#)
- [Port Status - Physical Widget Details | 71](#)

The Port Status - Physical widget provides summary and detailed information about the status of physical ports on managed devices.

This topic describes:

Port Status - Physical Widget Summary

The Port Status - Physical widget displays summary information about the status of physical ports on managed devices. It has the following pie charts:

- Admin Status pie chart—Shows the distribution of ports that are administratively up or down and states the total number of ports. Mouse over a chart segment to see more information about it.
- Free vs. Used pie chart—Shows the distribution of ports that are free or used and states the total number of ports. Mouse over a chart segment to see more information about it.

Port Status - Physical Widget Details

The Port Status - Physical widget details window has a table containing detailed information about the status of physical ports on managed devices. See [“Port Status Monitor” on page 740](#) for descriptions of the table columns.

RELATED DOCUMENTATION

[Understanding the Dashboard | 56](#)

[Using Dashboard Widgets | 57](#)

[Network Director Documentation home page](#)

Top Talker - Wired Devices Widget

IN THIS SECTION

- [Top Talker - Wired Devices Widget Summary | 72](#)
- [Top Talker - Wired Devices Widget Details | 72](#)

The Top Talker - Wired Devices widget provides summary and detailed information about the hosts that are using the most bandwidth. Hosts are endpoints that are directly connected to access ports of wired switches.

This topic describes:

Top Talker - Wired Devices Widget Summary

The Top Talker - Wired Devices widget has a bar chart that shows summary information about the hosts that are using the most bandwidth. Host names are listed on the vertical axis. Data usage in kilobytes is shown on the horizontal axis. Mouse over a bar to see more information about that host.

Top Talker - Wired Devices Widget Details

To open the Top Talker - Wired Devices widget details page, click the **Maximize** button in the widget's title bar. The Top Talker - Wired Devices widget details window has a table containing detailed information about the hosts that are using the most bandwidth. [Table 17 on page 72](#) describes the columns in the table. To close the details page, click the **Minimize** button in the title bar.

Table 17: Top Talker - Wired Devices Widget Details

Column	Description
Host Name	Host's host name.
MAC Address	Host's MAC address
Data Usage (KBytes)	Data used by the host, in kilobytes.
Device Serial Number	Device's serial number.

RELATED DOCUMENTATION

[Understanding the Dashboard | 56](#)

[Using Dashboard Widgets | 57](#)

[Network Director Documentation home page](#)

Top Overlay Networks Widget

IN THIS SECTION

- [Top Overlay Networks Widget Summary | 73](#)
- [Top Overlay Networks Widget Details | 74](#)

Virtual Extensible Local Area Network (VXLAN) represents a technology that enables you to segment your networks (as VLANs do), but that also solves the scaling limitation of VLANs and provides benefits that VLANs cannot. VXLAN is often described as an overlay technology because it enables you to stretch Layer 2 connections over an intervening Layer 3 network by encapsulating (tunneling) Ethernet frames in a VXLAN packet that includes IP addresses.

This topic describes:

Top Overlay Networks Widget Summary

The Top Overlay Networks widget displays a summary of the VXLANs in your network in a table. [Table 18 on page 73](#) displays the details that are displayed in this table.

Table 18: Top Overlay Networks Widget Summary page Field Descriptions

Column	Description
VXLAN	Unique ID of the VXLAN.
Tenant	Name of the tenant that uses the given VXLAN.
Network	IP address and the subnet mask of the network that is assigned to the tenant.
VMs	Number of virtual machines (VMs) that are active for the tenant.
Aggregate Bandwidth	Aggregated bandwidth used by the overlay network for the last 10 minutes.

Click **Show Details** corresponding to a VXLAN entry to view detailed information about that VXLAN. The Top Overlay Networks Widget Details page opens.

Top Overlay Networks Widget Details

To open the Top Overlay Networks Widget Details page, click **Show Details** in the Top Overlay Networks Widget table in the Summary View. Top Overlay Networks Widget Details page has two tabs—List View and the Topology View.

The list view displays detailed information about the VMs that are part of the selected overlay network or VXLAN. [Table 19 on page 74](#) describes the fields in this page.

Table 19: Top Overlay Details page Field Descriptions

Column	Description
VM Name	Name of the VM.
BW Utilization	Average bandwidth utilized by the VM for the last 10 minutes.
Host Name	The ESXi hostname of the VM.
Guest Operating System	Operating system running on the VM.
IP Address	IP address of the VM.
MAC Address	MAC address of the VM.

The Topology view highlights the VMs and the bare metal servers that are part of the selected VXLAN. You can mouse over each entity to view more details.

RELATED DOCUMENTATION

[Understanding the Dashboard | 56](#)

[Using Dashboard Widgets | 57](#)

[Network Director Documentation home page](#)

3

PART

Working in Build Mode

About Build Mode | **77**

Discovering Devices | **94**

Setting Up Sites and Locations Using the Location View | **105**

Building a Topology View of the Network | **125**

Creating Custom Device Groups | **152**

Configuring Quick Templates | **161**

Configuring Device Settings | **169**

Configuring Authentication, Authorization, and Access for Your Network | **202**

Configuring Interfaces and VLANs | **246**

Configuring Firewall Filters (ACLs) | **349**

Configuring Class of Service (CoS) | **396**

Configuring Media Access Control Security (MACsec) | **417**

Configuring Link Aggregation Groups (LAGs) | **425**

Creating and Managing Fabrics | **461**

Configuring Overlay Networks and Tenants | **527**

Configuring VRRP Profiles | **544**

Managing Network Devices | **550**

About Build Mode

IN THIS CHAPTER

- [Understanding Build Mode in Network Director | 77](#)
- [Understanding the Build Mode Tasks Pane | 82](#)
- [Understanding Network Configuration Profiles | 88](#)
- [Assigning Profiles to an Interface, Device, or a Group of Devices | 91](#)

Understanding Build Mode in Network Director

IN THIS SECTION

- [Discovering Devices | 77](#)
- [Building the Logical, Location, and Custom Views | 78](#)
- [Configuring Devices | 79](#)
- [Managing Devices | 81](#)

In Build mode, you build the network managed by Junos Space Network Director. It provides you with the ability to use device discovery to bring devices under Network Director management, to customize your view of the devices, to configure devices, and to perform some common device management tasks.

This topic describes:

Discovering Devices

Device discovery finds your network devices and brings them under Network Director management. You provide Network Director with identifying information about the devices you want Network Director to manage—an IP address or hostname, an IP address range, an IP subnetwork, or a CSV file that contains this information. Network Director uses the information to probe the devices by using either ping or SNMP get requests. If a device probe is successful, Network Director then attempts to make an SSH connection

to the device using the login credentials you supply. If the connection is successful and the device is a supported device, Network Director adds the device to its database of managed devices. Network Director uses Juniper Network's Device Management Interface (DMI), which is an extension to the NETCONF network configuration protocol, to connect to and configure its managed devices.

You can also discover devices using the device discovery feature provided by the Junos Space Network Management Platform. Devices you discover using Junos Space device discovery are brought under Network Director management if they are supported by Network Director.

Besides bringing your devices under Network Director management, device discovery:

- Reads the device configuration and saves it in the Junos Space configuration database. Network Director uses this record of the device configuration to determine what configuration commands it needs to send to a device when you deploy the configuration on the device. For this reason, it is important for the Junos Space configuration record to match, or be in sync with, the device configuration. For more information about how the Junos Space configuration record and device configuration are kept in sync, see [“Understanding Resynchronization of Device Configuration” on page 611](#).
- Imports the device configuration into the Build mode configuration. For more information about importing device configurations, see [“Importing Device Configurations” on page 80](#).

Building the Logical, Location, and Custom Views

When a device is discovered in the physical network mode, it is added to the network tree in the View pane. In Logical View, all switches are added to the Unassigned node under the switching network. You can then assign them to the Access, Aggregation, or Core nodes to complete the Logical View of the switching network.

Similarly, in Location View, all discovered devices are added to the Unassigned node. You can use Build mode to create the Location View—that is, create the sites, buildings, floors, closets, and outdoor areas that reflect the physical location of your network devices—and to assign the discovered devices to these locations.

NOTE: Network Director displays the Virtual Chassis and Virtual Chassis Fabric systems in the Location view network tree only if at least one of their member devices are *not* assigned to any location entity. If all the member devices are assigned to location entities, then the Virtual Chassis and Virtual Chassis Fabric systems are not displayed in the network tree.

The Custom Group View displays only the top level—My Network—until you create one or more custom groups. Custom group is another way of grouping your devices based on your business needs. You can create custom groups and add devices to each custom group. You can manually add devices to a custom group or you can define rules to add devices, that match the rule condition, to the custom group once they

are discovered by Network Director. You can view the custom groups and devices that are assigned to each group in the Custom Group view.

NOTE: This section does not apply to virtual devices that Network Director manages.

Configuring Devices

In Build mode, you can define the configuration of network devices in your Physical network. To support rapid, large-scale deployment of devices, you can define much of your Build mode configuration in a set of profiles. You can reference profiles in other profiles or apply them to multiple objects in your network—devices, ports, logical entities. For example, you can create a Port profile that sets up class-of-service (CoS), authentication, firewall filters, and Ethernet switching settings that are appropriate for access ports that connect to employee desktop VoIP phones and then assign that profile to access ports across multiple switches.

NOTE: This section does not apply to virtual devices that Network Director manages.

[Figure 8 on page 80](#) shows an example landing page for profile configuration, in this case VLAN profiles. This page lists all existing VLAN profiles. From this page you can create new profiles, modify or delete existing profiles, assign profiles to objects, and change or view assignments. For more information about profiles, see [“Understanding Network Configuration Profiles” on page 88](#).

Figure 8: Manage VLAN Profiles Page

Profile Name	VLAN Name	Family Type	VLAN ID	VLAN Range	VLAN ID List	Description	Assignment State
v300_152	v300	Campus Switching ELS	300	-	-	-	Assigned
v200_150	v200	Campus Switching ELS	200	-	-	-	Assigned
v123_148	v123	Campus Switching ELS	123	-	-	-	Assigned
v100_146	v100	Campus Switching ELS	100	-	-	-	Assigned
v1_144	v1	Campus Switching ELS	1	-	-	-	Assigned
v100_5	v100	Campus Switching ELS	1000	-	-	-	Assigned

In addition to creating configuration profiles, in Build mode you can configure network domains, Link Aggregation Groups (LAGs) on switches, and so on.

Deploying Device Configurations

After you build your device configurations in Build mode, you need to deploy the configurations on the devices. None of the configurations you create in Build mode affect your devices until the configurations are actually deployed on the devices.

To deploy the configuration on devices, use Deploy mode. When you change a device's configuration in Build mode, the device becomes available in Deploy mode for configuration deployment.

For more information about deploying configuration changes, see [“Deploying Configuration Changes” on page 573](#).

Importing Device Configurations

As part of device discovery, Network Director analyzes the configuration of a newly discovered device and automatically imports the configuration into the Build mode configuration for that device.

As it imports the device configuration, Network Director automatically creates profiles to match the configuration. It first determines whether any existing profiles match the configuration, and if so, assigns those profiles to the device. It then creates and assigns new profiles as needed. For example, if an access switch has some ports that match the configuration of an existing Port profile, Network Director assigns the existing Port profile to those ports. For the other ports, Network Director creates as many Port profiles as needed to match the port configurations and assigns them to the ports.

You can manage the profiles that Network Director creates as part of device discovery in the same way that you manage user-created profiles—that is, you can modify, delete, or assign them to other devices.

Out-of-Band Configuration Changes

Out-of-band configuration changes are configuration changes made to a device outside of Network Director. Examples include changes made by:

- Using the device CLI.
- Using the device Web-based management interface (the J-Web interface or Web View).
- Using the Junos Space Network Management Platform configuration editor.
- Using RingMaster software.
- Restoring or replacing device configuration files.

When an out-of-band change is made, the device configuration no longer matches the Build mode configuration, and the device configuration state changes to out of sync. You cannot deploy configuration on a device that is out of sync. Use the Resynchronize Device Configuration task in Deploy mode to resynchronize the device configuration. For more information about how Network Director resolves out-of-band configuration changes and synchronizes the Build mode configuration with the device configuration, see [“Understanding Resynchronization of Device Configuration” on page 611](#).

TIP: Before you make configuration changes in Build mode, make sure that devices that will be affected are in sync. Resynchronizing the device configuration can result in losing pending Build mode configuration changes for that device.

Managing Devices

In addition to the tasks that allow you to build your network, Build mode provides a number of tasks for day-to-day device management. For example, you can:

- View a device’s hardware component inventory or its installed licenses
- Reboot a device or groups of devices
- Connect to a device’s CLI through SSH or to its web-based management interface
- View the profiles assigned to a device

RELATED DOCUMENTATION

[Understanding the Build Mode Tasks Pane | 82](#)

[Understanding the Network Director User Interface | 3](#)

[Understanding Network Configuration Profiles | 88](#)

Understanding the Build Mode Tasks Pane

The Tasks pane in Build mode contains all the tasks you can do in Build mode. Click a specific task to begin that task.

The tasks listed in the Tasks pane depend on the scope you select in the View pane—that is, what view (Logical, Location, Device, Virtual, or Custom Group) you have selected and what object you have selected. Not all tasks are available in all scopes. As you change your selections in the View pane, the contents of the Tasks pane also change.

Build mode tasks are divided into the following categories in the Tasks pane.

Network Director enables you to perform the following tasks for devices in your physical network:

- **Device Discovery**—Before your devices can be managed by Network Director, you must use device discovery to discover them. As Network Director discovers devices, it adds them to your network view in the View pane. [Table 20 on page 83](#) describes the device discovery tasks.
- **Device Management**—After devices have been discovered, you can perform administrative tasks on them, such as viewing a list of the device's physical components, connecting to a device using SSH, rebooting a device, or assigning a switch to its logical role in the network. [Table 21 on page 83](#) describes the device management tasks.
- **Wired**—You can create configuration profiles and quick templates for the different wired devices—EX Series Ethernet Switches, EX Series Switches, QFX Series switches and MX Series routers.
- **Location Management**—You can build your Location view of the network by creating sites, buildings, floors, closets, and outdoor areas and assigning devices to these locations. [Table 22 on page 85](#) describes the location management tasks.
- **Connectivity**—For switches in your network that are connected to your virtual network, you can view the connectivity between a given switch and the corresponding virtual switch and between the virtual switch and the virtual machine. [Table 23 on page 86](#) describes the connectivity tasks.
- **Profile and Configuration Management**—Network Director provides a set of configuration profiles that you can create to provision multiple devices in your network. [Table 24 on page 86](#) describes the profile and configuration management tasks.
- **Key Tasks**—Network Director enables you to group the tasks that you perform frequently and create a list of key tasks. You can add any task from the Tasks pane to the Key Tasks list by selecting a task and clicking the plus (+) sign that appears adjacent to the task. For some modes, you can see that Network

Director has predefined some key tasks for you. You can modify this set of tasks to suit your requirements. This feature is available in Task pane irrespective of your current mode, scope, or view.

For more information about Build mode features, see [“Understanding Build Mode in Network Director” on page 77](#).

[Table 20 on page 83](#) through [Table 24 on page 86](#) describe the tasks that you can perform in the physical network category, including the scope in the View pane that you must select to access the task.

Table 20: Device Discovery Tasks

Task	Description	Scope
Discover Devices	Discovers supported switches in the network and brings them under Network Director management.	Any
View Discovery Status	Displays the status of device discovery jobs.	Any

Table 21: Device Management Tasks

Task	Description	Scope
Assign Device to Logical Category	Assigns switches to one of the following logical categories within the switching network: Core, Aggregation, Access.	View: Logical Object: Switching Network, Core, Aggregation, Access, Unassigned, or an individual switch
Change Location of Device	Changes where a device is located in Location view.	View: All Object: Individual switch
Create MC-LAG	Configures multichassis link aggregation groups (MC-LAG).	View: All Object: All
Create ESI-LAG	Configures Ethernet Segment Identifier link aggregation groups (ESI-LAG) in a campus environment.	View: All Object: All
Delete Devices	Deletes a switch as a managed device from Network Director. If you select a scope that contains more than one switch, you can choose which devices are deleted.	View: All Object: All
Launch Web View	Launches the Web-based management interface for the selected device in a separate window: the J-Web interface for EX Series switches .	View: All Object: Individual switch

Table 21: Device Management Tasks (*continued*)

Task	Description	Scope
Manage LAG	Creates and manages Link Aggregation Groups (LAGs).	View: All Object: Individual switch
Manage Port Groups	Enables you to create and manage port groups. Port groups enable you to configure a group of ports with a single configuration.	View: Logical Object: Switching Network, Core, Aggregation, Access, Unassigned, an individual switch (both EX Series and QFX Series)
Manage Virtual Chassis Fabric (VCF)	Manages a Virtual Chassis Fabric (VCF)	View: All Object: VCF
Reboot Devices	Reboots devices. If you select a scope that contains more than one switch, you can choose which devices get deleted.	View: All Object: All
Setup Virtual Chassis Fabric (VCF)	Sets up a Setup Virtual Chassis Fabric (VCF)	View: All Object: All
Show Current Configuration	Shows the running configuration on a switch.	View: All Object: Individual switch
SSH to Device	Launches an SSH connection to the selected device.	View: All Object: Individual switch
Validate Pending Configuration	Validates configuration changes that have not yet been deployed on devices.	View: All Object: All
View Assigned Profiles	Displays the profiles assigned to the selected device.	View: All Object: Individual switch
View Inventory	Displays information about all the devices in the currently selected object and all its child objects.	View: All Object: All
View License Information	View the licenses installed on the device and their status.	View: All Object: Individual switch
View Physical Inventory	Displays information about the selected device's hardware components.	View: All Object: Individual switch

Table 22: Location Management Tasks

Task	Description	Scope
Add Building	Creates a new building in the selected site. NOTE: Use this task only to create the building. Floors and closets in the building must be created separately.	View: Location Object: A site
Add Closet	Creates a new closet in the selected floor.	View: Location Object: A floor
Add Floor	Creates a new floor in the selected building. NOTE: Use this task only to create the floor. Closets in the building must be created separately.	View: Location Object: A building
Add Outdoor Area	Creates a new outdoor area in the selected site.	View: Location Object: A site
Add Site	Creates a new site in Location view. NOTE: Use this task only to create the site object. Buildings, floors, closets, and outdoor areas in the site must be created separately.	View: Location Object: My Network only
Delete Building/Edit Building	Deletes or modifies the selected building.	View: Location Object: A building
Delete Closet/Edit Closet	Deletes or modifies the selected closet.	View: Location Object: A closet
Delete Floor/Edit Floor	Deletes or modifies the selected floor.	View: Location Object: A floor
Delete Outdoor Area/Edit Outdoor Area	Deletes or modifies the selected outdoor area.	View: Location Object: An outdoor area
Delete Site/Edit Site	Deletes or modifies the selected site.	View: Location Object: A site

Table 22: Location Management Tasks (*continued*)

Task	Description	Scope
Assign Devices to Building	Assigns switches to a building.	View: Location Object: A building
Assign Devices to Closet	Assigns switches to a closet.	View: Location Object: A closet
Assign Devices to Floor	Assigns switches to a floor.	View: Location Object: A floor
Setup Locations	<p>Opens the page by using which you can create an entire site—that is, define buildings, floors, closets, outdoor areas and to assign devices to these locations.</p> <p>NOTE: Use this task only to create a site. Do not use it to modify an existing site.</p>	View: Location Object: My Network and any location node within an existing site.

Table 23: Connectivity Tasks

Task	Description	Scope
View Virtual Network Connectivity	<p>Pictorially displays the network connectivity between the selected switch and the virtual switch and between the virtual switch and the virtual machine.</p> <p>If the selected switch is not connected to a virtual network, Network Director displays the standalone switch.</p>	View: Logical, Location, Device Object: Individual switch
View Virtual Machines	Displays the virtual machines that are connected to the selected switch.	View: Logical, Location, Device Object: Individual switch

Table 24: Profile and Configuration Management Tasks

Task	Description	Device Family	Scope
Manage Quick Templates	Enables you to create and manage quick templates. Quick templates enable you to define your network configuration in the form of templates that you can apply to multiple devices in your network.	EX Series QFX Series MX Series	All

Table 24: Profile and Configuration Management Tasks (*continued*)

Task	Description	Device Family	Scope
View Deployed Templates	Enables you to view the list of quick templates that are deployed.	EX Series QFX Series MX Series	All
Access	Creates and manages Access profiles. Use Access profiles to configure authentication methods (RADIUS, LDAP, and local), accounting methods (RADIUS and LDAP), and authentication/accounting servers.	EX Series QFX Series	Any
Authentication	Creates and manages Authentication profiles. Use Authentication profiles to specify authentication method and authentication parameters for authenticating clients and users who connect to an access port on a switch.	EX Series QFX Series	Any
CoS	Creates and manages CoS profiles. Use CoS profiles to configure class-of-service (CoS) attributes to be applied to interfaces or to user traffic.	EX Series QFX Series	Any
Device Common Settings	Creates and manages Device Common Settings profiles. Use Device Common Settings profiles to configure basic system settings, such as users, time and time servers, SNMP, system logging, and so on.	EX Series QFX Series	Any
Fabric	Creates and manages Fabric profiles. Use Fabric profiles to configure gateway FC fabrics on QFX Series devices that act as a FCoE-FC gateway.	QFX Series	Any
Filter	Creates and manages Filter profiles. Use Filter profiles to define Layer 2 and Layer 3 firewall filters (ACLs).	EX Series QFX Series	Any
LDAP	Creates and manages LDAP profiles. Use LDAP profiles to specify details about your LDAP authentication and accounting server. LDAP profiles can then be linked to an access profile.	EX Series	Any
Port	Creates and manages Port profiles for EX Series switches. Use Port profiles to configure interface settings, such as PoE settings, protocol family, port mode, physical link settings, firewall filters, and port security settings for interfaces.	EX Series QFX Series	Any

Table 24: Profile and Configuration Management Tasks (*continued*)

Task	Description	Device Family	Scope
Radius	Creates and manages Radius profiles. Use Radius profiles to specify details about your RADIUS authentication and accounting server. Radius profiles can then be linked to an access profile.	EX Series	Any
VLANs	Creates and manages VLAN profiles. Use VLAN profiles to define VLANs, including the firewall filters to be applied to the VLANs and other settings.	EX Series QFX Series	Any

RELATED DOCUMENTATION

[Understanding Build Mode in Network Director | 77](#)

[Understanding Network Configuration Profiles | 88](#)

[Understanding the Network Director User Interface | 3](#)

[Understanding Quick Templates | 161](#)

[Network Director Documentation home page](#)

Understanding Network Configuration Profiles

To support rapid network deployment, Junos Space Network Director enables you to define your network configuration in a set of profiles that you can apply to multiple objects in your network. For example, you can define a Port profile to set up class-of-service (CoS), authentication, firewall filters, and Ethernet switching settings that are appropriate for all access ports in your network that connect to employee desktop VoIP phones.

After you have defined a profile, you can associate it with devices in one of two ways:

- By directly assigning it to a device (or to ports on the device). When you assign a profile to a device, you can configure certain device-specific parameters. For example, when you assign a VLAN profile to a device, you can configure the IP address for that VLAN on that device. Or when you assign a Port profile for a Layer 3 interface to the interface, you can configure the IP address for that interface.
- By referencing the profile in another profile. Some profiles are not assigned directly to network devices—instead they are referenced from other profiles that are, in turn, assigned to network devices. For example, the settings in the CoS, Filter, and Authentication profiles are assigned indirectly to a port by the profiles being included in the Port profile.

Because a child profile might be a required setting in its parent profile, you must create the child profiles before you create the parent profiles. For example, to create a Port profile, create the profiles in this order:

1. Access, VLAN, CoS, and Filter profiles
2. Authentication

Network Director also includes six predefined Port profiles and one predefined CoS profile for EX Series switches. You can choose to apply the Port profiles to one or more ports of a single device or a group of devices, and the CoS profile to a Port profile (using an Authorization profile).

After you have created and included the child profiles in to parent profiles, you can assign these parent profiles at various levels in your wired networks. [Table 25 on page 89](#) shows the levels at which you can assign each of these parent profiles.

Table 25: Profile Associations at Various Levels

Name of the Profile	EX Series Ethernet Switches	QFX Series Switches
Device Common Settings profile	Device	Device
VLAN profile	Device	Device
Port profile	Port	Port
Fabric profile	Not applicable	Port
FC Gateway Service	Not applicable	Port
Local Switching VLAN profile	Not applicable	Not applicable
mDNS profile	Not applicable	Not applicable
Remote Sites profile	Not applicable	Not applicable

Once you have assigned profiles to devices or ports, you can view the profile associations in the **Profiles Assigned to the Device** page. For more information, see [“Viewing Profiles Assigned to a Device” on page 558](#).

In addition to the profiles you create yourself, Network Director creates profiles for you from existing device configuration. Typically, you create profiles and associations manually when you are setting up a new network from scratch, adding a new device to your existing network, or when you want to make certain customized changes to the way your network is currently operating. Network Director creates profiles for you when:

- You discover existing devices in your network. As part of device discovery, Network Director examines the configurations present in the discovered device or devices. If configurations match existing profiles,

Network Director assigns the matching profiles to the appropriate levels on the devices. If configurations do not match existing profiles, Network Director creates the required profiles and associates them at the appropriate levels. For more information about device discovery, see [“Discovering Devices in a Physical Network” on page 94](#).

NOTE: If Network Director fails to read the configuration of one or more devices after the device discovery, such devices are not displayed in the Assign Profile page. You will not be able to assign profiles to such devices. The Manage Jobs page in System mode displays details of the device discovery jobs. Use the information displayed on this page to take appropriate corrective steps to enable Network Director to reread the configuration of the failed device. For more information, see [“Troubleshooting Device Discovery Error Messages” on page 102](#).

- You first install Network Director and supported devices are already being managed by Junos Space. In this case, Network Director imports the device configurations into profiles the same way it does when you discover devices with Network Director.
- You resynchronize the Network Director configuration with the device configuration in order to resolve out-of-band configuration changes—that is, configuration changes that are not made with Network Director. Out-of-band configuration changes result in the device configuration not matching or being in sync with the Network Director configuration for the device. When you resynchronize the Network Director configuration with the device configuration, Network Director creates and associates new profiles if none of the existing profiles match the changed configuration. For more information about resynchronization of device configuration, see [“Understanding Resynchronization of Device Configuration” on page 611](#).

After a profile is created, you can edit it from the Manage Profile page by selecting the profile and clicking Edit. The only exception is when the profile that you want to edit is part of a job that is scheduled for deployment. When you schedule a deployment job, that job and any profiles assigned to that job are locked. You cannot edit the job or any of its assigned profiles until the job is completed or gets cancelled. For more information, see [“Deploying Configuration Changes” on page 573](#).

When you delete a device in Network Director, the system deletes only the device and the profile associations. The profiles are retained in the system. If you rediscover the deleted devices into the system at a later stage, without making any configuration changes on the device, Network Director identifies this and reinstates the previous profile associations.

RELATED DOCUMENTATION

[Understanding Access Profiles | 217](#)

[Understanding Authentication Profiles | 235](#)

[Understanding Class of Service \(CoS\) Profiles | 396](#)

[Understanding Device Common Settings Profiles | 169](#)

[Understanding Filter Profiles | 349](#)

[Understanding Port Profiles | 246](#)


[Understanding VLAN Profiles | 327](#)

[Network Director Documentation home page](#)

Assigning Profiles to an Interface, Device, or a Group of Devices

After you create an authorization profile, CoS profile, device common settings profile, fabric profile, FC Gateway services profile, filter profile, port profile, VLAN profile, and Local Switching profile, you can assign each of these profile to an interface, device, or a group of devices.

To assign a profile:

1. Click  **Build** in the Network Director banner.
2. Under Select View, select either **Logical View**, **Location View**, **Device View** or **Custom Group View**.

TIP: Do not select **Dashboard View**, **Datacenter View** or **Topology View**.

3. From the Tasks pane, select the type of network (Wired), the appropriate functional area (System or AAA), and select the name of the profile that you want to create. For example, to create a port profile for a wired device, click **Wired > Profiles > PORT**. The Manage Profile page opens.
4. Select the profile that you want to assign and click **Assign**.
The Assign Authorization Profile page appears displaying a hierarchal list of network objects that is already defined or discovered for your network.
5. Select a level and click **Next** to view the objects available at that level.
6. Select one or more devices or groups from the list.

NOTE: If Network Director fails to read the configuration of one or more devices after the device discovery, such devices are not displayed in the Assign Profile page. You will not be able to assign profiles to such devices. The Manage Jobs page in System mode displays details of the device discovery jobs. Use the information displayed on this page to take appropriate corrective steps to enable Network Director to reread the configuration of the failed device. For more information, see [“Discovering Devices in a Physical Network” on page 94](#).

7. If you want to view the existing assignment of a device, select it and click **View Assignments**.

The Profile Details window opens displaying the device's current profile assignment.

Click **Close** to close the window.

8. If you want to remove an existing assignment from a device, select it and click **Remove**. The system removes the assignment from the selected device.

9. Do one of the following to assign the Authorization profile to a device:

- Click **Assign > Assign to Device** to assign the Authorization profile to the selected devices.

10. Click **Next** or **Review**.

The system displays the associations that you created. To modify any of these assignments, click **Edit** or **Profile Association**.

11. Click **Finish** to save the profile associations.

After you click Finish, the Create Profile Assignments Job Details window opens with a report on the status of the profile assignment job. If you have assigned the profile to a large number of objects, the profile assignment job can take some time to complete. Instead of waiting for the Job Details window to report job completion status, you can close it and check the details of the profile assignment job at a later time using the Manage Job task in System mode.

NOTE: If any assignment fails, the profile assignment job fails and none of the assignments are created. Check the details for the profile assignment job for information about why the assignment failed.

RELATED DOCUMENTATION

[Understanding Filter Profiles](#) | 349

[Understanding Class of Service \(CoS\) Profiles](#) | 396

[Network Director Documentation home page](#)

Discovering Devices

IN THIS CHAPTER

- Discovering Devices in a Physical Network | 94
- Understanding the Device Discovery Process | 101
- Troubleshooting Device Discovery Error Messages | 102

Discovering Devices in a Physical Network

IN THIS SECTION

- Specifying Target Devices | 95
- Specifying Discovery Options | 97
- Specifying Schedule Options | 99
- Reviewing Device Discovery Options | 99
- Viewing the Discovery Status | 99

You can discover and synchronize physical devices such as EX Series switches, QFX Series devices, MX Series routers, in your network that are managed by Network Director.

You can also discover and manage virtual devices in your virtual network from Network Director.

NOTE: To discover, EX Series switches, Network Director connects to port 22 (the default port) on the Junos Space JA2500 Appliance or the Junos Space Virtual Appliance by using SSH. You can configure port 22 on the Junos Space appliances through **Administration > Applications** on the Junos Space Platform page. Select **Network Application Platform** and click **Actions > Modify Application Settings**. Change the SSH port for device connection to 22.

Device discovery is a three-step process in which you specify the target devices, the discovery options, and the schedule options.

While in Build mode, from the Tasks pane, click **Discover Devices** from the Device Discovery menu. The Discover Devices page is displayed.

This topic describes:

Specifying Target Devices

You can add devices to Network Director for device discovery by clicking either **Import from CSV** or **Add**, or both together. Click **Import from CSV** to add devices in bulk. You can add a large number of devices to Network Director by using a CSV file that contains information extracted from an LDAP repository. During device discovery, you can associate the devices with logical, location, and custom groups. You can list all devices to be discovered in the CSV file along with their logical, location, and custom groups. This eliminates the need to make an explicit association later. If you do not assign groups to the devices, the devices are added to the Unassigned folder by default. You can also change the assignment later. Associating new devices with groups makes the network simpler to manage and maintain.

To specify the target devices that you want Network Director to discover:

1. Enter a name for the device discovery job.

The default name is ND Discovery.

2. To add devices in bulk, click **Import from CSV** from the Device Targets window.

The Upload CSV File dialog box is displayed.

3. Click **Browse**.

The File Upload dialog box is displayed.

4. Navigate to the target CSV file on your computer, select the file, and click **Open**.

The CSV File Upload dialog box reappears, this time displaying the name of the selected file.

NOTE: The selected CSV file must follow the same file format as that of the sample CSV file.

5. Click **Upload** to upload the selected CSV file.

6. To add individual devices by specifying the IP address credentials, click **Add** in the Device Targets table.

The Add Device Target dialog box appears.

7. In the Add Device Target dialog box, perform the following steps:

a. Choose one of the following options to specify target devices:

- Select the **IP** option and enter the IP address of the device.
- Select the **IP-Range** option and enter a range of IP addresses for the devices.

The maximum number of IP addresses for an IP range target is 1024.

- Select the **IP-Subnet** option and enter an IP subnet for the devices.
- Select the **HostName** option and enter the hostname of the device.

b. In the Assign To section, specify the following groups to which the newly discovered devices can be assigned:

- From the Logical Group drop-down menu, select **Core**, **Aggregation**, **Access**, or **Layer 3 Fabric** to specify the logical grouping of the device.

Select the Layer 3 Fabric option to discover a Layer 3 fabric that is not created using Network Director and OpenClos. You can discover devices in Network Director that belong to the same IP subnet. To discover a Layer 3 Fabric, specify the IP subnet range, as all the devices that belong to the same Layer 3 Fabric resides in the same subnet. Network Director expands the IP subnet range and reaches every single IP address that you have specified in the IP range.

Network Director initially discovers the Layer 3 Fabric based on the IP subnet and range. However, you can manually discover Layer 3 Fabric at a later stage by entering the host name of the Layer 3 Fabric.

- For the Location Group field, click **Select** to choose the location group for the device or input the location path for the association. To clear the selection, click **Clear**.

Use the following format for the location path for the respective associations:

- *site-name#S/building-name#B*
- *site-name#S/building-name#B/floor-name#F#floor-level*
- *site-name#S/building-name#B/floor-name#F#1 - 1st level*
- *site-name#S/building-name#B/floor-name#F#floor-level/closet-name#C*
- *site-name#S/building-name#B/floor-name#F#floor-level/aisle-name#A/rack-name#R*
- *site-name#S/outdoorarea-name#O*

If the location paths do not point to existing locations in Network Director, new location groups are created before the devices are assigned to them.

- For the Custom Group field, click **Select** to choose the custom group for the device or input the custom group path. To clear the selection, click **Clear**.

Use the following format for the custom group path:

- *customgroup1-name/customgroup2-name*

If the custom group paths do not point to existing custom groups, new groups are created before the devices are assigned to them.

- c. Click **Add** to save the target devices that you specified, or click **Add More** to add more target devices. When you have added all target devices that you want Network Director to discover, click **Add**.

The Discover Targets table displays the addresses of the configured target devices.

8. Following device discovery management options are available:

- To edit a target device, select a row from the Device Targets table and click **Edit**. Make the required changes and click **Add** to display the IP addresses in the Device Targets table
- To delete a target device, select a row from the Device Targets table and click **Delete**.
- To view and download a sample CSV file, click **CSV Sample**. The Opening Device_Discovery_CSV.csv file dialog box is displayed. You can open the sample CSV file or save the sample CSV file.

9. Click **Next** or click **Discovery Options** from the top wizard workflow to go to the Discovery Options page. Specify the options as described in [“Specifying Discovery Options” on page 97](#).

Specifying Discovery Options

To add the device credentials and specify the probes:

1. Add the device credentials. To add the credentials, click **Add** from the Device Credentials table.

The Add Device Credentials dialog box is displayed.

NOTE: If the credentials were specified in the CSV file, the Credentials table displays those values. If the credentials were not specified in the CSV file, then enter the values in the Add Device Credentials dialog box.

- Specify the administrator username and password, and confirm the password. The username and password must match the name and password configured on the device. The username is a mandatory field.
- Click **Add** to save the username and password that you specified or click **Add More** to add another username and password.

Click **Add** after you have finished adding all login credentials. The Device Credentials table displays the usernames that you configured.

2. Specify the probes from the Specify Probes table. Select a probe method to discover the target devices.

- Select **Use Ping** if SNMP is not configured for the device and clear the **Use SNMP** check box.

Network Director uses the Juniper Networks Device Management Interface (DMI) to directly connect to and discover devices. DMI is an extension to the NETCONF network management protocol.

- Select **Use SNMP** if SNMP is configured for the device, and clear the **Use Ping** check box.

Network Director uses the **SNMP GET** command to discover target devices.

- Select both the **Use Ping** and the **Use SNMP** check boxes, to enable Network Director for faster discovery of the target devices, provided the device is pingable and also SNMP is enabled on the device.

NOTE: Network Director uses the Juniper Networks Device Management Interface (DMI) adapter to manage devices that do not run Junos OS. However, if you enable Use SNMP, Network Director detects whether the device is running a DMI-complaint software or not.

3. Click **Add** if you have selected the Use SNMP check box.

The Add SNMP Settings dialog box is displayed.

Select either **SNMP V1/V2C** or **SNMP V3**. Based on the selection, you need to enter the details as follows:

- If you selected SNMP V1/V2C, specify a community string, which can be *public*, *private*, or a predefined string.

Click **Add** in the Add SNMP Settings dialog box or click **Add More** to add more strings to the community. If you click **Add More**, when you are done adding all the strings, click **Add** to save the SNMP settings for V1/V2C.

- If you selected SNMP V3:

- Enter a username
- Select the privacy type (AES 128, DES, or None).
- Enter the privacy password (if AES 128 or DES). If you specify none for the privacy type, the privacy function is disabled.
- Select the authentication type (MD5, SHA, or none).
- Enter the authentication password (if MD5 or SHA). If you specify none for the authentication type, the authentication function is disabled.
- Click **Add** to save the SNMP settings and close the dialog box, or click **Add More** to add additional configurations. If you clicked Add More, click **Add** to save the settings and close the dialog box.

The Specify Probes table displays the configured SNMP settings.

4. Click **Next** or click **Schedule Options** from the top wizard workflow to go to the Discovery Schedule Options page. Specify the options as described in the [“Specifying Schedule Options” on page 99](#).

Specifying Schedule Options

To specify the scheduler details:

1. Click **Run Now** if you want to discover the devices immediately or Click **Schedule at a later time** if you want to schedule the device discovery for a future time.

If you select **Schedule at a later time**, specify the date and time to run the device discovery.

2. Click **Next** or click **Review** from the top wizard workflow to view the configuration. See [“Reviewing Device Discovery Options” on page 99](#).

Reviewing Device Discovery Options

From this page, you can save or make changes to the device discovery options.

- To make changes to the device discovery options, click the **Edit** button associated with the configuration you want to change.

Alternatively, you can click the appropriate buttons in the profile workflow at the top of the page that corresponds to the configuration you want to change.

When you are finished with your modifications, click **Review** to return to this page.

- Click **Finish** when you are done with the configurations.

A message window opens, displaying the status of the device discovery job name and job ID. Click **OK**.

The Device Discovery Jobs page is displayed with the list of jobs scheduled.

Viewing the Discovery Status

After you have configured the device discovery options, you can view the device discovery status from the **View Discovery Status** option from the **Device Discovery** menu.

The **Device Discovery Jobs** page displays all the scheduled device discovery jobs. You can view the following details from the Device Discovery Jobs page as described in [Table 26 on page 100](#).

Table 26: Viewing Device Discover Jobs

Field	Description
Job ID	An identifier assigned to the job.
Job Name	The name of the job (user-created).
Percent	Percentage of the job that is complete.
Status	<p>Job status. The possible states are:</p> <ul style="list-style-type: none"> • CANCELLED—The job was cancelled by a user. • FAILURE—The job failed. This state is displayed if any of the devices in the job failed. But some of the devices might have completed successfully. View the job details for the status of each device. • INPROGRESS—The job is running. • SCHEDULED—The job is scheduled but has not run yet. • SUCCESS—The job completed successfully. This state is displayed if all of the devices in the job completed successfully.
Summary	Summary of the job scheduled and executed with status.
Scheduled Start Time	The UTC time on the client computer when the job is scheduled to start.
Actual Start Time	The actual time when the job started.
End Time	The time when the job was completed.
User	The login ID of the user that initiated the job.
Recurrence	The recurrent time when the job will be restarted.

To view the details of a job, select the check box against Job ID or Job Name and click **Show Details**. The Discover Network Elements window displays details of the device discovery job.

NOTE: During device discovery, if Network Director is unable to read the device configurations, then the status displays Failed state. For such failures, you can check the reason for failure from the Manage Jobs page in System mode. You must make the required changes to the device configuration using the CLI so that Network Director can read the configuration. Network Director automatically resynchronizes once you enable a device discovery job. If Network Director cannot discover the device even after resynchronization, then you must rediscover the device after making the appropriate changes in the device configurations by using the CLI.

RELATED DOCUMENTATION

[Viewing the Device Inventory Page | 551](#)

[Troubleshooting Device Discovery Error Messages | 102](#)

[Network Director Documentation home page](#)

Understanding the Device Discovery Process

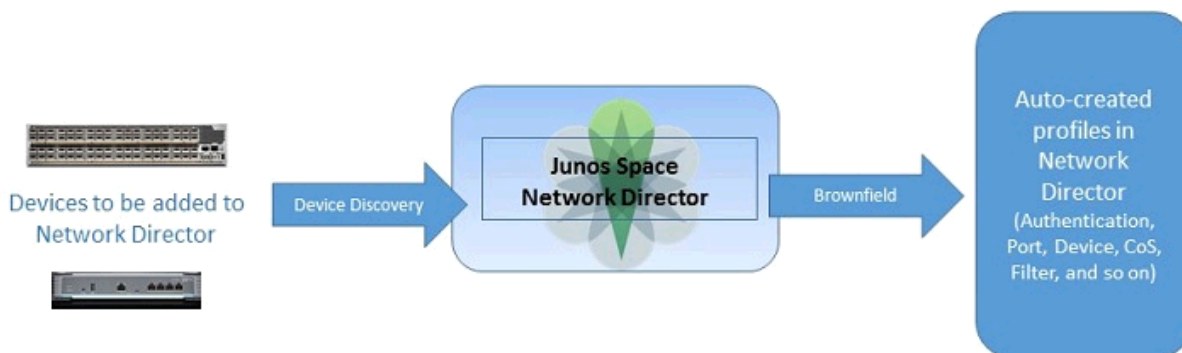
When a new device with network configurations is added to the network, Network Director runs a job to discover the device. Two minutes after device discovery, Network Director initiates another job called the brownfield process. The brownfield process ensures that the new device is ready to be used in the network by deploying the required configurations to the device.

To support rapid network deployment, Junos Space Network Director enables you to define your network configuration in a set of profiles that you can apply to multiple objects in your network. For example, you can define a Port profile to set up class-of-service (CoS), authentication, firewall filters, and Ethernet switching settings that are appropriate for all access ports in your network that connect to employee desktop VoIP phones.

You can manually create Profiles from the Network Director user interface or the profiles may be created automatically by Network Director when you discover a device. Once a device, that has network configurations, is discovered, Network Director initiates a Brownfield process to read the configuration and create the necessary profiles for all the supported configuration from the discovered device.

Figure 9 on page 101 displays how the brownfield process works in Network Director.

Figure 9: Brownfield Process



Using the brownfield process, Network Director completes the following actions:

- Fetching the complete configuration file of the newly discovered device and looking for matching profiles in the Network Director database.

- Using basic configurations from the already existing matching profiles in the database (such as VLAN IDs, ports, authentication protocols, class of service, firewalls, and so on) to deploy on the newly discovered device. If a matching profile does not exist, Network Director uses the configuration in the newly discovered device to create a new profile that can then be associated with other devices added to the network in the future.

Benefits of the Device Discovery Process

- A newly discovered device becomes functional in a matter of minutes after it is brought into the network because Network Director automatically assigns an existing profile to the device or creates a new profile without manual intervention.
- Device configurations are reused so you do not need to configure basic features (for example, the VLAN ID) for every newly discovered device added to the network.
- Bulk provisioning of profiles on devices means you can change any parameter (for example, VLAN) on the profile to effect the changes on multiple devices simultaneously.

RELATED DOCUMENTATION

[Understanding Network Configuration Profiles](#) | 88

Troubleshooting Device Discovery Error Messages

While you are discovering devices by using Network Director, you might encounter some issues. Network Director enables you to detect the errors and provide solutions to the potential errors that you encounter.

Error Message	Solution
Error Messages Displayed During Discovery of EX Series Switches	

Error Message	Solution
SSH connection failed. Device might not be reachable.	<p>For EX Series switches, Network Director connects to port 22 (default port) on the JA2500 Appliance or the Junos Space Virtual Appliance by using SSH. Ensure that you have configured port 22 on the Space appliance through Administration > Applications in the Junos Space Platform page. To do this, select Network Application Platform and click Actions > Modify Application Settings. Change SSH port for device connection field to 22.</p> <p>If port 22 is open on the Junos Space Appliance, and you still get the error, then check if port 22 is open on the switch and if the switch is accepting SSH connections on port 22.</p>
User Authentication failed.	Check the read and write credentials used during device discovery.
Device is not reachable.	If ping is enabled during device discovery, then check whether the switch is reachable using the CLI command ping .
Junos Space is unable to query the device information through SNMP. Check the SNMP settings on the device to verify SNMP is not blocked and the SNMP settings specified in Junos Space match the device SNMP settings.	If the SNMP option is enabled in Network Director during device discovery, check and ensure that SNMP is enabled on the switch. Also, check and ensure that the SNMP settings on Network Director and Junos Space match with the SNMP settings on the switch.
General Error Messages	

Error Message	Solution
Device Failed to return System information.	This message is displayed if the switch is too busy to respond to operational commands. Try discovering the device again.
Failed to configure device, Check Device state.	Check whether the Edit lock is open on the switch and close it if it is open. The configuration commit fails if the Edit lock is open.
Device has been added, but failed to synchronize. Please try manual re-synchronization. Error while reading config from device: device_name, Detail - Fail while executing following RPC: <get-configuration database=committed><configuration></configuration></get-configuration>	Try to resynchronize the devices manually. For details, see "Resynchronizing Device Configuration" on page 616 .
Error while reading config from device: device-name Failed while executing the following RPC: <get-hardware-inventory/>	<p>Check the hardware details of the switch using the CLI command show chassis hardware detail.</p> <p>If the output displays a message error: command is not valid, then the Junos OS image on the specified switch is corrupted and you need to upgrade to the latest version of Junos OS.</p>

RELATED DOCUMENTATION

[Discovering Devices in a Physical Network | 94](#)

[Resynchronizing Device Configuration | 616](#)

[Network Director Documentation home page](#)

Setting Up Sites and Locations Using the Location View

IN THIS CHAPTER

- [Understanding the Location View | 105](#)
- [Setting Up the Location View | 107](#)
- [Creating a Site | 111](#)
- [Configuring Buildings | 112](#)
- [Configuring Floors | 114](#)
- [Setting Up Closets | 116](#)
- [Assigning and Unassigning Devices to a Location | 117](#)
- [Changing the Location of a Device | 119](#)
- [Deleting Sites, Buildings, Floors, Wiring Closets, and Devices | 121](#)
- [Configuring Outdoor Areas | 123](#)

Understanding the Location View

The Location View is one of the perspectives that Network Director enables you to view and analyze your network. Using this view, you can view devices and data based on their physical location and proximity in the network. By physical location, we mean the buildings, floors, aisles, racks, wiring closets, and outdoor areas where the devices reside. After these locations are defined and devices assigned, the Location View gives you a visual representation of your devices based on where they reside.

You can define the physical location where the devices in the network are deployed in a hierarchical way, and define location entities from a site down to the wiring closet. When in the Location View, the network tree shows the network in terms of buildings, floors, aisles, racks, wiring closets, and outdoor areas nested beneath the building. The hierarchy of the locations is:

- **Site**—Your campus; the highest node in your location.
- **Building**—One entry for every building at your site. Buildings are listed in alphabetical order, not by address or the order in which you identified them to the system.

- Floors—One entry for each floor within the building; Floors are nested within the building.
- Aisles—One entry for each aisle in a floor. Aisles are nested within the floor.
- Racks—One entry for each rack in an aisle. Racks are nested within the aisle.
- Outdoor Area—One entry for each named area; Outdoor areas are associated with buildings.
- Devices—Most are assigned to buildings, floors, outdoor areas, or racks. Devices are not assigned at the site level; those devices are considered unassigned.

The hierarchical model enables you to define a location by using either of these methods:

- Using the Location Setup wizard to set up a location in a single process, starting at the site level and progressing to the racks and outdoor areas. The wizard also provides an option to create part of the location, such as defining the site and building, then to use the individual procedures to create floors and wiring closets for the building you created.
- Using separate tasks to create location entities in sequence in a top-to-bottom order. You can create the higher level entities such as a site or building first and save them. Later, you can add floors and wiring closets when information about them becomes available.

RELATED DOCUMENTATION

[Setting Up the Location View | 107](#)

[Creating a Site | 111](#)

[Network Director Documentation home page](#)

Setting Up the Location View

You can build a new location site by the individual nodes, or you can use the Location Setup page. The wizard guides you through the top-down process from the site node down to the assignment of devices.

NOTE: Use the Location Setup page only to design new sites; it is not meant for editing existing sites. If you enter data for an existing site, it is rejected when you attempt to commit the data.

A site is the cornerstone of the location-based view of your network. Until you define a site, the default view of your network tree only shows you a list of your unassigned devices. After you define a site, you can build a tree structure of buildings, floors, wiring closets, aisles, and outdoor areas. As you define your network, you can assign devices to the various components of your network. [Table 27 on page 107](#) describes the devices that you can assign to each of the location component.

Table 27: Devices that can be Assigned to each Location Component

Component	Devices that can be assigned
Site	None
Building	EX Series switches and QFX Series switches
Floor	EX Series switches and QFX Series switches
Closet	EX Series switches and QFX Series switches
Aisle	None
Rack	EX Series switches and QFX Series switches
Outdoor Area	EX Series switches and QFX Series switches

The Location Setup page displays the network tree as you add components to your network. Use the buttons on this page to add various components—such as, buildings, outdoor areas, floors, aisles, racks—to your network. These buttons change depending on the component that you select in the network tree.

After the location is set up, you can view the devices in the network by expanding and collapsing these location nodes in the Location view.

To set up your Location view:

1. Ensure you are in the Build mode and Location or Topology view. Click **Build** in the Network Director banner to enter Build mode; select **Location** view or **Topology** view from the View selector.

2. If you are accessing the Location Setup page from the Location view, select the root node (for example, My Network) in the View pane.
3. Do one of the following depending on the view you are in:
 - From the Tasks pane in the Location view, select **Location Management > Setup Locations**.
 - From the Tasks pane in the Topology view, select **Location > Setup Locations**.

The Location Setup page opens.

4. Click **Add Site** to add a new site.

Network Director adds a new site under the root node and names it as **Site-1**.

5. Select the new site and perform any of the following actions:

- Click **Edit Site** to modify the name of the site and specify the site address. The Edit Site window opens.

Topology view uses this address to place the devices assigned to this site on the topology map. For more details on editing a site, see [“Creating a Site” on page 111](#).

- Click **Add Building** to add a building to your site.

Network Director adds a new building under the site and names it as **Building-1**.

- Click **Outdoor Area** to add an outdoor area to your site. Network Director adds a new outdoor area under the site and names it as **Outdoor Area-1**. You can associate an outdoor area to a site or a building and upload an image or map of that area. After you designate an outdoor area, you can edit or view the map using the Edit Outdoor Area task.

- Click **Delete** to delete the site.

6. If you added a building, select the building and perform any of the following actions to continue building your network:

- Click **Add Floor** to add floors to the building.
- Click **Assign Device** to assign devices to the selected building. The Associate Devices to Building window opens displaying all the unassigned devices in your network. Select the devices that you want to add to the building and click **Add**.

Network Director adds the selected devices to the network tree.

- Click **Edit Building** to edit the name and address of the building. For more details on editing a building, see [“Configuring Buildings” on page 112](#).
- Click **Delete** to delete the building.

7. If you added an outdoor area, select the outdoor area and perform any of the following actions to continue building your network:

- Click **Assign Device** to assign devices to the selected outdoor area. The Associate Devices to Outdoor window opens displaying all the unassigned devices in your network. Select the devices that you want to add to the building and click **Add**.
 - Click **Edit Outdoor Area** to edit the name of the outdoor area and to upload the image of the outdoor area. For more details on editing an outdoor area, see [“Configuring Outdoor Areas” on page 123](#).
 - Click **Delete** to delete the building.
8. If you added a floor to the building, select the floor and perform any of the following actions to continue building your network:

NOTE: You can add aisles and racks to a floor only from the Location view. However, you can view aisles, racks, and devices that you have assigned to these components from the Topology view.

- Click **Add Closet** to add a wiring closet to the floor.
 - Click **Add Aisle** to add an aisle to the floor.
 - Click **Assign Device** to assign devices to the selected floor. The Associate Devices to Floor window opens displaying all the unassigned devices in your network. Select the devices that you want to add to the floor and click **Add**.
 - Click **Edit Floor** to modify the name of the floor, floor level and upload the floor plan. For more details on editing a floor, see [“Configuring Floors” on page 114](#).
 - Click **Delete** to delete the floor.
9. If you added a wiring closet, select the wiring closet and perform any of the following actions:
- Click **Assign Device** to assign devices to the selected closet. The Associate Devices to Closet window opens displaying all the unassigned devices in your network. Select the devices that you want to add to the closet and click **Add**.
 - Click **Edit Closet** to modify the name of the wiring closet. In the Edit Closet window, modify the wiring closet name and click **Done**.
 - Click **Delete** to delete the wiring closet.
10. If you added an aisle, select the aisle and perform any of the following actions:

NOTE: You can add aisles and racks to a floor only from the Location view. However, you can view aisles, racks, and devices that you have assigned to these components from the Topology view.

- Click **Add Rack** to add a rack to the aisle.

- Click **Edit Aisle** to modify the name of the aisle. In the Edit Aisle window, modify the name and click **Done**.
- Click **Delete** to delete the aisle.

11. If you added a rack, select the rack and perform any of the following actions:

NOTE: You can add aisles and racks to a floor only from the Location view. However, you can view aisles, racks, and devices that you have assigned to these components from the Topology view.

- Click **Assign Device** to assign devices to the selected rack. The Associate Devices to Rack window opens displaying all the unassigned devices in your network. Select the devices that you want to add to the rack and click **Add**.
- Click **Edit Closet** to modify the name of the rack. In the Edit Rack window, modify the name and click **Done**.
- Click **Delete** to delete the rack.

12. Click **Done** to save the location details.

Network Director displays the location details along with the assigned devices in Location view.

RELATED DOCUMENTATION

[Understanding the Location View | 105](#)

[Deleting Sites, Buildings, Floors, Wiring Closets, and Devices | 121](#)

[Changing the Location of a Device | 119](#)

[Configuring Buildings | 112](#)

[Configuring Floors | 114](#)

[Setting Up Closets | 116](#)

[Network Director Documentation home page](#)

Creating a Site


IN THIS SECTION

- [How to Add or Edit a Location Site | 111](#)
- [Creating or Editing a Site | 111](#)

A site is the cornerstone of the location-based view of your network. Until you define a site, the default view of your network tree merely shows you a list of your unassigned devices. After you define a location site, you can build a tree structure of buildings, floors, wiring closets, and outdoor areas that can each be assigned devices. You are able to view the devices in the network by expanding and collapsing these location nodes. To setup a location in Network Director, the first step is to create a site.

This topic describes:

How to Add or Edit a Location Site

1. Click the Build Mode icon  in the Network Director banner.
2. Select **Location View** from the list in the View pane.
3. Click **Add Site** to add a new site or click **Edit Site** in the Tasks pane.
4. Fill in or change the fields on the page that opens.
5. Click **Done** to define the site and to save the configuration.

Creating or Editing a Site

Only a few fields are required to establish a site as shown in [Table 28 on page 111](#).

Table 28: Site Creation Fields

Site Name	A descriptive name for the site. This field is mandatory.
City	The city where the site is located.
State	The state where the site is located.

Table 28: Site Creation Fields (continued)

Country	<p>The country where the site is located. Select the country from the list.</p> <p>This field is mandatory. Network Director validates the country code against the country codes in the network’s controllers. If the codes do not match, a warning message is sent.</p>
---------	---

RELATED DOCUMENTATION

Understanding the Location View 105
Configuring Buildings 112
Network Director Documentation home page

Configuring Buildings

IN THIS SECTION

- [How to Add or Edit a Building | 112](#)
- [Adding or Editing a Building for a Location | 113](#)

At any time after you create a site, you can grow your location by adding buildings. You add a building to a site either from within the Location wizard or independently from the Add Building page.

This topic describes:

How to Add or Edit a Building

To add or change a building definition:

1. Ensure you are in the Build mode and Location view. Click **Build** in the Network Director banner to enter Build mode; select **Location View** from the list in the View pane.
2. If you want to add a building to a site:
 - a. Select the site in the Tasks pane , for example, Main Campus.

The Tasks pane refreshes to show your selected site and the tasks available at the site node.

- b. Click **Add Building** in the Tasks pane to open the **Add Building** page.
- 3. If you want to edit an existing building definition:
 - a. Select the building within the site, for example, Headquarters Building.

The Tasks pane refreshes to show your selected building and the available tasks that you can perform at the building node.
 - b. Click **Edit Building** in the Tasks pane to open the **Edit Building** page.
- 4. Fill in the fields and click **Done** to submit the information and to refresh the network tree.

Adding or Editing a Building for a Location

Table 29 on page 113 describes the fields needed to establish a building.

Table 29: Add or Edit Building Fields

Field	Description
Building Name	Type a representative name for the building. The Building Name is a required field.
Address	Type an address. The address can be the street address, building number, or any other identification that helps distinguish it from other buildings.
Done	Click to submit the information. Your view updates to reflect the building change under the site name in the network tree.
Cancel	Click to close the window without changes.

RELATED DOCUMENTATION

Understanding the Location View 105
Configuring Floors 114
Assigning and Unassigning Devices to a Location 117
Network Director Documentation home page

Configuring Floors

IN THIS SECTION

- [How to Add or Edit a Floor | 114](#)
- [Adding or Editing a Building Floor for a Location | 115](#)


You can refine the a building location and designate floors within the building. Use the Add Floor page to:

- Name a floor
- Note the floor level
- Upload a floor plan for viewing
- View an uploaded floor plan

This topic describes:

How to Add or Edit a Floor

Within each building you can define the number of floors and attach the floor plan for online viewing.

1. Click the Build Mode icon  in the Network Director banner.
2. Select **Location View** from the list in the View pane.
3. If you want to add a floor to a building:
 - a. Select the building in the network tree to which you want to add floors, for example, Headquarters.
The Tasks pane refreshes to show your selected building and the available tasks for the building.
 - b. Click **Add Floor** in the Tasks pane to add a new floor to the building.
4. If you want to edit an existing floor definition:
 - a. Select the floor within the building, for example, Lobby-Floor 1.
The Tasks pane refreshes to display the selected building floor and the available tasks that you can perform at the floor node.

- b. Click **Edit Floor** in the Tasks pane to open the Edit Floor page.
5. Fill in the fields for the floor name and level.
6. (Optional) Upload an image of the floor plan.
7. (Optional) View the floor plan, if available.
8. Click **Done** to submit the information and to refresh the network tree.

Adding or Editing a Building Floor for a Location

To add or change information about a building floor, use the fields in [Table 30 on page 115](#).

Table 30: Floor Field Descriptions

Field	Description
Floor Name	Type the name of the floor. This field is required.
Floor Level	Use the arrow keys to set the floor number.
Add/Update	Upload a image of the floor plan.
View	View an existing floor plan.
Done	Saves the floor configuration information, and returns you to Device Inventory page in the default view.
Cancel	Discards any configuration changes.

RELATED DOCUMENTATION

Understanding the Location View 105
Setting Up Closets 116
Configuring Buildings 112
Assigning and Unassigning Devices to a Location 117
Network Director Documentation home page

Setting Up Closets

IN THIS SECTION


- [How to Add or Edit a Closet | 116](#)
- [Adding or Editing a Wiring Closet | 117](#)

Use the Add Closet or Edit Closet tasks to create or change the name of a wiring closet. These tasks are visible only from a floor node in a building.

This topic describes:

How to Add or Edit a Closet

To add or change a wiring closet:

1. Click the Build Mode icon  in the Network Director banner.
2. Select **Location View** from the list in the View pane.
3. Navigate to the building and floor where you are adding or changing the closet.
4. If you are adding a wiring closet:
 - a. Select a building floor in the network tree to which you want to add a wiring closet.
The Tasks pane refreshes to show your selected floor and the available tasks for the floor.
 - b. Click **Add Closet** in the Tasks pane.
5. If you are changing a wiring closet, click **Edit Closet** in the Tasks pane.
6. Type the closet name and click **Done** to save the configuration.
The closet appears with the change in the network tree.

Adding or Editing a Wiring Closet

The Add Wiring Closet or Edit Wiring Closet pages allow you to name a wiring closet. Simply type the name of the new or changed wiring closet and click **Done** to submit the information to the system. Your network tree refreshes to show the wiring closet.

RELATED DOCUMENTATION

[Understanding the Location View | 105](#)

[Configuring Floors | 114](#)

[Assigning and Unassigning Devices to a Location | 117](#)

[Network Director Documentation home page](#)

Assigning and Unassigning Devices to a Location

IN THIS SECTION

- [How to Assign or Unassign Devices | 118](#)
- [Assigning Devices | 118](#)

You can assign devices or remove assignments from devices by their location. Your choices for device assignment are dependent upon the type of device and your position in the site. For details on which devices can be assigned to a location node, see the Devices that can be Assigned to each Location Component table from the [“Setting Up the Location View” on page 107](#).

This topic describes:

How to Assign or Unassign Devices

To assign devices to a specific location:

While in Build mode,

1. Select **Location View** from the list in the View pane.

The network tree displays discovered devices under the physical locations already defined in Network Director. The root node (for example, My Network) is selected by default. The devices that are assigned to the locations are displayed under the nodes for respective locations, such as buildings and floors. All devices that are not assigned to any location are displayed under the Unassigned node.

2. Navigate the network tree to the location where you want to add a device.

Both the Tasks pane and Device Inventory page update to reflect the location's current configuration.

3. Select one of the following tasks in the pane to open Add/Remove Devices for Selected Location.

- Assign Devices to Building
- Assign Device to a Floor
- Assign Devices to a Wiring Closet
- Assign Devices to an Outdoor Location

4. Navigate the tree to find an available device under Unassigned in the left portion of the page.

5. Select the device and click the double right arrows to assign it to the target location on the right. To unassign a device, select the device in the Assigned Devices to Selected Location column and click the double left arrows. Repeat this step until you have finished assigning and unassigning devices.

6. Click **OK** at the bottom of the page to save the assignment. The network tree refreshes to display the device in the new location.

Assigning Devices

Use the Add/Remove Devices for Selected Location to find a device and assign it to a location within a site. Locate the device in the Available Devices column and assign it by clicking the double right arrows. Use the same method to unassign a device by selecting it in the Assigned Devices to Selected Location column and double clicking the double left arrows.

You can assign switches, Virtual Chassis devices and members, Virtual Chassis Fabrics and corresponding member devices to buildings, floors, aisles, and closets.

While assigning Virtual Chassis devices, or Virtual Chassis Fabric devices to a location within a site, you can either assign the logical device—Virtual Chassis—as a single device or one or more member devices that belong to these logical devices, but not both.

NOTE: Network Director displays the Virtual Chassis and Virtual Chassis Fabric in the Location view network tree only if the following conditions are met:

- Virtual Chassis or the Virtual Chassis Fabric is assigned to a location.
- At least one of their member devices are *not* assigned to any location entity.

If all the member devices are assigned to location entities, then the Virtual Chassis and Virtual Chassis Fabric is not displayed in the network tree.

RELATED DOCUMENTATION

[Understanding the Location View | 105](#)

[Configuring Buildings | 112](#)

[Configuring Floors | 114](#)

[Setting Up Closets | 116](#)

[Configuring Outdoor Areas | 123](#)

[Network Director Documentation home page](#)

Changing the Location of a Device

IN THIS SECTION

- [How to Move a Device to a New Location | 120](#)
- [Changing the Location of a Device | 120](#)

The Change Location of Device task is an easy way to move a device address to another building, floor, or wiring closet location within the site. The Change Location of Device task is available whenever you select an assigned device in the Location or Logical views.

This topic describes:

How to Move a Device to a New Location

To move a device address to another location:

- 1. Select a device in the network tree that is currently assigned to a building, floor, or closet.
- 2. Click **Change Location of Device** to open the Change Location of Device page.
- 3. Using the Location View, navigate the tree and select the new location for the device.
- 4. Click **OK** to move the device assignment and to save the new configuration.

Changing the Location of a Device

The Change Location of Device page consists of two components: Selected Device Details and Location View. Use the Selected Device Details portion of the page to review information about the device and its current location. The fields in Selected Device Details page are described in [Table 31 on page 120](#).

Table 31: Contents of Selected Device Details

Field	Description
Device Name	Hostname
Device IP	Device Address
Device Family	Hardware family of products, for example, Junos-QFX.
Location	Gives the current location of the device in the format of site/building/floor/cabinet

Location View is a copy of the network tree for you to navigate and designate the new location for the device.

RELATED DOCUMENTATION

Understanding the Location View 105
Assigning and Unassigning Devices to a Location 117
Network Director Documentation home page

Deleting Sites, Buildings, Floors, Wiring Closets, and Devices

IN THIS SECTION

- [How to Delete a Location Object | 121](#)
- [Deleting Sites | 121](#)
- [Deleting Buildings | 122](#)
- [Deleting Floors | 122](#)
- [Deleting Closets | 122](#)
- [Deleting Devices | 122](#)

From the Build mode Tasks pane, you can delete any sites, buildings, floors, wiring closets and their associated devices. When you delete one of these objects, it removes not only that item but all child objects within the node. All associations related to the node and below are also removed. Devices are moved to the Unassigned node in the network tree. Be sure you understand what is being deleted on the node when you choose to delete a node.

For example, if you delete a building, it deletes the building, all floors, all wiring closets in that building. All of the devices in the building are moved to Unassigned in the network tree. When you delete a building, the site and any other buildings and their associations remain.

How to Delete a Location Object

1. Ensure you are in the Build mode and Location view. Click **Build** in the Network Director banner to enter Build mode; select **Location View** from the list in the View pane.
2. Select any object within the site. The option to delete the object appears in the Tasks pane.
3. Confirm the deletion of the object.

Deleting Sites

There is only one method of deleting a site: select the site in the Tasks pane and click **Delete Site**. Use caution with this selection. When you click **Delete Site** you are given the opportunity to confirm or cancel

the deletion. If you confirm the deletion, you remove the site and everything in the site. All devices become unassigned and are not associated with any buildings, floors, or wiring closets.

Deleting Buildings

When you delete a building, it removes the building, all floors, and all wiring closets within that building. All devices become unassigned and are not associated with the building, its floors, or its wiring closets. Only one building can be deleted at a time. To delete a building, select the building in the network tree and click **Delete Building**. Confirm the deletion to remove the objects and to disassociate the devices. If a site is deleted, all of the buildings within the site are also deleted.

Deleting Floors

When you delete a floor, it removes the selected floor and all wiring closets on that floor. All devices assigned to the floor or to the closets on that floor become unassigned and become available for reassignment. To delete a floor, select the floor in the network tree and click **Delete Floor**. Confirm the deletion to remove the objects and to disassociate the devices. If a site or building is deleted, the floors are also deleted.

Deleting Closets

When you delete a closet, it removes the selected closet and unassigns the devices in the closet. Those devices then become available for reassignment. To delete a closet, select the closet in the network tree and click **Delete Closet**. Confirm the deletion to remove the objects and to disassociate the devices. If a site, building, or floor is deleted, the associated closets are also deleted.

Deleting Devices

At every node in the network tree, you can choose to delete devices directly.

BEST PRACTICE: However, it is usually best to select the node directly above the device so that you do not accidentally unassign more devices than desired.

- Select the node (site, building, floor, or closet) directly above the device.
- Click **Delete Devices** to open the Delete Devices page.
- Click the plus signs to expand the node until you locate the device.

- Click one or more boxes to select the devices. If you do not navigate down to the device level and select a node at a higher level (such a closet or floor), the system selects all devices at and below the node.
- Click **OK** and confirm your selection to remove the assignment. The devices are moved to the Unassigned node of the network tree.

RELATED DOCUMENTATION

[Understanding the Location View | 105](#)

[Network Director Documentation home page](#)

Configuring Outdoor Areas

IN THIS SECTION

• [How to Configure an Outdoor Area | 123](#)

• [Configuring an Outdoor Area | 124](#)

You can associate an outdoor area to a site or a building for wired network and upload an image or map of that area. After you designate an outdoor area, you can edit or view the map using the Edit Outdoor Area task.

This topic describes:

How to Configure an Outdoor Area

To create an outdoor area without using the wizard:

- Ensure you are in Build mode and Location view. Click **Build** in the Network Director banner to enter Build mode; select **Location View** from the list in the View pane.
- Click **Add Outdoor Area** in the Tasks pane. The Add Outdoor Area page opens.
- Fill in the name and upload the optional map.
- Click **Done** to save the data and to return to the default view.

Configuring an Outdoor Area

Table 32 on page 124 describes the fields and buttons necessary to create or change an outdoor area.

Table 32: Outdoor Area Fields

Field	Description
Outdoor Area Name	Type the name of the outdoor area. Network Director associates the outdoor area with the building.
Upload	Optional step to upload an image of the outdoor area. Use the Upload Map window to navigate to the image file location.
Done	Click to save the configuration. The network tree is updated to reflect the change.
Add/Update	Click to add a map or overlay an existing map of the area.

RELATED DOCUMENTATION

Understanding the Location View 105
Setting Up the Location View 107
Assigning and Unassigning Devices to a Location 117
Network Director Documentation home page

Building a Topology View of the Network

IN THIS CHAPTER

- Understanding the Network Topology in Network Director | 126
- Understanding the Topology View Tasks pane | 130
- Setting Up the Topology View | 133
- Managing the Topology View | 134
- Adding and Managing OUI Data in Network Director | 150

Understanding the Network Topology in Network Director

Junos Space Network Director provides features for monitoring and managing Juniper Networks EX Series Ethernet Switches, and QFX Series devices besides enabling connectivity visualization between discovered and managed devices such as routers and switches. Connectivity between devices and their association with their location provide the foundation for rendering topology in a complete manner.

As a network administrator, you must have a clear understanding of the various networking devices in your network, their physical locations, and how these devices are interconnected in your network. The network topology represents the interconnection between various devices in your network, which are managed by Network Director, based on their connectivity and association to their physical surroundings. The network topology provides a visual insight into the network, which is useful for debugging, troubleshooting, planning, and executing administrative actions.

Before you access the topological view of your network, you must:

- Connect your Network Director system to the Internet before accessing Topology View in Network Director as this feature works only while the system is connected to the Internet.

NOTE: Ensure that Internet connection is available for both the Network Director and Network Director client systems.

- Discover the devices managed by Network Director in your network. For details about discovering devices, see [“Discovering Devices in a Physical Network” on page 94](#).

NOTE: You must specify the SNMP parameters during device discovery to have all the devices discovered and managed by Network Director available in Topology View. However, you can specify the SNMP parameters in the **Refresh Topology** task from the Topology View also.

NOTE: Ensure that you have enabled the LLDP, STP, or RSTP protocols on the devices as Network Director uses these protocols to determine the connectivity of devices with their neighbors in the network. LLDP and RSTP protocols are enabled by default on all EX Series switches and QFX Series devices.

- Set up the physical location of the devices in your network based on the geographical location of the devices. Some examples of location nodes are: sites, buildings, floors, closets, and aisles. You can set up one location node within another location node as in buildings within a site or floors within a building. You can then assign the network devices based on their location in buildings, floors, outdoor areas, closets, and racks. Apart from using the Location Management tasks in the Location View, you can set up the location details of the network devices managed by Network Director by using the *Setup Locations* task from the Task menu in the Topology View. For details, see [“Setting Up the Location View” on](#)

[page 107](#). The device-to-topology-group location relationship is established when a device is placed at a specific location on a topology map.

Network topology enables you to view all the discovered devices in your network, overlaid on a map where the devices are located across sites, buildings, floors, outdoor area, closets, and racks along with their physical interconnection with other devices in your network. Topology also provides visualization around physical connectivity between various discovered interconnected devices.

You can use the Topology View to zoom in or zoom out of a site to a building and a building to a site. In the Topology View, you can also double-click a node such site, buildings, and so on to navigate to the next node. You can also see the connectivity between a device and its immediate neighbors, alarms details, and so on. Network Director also enables you to assign devices to buildings, floors, closets, and outdoor areas on the map.

Network topology also provides visualization around physical connectivity between various discovered interconnected devices. You can upload a floor plan at the floor level or a map at the outdoor area if you already have the floor plan or map available. You can then move the nodes in the Topology View page based on the floor plan.

The topology display is created by layering the device images on top of the imported floor plan images as shown in [Figure 10 on page 128](#) and [Figure 11 on page 129](#).

Figure 10: Typical Floor Plan Displaying the Closets and Devices

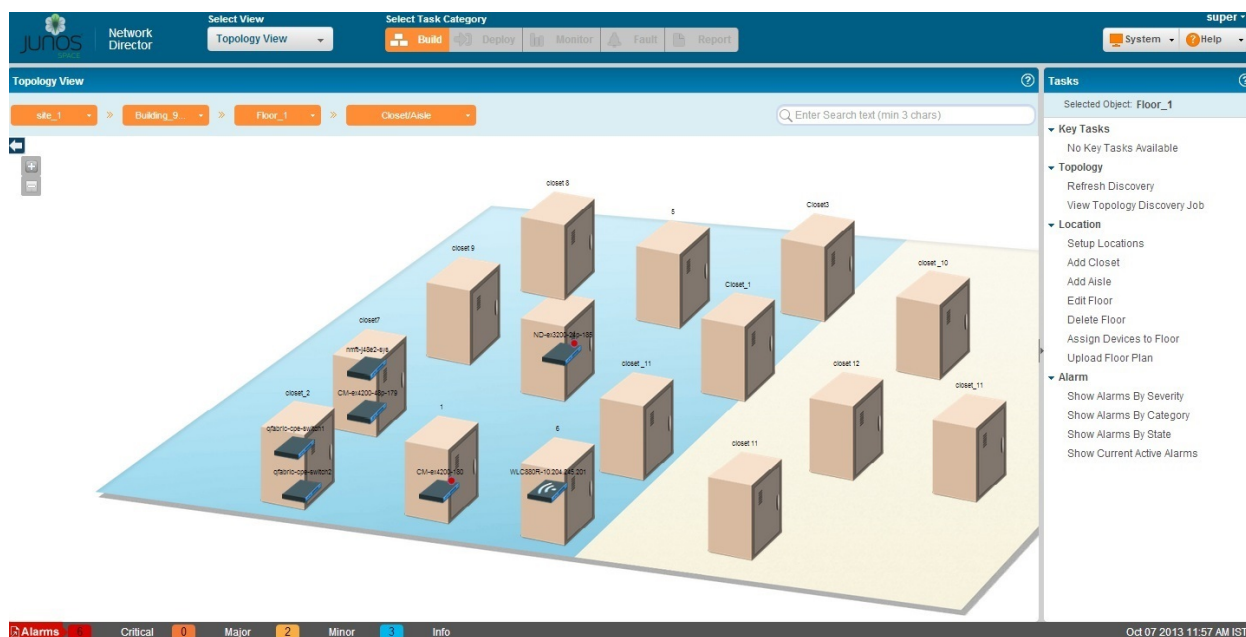



Figure 11: Typical Floor Plan Displaying the Racks and QFX Devices



RELATED DOCUMENTATION

[Discovering Devices in a Physical Network | 94](#)

[Setting Up the Location View | 107](#)

[Managing the Topology View | 134](#)

[Network Director Documentation home page](#)

Understanding the Topology View Tasks pane

The Tasks pane in Topology View contains all the tasks you can do in the Topology View. Click a specific task to begin that task.

Not all tasks are available by default in the Topology View. As you change your selections in the Topology map pane, the tasks in the Tasks pane also change—for example, tasks such as View Virtual Chassis Fabric are visible only after you select a device that is part of a Virtual Chassis Fabric. Similarly, Alarm tasks are available only after you select a site, building, floor, closet, or device.

Topology View tasks are divided into the following categories in the Tasks pane.

- **Key Tasks**—The most preferred tasks that you want to do while you are using the Topology View. You can add important tasks from the topology task menu and hence the Key Tasks are a duplicate of select tasks from the Topology and Location tasks menu.
- **Discovery**—The tasks you do to create the topology of your network.
- **Location**—The tasks you do to create a location such as a site, building, floor, closet, aisle, rack, and outdoor area.
- **Alarm**—The tasks that enable you to monitor the fault alarm details of the devices.

Table 33 on page 130 describes the Topology View tasks.

Table 33: Topology View Tasks

Task	Description
Key Tasks	A duplicate of the most important tasks from other tasks menu. You can add your frequently used tasks to the key tasks menu.
Connectivity	
Displays the device-level connectivity	
View Virtual Network Connectivity	Displays the connectivity of the selected device with the virtual network. NOTE: This task is available only after you select a device that is connected to a virtual network.
View Virtual Machines	Displays virtual machines in the grid format that are connected through the switch.
View Device Connectivity	Displays the connection details of a device with its neighbors in graphical and grid views. If the selected device is connected to more than 60 devices, then the connection details are displayed only in grid view.

Table 33: Topology View Tasks (continued)

Task	Description
View VC/VCF Connectivity	Displays the Virtual Chassis (VC) or Virtual Chassis Fabric (VCF) connectivity from the selected device. NOTE: This task is available only after you select a device that is part of the VC or the VCF.
View Layer 3 Fabric Connectivity	Displays the devices and their physical connectivity in the spine and leaf topology.

Discovery

Displays the device discovery tasks

Refresh Topology	Refreshes the devices discovered and managed by Network Director earlier from Build mode. This task also refreshes the device connectivity.
View Topology Discovery Job	Displays the discovery jobs in the Topology View pane.

Location

Displays the location related tasks

Setup Locations	Creates a new location. This task has sub tasks to add sites, buildings, floors, outdoor area, and closets. NOTE: You cannot create aisles and racks from Topology View. You must create aisles and racks from the Location View work flow.
Add Site	Creates a new site in Location View. NOTE: Use this task only to create the site object. Buildings, floors, closets, and outdoor areas in the site must be created separately.
Add Building	Creates a new building in the selected site. NOTE: Use this task only to create the building. Floors and closets in the building must be created separately.
Add Outdoor Area	Creates a new outdoor area in the selected site.
Upload Floor Plans	Enables you to upload a floor plan for a floor in a building, if you already have a floor plan.
Upload map	Enables you to upload a map to an outdoor area.

Table 33: Topology View Tasks (*continued*)

Task	Description
Delete Site/Edit Site	Deletes or modifies the selected site.
Alarm	
Displays the alarm related tasks	
Show Alarms by Severity	Displays the fault alarm details sorted based on the severity; that is from critical, major, minor, and info.
Show Alarms by Category	Displays the fault alarm details sorted based on the category; that is from active, acknowledged, and cleared.
Show Alarms by State	Displays the fault alarm details sorted based on the state; that is from active, acknowledged, and cleared.
Show Current Active Alarms	Displays any active alarm that has not yet been cleared.
Device Management	
Displays the various device management tasks	
SSH To Device	Launches the SSH connection to the device. You can launch the SSH connection for a device from a location such as a building, floor, rack, or closet. This task is available when you click the device in a particular location. For more details, see “Accessing a Device’s CLI from Network Director” on page 565 .
Launch Web View	<p>Launches the Web user interface connection to the device. You can launch the Web user interface connection for a device from a location such as a building, floor, rack, or closet. This task is available when you click the device in a particular location. For more details, see “Accessing a Device’s Web-Based Interface from Network Director” on page 566.</p> <p>NOTE: This task is applicable to only to those devices that support Web user interface access.</p>
Manage Port Admin State	Enables or disables one or more ports of the selected device. For more details, see “Enabling or Disabling Network Ports on Switches” on page 645 .

RELATED DOCUMENTATION

[Understanding the Network Topology in Network Director](#) | 126

Setting Up the Topology View

Topology View enables you to view all the discovered devices in your network, overlaid on a map. You can create sites, buildings, floors, outdoor areas, closets, and racks by using the Location wizard in the Topology View. You can use the Topology View to zoom in or zoom out of a site or a building. You can also see the connectivity between a device and its immediate neighbors, alarms details, port details, and so on.

Before you start, ensure that:

- The devices are discovered using the Device Discovery task. For detailed steps, see [“Discovering Devices in a Physical Network” on page 94](#).
- You have specified SNMP parameters during device discovery.
- Ensure that you have enabled the LLDP, STP, or RSTP protocols on the devices as Network Director uses these protocols to determine the connectivity of devices with their neighbors in the network. LLDP and RSTP protocols are enabled by default on all EX Series and QFX Series devices.

Perform the following steps to set up your network in the Topology View:

1. After you have discovered the devices using the Device Discovery task, open Location View or the Topology View. You can set up sites, buildings, floors, outdoor areas, closets, aisles, and racks using Location View. The Topology View allows you to create most of these components except racks and aisles. For detailed steps on building your location view, see [“Setting Up the Location View” on page 107](#).
2. Assign devices to the various entities in the Location View. Network Director enables you to assign devices to the various components of your network. [Table 27 on page 107](#) describes the devices that you can assign to each of the location component.
3. Save the location settings.
4. Select **Topology** from the Network View Selector.

Network Director lays out the sites that you created in Location View, in the topology map. Network Director uses the location that you specified for each site and building to place them on the map. A site that does not have the location specified, will be placed in the default location—United States.

5. Click **Discovery > Refresh Discovery** from the Tasks menu to refresh the topology of devices. Network Director refreshes the Topology View for all the devices that you have added to Topology View. After

a successful discovery, you can select and view the connectivity of a particular device with their immediate neighbors using the tasks in the Connectivity section of the Tasks pane.

6. Follow the steps outlined in [“Managing the Topology View” on page 134](#) to perform various tasks from Topology View.

RELATED DOCUMENTATION

[Understanding the Network Topology in Network Director | 126](#)

[Managing the Topology View | 134](#)

[Network Director Documentation home page](#)

Managing the Topology View

IN THIS SECTION

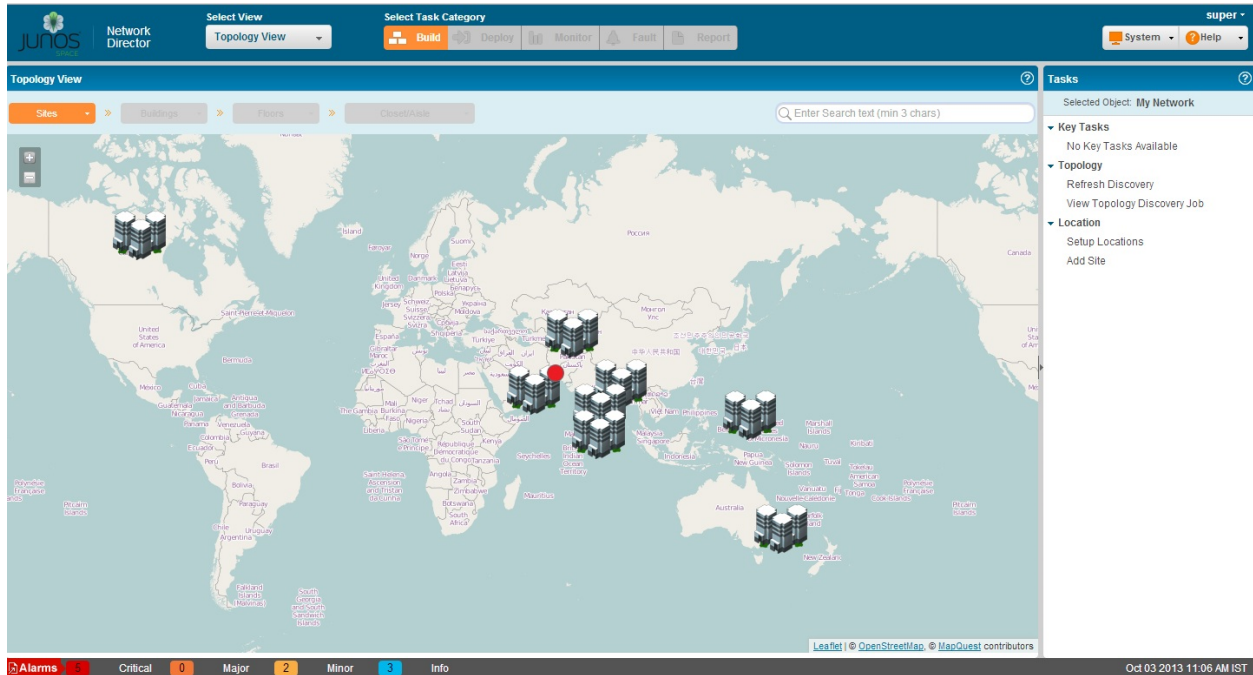
- [Viewing the Network Topology | 134](#)
- [Refreshing the Topology | 137](#)
- [Viewing Topology | 138](#)
- [Viewing Topology Discovery Job | 139](#)
- [Setting Up Locations | 139](#)
- [Viewing the Alarm Details | 140](#)
- [Discovering the Linux Hosts | 140](#)
- [Displaying Device Connectivity | 140](#)
- [Displaying Virtual Chassis and Virtual Chassis Fabric Connectivity | 145](#)
- [Uploading Floor Plans | 148](#)
- [Uploading Topology Map | 149](#)

Viewing the Network Topology

Topology View enables you to view all the discovered devices in your network, overlaid on a map. You can create sites, buildings, floors, outdoor areas, closets, and racks by using the Location wizard in Topology View. You can use the Topology View to zoom in or zoom out of a site or a building. You can also see the

connectivity between a device and its immediate neighbors, alarms details, port details, and so on. An example of how the topology map looks like after you have added the location details is shown in [Figure 12 on page 135](#).

Figure 12: Main Topology Window



The topology display is created by layering the device images on top of the imported floor plan images.

In addition to the tasks outlined in the [Table 33 on page 130](#), you can perform the following tasks from the Topology View pane:

- **Zoom In and Zoom Out**—You can use the zoom in () or zoom out () buttons to get a detailed or high-level view. Network Director enables you to zoom in and view details up to the rack level, if you have defined racks and assigned devices to the rack.
- **Pan**—Network Director enables you to pan the topolog. You can move the devices to different parts on the topology by holding the mouse button down and dragging to a specific part of the map.





Network Director displays the sites, buildings, and geographical coordinates on the topology map based on the address specified while setting up the locations. However, you can move the devices around to another location by selecting, dragging, and dropping the device to the correct location. For example, you can move a site from a US site to a Bangalore site on the topology by using the Pan feature. This helps to determine the correct location of a site because while setting up the location details, it is not mandatory to provide the address information.


- **Search and Locate**—You can search for all nodes such as sites, building, floors, specific devices, all devices in a floor, and so on by entering a search keyword or the complete name of the device in the Search field. On entering the search criteria, you might see the list of nodes or just one node based on the search


criteria. When you select the node, Network Director locates the node and Topology View is panned to ensure the selected node is centered on the page. In cases where a device is at the edge of the map the corner is aligned accordingly to bring device in view.

- **View details**—If you mouse over a entity (site, building, device), the entity gets highlighted and Network Director displays details such as the building name or IP address and the number of active alarms on that entity. A colored dot that appears on the upper right corner of the entity identifies whether there are alarms on the entity and the color of the dot indicates the severity level of the alarm. See [Table 34 on page 136](#) to know more about the alarm severity indicator for each alarm severity.

Table 34: Alarm Severity Indicator

Alarm severity	Indicator
Critical	
Major	
Minor	
Informational	

NOTE: Each of these entities might have alarms of different severities. Network Director displays the indicator based on the most severe alarm on an entity. For example, if a device has 3 informational alarms and one major alarm, Network Director displays the indicator for the major alarm ().

- **Highlight and select View Device Connectivity**—Select a device and click **View Device Connectivity**, **View VC/VCF Connectivity**, or **View Virtual Network Connectivity** to view the selected device's connectivity with other devices. Each of these tasks, except the **View Device Connectivity**, are visible only when you select a device that is part of a Virtual Chassis, Virtual Chassis Fabric, or a Virtual network. The device images are displayed along with details such as name, IP address, and the connectivity link between the devices.
- **Navigation**—Use the navigation breadcrumbs at the top of the page to navigate through sites, buildings, floors, closets, or aisles. For there are more than one entities at any give level, you can use the Down arrow in the breadcrumb to navigate to that entity. For example, if you want to navigate from floor-1 in building-1 to floor-3 in building-2, you can use the down arrow in the building breadcrumb to select building-2 and the down arrow in the floor breadcrumb to select floor-3.
- **Host Information**—You can use the  button to expand the members to view the host details. You can also view the virtual machines if the host is a hypervisor and managed by Network Director.

Refreshing the Topology

You can refresh the devices discovered and managed by Network Director from the Tasks pane in the Topology View and the Datacenter View. You can add the Refresh Topology task to the Key tasks in both the views if you will use this task frequently.

To refresh the device discovery process:

1. Click **Refresh Topology** from Key Tasks or the Topology View or from the Datacenter Discovery panes.

The Refresh Discovery window is displayed along with the SNMP details that you specified in the discovery options in the Device Discovery task page.

If you have not specified the SNMP details while discovering the devices, proceed to Step 3.

2. Select the **SNMP version** from the SNMP version table and click **Discover**. The Refresh Topology window appears displaying the progress of the topology discovery.
3. If SNMP version details are not displayed in the SNMP version table, click **Add** on the Refresh Discovery window.

The Add SNMP Settings dialog box is displayed.

Select either **SNMP V1/V2C** or **SNMP V3**. Based on the selection, you need to enter the details as follows:

- If you selected SNMP V1/V2C, specify a community string, which can be *public*, *private*, or a predefined string.

Click **Add** in the Add SNMP Settings dialog box or click **Add More** to add more strings to the community. If you click **Add More**, when you are done adding all the strings, click **Add** to save the SNMP settings for V1/V2C.

- If you selected SNMP V3, specify the following:
 - Enter a username.
 - Select the privacy type (AES 128, DES, or None).
 - Enter the privacy password (if AES 128 or DES). If you specify none for the privacy type, the privacy function is disabled.
 - Select the authentication type (MD5, SHA, or none).
 - Enter the authentication password (if MD5 or SHA). If you specify none for the authentication type, the authentication function is disabled.
 - Click **Add** to save the SNMP settings and close the dialog box, or click **Add More** to add additional configurations. If you clicked Add More, click **Add** to save the settings and close the dialog box.

The specified details are displayed in the SNMP version table.

4. Click **Discover**.

The Refresh Topology window is displayed, showing details of the device discovery.

Viewing Topology

The Refresh Topology window displays the refresh device discovery job details as described in [Table 35 on page 138](#).

Table 35: Refresh Topology Job Details

Field	Description
Job Name	The Refresh topology job name along with the Job ID.
Start Time	The time at which the refresh discovery job is initiated.
End Time	The time at which the refresh discovery job is completed.
Percentage Progress	Displays the progress of the job in percentage. When the job is completed, displays 100 percentage.
Status	The status of the job. The status is <i>In Progress</i> until the job is completed.
Target Devices Count	The total number of targets for the discovery of devices.
Discovered Devices Count	<p>The total number of discovered devices for which the SNMP parameters are specified.</p> <p>NOTE: In the Topology View, you can view the network connections of only those discovered devices for which LLDP, SNMP, and STP parameters are set.</p>
Discovered Subnets Count	The total number of subnets discovered.
Targets	The Targets table displays the Management IP addresses of the devices discovered along with the status of the discovery process.

Use the right and left arrows to navigate through the discovered pages. You can specify the details to be displayed in a page by selecting the **show items** list box and specifying the number of items to be displayed in one page.

Click **Close** to close the Refresh Topology page and return to the main Topology page.

Viewing Topology Discovery Job

To view the discovery jobs in the Topology View and the Datacenter View, click **View Topology Discovery Job**. The Topology Discovery Jobs window opens displaying details of the topology related jobs as described in [Table 36 on page 139](#). To view any hidden column, mouse over column heading, select the down arrow, and then click **Columns**. Select the check box to display the hidden columns.

Table 36: Job Details for Topology Discovery

Field	Description
Job ID	For each job-based task, the audit log includes a job ID.
Job Name	The name of the job.
Percent	The percentage of completion of the job.
State	<p>The status of the job:</p> <ul style="list-style-type: none"> • Success—Job completed successfully • Failure—Job failed and was terminated • Job Scheduled—Job is scheduled but has not yet started • In progress—Job is has started, but not completed • Cancelled—Job is cancelled
Summary	Summary of the job scheduled and executed with status.
Scheduled Start Time	The UTC time on the client computer when the job is scheduled to start.
Actual Start Time	The actual time when the job started.
End Time	The time when the job was completed.
User	The login ID of the user that initiated the task.
Recurrence	The recurrent time when the job will be restarted.

To view the details of a topology discovery job, select a row and click **Show Details**. To cancel a scheduled job, select a job that is scheduled for a later time or a job that is in progress and click **Cancel**.

Setting Up Locations

You can set up locations and assign devices to these locations by setting up the Location View. To set up Location View, see [“Setting Up the Location View” on page 107](#).

Viewing the Alarm Details

You can view the alarm details at site, building, floor, closet, aisle and rack levels. Based on the location node, the alarms are aggregated and displayed. That is, all the alarms for a particular building is aggregated and displayed at the building level. The alarms display as red, orange, yellow, and blue dots indicating critical, major, minor, and info alarms. Network Director updates and displays the alarm status changes in real time in the Topology view. To view the Alarm details for a device, navigate to the device level, select a device, and click one of the following options from Alarm in the Tasks pane:

- Show Alarms By Severity, see [“Alarms by Severity Monitor” on page 778](#)
- Show Alarms By Category, see [“Alarms by Category Monitor” on page 778](#)
- Show Alarms By State, see [“Alarms by State Monitor” on page 779](#)
- Show Current Active Alarms, see [“Current Active Alarms Monitor” on page 776](#)

Discovering the Linux Hosts

Based on the LLDP discovery method, you can discover the hosts for various Linux platforms such as Ubuntu, CentOS, and Red Hat by clicking **Refresh Topology**. You can also view the additional information in the tool tip that appears when you mouse over a host. These hosts also show icons that identify if a server is based on its respective platform. For example, if a server is a generic Linux server, a Linux icon is shown in the topology.

NOTE: You must enable LLDP both on the switch and on the host.

Displaying Device Connectivity

At the device level, you can view the connectivity details of a device and the details of all the devices that are connected to the specified device by using Topology View in Network Director. The Device Connectivity View also displays various details about the selected device and the immediate neighbors. The level of detail that Network Director displays in the Device Connectivity View differs based on the type of device that you select.

To view the connectivity details of a device:

1. Do one of the following:
 - While in the Logical, Location, Device, or Custom View, select the device for which you want to view connectivity from the View pane and click **Connectivity > View Device Connectivity** from the Tasks pane.

- In the Topology View, navigate to a device in a building, floor, outdoor area, closet or rack and select the device for which you want to view the connectivity details and click **Connectivity > View Device Connectivity** from the Topology Tasks pane.


NOTE: The Connectivity task container is available only after you select a device.

The Device Connectivity page opens. You can view the device connectivity details either in graphical view or in grid view. The default view is the graphical view.

In the graphical view, the device is displayed in the center and its network connectivity to all the connected devices are displayed as in [Figure 13 on page 141](#). Mouse over a device to view details of the highlighted device.

NOTE: If the selected device is connected to a device that is not a Juniper Networks device, the latter appears dimmed in the Device Connectivity page indicating that the device is not managed by Network Director.

Figure 13: Displaying the Connection Details in Graphical View



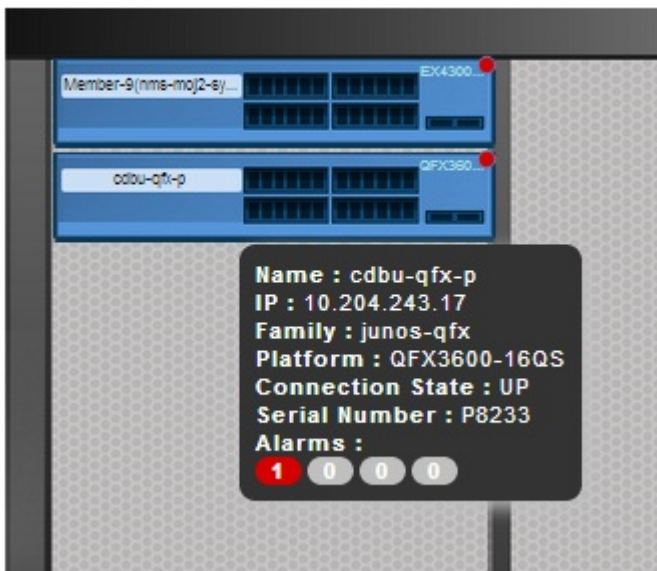
If the selected device is connected to more than sixty devices, then all the connected devices are highlighted in a circular form or a grid form. If the selected device is connected to less than 60 devices, then the links between the interconnected devices are displayed.

The device images are displayed along with details such as name, IP address, and alarm state information in colored labels that provide health and reachability information. You can also view the details of the hosts or virtual machines that are connected to the switches.

You can view the following details in the Device Connectivity - Graphical view:

- Name—The name of the device provided while configuring the device. The device name is displayed as a label.
- IP—The IP address of the device.
- Family—The family or platform to which the device belongs to. Family can be an JUNOS-EX, JUNOS-QFX, JUNOS-QF, and MSSOS.
- Serial Number—Serial number of the device.
- Alarms—The alarm details displaying the number of critical, major, minor alarms, or info for the device. Alarms details are color coded to indicate their severity level as shown in [Figure 14 on page 142](#). Network Director updates and displays the alarm status changes in real time in the Topology view.


Figure 14: Alarm Indicator



- Connection State—Connection status of the device. Connection state can be UP, DOWN or N/A. Network Director updates and displays the connection state changes in real time in the Topology view.

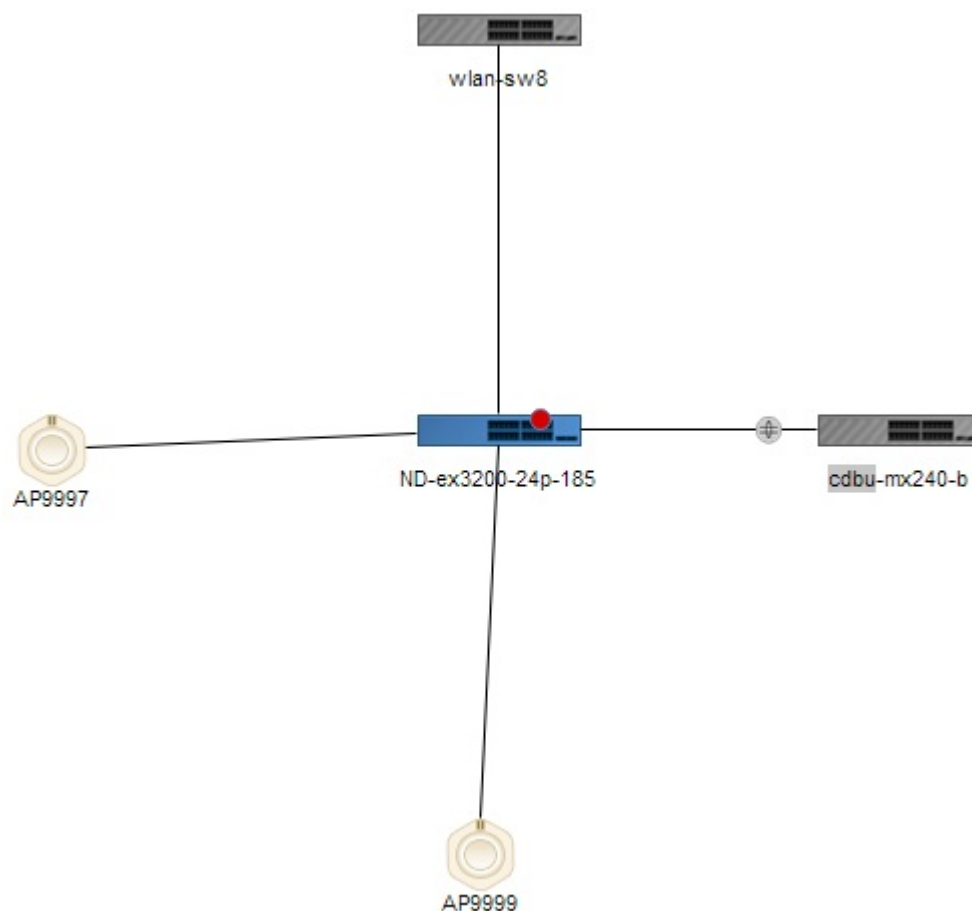
- Link status—Indicates whether the link between two devices is UP or DOWN as shown in [Figure 15 on page 143](#). Network Director updates and displays the link status changes in real time in the Topology view.

Figure 15: Link Status Indicator



- LAG—Identifies connections that are configured as LAGs as shown in [Figure 16 on page 144](#).

Figure 16: LAGs in the Device Connectivity view



You can view the following details of the virtual machine that are connected to hosts:

- Virtual Machine—Name of the virtual machine.
- Host Name—Name of the host to which the virtual machine is connected to.
- VNetwork—Name of the virtual network.
- OS—Name of the operating system on which the virtual machine is running.
- Connection State—Connection status of the virtual machine. Connection state can be UP, DOWN or N/A.
- Power State—State of the power supply: Powered On or Powered Off.

2. Click **Show Grid View** to view the device connectivity details in a tabular format as displayed in [Figure 17 on page 145](#).

Figure 17: Displaying the Connection Details in Grid View

Device Connectivity : nd-72q1-elit						
External Links Fabric Links						
Show Graph View						
Source Device	Source Port	Source Port Bandw...	Destination Device	Destination Port	Destination Port Bandw...	Link Status
nd-72q1-elit	[LAG] ae0	NA	nd-36q1-elit	[LAG] ae0	NA	UP
nd-72q1-elit	[LAG] ae1	NA	nd-36q1-elit	[LAG] ae1	NA	DOWN
nd-72q1-elit	et-0/0/0 (ae0)	0	nd-36q1-elit	et-0/0/0 (ae0)	0	UP
nd-72q1-elit	et-0/0/1 (ae1)	0	nd-36q1-elit	et-0/0/1 (ae1)	0	DOWN
nd-72q1-elit	et-0/0/2 (ae2)	0	nd-opus-48s4	et-0/0/48	0	UP

The following details are displayed in the table:

- Source Node—Name of the device specified while configuring the device.
- Source Port—Source port of the device.
- Source Port Bandwidth %—Realtime percentage of bandwidth utilized at the source port.
- Destination Node—Name of the device or devices the device is connected to.
- Destination port—The port number on the destination device to which the source device is connected to.
- Destination Port Bandwidth %—Realtime percentage of bandwidth utilized at the destination port.
- Link Status—Indicates if the link to the device is UP or DOWN.

You can sort the details in the table in the ascending order or descending order for each column. You can also use filters to display only the desired device connectivity details.

Displaying Virtual Chassis and Virtual Chassis Fabric Connectivity

You can view the connectivity between the components of Virtual Chassis and Virtual Chassis Fabric using the View VC/VCF Connectivity task. You can access this tasks from all views except Dashboard View and Datacenter View.

To view the connectivity details for a Virtual Chassis or a Virtual Chassis Fabric:

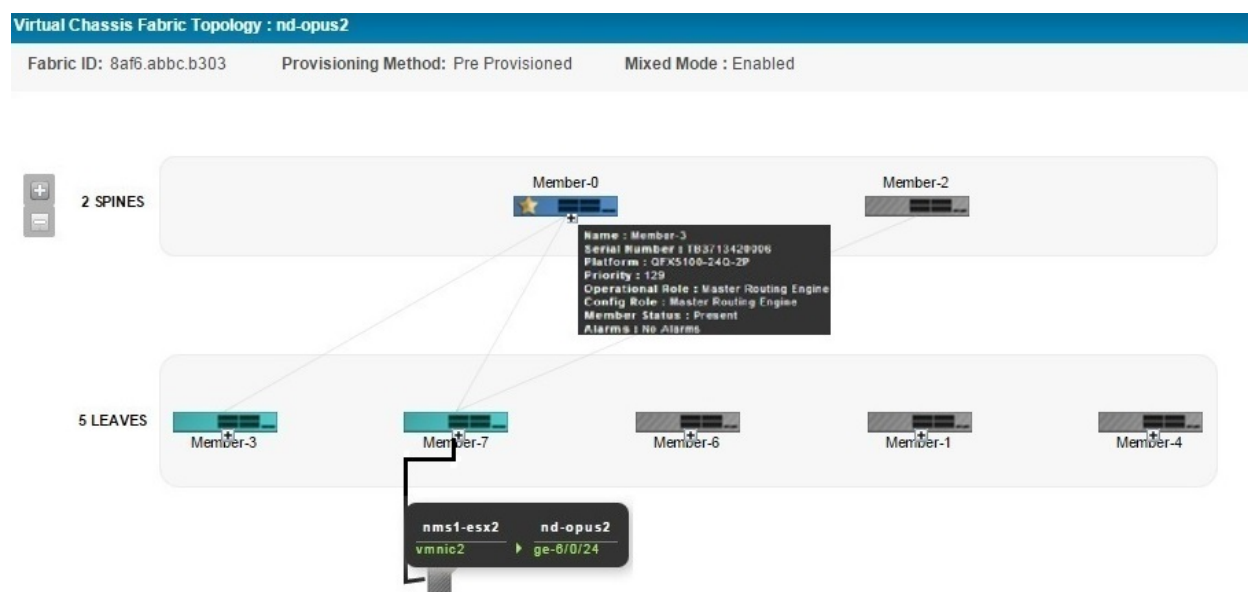
1. Do one of the following:

- While in the Logical, Location, Device, or Custom View, select the device for which you want to view the Virtual Chassis or Virtual Chassis Fabric connectivity from the View pane and click **Connectivity** > **View VC/VCF Connectivity** from the Tasks pane.
- In the Topology View, navigate to a device in a building, floor, outdoor area, closet or rack and select the device for which you want to view the Virtual Chassis or Virtual Chassis Fabric connectivity details and click **Connectivity** > **View VC/VCF Connectivity** from the Topology Tasks pane.

NOTE: The Connectivity task container is available only after you select a device.

- Click **Connectivity > View VC/VCF Connectivity** from the Tasks pane. Network Director displays the connectivity between the members of the selected Virtual Chassis as shown in [Figure 18 on page 146](#). The inactive members and any members having alarms in down state are shown as grey icons in the topology identifying the state of the member. Mouse over a member to view details of that member.

Figure 18: Displaying the Connectivity for a Virtual Chassis Fabric



In the [Figure 18 on page 146](#), the connection details are represented by green, yellow, and red lines.

- The green line indicates that the node is connected to all the Interconnects properly and all functions are normal.
- The yellow line indicates that the node is only partially connected to the Interconnect. That is, the node may be connected to one Interconnect, but not all the Interconnects.
- The red line indicates that the node is not connected to any of the interconnects.

The following details are displayed in the top panel of the Virtual Chassis Topology View as shown in [Table 37 on page 146](#).

Table 37: Common Details for the Virtual Chassis and Virtual Chassis Fabric

Details	Description
VC ID	All members of a Virtual Chassis configuration share one Virtual Chassis identifier (VCID). This identifier is derived from internal parameters.

Table 37: Common Details for the Virtual Chassis and Virtual Chassis Fabric (*continued*)

Details	Description
Provisioning Method	<p>The provisioning mode of the member. Provision mode can be <i>autoprovisioned</i>, <i>preprovisioned</i> or <i>not preprovisioned</i>.</p> <p>Autoprovisioning a Virtual Chassis Fabric (VCF) enables you to “plug and play” devices into your VCF after minimal initial configuration.</p> <p>In a VCF, you can have two to four members configured in the Routing Engine role. Of this, one member acts as the Master Routing Engine and another member acts as the backup Routing Engine. In a preprovisioned configuration, the selection of which member functions as the Master Routing Engine and which as the backup Routing Engine is determined by the software based on the primary-role election algorithm.</p> <p>In a configuration that is not preprovisioned, the selection of the primary and backup is determined by the primary-role priority value and secondary factors in the primary-role election algorithm.</p>
VC Mode	Indicates whether the Virtual Chassis is mixed or not.

The following details are displayed in the Virtual Chassis Topology view for each member depending on the role of the member as shown in [Table 38 on page 147](#).

Table 38: Details of VC/VCF Members

Details	Description	Role
Name	Name of the member switch provided while configuring the device. The device name is displayed as a label.	Primary Backup Line Card
Serial Number	Serial number of the member switch.	Primary Backup Line Card
Platform	Platform of the device. Platform can be QFX5100, QFX5110, or QFX10002.	Primary Backup Line Card
Priority	The primary-role priority value. This is the most important factor in determining the role of the member switch within the Virtual Chassis configuration.	Primary Backup Line Card

Table 38: Details of VC/VCF Members (*continued*)

Details	Description	Role
Operational Role	Operational role of the device. A device might be configured for a particular role, but can operate in the same or a different role. For example, a spine device configured with a Routing Engine role might operate as a line card. Therefore, the operational role of this device is Line Card. Operational role can be Routing Engine or Line Card.	Primary Backup Line Card
Config Role	The configured role of the device. This can be Routing Engine or Line card.	Primary Backup Line Card
Member Status	Displays the status of each member device: <ul style="list-style-type: none"> • Present—The device is connected and working fine. • Not Present—The device is not connected to the VC or VCF. • Inactive—The device is connected, but not running. • Non Provisioned—A configuration in which the roles of the members are assigned automatically; not configured statically (preprovisioned). • Pre Provisioned—A configuration that allows you to deterministically control the member ID and role assigned to a member by associating the member with its serial number. 	Primary Backup Line Card

In addition to the above details, when you expand the host details, you can also view the details of the member and the link connected to the virtual machine.

Uploading Floor Plans

You can upload the floor plan from the Topology View if you already have a floor plan for a specified building in a site.

To upload the floor plan:

1. Select a **Site > Building > Floor**. Alternatively, create a site, building, and floor. Click the **Upload Floor Plan** task from the Location task in the Tasks pane.

The Upload Floor Plan dialog box is displayed.

2. Click **Browse** next to Image File to choose a floor plan file.
3. Navigate to the folder where you have saved the floor plan on your system and click **Open**.

4. Click **Upload** to upload the floor plan image.
5. Click **Cancel** if you do not want to upload the floor plan and quit the Upload Floor Plan dialog box.

Uploading Topology Map

You can upload a topology map for an outdoor area from the Topology View.

To upload the map:

1. Select a **Site > Outdoor area**. Alternatively, create an outdoor area within a site. Click the **Upload map** task from the Location task in the Tasks pane.
The Upload Map dialog box is displayed.
2. Click **Browse** next to Image File to choose a map file.
3. Navigate to the folder where you have saved the topology map for an outdoor area on your system and click **Open**.
4. Click **Upload** to upload the topology map image.
5. Click **Cancel** if you do not want to upload the map and quit the Upload Map dialog box.

RELATED DOCUMENTATION

[Discovering Devices in a Physical Network | 94](#)

[Setting Up the Location View | 107](#)

[Understanding the Network Topology in Network Director | 126](#)

[Network Director Documentation home page](#)

Adding and Managing OUI Data in Network Director

Network Director uses Link Layer Discovery Protocol (LLDP) to detect the type of network device (such as desktop computer, VoIP phone, or network servers) in a campus network. However, network printers are an exception as most printers do not use LLDP. As a result, Network Director might not be able to identify the device as a printer. This is where the organizationally unique identifier (OUI) of printers come into play. OUI is formed using the first three octets of the device MAC address. OUI is unique to a vendor or manufacturer and can be identified globally.

You can build a database of OUIs that Network Director can use to identify the device type. During the topology discovery, if the LLDP-based discovery does not identify the device type, Network Director looks up in the OUI database to see whether there is a match. If Network Director finds a matching entry, the device is notated as a printer in the network topology. If there is no match, the device is marked as an unknown device.

To add and manage OUI data:

1. While in the Build mode with Logical View selected, click **Connectivity > Manage OUI** from the Tasks pane.

The Manage OUI page opens.

NOTE: Network Director displays a list of well-known printer manufacturers and their OUI details in the Manage OUI page. Make sure that the details that you want to add are not already listed before you proceed with adding OUI data for your device.

2. Click **Add** to add new OUI details to Network Director.

The Add MAC OUI window opens.

3. Enter the MAC OUI for the device that you want to add. The first three octets of a MAC address of the device forms the OUI. OUI is unique to a vendor or manufacturer and can be identified globally.
4. Enter the name of the vendor and select the type of device.
5. Click **Save** to save the OUI details and return to the Manage OUI page.
6. To delete one or more OUI details, select the rows that you want to delete and click **Delete**.

RELATED DOCUMENTATION

Creating Custom Device Groups

IN THIS CHAPTER

- Understanding Custom Device Groups | 152
- Creating Custom Device Groups | 155

Understanding Custom Device Groups

IN THIS SECTION

- Where Is the Custom Group Function Located in Network Director? | 153
- How Do Custom Group Rules Work? | 153
- What Happens When I Edit a Custom Group Rule? | 154
- When Are Rules Executed? | 154

Custom group is way of grouping your devices based on your business needs. You can create custom groups and add devices to each custom group. You can manually add devices to a custom group or you can define rules to add devices, that match the rule condition, to the custom group once they are discovered by Network Director. You can view the custom groups and devices that are assigned to each group in the Custom Group view.

A custom group can include devices such as switches. Creating custom device groups enables the configuration of multiple devices simultaneously—you can create multiple custom groups and directly associate devices at any level. Up to this point, Custom Groups are the same as selecting related items in the location view tree. What makes Custom Groups unique is that you can also configure a custom group to automatically add devices after discovery. You indicate the criteria for additional devices by editing rules. Custom groups can then be created in a hierarchy up to eight levels deep. Each layer can contain up to 32 peer containers under a single parent container.

Where Is the Custom Group Function Located in Network Director?

Network Director has different views that you select to see different aspects of your data. You select one of these views at a time from the Select View option in the Network Director banner. The options are Logical View, Location View, Device View, Custom Group View, Datacenter View, and Topology View. To create a Custom Group, Network Director must be in Custom Group View. Custom Groups are created at the top level of the network—My Network.

Once Custom Groups are created, they appear in all views as options for profile assignment—assigning a profile to a Custom Group assigns that profile to all members of the group.

How Do Custom Group Rules Work?

Adding rules to a Custom group consists of creating a three part rule statement, with a rule basis, an operator, and matching criteria. Possible combinations are shown in [Table 39 on page 153](#).

Table 39: Three Options of a Rule Statement

Rule Basis	Operator	Matching Criteria
Device Role	Equals	Access
		Aggregation
		Core
		Unassigned (available only for EX Series switches for which logical category can be defined)

Table 39: Three Options of a Rule Statement (*continued*)

Rule Basis	Operator	Matching Criteria
Device Type	Equals or Not Equals	Switch
		Virtual Chassis
Serial Number	Equals or Contains	<i>You provide serial numbers or letters</i>
SKU or Model	Equals or Contains	<i>You provide model numbers or letters</i>
Management IP Address	Equals or Regex	<i>You provide IP address</i>
Location	Select a previously configured location: NOTE: For location directions, see “Setting Up the Location View” on page 107 . 1. Click Please select . 2. From the Select Location window, select a location. 3. Click OK .	
Device Role	Equals	<i>You provide preconfigured device role</i>
Device Type	Equals or Not Equals	Switch
		Virtual Chassis
Firmware Version	Equals or Contains	<i>You provide a full or partial firmware version for devices</i>

What Happens When I Edit a Custom Group Rule?

When you edit a rule, devices that were added to the group but no longer qualify because of the rule edit are not automatically removed from the group. You must remove those devices manually. If more devices are now qualified to join the group because of your rule edit, the devices are added to the group on the next device notification change to the network.

When Are Rules Executed?

The option **Associate devices based on the rules for the custom groups while saving group information** is enabled by default. If the option is disabled, the rule engine will be activated only when there is some

change in the device property. When a device property change occurs, rules are processed and devices are added to the group, if the group has a rule for those actions.

RELATED DOCUMENTATION

[Creating Custom Device Groups | 155](#)

[Network Director Documentation home page](#)

Creating Custom Device Groups

IN THIS SECTION

● [Creating Custom Groups | 155](#)

● [Creating a Custom Group | 156](#)


From Network Director, you can create a custom group, then add devices such as switches to the group. Creating custom device groups enables the configuration of multiple devices simultaneously—you can also create multiple custom groups and directly associate devices at any level. Up to this point, Custom Groups behave the same way as selecting related items in the location view tree. What makes Custom Groups unique is that you can also configure a custom group to automatically add devices after discovery. You indicate the criteria for additional devices with rules. Custom groups can then be created in a hierarchy up to eight levels deep. Each layer can contain up to 32 peer containers under a single parent container.

NOTE: A device can be part of a group at only one level in a hierarchy.

This topic describes:

Creating Custom Groups

To create custom groups:

1. In the top banner, under **Views**, select **Custom Group View**.
2. Click  in the Network Director banner.

3. Click **Set Up Custom Group** under Key Tasks in the Tasks pane.

The Set Up Custom Group page opens, displaying a list of currently configured Custom Groups.

4. Configure the custom group, following the directions [“Creating a Custom Group” on page 156](#).

5. Click **Done**.

The new custom group appears in the Groups List.

Creating a Custom Group

Use the Set Up Custom Group page to define a group of devices that you can configure simultaneously.

To add a new custom group:

1. Type a Custom Group Name for the new group and then click **Add**.

The Custom Group tree is displayed with your new group added.

2. Click **Done** now to create the group with no child groups, devices, or rules. The *Message Data Saved Successfully* is displayed. Click **OK**.

For additional configuration, select your new group.

The options **Add Child Group**, **Assign Devices**, and **Add/Edit Rule** appear.

3. To add a child group under the new custom group:

- a. Be sure the correct custom group is selected—this group will become the parent group.

- b. Click **Add Child Group**.

The Add Child Group window opens, displaying a default child group name such as Group-0.

- c. Replace the default child group name.

- d. Click **Add**.

The new child group appears in the Custom Group list tree under the parent group.

TIP: Custom groups can be created in a hierarchy up to eight levels deep. Each layer can contain up to 32 peer containers under a single parent container.

4. To assign devices to a custom group:

- a. Select a custom group, either a parent or child group, and then click **Assign Devices**.

The Assign Devices To Custom Group window opens, displaying a list of discovered network devices, their IP addresses, and their platforms. Platforms include junos-ex, junos-qfx, junos-qf, and mssos. These are devices that can be added to the group.

- b. Select one or more devices by adding a check mark and then click **Add**.

The devices are listed under the appropriate group in the Custom Groups List.

NOTE: A device can be part of a group at only one level in a hierarchy.

5. To add a rule that will automatically add devices to a parent or child custom group:

- a. Select a custom group, either a parent or child group, that will have devices added to it automatically when a specific rule has been met.

- b. Click **Add/Edit Rule(s)**.

The Add/Edit Rules window opens.

- c. Click **Add Rule**.

A rule statement is displayed with three columns—two columns display the words *Please select...*. The third column is blank.

- d. From the first *Please select...* option in the rule statement, select the basis for the rule. You are indicating that automatic additions to the list will be based on either **Device Type**, **Firmware Version**, **Device Role**, **Serial Number**, **SKU/Model**, **Management IP**, or **Location**.
- e. From the second *Please select...* option in the rule statement, select an available operator, either **Equals**, **Not Equals**, **Like**, **Regex**, or **Contains**—the operators presented depend on the basis you selected in the first column. For example, if the basis for the rule is **SKU/Model**, then the only operator options are **Equals** and **Not Equals**. If the basis for the rule is **Location**, then your only option is to click **Select** for a list of locations.

TIP: The **Equals** operation matches all characters of the matching criteria. The **Like** operation matches the first few characters of the matching criteria.

- f. For the third option in the rule statement, provide a matching criteria. Matching criteria are indicated in the third column of the list shown in [Table 40 on page 158](#).

TIP: Some rules have no third option.

Table 40: Three Options of a Rule Statement

Rule Basis	Operator	Matching Criteria
Device Role	Equals	Access
		Aggregation
		Core
		Unassigned (available only for EX Series switches for which logical category can be defined)
Device Type	Equals or Not Equals	
		Switch
		Virtual Chassis
Serial Number	Equals or Contains	<i>You provide serial numbers or letters</i>
SKU or Model	Equals or Contains	<i>You provide model numbers or letters</i>
Management IP	Equals or Regex	<i>You provide IP address or regular expression</i>
	TIP: Regex, a regular expression, consists of a sequence of characters that forms a search pattern.	TIP: For example, <code>(?<=\.){2}(?=[A-Z])</code> is a regular expression.

Table 40: Three Options of a Rule Statement (*continued*)

Rule Basis	Operator	Matching Criteria
Location	Select a previously configured location: NOTE: For directions to configure locations, see “Setting Up the Location View” on page 107 . i. Click <i>Please Select</i> . ii. From the Select Location window, select a location. iii. Click OK .	
Firmware Version	Equals or Contains	<i>You provide a full or partial firmware version for devices.</i>

- g. Click **OK**.

Rules are executed when new devices are discovered. Devices that match the defined rules are added to the group dynamically once discovery is complete.

TIP: If you add more than one rule to a Custom Group, then all rules must be met for a device to join the group.

- The option **Associate devices based on the rules for the custom groups while saving group information** is enabled by default. When a device property change occurs, rules are processed and devices are added to the group, if the group has a rule for those actions. If you disable the option, the rule engine will be activated only when there is some change in the device property.
- Click **Done**.
A status window opens with either the message *Data saved successfully* or with an error message. Click **OK**.
- To edit a rule, select the appropriate custom group and then click **Add/Edit Rule**. When you edit a rule, devices in the group that no longer qualify because of the rule change are not automatically removed from the group. You must remove those devices manually. If more devices are now qualified to join the group because of your rule edit, the devices are added to the group on the next device notification change to the network.

TIP: To delete a device from the group, select the device and then click **Delete**. To delete an entire Custom Group, select the group and then click **Delete**. You are asked to confirm the deletion—click **OK**.

RELATED DOCUMENTATION

[Creating and Managing Port Groups | 323](#)

[Assigning a VLAN Profile to Devices or Ports | 345](#)

[Assigning Device Common Settings to Devices | 198](#)

[Assigning and Unassigning Devices to a Location | 117](#)

[Understanding Custom Device Groups | 152](#)

[Network Director Documentation home page](#)

Configuring Quick Templates

IN THIS CHAPTER

- Understanding Quick Templates | 161
- Configuring and Managing Quick Templates | 163

Understanding Quick Templates

Quick templates is a way to create a base build for devices. This feature enables you to use a CLI-based text editor to define your network configuration in the form of a template that you can apply to multiple devices in your network in addition to the profile assignment feature. Because quick templates are driven by Device Management Interface (DMI) schema, you can use them to set all the configuration parameters for any supported device.

By using these quick templates, you can configure, for example, routing protocols such as BGP, OSPF, ISIS, or even static routes by specifying the device configuration. You can append or add the system commands or the user-defined commands in the form of the variables in the CLI-based text editor. The user-defined commands support variables in the format `$(variable_name)`, which must be populated with data when you apply a template to a device.

The variable name defined for each CLI must be unique. Otherwise, you cannot assign different values to those variables even though they are used in different CLIs. For example, if a variable say `$(description)` is used in two CLIs `set vlans $(name) description $(description)` and `set snmp description $(description)`, you will not be able to define different values to the descriptions. To define different values, you must change the variable name for one of the commands.

The [Table 41 on page 161](#) shows data types supported for the values entered for variables.

Table 41: Variable Data Types

Data Type	Description
Container	Holds other data types.
String	Contains character strings.

Table 41: Variable Data Types (continued)

Data Type	Description
Integer [Number]	Specifies a numeric value without a fractional component.
Boolean	Has two possible values: true and false. True if checked and False if unchecked.
Enumeration	Defines a variable to be a set of predefined constants. The variable is equal to one of the values that have been predefined for it.
Choice	Provides a radio button. Check the radio button to use the configuration option in the template.
String - Key [column in a table]	Identifies the uniqueness of the record in the table. If the table has a key specified , only one record with the given key could exist.

The Save option in the Create Quick Templates page enables you to save and also validate a template. If there are any conflicts in the configuration, you must resolve the conflicting variables in the configuration elements manually, before you deploy the configuration to the devices. Upon successful validation (and after you apply a template to a device), you can deploy the configurations (specified in the templates) to the devices. You can choose to deploy the configuration immediately, or at a later time. Depending upon the approval mode selected for your deployment, you can either deploy the changes directly or you can get an approval from the approver before deploying the changes. For more information about types of approval modes supported for deployments in Network Director, see [“Setting Up User and System Preferences” on page 26](#).

Benefits of Quick Templates

- Configuring a large number of devices can be tedious and time-consuming. Quick templates can apply necessary configurations on multiple devices at the same time, helping you save time and effort.
- Modifying configurations across multiple devices may lead to configuration errors. Deploying configuration using quick templates simplifies device configurations and reduces configuration errors.

RELATED DOCUMENTATION

[Setting Up User and System Preferences | 26](#)

[Deploying Configuration to Devices | 580](#)

Configuring and Managing Quick Templates

IN THIS SECTION

- [Creating a Quick Template | 164](#)
- [Applying Templates to Devices | 165](#)
- [Editing a Quick Template | 166](#)
- [Deleting a Quick Template | 166](#)
- [Cloning a Quick Template | 166](#)
- [Using the Quick Template Details Window | 167](#)
- [Viewing Deployed Quick Templates | 167](#)

You can create and manage custom templates for your device configurations that are deployable through Network Director. Unlike other features that support implementation of only some of the device configurations, quick templates enables you to set up all the configuration parameters for any supported device because it is Device Management Interface (DMI) schema-driven.

Each device type is described by a unique data model that contains all the configuration data for that device. The Schema window shows the device family that you select while you create a template and the DMI schema that lists all the possible fields and attributes for a type of device. The latest schema describe the new features associated with recent device releases. After you create a quick template, you can add or delete device configuration details to and from quick templates by loading the configuration data from the schema. You need to apply these templates to devices manually.

If you click the **More tips** link you are guided on the variable and the command syntax usages. It also provides instructions on how to issue sub-commands. When defining your network configuration in quick templates by using a particular command, ensure that you define the sub-commands individually. Stating sub-commands as a single command causes errors. For example, the commands **set snmp location sunnyvale** and **set snmp contact admin@example.com** are valid when defined individually. However, if you combine these commands into the single command **set snmp location sunnyvale contact admin@example.com** schema validation treats the end command **contact** as an extra entry and throws an error.

To avoid any conflicts with the profile configurations while creating the template, a warning message **Please don't create any Profile conflict configuration** is displayed to indicate that you must not create a configuration as part of the template if the same configuration is available as part of the profile configuration.

The Templates page in the Quick Templates workspace lists the device templates created, in a tabular view. The [Table 42 on page 164](#) lists the columns in the table along with a description:

Table 42: Quick Templates

Column	Description
Creation Time	Date and time when the template was created.
Template Name	Name of the quick template.
Device Family	Name of the device family for which the template is created. Selecting the option Common indicates that the template is applicable for all the device families.
OS Version	Junos OS version of the device family selected.
Description	Description of the quick template.
Last Updated Time	Date and time when the template was last modified.
Last Updated By	User name of the person who created the template.

This topic describes:

Creating a Quick Template

Quick templates enable you create a template to define configurations for your devices. You can create and deploy quick templates from the Wired workspace.

To create a quick template:

1. Click the Build Mode icon in the Network Director banner.
2. Select **Wired > Tasks > Manage Quick Templates** in the Tasks pane.

The Manage Quick Template page appears.

3. Click **Add**.

The Create Quick Template page opens.

4. Specify the following details:

- **Name**—Type a name for the quick template. The quick template name is required. The quick template name must be unique and limited to 63 characters.
- **Description**—Type a description for the quick template. The description is optional and limited to 255 characters.

- **Device Family**—From the Device Family list, select an appropriate device family. Selecting the option **Common** in device family creates a generic template, which can be applied to any device family. Therefore, specify only the most common settings such as system, SNMP, or track group settings that are applicable to all the platforms. If you want to define the settings that are specific to a platform select the appropriate platform from the device family instead of the Common option. For the list of device families supported by Network Director, see the latest [Network Director Release Notes](#).
- **OS Version**—From the OS Version list, select an appropriate DMI Schema version running on that platform. If you are unable to locate the DMI schema for a device family, you can update the DMI schema version on the Junos Space server. For more information about updating the DMI schema on the Junos Space server, see Junos Space documentation.

The Schema window displays the device family and the OS version selected in this step.

5. Type or paste the Junos commands in the form of variables in the CLI-based text editor provided in the CLI Commands section. For information on the type of supported variables, see [“Understanding Quick Templates” on page 161](#). Alternatively, you can navigate through the configuration option levels (at the left side) in Schema and double-click the configuration option you want to add to the quick template. The selected configuration option is displayed in the CLI Commands CLI-based text editor. The configuration options available here depend on the device family you selected.
6. Optionally, you can modify the configuration in the CLI Commands text area by using the tool bar functionalities such as undo, redo, cut, copy, paste, and find.
7. Click **Save**.

The template you created is displayed in the quick templates table.

Applying Templates to Devices

After you create a template, you can define your device configuration to be managed by using the quick templates, and apply these templates to the multiple devices.

To assign a template to a device:

1. Select the check box against the quick template for which you want to assign the profile.
2. Click **Assign**.

The Assign Quick Template : template names page opens.

3. Choose at least one device to which the profile needs to be assigned.
4. Click **Next**.

5. Choose a device and specify the quick template variables in Configure attributes page and click **Save**.
For example, when you configure a VLAN interface in a quick template, you can specify the variables VLAN and interface names for that template for a selected device.
6. Optionally, you can apply the settings specified here to all the selected devices of a device family by selecting the check box against the option **Apply above settings to all other selected devices**.
7. Click **Next** and then click **Finish**.
8. Review the profile association with the quick template and then click **Finish**.

Editing a Quick Template

You can edit a quick template to modify configurations for your devices.

To edit a quick template:

1. Select the check box against the quick template that you want to modify.
2. Click **Edit**.

The Edit Quick Template : template name page opens.

3. Make the required changes to the quick template and click **Save**.

Deleting a Quick Template

To delete a quick template:

1. Select the check box against the quick template that you want to delete.
2. Click **Delete**.

The Delete Quick Templates window opens.

3. Click **Yes** to delete the quick template; else click **No**.

Cloning a Quick Template

A cloned quick template is a copy of an existing quick template. You can use the quick template as a primary copy to create clone of that template. When you clone a quick template, you create a copy of the entire device configuration, including its settings, and other contents. Cloning a quick template saves time if you

are deploying device configuration that are similar to the primary copy, rather than creating a template and defining configurations multiple times.

To create a copy of an existing template:

1. Select the check box against the quick template you want to clone.
2. Click **Clone**.

The cloned template named primary template-clone is shown in the list of templates.

Using the Quick Template Details Window

Use the Quick Template Details window to view the details of the quick template. [Table 43 on page 167](#) describes the fields in this window.

Table 43: Quick Template Details

Field	Description
Name	Displays the name of the quick template.
Description	Provides a description of the quick template.
Device Family	Displays the device family for which quick template is created.
OS Version	Displays the Junos OS version for the selected device family.
CLI Commands	Displays the CLI commands configured for the device family.

Viewing Deployed Quick Templates

You deploy the device configurations defined in a quick template after you have applied the template to a device. The View Deployed Templates option enables an administrator or an operator to view the list of templates that are deployed to the devices.

You can mouse over the template name to view the date and time when the template was created and last modified.

The View Deployed Templates page lists the deployed templates device in a tabular view. The [Table 44 on page 168](#) lists the columns in the table along with a description.

Table 44: View Deployed Template

Column	Description
Template Name	Indicates the name of the template whose configuration is deployed to the system.
Creation Time	Indicates the date and time when the template was created.
Last Updated Time	Indicates the date and time when the template was last modified.
User Name	Indicates the user name of the person who created the template.

Depending upon the type of approval mode configured—Manual Approval or Auto Approval mode— you can either deploy the device configurations defined in the template directly or by pursuing an approval from a configuration approver for the device changes.

To view the list of quick templates that are deployed to a device:

1. Click the Build Mode icon in the Network Director banner.

2. Select a device in the View pane.

The View Deployed Templates option appears under Wired>Tasks.

3. Click **View Deployed Templates**.

The Deployed Templates For Device: device name page displays listing the templates applied for that device.

RELATED DOCUMENTATION

[Understanding Quick Templates | 161](#)

[Deploying Configuration to Devices | 580](#)

[Network Director Documentation home page](#)

Configuring Device Settings

IN THIS CHAPTER

- [Understanding Device Common Settings Profiles | 169](#)
- [Creating and Managing Device Common Settings | 170](#)
- [Assigning Device Common Settings to Devices | 198](#)

Understanding Device Common Settings Profiles

Network Director enables you to configure device-level settings for switches in the Device Common Settings profile. Once you create the profiles, you can assign the profiles to a switch and you can deploy the profiles using the Deploy mode tasks.

Network Director also creates Device Common Settings profiles when it discovers devices. It creates a Device Common Settings profile for each device it discovers, importing the device-level settings from the device into the profile.

While configuring the profiles, you can specify the basic settings, which includes the profile name, device user list, and time settings. Apart from the basic settings, you can optionally specify the management and protocol settings too.

RELATED DOCUMENTATION

[Creating and Managing Device Common Settings | 170](#)

[Assigning Device Common Settings to Devices | 198](#)

[Understanding Network Configuration Profiles | 88](#)

[Network Director Documentation home page](#)

Creating and Managing Device Common Settings

IN THIS SECTION

- [Managing Device Common Settings | 170](#)
- [Creating a Device Common Settings Profile | 172](#)
- [Specifying Basic Settings for Device Common Settings | 174](#)
- [Specifying Management Settings for EX Switching Device Common Settings | 177](#)
- [Specifying Management Settings for Campus Switching ELS Device Common Settings | 181](#)
- [Specifying Management Settings for Data Center ELS Device Common Settings | 185](#)
- [Specifying Protocol Settings for EX Switching Device Common Settings | 188](#)
- [Specifying DHCP Relay/DNS Settings for Campus Switching ELS Device Common Settings | 192](#)
- [Specifying Protocol Settings for Data Center Switching ELS Device Common Settings | 194](#)
- [Reviewing and Saving a Device Common Settings Configuration | 197](#)
- [What to Do Next | 197](#)

Use the Manage Device Common Settings page to create new device common settings for switching devices and to manage the existing device common settings.

This topic describes:

Managing Device Common Settings

From the Manage Device Common Settings page, you can:

- Create a new Device Common Settings profile by clicking **Add**. For directions, see [“Creating a Device Common Settings Profile” on page 172](#).
- Modify an existing Device Common Settings profile by selecting it and clicking **Edit**.
- Assign a Device Common Settings profile to a device by selecting a profile and clicking **Assign**. For directions, see [“Assigning Device Common Settings to Devices” on page 198](#).
- Modify an existing assignment of a Device Common Settings profile by selecting the profile and clicking **Edit Assignment**.
- View information about a Device Common Settings profile by either double-clicking the profile name or by selecting the profile and clicking **Details**.
- Delete a Device Common Settings profile by selecting a profile and clicking **Delete**.

TIP: You cannot delete common settings profiles that are in use—that is, assigned to devices or used by other profiles. To see the current assignments for a profile, select the profile and click **Details**.

- Clone a Device Common Settings profile by selecting a profile and clicking **Clone**.

Table 45 on page 171 describes the device information available on the Manage Device Common Settings page. This page lists all Device profiles defined for your network, regardless of your current selected scope in the network view.

Table 45: Manage Device Common Settings Settings

Field Name	Action
Profile Name	Name given to the profile when the profile was created.
Family Type	The device family; EX Series switch, Campus Switching ELS.
Description	Description of the Device profile entered when the profile was created.
Assignment State	Displays the assignment state of the profile. A profile can be: <ul style="list-style-type: none"> • Unassigned—When the profile is not assigned to any device • Deployed—When the profile is assigned to a device and is deployed from Deploy mode • Pending Deployment—When the profile is assigned to a device, but not yet deployed in the network. For deployment directions, see “Deploying Configuration to Devices” on page 580.
Assigned to	Displays the number of devices to which the profile assignment is done.
Creation Time	Date and time when the profile was created.
Last Updated Time	Date and time when the profile was last modified.
User Name	The username of the person who created or modified the profile.

TIP: All columns might not be displayed. To show or hide fields listed in the Manage Authorization Profiles table, click the down arrow on the field header, select **Columns**, and select or clear the check box adjacent to the field that you want to show or hide.

Creating a Device Common Settings Profile

In Network Director, as an administrator, you can configure Device Common Settings profiles by using the Create Device Profile page for switches. You can view the summary of the configurations before saving the Device profile.

At minimum, you must specify the Device profile and profile name in the workflow. You can include additional configuration such as:

- Device users
- Management services
- Multicast, spanning-tree protocol (STP)
- Domain Name Server
- DHCP servers, DHCP Relay servers, Login Banner, and Global PoE settings for switches


You can create profiles on the basis of the device family and each Device profile is specific to a device family. After you create a Device profile, you assign the profiles to different devices.

NOTE: You can assign only one profile to a device. However, you can assign the same profile to multiple devices.

To create a Device profile:

1. Under Views, select one of these options: **Logical View**, **Location View**, **Device View** or **Custom Group View**.

TIP: Do not select **Dashboard View**, **Datacenter View**, or **Topology View**.

2. Click  in the Network Director banner.
3. From the Tasks pane, select the type of network (Wired), the appropriate functional area (Wired), and select the name of the profile that you want to create.

4. Click **Add** to add a new profile.

If you chose to create a profile for the wired network, Network Director opens the Device Family Chooser window.

- a. From the Device Family Chooser, select the device family for which you want to create a profile. The available device families are **Switching (EX)**, **Campus Switching ELS** (Enhanced Layer 2 Software) and **Data Center Switching ELS**.

- b. Click **OK**.

The Create Device Common Settings wizard for the selected device family is displayed. It consists of four sections, Basic Settings, Management Settings, Protocol Settings, and Review.

5. Specify the basic settings. Complete the Basic Setting wizard page as described in both the online help and in [“Specifying Basic Settings for Device Common Settings” on page 174](#).
6. When you have completed the basic settings, either click **Next** or click **Management Settings** at the top of the wizard window.
7. Complete the Management Settings described in both the online help and in the sections [“Specifying Management Settings for EX Switching Device Common Settings” on page 177](#) and [“Specifying Management Settings for Campus Switching ELS Device Common Settings” on page 181](#).
8. When you have completed the management settings, click **Next**.
9. Complete the protocol settings as described both online help and in the sections [“Specifying Protocol Settings for EX Switching Device Common Settings” on page 188](#) and [“Specifying DHCP Relay/DNS Settings for Campus Switching ELS Device Common Settings” on page 192](#).
10. When you have completed the protocol settings, either click **Next** or click **Review** at the top of the wizard window.
11. You can either save your profile or make changes to your profile from the Review page. For more information, see [“Reviewing and Saving a Device Common Settings Configuration” on page 197](#).
12. Click **Finish** to save the Device profile configuration.

The system saves the Device profile and displays the Manage Device Common Settings page. Your new or modified Device profile is listed in the table.

Specifying Basic Settings for Device Common Settings

To configure the basic settings for any Device Common Settings profile, enter the settings described in [Table 46 on page 174](#). Mandatory settings are indicated by a red asterisk (*) that appears next to the field label in the user interface.

Table 46: Device Profile Basic Settings

Field	Action
Profile Name	Type a name for the profile. Profile name must not contain special characters or spaces. Note that profiles that are automatically created by Network Director as part of device discovery or out-of-band changes might contain the underscore (_) character.
Description	Type a description of the profile containing up to 256 characters.
Login Banner for EX Series switches and Campus Switching ELS,	Enter the banner text—this text is displayed in the banner when you log in to the device.
Device Users	

Table 46: Device Profile Basic Settings (*continued*)

Field	Action
Task: Add a Device User	<p>To add a device user:</p> <ol style="list-style-type: none"> Click Add under Device Users. The Add User window opens. Provide a username and password. Confirm the password. Enter a combination of 6 through 128 alphanumeric characters and special characters. The password is case sensitive and must be a combination of at least two different types of characters or a combination of upper case and lower case letters. TIP: Do not create a user with the name <i>root</i>. Select a role for the user: <ul style="list-style-type: none"> For switches, the role options are: Operator, Read-only, Super-user, or Unauthorized. Operators have clear, network, reset, trace, and view privileges. Super-Users have all privileges. Click OK. The user is added to the list of Device Users. <p>TIP: To edit an entry, select a row from the Device Users table and click Edit to modify the information. To delete an entry select a row from the Device Users table and click Delete to delete the user.</p>
Time Settings	
Time settings apply to all platforms.	
Time Zone	Select a country and time zone from the list.

Table 46: Device Profile Basic Settings (*continued*)

Field	Action
Add a Time Server	<p>To add a time server:</p> <ol style="list-style-type: none"> 1. Click Add under Time Server. The Add Time Server window opens. 2. Provide an IP address and, optionally for switches only, mark the corresponding time server as Preferred. 3. Click OK. The server is added to the list of Time Servers. <p>TIP: To edit the settings of a time server, select it and then click Edit.</p>

To configure management settings, click **Next** or click **Management Settings** at the top of the wizard window. To skip the management settings and protocol settings, click **Review** at the top of the wizard window.

Management Settings are described in both the online help and in the sections [“Specifying Management Settings for EX Switching Device Common Settings” on page 177](#) and [“Specifying Management Settings for Campus Switching ELS Device Common Settings” on page 181](#).

Specifying Management Settings for EX Switching Device Common Settings

To configure the management settings for an EX switching Device profile:

1. Enter the settings described in [Table 47 on page 178](#). All settings are optional. Default values are applied to the configuration if you skip the management settings configuration.

Table 47: Device Profile Management Settings for EX Switching

Task	Action
Enable Services	<p>You can enable one or more network protocol services for this Device profile: FTP, TELNET, HTTPS, or HTTP.</p> <p>NOTE: HTTP and HTTPS are not available for EX9200 Series switches.</p>
Configure PoE	<p>To add Power over Ethernet (PoE) configuration for EX switching, enable Configure PoE and provide these settings:</p> <p>NOTE: PoE configuration will be added only to switches that support PoE.</p> <ol style="list-style-type: none"> Using the arrows, adjust the Guard Band value from 0 through 19 watts. A guard band reserves a specified amount of power from the PoE power budget for the switch or line card in case of a spike in PoE consumption. For switches with multiple PoE line cards, such as the EX2300 switch, the guard band wattage is set to the specified value on all line cards, unless a line card has been explicitly configured with a different value. <p>TIP: The valid guard band range (in watts) for EX2300 and EX3400 switches is 0 through 15. Any value outside this range causes the profile deployment to fail.</p> <ol style="list-style-type: none"> Select a Management Mode for PoE, either Class or Static: <ul style="list-style-type: none"> • Class Management—In class PoE management mode, the maximum power for an interface is determined by the class of the connected powered device. • Static Management—In the static PoE management mode, you specify the maximum power for each PoE interface. The PoE controller then allocates this amount of power to the interface from its total budget. For PoE Global, you can indicate Enable All, Disable All, or None. <p>NOTE: If you deselect Configure PoE, PoE is disabled and the global PoE settings supported by this profile (poe guard-band, poe fpc all guard-band, poe management, poe fpc all management, and poe interface all) are deleted from the switch when the profile is deployed on the switch.</p>

Syslog Settings

Optionally, expand the Syslog Settings and provide the following system logging settings.

Table 47: Device Profile Management Settings for EX Switching (continued)

Task	Action
Enable Device Logging for Switches	<p>To enable device logging for switches:</p> <ol style="list-style-type: none"> Under Enable Device Log, click Add. The Add Log window opens. Select the log type for switching, either Console, File, User, or Host. <ul style="list-style-type: none"> Console logging sends system log messages to the console. File logging sends system log messages to the file you specify in File Name. User logging sends system log messages to the terminal session of the user specified in User Name. You will also need to provide the name of the user. Host logging sends system log messages to the server specified in Host. Host can be either an IP address or host name. Under Services, click Add. The phrase <i>Click to enter value</i> appears in both the Service column and Severity Filter column. Click the phrase <i>Click to enter value</i> in the Service column. A list box replaces the phrase in the Service column. From the Service list, select a logging service: Any, Authorization, Change-log, Conflict-log, Daemon, DFC, External, Firewall, FTP, Interactive-commands, Kernel, NTP, PFE, Security or User. Click the phrase <i>Click to enter value</i> in the Severity Filter column. A list box replaces the phrase in the Severity Filter column. Select an available severity filter from the list, either Alert, Any, Critical, Emergency, Error, Info, None, Notice, or Warning. The filter is added to the list of Severity Filters. The filter is activated when the corresponding service is triggered. Click OK. The log is added to the Enable Device Log list.
Edit Logging Settings	Select a Log Type from the Enable Device Log list and click Edit to change the configuration.

Table 47: Device Profile Management Settings for EX Switching *(continued)*

Task	Action
Delete Logging Settings	Select a Log Type from the Enable Device Log list and click Delete to remove the server configuration.

To configure protocol settings, either click **Next** or click **Protocol Settings**. To use the default protocol settings, skip to final review by clicking **Review** at the top of the wizard window.

Protocol Settings options are described in the section [“Specifying Protocol Settings for EX Switching Device Common Settings” on page 188](#),

Specifying Management Settings for Campus Switching ELS Device Common Settings

To configure the management settings for an ELS campus switching device common setting profile:

1. Enter the settings described in [Table 48 on page 182](#). All settings are optional—default values are applied to the configuration if you skip the management settings.

Table 48: Management Settings for ELS Switching Device Profile

Task	Action
Enable Services	You can enable one or more network protocol services for this Device profile: FTP , Telnet , HTTPS , or HTTP . By default, none are selected.
Configure PoE	<p>To add Power over Ethernet (PoE) configuration for ELS Switching, enable Configure PoE and provide these settings:</p> <p>NOTE: PoE configuration will be added only to switches that support PoE.</p> <ol style="list-style-type: none"> Using the arrows, adjust the Guard Band value from 0 through 19 watts. A guard band reserves a specified amount of power from the PoE power budget for the switch or line card in case of a spike in PoE consumption. For switches with multiple PoE line cards, such as the EX2300 switch, the guard band wattage is set to the specified value on all line cards, unless a line card has been explicitly configured with a different value. <p>TIP: The valid guard band rang (in watts) for EX2300 and EX3400 switches is 0 through 15. Any value outside this range causes the profile deployment to fail.</p> <ol style="list-style-type: none"> Select a Management Mode for PoE, either Class or Static: <ul style="list-style-type: none"> • Class Management—In class PoE management mode, the maximum power for an interface is determined by the class of the connected powered device. • Static Management—In the static PoE management mode, you specify the maximum power for each PoE interface. The PoE controller then allocates this amount of power to the interface from its total budget. For PoE Global, you can indicate Enable All, Disable All, or None. <p>NOTE: If you deselect Configure PoE, PoE is disabled and the global PoE settings supported by this profile (poe guard-band, poe fpc all guard-band, poe management, poe fpc all management, and poe interface all) are deleted from the switch when the profile is deployed on the switch.</p>

Syslog Settings

Optionally, expand the system logging section and configure device logging.

Table 48: Management Settings for ELS Switching Device Profile (continued)

Task	Action
Enable Device Logging for ELS Switches	<p>To enable device logging for ELS switches:</p> <ol style="list-style-type: none"> Under Enable Device Log, click Add. The Add Log window opens. Select the log type for ELS switching, either Console, File, User, or Host (default). <ul style="list-style-type: none"> Console logging sends system log messages to the console. File logging sends system log messages to the file you specify for File Name. User logging sends system log messages to the terminal session of the user you specify for User Name. You will also need to provide the name of the user. Host logging sends system log messages to the server you specify for Host. Host can be either an IP address or host name. Under Services, click Add. The phrase <i>Click to enter value</i> appears in both the Service column and Severity Filter column. Click on the phrase <i>Click to enter value</i> in the Service column. A list box replaces the phrase in the Service column. From the Service list, select a logging service: Any, Authorization, Change-log, Conflict-log, Daemon, DFC, External, Firewall, FTP, Interactive-commands, Kernel, NTP, PFE Security or User. Click on the phrase <i>Click to enter value</i> in the Severity Filter column. A list box replaces the phrase in the Severity Filter column. Select a Severity Filter from the list, either Alert, Any, Critical, Emergency, Error, Info, Notice, or Warning. Click OK. The filter is added to the list of Enabled Device Logs with entries in the Log Type column and filter name column. The filter will be activated when the corresponding log type is triggered.
Task: Edit Logging Settings	Select an entry from the Enable Device Log table and click Edit to change the settings.

Table 48: Management Settings for ELS Switching Device Profile *(continued)*

Task	Action
Task: Delete Logging Settings	Select an entry from the Enable Device Log table and click Delete to remove the server settings.

To configure DHCP Relay and DNS, either click **Next** or click **DHCP Relay/DNS Settings**. To skip the protocol settings, click **Review** at the top of the wizard window.

DHCP Relay and DNS options are described in the section [“Specifying DHCP Relay/DNS Settings for Campus Switching ELS Device Common Settings”](#) on page 192.

Specifying Management Settings for Data Center ELS Device Common Settings

To configure the management settings for a Data Center ELS Device profile:

1. Enter the settings described in [Table 49 on page 186](#). All settings are optional—default values are applied to the configuration if you skip the management settings.

Table 49: Device Profile Management Settings for Data Center ELS

Field	Action
Enable Services	You can enable one or more network protocol services for this Device profile:

System Logging Settings

Optionally, expand the Syslog Settings and provide the following system logging settings:

Table 49: Device Profile Management Settings for Data Center ELS (continued)

Field	Action
Enable Device Logging for Switches	<p>To enable device logging for switches:</p> <ol style="list-style-type: none"> Under Enable Device Log, click Add. The Add Log window opens. Select the log type for switching: Console, File, User, or Host. <ul style="list-style-type: none"> Console logging sends system log messages to the console. File logging sends system log messages to the file you specify in File Name. User logging sends system log messages to the terminal session of the user specified in User Name. You will also need to provide the name of the user. Host logging sends system log messages to the server specified in Host. Host can be either an IP address or host name. Under Services, click Add. The phrase <i>Click to enter value</i> appears in both the Service column and Severity Filter column. Click on the phrase <i>Click to enter value</i> in the Service column. A list box replaces the phrase in the Services column. From the Services list, select a logging service: any, Authorization, Change-log, Conflict-log, Daemon, DFC, External, Firewall, FTP, Interactive-commands, Kernel, NTP, PFE or Security. Click on the phrase <i>Click to enter value</i> in the Severity Filter column. A list box replaces the phrase in the Severity Filter column. Select a Severity Filter from the list: Alert, Critical, Debug All, Emergency, Error, Info, Notice, or Warning. The filter is added to the list of Severity Filters. The filter is activated when the corresponding service is triggered. Click OK.
Edit Logging Settings	Select a Log Type from the Enable Device Log table and click Edit to change the information.

Table 49: Device Profile Management Settings for Data Center ELS (*continued*)

Field	Action
Delete Logging Settings	Select a Log Type from the Enable Device Log table and click Delete to remove the server information.

To configure protocol settings, either click **Next** or click **DHCP/DNS Settings**. To skip the DHCP/DNS settings, click **Review** at the top of the wizard window.

DHCP/DNS Settings options are described in the section [“Specifying Protocol Settings for Data Center Switching ELS Device Common Settings” on page 194](#).

Specifying Protocol Settings for EX Switching Device Common Settings

To configure the protocol settings for an EX Switching Device profile, enter the settings described in [Table 50 on page 188](#). All settings are optional.

Table 50: Device Profile Protocol Settings for EX Switching

Field	Action
Enable Storm Control	
	Select this option to enable storm control on a switch.
Spanning Tree Settings	

Table 50: Device Profile Protocol Settings for EX Switching (*continued*)

Field	Action
Spanning Tree Protocol Settings for switches only	<p>Select one of spanning-tree protocol (STP) settings for switches: STP, RSTP (default), MSTP, or None of these.</p> <ul style="list-style-type: none"> Spanning Tree Protocol—With STP configured, the switches use the IEEE 802.1D 2004 specification, force version 0. This configuration runs a version of RSTP that is compatible with classic, basic STP as defined in the 802.1D 1998 specification. Rapid Spanning Tree Protocol—RSTP provides faster reconvergence time than the original STP both by identifying certain links as point-to-point and by using protocol handshake messages rather than fixed timeouts. VLAN Spanning Tree Protocol (VSTP) and RSTP can be configured concurrently. You can selectively configure up to 253 VLANs by using VSTP; the remaining VLANs will be configured by using RSTP. VSTP and RSTP are the only spanning-tree protocols that can be configured concurrently on a switch. Multiple Spanning Tree Protocol—MSTP enables multiple VLANs to be mapped to the same spanning-tree instance, reducing the number of spanning-tree instances needed to support a large number of VLANs. MSTP provides multiple forwarding paths for data traffic and enables load-balancing. It improves the fault tolerance of the network because a failure in one instance, or forwarding path, does not affect other instances. <p>You can also select the Enable VSTP check box to enable VSTP.</p>
Multicast Settings	
Enable IGMP	Selecting this option enables Internet Group Management Protocol (IGMP) on all the interfaces for the selected device. Default is disabled. IGMP is a communications protocol used by both hosts and adjacent routers on IP networks to establish multicast group memberships.
Enable IGMP Snooping	Enables IGMP snooping on all VLANs. Default is enabled.
Enable DHCP Relay	
Select this option to display the DHCP Relay settings.	

Table 50: Device Profile Protocol Settings for EX Switching *(continued)*

Field	Action
Add DHCP Relay to Device Profile	<p>To add DHCP Relay to this Device profile:</p> <ol style="list-style-type: none">1. Select Legacy DHCP Relay (default).2. Add one or more DHCP servers to the Device Common Settings profile:<ol style="list-style-type: none">a. Click Add under DHCP Servers. The Add Server window opens.b. Type an IP Address.c. Click OK. The server is added to the list of DHCP Servers.

Table 50: Device Profile Protocol Settings for EX Switching (*continued*)

Field	Action
Add Extended DHCP Relay to a Device Profile	<p>To add Extended DHCP Relay to this Device profile:</p> <ol style="list-style-type: none"> 1. Select Extended DHCP Relay instead of Legacy DHCP Relay. 2. Add one or more DHCP Server Groups to the Device Common Settings profile: <ol style="list-style-type: none"> a. Click Add under Add DHCP Servers Group. The Add Server Group window opens. b. Provide a name for the server group. c. Optionally, make this an active server group by checking Active Group. d. Add servers to the group by clicking Add under DHCP Servers. The phrase <i>Click to enter value</i> appears in the IP Address column. e. Select <i>Click to enter value</i> and then enter an IP Address. f. Click OK. The server is added to the DHCP server group list. g. Add a relay interface group by clicking Add under Add Relay Interface Group. The Add DHCP Relay Interface window opens. h. Type a DHCP interface group name. i. Select a server group from the Server Group list. j. Click OK. The group is added to the Relay Interface Group list.

Click either **Next** or **Review**, to see the Review page. For review directions, see [“Reviewing and Saving a Device Common Settings Configuration” on page 197](#).

Specifying DHCP Relay/DNS Settings for Campus Switching ELS Device Common Settings

To configure the DHCP relay and DNS settings for a Campus Switching ELS Device profile, enter the settings described in [Table 51 on page 192](#). All settings are optional.

Table 51: Device Profile Protocol Settings for ELS Switching

Task	Action
DHCP Relay	
DHCP relay enables a switch to relay DHCP requests from a client to a DHCP server when the client and server do not reside on the same VLAN. You define the client interfaces for DHCP relay as part of the process of assigning the profile to a device. Select Enable DHCP Relay to enable DHCP Relay and view the DHCP Relay configuration.	
Configure Legacy DHCP Relay for ELS Switches	<div>To configure Legacy DHCP Relay for Campus Switching ELS:</div> <div><div>1. Select Legacy DHCP Relay (default).</div><div>2. Add DHCP servers for Legacy DHCP Relay:<div><div>a. Click Add under DHCP Servers.<div>The Add Server window opens.</div></div><div>b. Enter the IP address of the DHCP server and then click OK.<div>The DHCP server name appears in the list of DHCP servers.</div></div></div></div><div><div>TIP:</div> You can add more than one DHCP server.</div></div>

Table 51: Device Profile Protocol Settings for ELS Switching (continued)

Task	Action
Configure Extended DHCP Relay for Campus Switching ELS	<p>To add Extended DHCP Relay to this Campus Switching ELS Device profile:</p> <ol style="list-style-type: none"> 1. Select Extended DHCP Relay. 2. Add one or more DHCP Servers Groups to the Device Common Settings profile: <ol style="list-style-type: none"> a. Click Add under Add DHCP Servers Group. The Add Server Group window opens. b. Provide a name for the server group. c. Optionally, make this an active server group by checking Active Group. d. Add servers to the group by clicking Add under DHCP Servers. The phrase <i>Click to enter value</i> appears in the IP Address column. e. Select <i>Click to enter value</i> and then enter an IP Address. f. Click OK. The server is added to the DHCP server group list. g. Add a relay interface group by clicking Add under Add Relay Interface Group. The Add DHCP Relay Interface window opens. h. Type a group name for the DHCP interface. i. Select a server group from the Server Group list. j. Click OK. The group is added to the Relay Interface Group list.
DNS Settings	

Table 51: Device Profile Protocol Settings for ELS Switching (*continued*)

Task	Action
Add a domain name server	<p>To add a domain name server to the Campus Switching ELS Common Settings:</p> <ol style="list-style-type: none"> 1. Click Add under Domain Name Servers. The Add Server window opens. 2. Provide an IP address for the DNS server. 3. Click OK. The server is added to the Domain Name Servers list. <p>TIP: To edit a DNS server's settings, select it and then click Edit. To delete a DNS server, select it and then click Delete.</p>

Click either **Next** or **Review**, to see the Review page. For review directions, see ["Reviewing and Saving a Device Common Settings Configuration" on page 197](#).

Specifying Protocol Settings for Data Center Switching ELS Device Common Settings

To configure the protocol settings for a Data Center Switching ELS Device profile, enter the settings described in [Table 52 on page 194](#). All settings are optional.

Table 52: Device Profile Protocol Settings for Data Center Switching ELS

Task	Action
DCBX Settings	
Select the Data Center Bridging Capability Exchange (DCBX) protocol features that you want to enable	
Enable DCBX	Select Enable DCBX . DCBX is a discovery and exchange protocol for conveying configuration and capabilities among neighbors to ensure consistent configuration across the network. It is an extension of Link Layer Discovery Protocol (LLDP).
Enable LLDP	Select Enable LLDP . LLDP is a discovery and exchange protocol for conveying configuration and capabilities among neighbors to ensure consistent configuration across the network.

Table 52: Device Profile Protocol Settings for Data Center Switching ELS (continued)

Task	Action
DHCP Relay	
DHCP relay enables a switch to relay DHCP requests from a client to a DHCP server when the client and server do not reside on the same VLAN. You define the client interfaces for DHCP relay as part of the process of assigning the profile to a device. Select Enable DHCP Relay to enable DHCP Relay and view the DHCP Relay configuration.	
Configure Legacy DHCP Relay for ELS Switches	<div>To configure Legacy DHCP Relay for Campus Switching ELS:</div> <div><div>1. Select Legacy DHCP Relay (default).</div><div>2. Add DHCP servers for Legacy DHCP Relay:<div><div>a. Click Add under DHCP Servers.<div>The Add Server window opens.</div></div><div>b. Enter the IP address of the DHCP server and then click OK.<div>The DHCP server name appears in the list of DHCP servers.</div></div></div></div><div>TIP: You can add more than one DHCP server.</div></div>

Table 52: Device Profile Protocol Settings for Data Center Switching ELS (continued)

Task	Action
Configure Extended DHCP Relay for Campus Switching ELS	<p>To add Extended DHCP Relay to this Campus Switching ELS Device profile:</p> <ol style="list-style-type: none"> 1. Select Extended DHCP Relay. 2. Add one or more DHCP Servers Groups to the Device Common Settings profile: <ol style="list-style-type: none"> a. Click Add under Add DHCP Servers Group. The Add Server Group window opens. b. Provide a name for the server group. c. Optionally, make this an active server group by checking Active Group. d. Add servers to the group by clicking Add under DHCP Servers. The phrase <i>Click to enter value</i> appears in the IP Address column. e. Select <i>Click to enter value</i> and then enter an IP Address. f. Click OK. The server is added to the DHCP server group list. g. Add a relay interface group by clicking Add under Add Relay Interface Group. The Add DHCP Relay Interface window opens. h. Type a group name for the DHCP interface. i. Select a server group from the Server Group list. j. Click OK. The group is added to the Relay Interface Group list.
DNS Settings	

Table 52: Device Profile Protocol Settings for Data Center Switching ELS (*continued*)

Task	Action
Add a domain name server	<p>To add a domain name server to the Campus Switching ELS Common Settings:</p> <ol style="list-style-type: none"> 1. Click Add under Domain Name Servers. The Add Server window opens. 2. Provide an IP address for the DNS server. 3. Click OK. The server is added to the Domain Name Servers list. <p>TIP: To edit a DNS server's settings, select it and then click Edit. To delete a DNS server, select it and then click Delete.</p>

Click either **Next** or **Review**, to see the Review page. For review directions, see [“Reviewing and Saving a Device Common Settings Configuration” on page 197](#).

Reviewing and Saving a Device Common Settings Configuration

From this page, you can save or make changes to Device Common Settings:

- To make changes to the settings, click the **Edit** associated with the configuration you want to change.

Alternatively, you can also click appropriate sections of the workflow at the top of the page that corresponds to the configuration you want to change.

When you have completed your modifications, click **Review** to return to this page.

- To save a new profile or to save modified settings to an existing profile, click **Finish**.

The Manage Device Common Settings page is displayed with the new or modified profile listed

What to Do Next

Once the Device Common Settings profile is created, you must assign the profile to the required device by using the Manage Device Profile page and then deploy the Device profile by using the **Deploy** mode. To assign a Device Common Settings profile to a device, see [“Assigning Device Common Settings to Devices” on page 198](#). For information about deploying your configurations, see [“Deploying Configuration to Devices” on page 580](#).

NOTE: A device can have only one Device profile assigned to it. However, you can assign the same Device profile to multiple devices.

RELATED DOCUMENTATION

[Assigning Device Common Settings to Devices | 198](#)

[Deploying Configuration to Devices | 580](#)

[Network Director Documentation home page](#)

Assigning Device Common Settings to Devices

IN THIS SECTION

- [Assigning Device Common Settings | 199](#)
- [Editing the Assignments of the Device Common Setting | 201](#)

Once a Device Common Settings profile is created or discovered (system-created profile), you must assign it to devices using the steps described in this topic. You can assign a Device profile to a either single device, a series of single devices, or a Custom Group of devices (see [“Creating Custom Device Groups” on page 155](#)).


NOTE: A device can have only one Device Common Settings profile assigned to it.

You must have one or more device profiles created or discovered before you can assign a device profile to a device. When you deploy an assigned device profile, the configuration is pushed onto the device.

This topic describes:

Assigning Device Common Settings

To assign device common settings to either a single device, a series of single devices, or members of a Custom Group:

1. Click  in the Network Director banner.
2. Select **Device Common Settings** from the Profile and Configuration Management menu in the Tasks pane.

The Manage Device Common Settings page is displayed. The page displays all the device profiles that you configured as well as the system-created profiles detected during device discovery.

3. Select an undeployed profile from the list of profiles and then click **Assign**.

The Assign Device Profile page for the selected device family appears with a wizard consisting of three parts, Device Selection, Profile Assignment, and Review. Device Selection is displayed.

4. Expand the Device Selection object tree and select one or more objects to receive the device profile. You must place a check next to a device to select it—simply highlighting the device does not select it.

NOTE: If Network Director fails to read the configuration of one or more devices after device discovery, those devices are not displayed in the Device Selection list. You will not be able to assign profiles to those devices. The Manage Jobs page in System mode displays details of the device discovery jobs. Use the information displayed on this page to take appropriate corrective steps to enable Network Director to reread the configuration of the failed device. For more information, see [“Discovering Devices in a Physical Network” on page 94](#).

5. Click either **Next** or click **Profile Assignment** from the wizard workflow.

The Profile Assignment page opens, displaying your selections, including their Device (name), Type, Assigned To, and Attributes. The Assigned To column now has the entry DEVICE and the Attributes column has the entry Undefined.

6. Click **Define** in the **Attributes** column in the Assignments table to configure the attributes.

The Configure Attributes window opens, listing all the Layer 3 interfaces available on the device.

- a. Select the Layer 3 interfaces that are required for DHCP relay from the Available list and using the right arrow, move them to the Selected list. You can reorder the interfaces using the UP and DOWN arrows.

b. Click **Save** to save the interface list and close the Configure Attributes window.

7. You can view the assignment details for the selected device and also remove any assignments:

- To view the assignment details, select the device and click **View Assignments**.

The Profile Details page for selected device appears. Expand the **Device** name to view the details of the assignment. The assignment status displays the status whether the device is deployed or is pending device update, and so on.

- To delete a device common setting assignment for a device, select the device from the Assignments table and click **Remove**.

8. Click **Next** or click **Review** from the wizard workflow to review the assignments. On the Review page, click **Edit** to edit the profile assignment.

9. Click **Finish** once you are done reviewing the profile assignment.

The Create Profile Assignments Job Details window appears with a status report for the profile assignment job—click **OK** to close this window. If you have assigned the profile to a large number of objects, the profile assignment job can take some time to complete. Instead of waiting for the Job Details dialog box to report job completion status, you can close it and check the details of the profile assignment job at a later time using the Manage Job task in System mode.

NOTE: If any assignment fails, the profile assignment job fails and none of the assignments are created. Check the details for the profile assignment job for information about why the assignment failed.

An assigned Device profile has the Assignment State *Pending Deployment* in the Manage Device Common Settings list. Deploy any device profile in this state by following the directions “[Deploying Configuration to Devices](#)” on page 580.

To view the details of a profile, select the profile from the Manage Device Common Settings page and then click **Details**.

SEE ALSO

[Creating and Managing Device Common Settings](#) | 170

[Deploying Configuration to Devices](#) | 580

[Network Director Documentation home page](#)

Editing the Assignments of the Device Common Setting

Use the Edit Assignments page to change device common setting assignments. To edit an existing assignment:

1. Select a profile from the **Manage Device Common Settings** page and click **Edit Assignment**.

The Edit Assignments page for the selected device appears.

2. Expand the **Devices** cabinet and make the desired change from the **Operation** column of the table.

3. Click **Define** from the **Attributes** column of the table to modify the attributes.

The Configure attributes page is displayed listing all the Layer 3 interfaces available on the device.

- Select the Layer 3 interfaces that are required for DHCP relay from the Available box and using the right arrow, move them to the Selected box.

You can rearrange the order of the interfaces using the Up and Down arrows.

- Click **Save** after you are done with selecting the interfaces.

4. Click **Apply** once you are done with the changes.

The Manage Device Common Settings page is displayed.

RELATED DOCUMENTATION

[Creating and Managing Device Common Settings | 170](#)

[Creating Custom Device Groups | 155](#)

[Network Director Documentation home page](#)

Configuring Authentication, Authorization, and Access for Your Network

IN THIS CHAPTER

- Understanding Central Network Access Using RADIUS and TACACS+ | 202
- Creating and Managing RADIUS Profiles | 206
- Creating and Managing LDAP Profiles | 212
- Understanding Access Profiles | 217
- Creating and Managing Access Profiles | 218
- Understanding Authentication Profiles | 235
- Creating and Managing Authentication Profiles | 237

Understanding Central Network Access Using RADIUS and TACACS+

IN THIS SECTION

- Why Do I Want Remote Authentication ? | 203
- Where Is RADIUS Installed on the Network? | 204
- How Is TACACS+ Installed on the Network? | 204
- A Comparison of RADIUS and TACACS+ | 205

Remote Access Dial In User Service (RADIUS) and Terminal Access Controller Access-Control System Plus (TACACS+) are two common security protocols used to provide centralized access into networks. RADIUS was designed to authenticate and log remote network users, while TACACS+ is most commonly used for administrator access to network devices like routers and switches. Both protocols provide centralized Authentication, Authorization, and Accounting (AAA) management for computers that connect and use a network service.

- *Authentication* - Who is allowed to gain access to the network? Traditionally authorized users provide a username and password to verify their identity for both RADIUS and TACACS+.
- *Authorization* - What services can a user access once they are authenticated? It is unlikely that you want your finance people to have access to the developer database. Visitors may have access only to the Internet, while only IT staff can access the entire passwords database.
- *Accounting* - What services did each user access and for how long? Accounting records record the user's identification, network address, point of attachment and a unique session identifier—these statistics are tracked and added to the user's record. This is useful when time on the system is billed to individuals or departments.

Why Do I Want Remote Authentication ?

Remote authentication enables you to keep your username and passwords in one place, on a central server. The advantage to using RADIUS or TACACS+ on this central server is that you don't configure changes on each separate network device when a user is added or deleted, or when a user changes a password. You only make one change to the configuration on the server and then devices continue to access the server for authentication. Although authentication is the most well known function of RADIUS and TACACS+, there are two additional functions provided, authorization and accounting.

NOTE: Instead of using a flat database on the RADIUS server, you can refer to external sources such as SQL, Kerberos, LDAP, or Active Directory servers to verify user credentials.

Why Not Just Rely on Firewalls and Filters for Access Control?

Routers and firewalls usually control access to services using filters based on source and/or destination IP addresses and ports. This means that restrictions are applied to devices and not to individual clients. For example if I enable traffic from 10.1.0.255 to access a particular web server, then anyone who is sitting at the machine with the address of 10.1.0.255 automatically has access to this server. Using RADIUS or TACACS+, that same person sitting at the machine with the address of 10.1.0.255 also has to provide a username and password to access a service.

What About Using LDAP For Authentication?

Lightweight Directory Access Protocol (LDAP) is a client/server protocol used to access and manage directory information. It reads and edits directories over IP networks and runs directly over TCP/IP using

simple string formats for data transfer. Directory servers include information about various entities on your network, such as user names, passwords, rights associated with user names, metadata associated with user names, devices connected to the network, and device configuration.

Use LDAP to obtain directory information, such as email addresses and public keys. If you want to make directory information available over the Internet, this is the way to do it. LDAP works well for captive portal authentication. However, LDAP does not implement 802.1X security easily. 802.1X was essentially designed with RADIUS in mind, so 802.1X challenge/response protocols like MSCHAPv2 work well with RADIUS.

Where Is RADIUS Installed on the Network?

RADIUS includes three components: an authentication server, client protocols, and an accounting server. The RADIUS server portion of the protocol is usually a background process running on a UNIX or Microsoft Windows server.

With RADIUS, the term client refers to a network access device (NAD) that provides the client part of the RADIUS service—a modem pool, a switch, a network firewall, or any other device that needs to authenticate users can be configured as a NAD to recognize and process connection requests from outside the network edge. When a NAD receives a user's connection request, it may perform an initial access negotiation with the user to obtain identity/password information. Then the NAD passes this information to the RADIUS server as part of an authentication/authorization request.

NOTE: RADIUS requires that each network client device be configured.

How Is TACACS+ Installed on the Network?

TACACS+ logon authentication protocol uses software running on a central server to control access by TACACS-aware devices on the network. The server communicates with switches or other TACACS-aware devices automatically—these devices do not require further configuration if they are TACACS-aware. The TACACS+ protocol is supported by most enterprise and carrier-grade devices.

Install the TACACS+ Service as close as possible to the user database, preferably on the same server. TACACS+ needs to be closely synchronized with your Domain, and any network connection issues, DNS problems, or even time discrepancies can cause a critical service failure. Installing TACACS+ on the same server as the user database can also improve performance.

TACACS+ servers should be deployed in a fully trusted internal network. If you keep your TACACS+ service within your trusted network, you need to open only one port, TCP 49. There should not be any direct access from untrusted or semi-trusted networks.

NOTE: RADIUS is typically deployed in a semi-trusted network, and TACACS+ uses internal administrative logins, so combining these services on the same server could potentially compromise your network security.

A Comparison of RADIUS and TACACS+

Table 53: RADIUS and TACACS+

	RADIUS	TACACS+
Primary Use	Authenticate and log remote network users	Provide administrator access to network devices like routers and switches
Authentication and Authorization	Authentication and Authorization checking are bundled together. When the client device requests authentication from the server, the server replies with both authentication attributes and authorization attributes. These functions can not be performed separately.	All three AAA functions (authentication, authorization, and accounting) can be used independently. Therefore, one method such as kerberos can be used for authentication, and a separate method such as TACACS+ can be used for authorization.
Accounting	The accounting features of the RADIUS protocol can be used independently of RADIUS authentication or authorization.	
Protocol	User Datagram Protocol (UDP)/IP with best-effort is used for delivery on ports 1645/1646, 1812/1813	TCP used for delivery on port 49. Also has multiprotocol support for AppleTalk Remote Access (ARA) protocol, NetBIOS Frame Protocol Control protocol, Novell Asynchronous Services Interface (NASI), and X.25 PAD connection.
Encryption applied to	Password	Username and password
802.1X Security	If you want to use 802.1x port-based network access control, you have to use the RADIUS client because the TACACS+ client does not support that feature.	
Model	client/server	
Recommended Environment	semi-trusted	trusted

RELATED DOCUMENTATION

[Creating and Managing RADIUS Profiles | 206](#)[Network Director Documentation home page](#)

Creating and Managing RADIUS Profiles

IN THIS SECTION

- [Managing RADIUS Profiles | 206](#)
- [Creating RADIUS Profiles | 208](#)
- [Specifying Settings for a RADIUS Profile | 208](#)
- [What to Do Next | 211](#)

Remote Authentication Dial In User Service (RADIUS) is a networking protocol that provides centralized authentication, authorization, and accounting (AAA) management for computers to connect and use a network service. By default, RADIUS servers are used for both accounting and authentication. From Network Director, you can create and manage RADIUS profiles that configure RADIUS server settings.

TIP: In addition to your RADIUS server, you can configure an LDAP server for EX Series ELS switch authentication also—for directions, see [“Creating and Managing LDAP Profiles” on page 212](#).

This topic describes:

Managing RADIUS Profiles

From the Manage RADIUS Profiles page, you can:

- Create a new profile by clicking **Add**. For directions, see [“Creating RADIUS Profiles” on page 208](#).
- Modify an existing profile by selecting it and clicking **Edit**.
- View information about a profile by selecting the group and clicking **Details** or by clicking the profile name.
- Delete profiles by selecting the profile and clicking **Delete**.

TIP: You cannot delete profiles that are in use—that is, assigned to objects or used by other profiles. To see the current assignments for a profile, select the profile and click **Details**.

- Clone a profile by selecting the profile and clicking **Clone**.

Table 54 on page 207 describes the information provided about RADIUS profiles on the Manage RADIUS Profiles page. This page lists all RADIUS profiles defined for your network, regardless of your current selected scope in the network view.

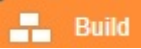
Table 54: RADIUS Profile Information

Field	Description
RADIUS Profile Name	Name given to the RADIUS profile when it was created.
Server Address	IP address of the RADIUS server.
Server Port	UDP port being used by the RADIUS server.
Creation Time	Date and time when this profile was created.
Update Time	Date and time when this profile was last modified.
User Name	The username of the user who created or modified the profile.

TIP: All columns may not be currently displayed. To show or hide fields in the table, click the down arrow on the field header, select **Columns**, and select or clear the check box adjacent to the field that you want to show or hide.

Creating RADIUS Profiles

To create a RADIUS profile:

1. Click  in the Network Director banner.
2. Under Select View, select either **Logical View**, **Location View**, **Device View** or **Custom Group View**.

TIP: Do not select **Dashboard View**, **Datacenter View** or **Topology View**.

3. From the Tasks pane, select the type of network (Wired), the appropriate functional area (System or AAA), and select the name of the profile that you want to create. For example, to create a port profile for a wired device, click **Wired > Profiles > PORT**. The Manage Profile page opens.
4. Click **Add** on the Manage RADIUS Profiles page.
The Create RADIUS Profile page appears.
5. Enter settings for the RADIUS profile on the Create RADIUS Profile page as described in [“Specifying Settings for a RADIUS Profile” on page 208](#).
6. Click **Done**.

Specifying Settings for a RADIUS Profile

Use the Create RADIUS Profile page to define authentication, authorization, and accounting settings for a RADIUS server.

[Table 55 on page 208](#) describes the RADIUS profile settings.

Table 55: RADIUS Profile Settings

Field	Action
Server Name	Type a name for the server, using up to 64 alphanumeric characters and no special characters other than the underscore. The name must be unique among servers.
Server Address	Type the IP address of the RADIUS server.

Table 55: RADIUS Profile Settings (*continued*)

Field	Action
Authentication Port (default is 1812)	Using the arrows, adjust the number of the UDP port to use for RADIUS authentication messages. The default UDP port is 1812, and the range is from 0 to 65535.
Secret	Provide a password for the RADIUS server.

Advanced Settings

You can change the advanced settings for a RADIUS server, or you can use the default settings.

Accounting Port (default is 1813)	Using the arrows, adjust the number of the UDP port to use for RADIUS accounting messages. The default UDP port is 1813, and the range is from 0 to 65535.
Retry Count (default is 3)	Using the arrows, adjust the retry count until it reflects the number of times Network Director retries connecting to the RADIUS server when the RADIUS server is unavailable.
Timeout (default is 5 seconds)	Using the arrows, adjust the timeout value. Timeout indicates how many seconds Network Director allows for RADIUS server connection before giving an unreachable error.
Dead Time (default is 5 seconds)	Using the arrows, adjust the number of seconds before Network Director checks a RADIUS server that was previously unresponsive. The default value is 5 seconds.
Use MAC as Password	Enable this option if you want each client device to use its MAC address as its password for the RADIUS server. If you enable Use MAC As Password, then the Authorization Password field becomes unavailable.
Authorization Password	If you are not using MAC addresses as passwords for the RADIUS server, provide a common password here.

Table 55: RADIUS Profile Settings (*continued*)

Field	Action
MAC Address Format	<p>Select None, Hyphens, Colons, One-Hyphen, or Raw to determine the MAC address format used with the RADIUS server. For example:</p> <ul style="list-style-type: none"> • None—For unicast IPv4, an example MAC address is 0123456789ab. For unicast IPv6, an example MAC address is 20010db8000000000000ff0000428329. • Hyphens—For unicast IPv4, an example MAC address using hyphens is 01-23-45-67-89-ab. For unicast IPv6, an example MAC address using hyphens is 2001-0db8-0000-0000-0000-ff00-0042-8329. • Colons—For unicast IPv4, an example MAC address using colons is 01:23:45:67:89:ab. For unicast IPv6, an example MAC address using colons is 2001:0db8:0000:0000:0000:ff00:0042:8329. • One-Hyphen: IPv6 unicast addresses other than those that start with binary 000 are logically divided into two parts: a 64-bit (sub-)network prefix, and a 64-bit interface identifier used to identify a host's network interface. The hyphen is placed between the two parts. • Raw: The IPv6 address is represented by all numbers—no sections containing all zeros are skipped and then represented by a double colon. For example, this is a raw IPv6 address: 2001:0000:0234:C1AB:0000:00A0:AABC:003F.

Table 55: RADIUS Profile Settings (*continued*)

Field	Action
Authentication Protocol (Default is PAP)	<p>Select PAP, CHAP, MSCHAP-V2, or None to determine an authentication protocol for the RADIUS server. These authentication protocols work as follows:</p> <ul style="list-style-type: none"> • PAP: stands for Password Authentication Protocol and is used by Point to Point Protocols to validate users before allowing them access to server resources. Almost all network operating system remote servers support PAP. However, PAP transmits unencrypted ASCII passwords over the network and is therefore not secure. Use it as a last resort when the remote server does not support the stronger authentication. • CHAP: stands for Challenge Handshake Authentication Protocol and authenticates a user or network host to an authenticating entity. CHAP provides protection against replay attacks by the peer through the use of an incrementally changing identifier and of a variable challenge-value. CHAP requires that both the client and server know the plaintext of the secret password—it is never sent over the network. CHAP provides better security than PAP does. • MSCHAP-V2: stands for Microsoft's implementation of the Challenge Handshake Authentication Protocol version 2 on the router for password-change support. This feature provides users accessing a router the option of changing the password when the password expires, is reset, or is configured to be changed at the next login. The MS-CHAP variant does not require either peer to know the plaintext of the secret password. MSCHAP-V2 is used as an authentication option with RADIUS servers used for Wi-Fi security using the WPA-Enterprise protocol.
Server Priority (default is 1)	<p>Enter a server priority to indicate the order in which RADIUS servers are accessed. Entering a one means that this server is checked first.</p>

Click **OK** to add the RADIUS server to the EX Switching Access profile. You can add more RADIUS servers if needed.

If you have multiple RADIUS servers, you can prioritize them in the Authentication Server Order section, using the arrows.

Click **Done** to create the RADIUS server profile.

The RADIUS server name appears in the list of RADIUS servers on the Manage RADIUS Profiles page.

What to Do Next

Link the RADIUS server to an Access profile. For directions, see [“Creating and Managing Access Profiles” on page 218](#).

RELATED DOCUMENTATION

[Creating and Managing Access Profiles | 218](#)

[Creating and Managing LDAP Profiles | 212](#)

[Understanding Central Network Access Using RADIUS and TACACS+ | 202](#)

[Network Director Documentation home page](#)

Creating and Managing LDAP Profiles

IN THIS SECTION

- [Managing LDAP Profiles | 212](#)
- [Creating LDAP Profiles | 214](#)
- [Specifying Settings for an LDAP Profile | 214](#)
- [What to Do Next | 217](#)

Lightweight Directory Access Protocol (LDAP) is an Internet protocol that email and other programs use to look up information from a server. Use LDAP to look up encryption certificates, pointers to printers and other services on a network, in addition to providing a single logon where one user password is used for different services. LDAP authentication is appropriate for any kind of directory-like information where fast lookups and infrequent updates are used. From Network Director, you can create and manage LDAP profiles for EX Switching ELS.

TIP: In addition to an LDAP server, you can configure a RADIUS server for both authentication and accounting purposes—for directions, see [“Creating and Managing RADIUS Profiles” on page 206](#).

This topic describes:

Managing LDAP Profiles

From the Manage LDAP Profiles page, you can:

- Create a new LDAP profile by clicking **Add**. For directions to add an LDAP profile, see [“Creating LDAP Profiles” on page 214](#).
- Modify an existing LDAP profile by selecting it and clicking **Edit**.
- View information about a profile by selecting the group and clicking **Details** or by clicking the profile name.
- Delete LDAP profiles by selecting the profile and clicking **Delete**.

TIP: You cannot delete profiles that are in use—that is, assigned to objects or used by other profiles. To see the current assignments for a profile, select the profile and click **Details**.

- Clone an LDAP profile by selecting a profile and clicking **Clone**.

[Table 56 on page 213](#) describes the information provided about LDAP profiles on the Manage LDAP Profiles page. This page lists all LDAP profiles defined for your network, regardless of your current selected scope in the network view.


Table 56: LDAP Profile Information

Field	Description
LDAP Name	Name given to the LDAP server profile when it was created.
Server Address	IP address of the LDAP server.
Server Port	UDP port being used by the LDAP server.
Domain Name	Domain using the LDAP server.
Creation Time	Date and time when this profile was created.
Update Time	Date and time when this profile was last modified.
User Name	The username of the user who created or modified the profile.

TIP: All columns of information may not be displayed. To show or hide fields in the table, click the down arrow on the field header, select **Columns**, and select or clear the check box adjacent to the field that you want to show or hide.

Creating LDAP Profiles

To create an LDAP profile:

1. Click  **Build** in the Network Director banner.
2. Under Select View, select either **Logical View**, **Location View**, **Device View** or **Custom Group View**.

TIP: Do not select **Dashboard View**, **Datacenter View** or **Topology View**.

3. From the Tasks pane, select the type of network (Wired), the appropriate functional area (System, or AAA), and select the name of the profile that you want to create. For example, to create a radius profile for a wired device, click **Wired > Profiles > PORT**. The Manage Profile page opens.
4. Click **Add** to add a new profile.
The Create LDAP Profile page for the selected device family is displayed.
5. Enter settings for the LDAP profile as described in [“Specifying Settings for an LDAP Profile” on page 214](#).
6. Click **Done**.

Specifying Settings for an LDAP Profile

Use the Create LDAP Profile page to define LDAP directory information services over an IP network.

[Table 57 on page 214](#) describes the LDAP settings.

Table 57: LDAP Profile Settings

Field	Action
Server Name	Type a name for the server, using up to 64 alphanumeric characters and no special characters other than the underscore. The name must be unique among servers.
Server Address	Type the IP address of the LDAP server.
Server Port (default is 389)	Using the arrows, adjust the number of the UDP port to use for LDAP authentication messages. The default port is 389 for unencrypted LDAP servers and 636 for unencrypted LDAP servers.

Table 57: LDAP Profile Settings (*continued*)

Field	Action
Advanced LDAP Settings	
Fully Qualified Domain Name	<p>Type a fully qualified domain name (FQDN)—this is the complete domain name for a specific computer, or host, on the Internet. The FQDN consists of two parts: the host name and the domain name. For example, an FQDN for a server might be ldap12.example.com. The host name is ldap12, and the host is located within the domain example.com. This domain name specifies all domain levels, including the top-level domain and the root zone. A fully qualified domain name is specified with a trailing dot, for example:</p> <p>ldap12.example.com.</p>
Dead Time (default is 5 seconds)	Using the arrows, adjust the number of seconds before Network Director checks an LDAP server that was previously unresponsive. The default value is 5 seconds.
Timeout (default is 5 seconds)	Using the arrows, adjust the number of seconds Network Director tries to establish connection with RADIUS server before giving an unreachable error.
Bind Mode (default is SASL-MD5)	<p>Select either SASL-MD5 or SIMPLE-AUTH to establish authentication for an LDAP session.</p> <p>Simple Authentication and Security Layer (SASL) is a framework for authentication and data security in Internet protocols. It decouples authentication mechanisms from application protocols, in theory enabling any authentication mechanism supported by SASL to be used in any application protocol that uses SASL.</p> <p>SIMPLE-AUTH sends the user's domain name and password in plain text. The server then checks the password against the password attribute in the named entry.</p> <p>TIP: We recommend that connections using SIMPLE-AUTH be encrypted using Transport Layer Security (TLS).</p>

Table 57: LDAP Profile Settings (continued)

Field	Action
MAC Address Format (default is Hyphens)	<p>Select None, Hyphens, Colons, One-Hyphen, or Raw to determine the MAC address format used with the LDAP server. For example:</p> <ul style="list-style-type: none"> • None: For unicast IPv4, an example MAC address is 0123456789ab. For unicast IPv6, an example MAC address is 20010db8000000000000ff0000428329. • Hyphens: For unicast IPv4, an example MAC address with hyphens is 01-23-45-67-89-ab. For unicast IPv6, an example MAC address with hyphens is 2001-0db8-0000-0000-0000-ff00-0042-8329. • Colons: For unicast IPv4, an example MAC address with colons is 01:23:45:67:89:ab. For unicast IPv6, an example MAC address with colons is 2001:0db8:0000:0000:0000:ff00:0042:8329. • One-Hyphen: IPv6 unicast addresses other than those that start with binary 000 are logically divided into two parts: a 64-bit (sub-)network prefix, and a 64-bit interface identifier used to identify a host's network interface. The hyphen is placed between the two parts. • Raw: The IPv6 address is represented by all numbers—no sections containing all zeros are skipped and then represented by a double colon. For example, this is a raw IPv6 address: 2001:0000:0234:C1AB:0000:00A0:AABC:003F. <p>TIP: A media access control address (MAC address) is a unique identifier assigned to network interfaces for communications on the physical network segment. MAC addresses are used as a network address for most IEEE 802 network technologies, including Ethernet.</p>
Base Domain	Base domains contain no extra dots. For example, example.com is a base domain, but www.example.com is not because it contains an extra dot.
Domain Prefix (default is cn)	Enter a domain prefix to identify a subdomain. The subdomain name can be used to identify services, devices, or regions.
Use MAC as Password (default is unchecked)	Check this option if you want each client device to use its MAC address as its password for the LDAP server.
Authorization Password	If you are not using individual MAC addresses as passwords for the LDAP server, provide a common password here.

Click **Done** to create the LDAP Server profile. The profile appears on the list on the Manage LDAP Profiles page.

What to Do Next

Link the LDAP server to an Access profile for Campus Switching with ELS. For directions, see [“Creating and Managing Access Profiles” on page 218](#).

RELATED DOCUMENTATION

[Creating and Managing Access Profiles | 218](#)

[Creating and Managing RADIUS Profiles | 206](#)

[Network Director Documentation home page](#)

Understanding Access Profiles

Access profiles enable access configuration on the network—this consists of authentication configuration and accounting configuration. Network Director supports RADIUS, Lightweight Directory Access Protocol (LDAP), and local authentication as authentication methods, and RADIUS for accounting.

Authentication prevents unauthorized devices and users from gaining access to your network. Authentication controls access to your network using authentication methods such as 802.1X, MAC RADIUS, or captive portal. For 802.1X and MAC RADIUS authentication, end devices or users must be authenticated before they receive an IP address from a DHCP server. For captive portal authentication, the switch enables the end devices to obtain an IP address, after which these devices can forward packets such as DHCP, DNS, and ARP.

Accounting servers collect and send information used for billing, auditing, and reporting, such as:

- User identity
- Connection start and stop times
- Number of packets received and sent
- Number of transferred bytes

The accounting information is stored locally or on a remote RADIUS server. You can track sessions by using this information. As network users roam through a Network , accounting records can be used to track their network usage.

RADIUS is an authentication and accounting server used for validating users who attempt to access the switch. RADIUS is a distributed client-server system—the RADIUS client runs on the switch, and the server runs on a remote network system.

LDAP is an Internet protocol for accessing and updating information in an X.500-compliant directory. Network administrators for LDAP clients can connect to X.500 directory service and add, delete, modify, or search for information if they have the required access rights to the directory. LDAP is designed to run over TCP/IP and can access information in both X.500 directories and many non-X.500 directories.

NOTE: LDAP is supported as an authentication and accounting method for Campus Switching ELS devices.

With local authentication, you configure a password for each user allowed to log in to the switch.

You can define one or more Access profiles. Each Access profile is specific to a device family. Use the Manage Access Profiles page to create, modify, view, and delete existing Access profiles.

RELATED DOCUMENTATION

[Creating and Managing Access Profiles | 218](#)

[Creating and Managing RADIUS Profiles | 206](#)

[Creating and Managing LDAP Profiles | 212](#)

[Network Director Documentation home page](#)

Creating and Managing Access Profiles

IN THIS SECTION

- [Managing Access Profiles | 219](#)
- [Creating an Access Profile | 220](#)
- [Specifying Basic Settings for an EX Series Switching Access Profile | 222](#)
- [Specifying RADIUS Accounting Settings for an EX Switching Access Profile | 224](#)
- [Specifying Basic Settings for a Campus Switching ELS Access Profile | 227](#)
- [Specifying RADIUS and LDAP Settings for Campus Switching ELS | 227](#)
- [Reviewing and Modifying the Access Profile Settings | 234](#)
- [What To Do Next | 235](#)

Access profiles enable authentication configuration for both methods and servers. Network Director supports the configuration of RADIUS, Lightweight Directory Access Protocol (LDAP), and local authentication as authentication methods, and RADIUS as an accounting method.

Use the Manage Access Profiles page to create new Access profiles and manage existing Access profiles.

This topic describes:

Managing Access Profiles

From the Manage Access Profiles page, you can:

- Create a new Access profile by clicking **Add**. For directions, see [“Creating an Access Profile” on page 220](#).
- Modify an existing profile by selecting it and clicking **Edit**.
- View information about an Access profile, including the interfaces it is associated with, by either clicking the profile name or by selecting the profile and clicking **Details**.
- Delete an Access profile by selecting the Access profile and clicking **Delete**.

TIP: You cannot delete profiles that are in use—that is, assigned to objects or used by other profiles. To see the current assignments for an Access profile, select the Access profile and click **Details**.

- Clone a profile by selecting a profile and clicking **Clone**.

TIP: The default Access profile named *Juniper Networks-access-profile* is always available.

[Table 58 on page 219](#) describes the information provided about Access profiles on the Manage Access Profiles page. This page lists all Access profiles defined for your network, regardless of the scope you selected in the network view.

Table 58: Manage Access Profile Fields

Field	Description
Profile Name	Name given to the profile when the profile was created.
Description	<p>Description of the profile that was entered when the profile was created.</p> <p>TIP: To display the entire description, you might need to resize the Description column by clicking the column border in the heading and dragging it.</p>

Table 58: Manage Access Profile Fields (*continued*)

Field	Description
Family Type	The device family on which the profile was created: EX Switching or Campus Switching ELS.
Creation Time	Date and time when the profile was created.
Last Updated Time	Date and time when the profile was last modified.
User Name	The username of the person who created or modified the profile.

TIP: All columns might not be displayed. To show or hide fields listed in the table, click the down arrow on the field header, select **Columns**, and select or clear the check box adjacent to the field that you want to show or hide.

Creating an Access Profile


In Network Director, you create an Access profile that is then used to authenticate network users. You can also specify servers to be used for user accounting purposes. You can create Access profiles for these kinds of hardware devices:

- EX Series Switches—configure Basic Settings and optional Accounting Settings.
- EX Series switches with ELS—configure Basic Settings and Server Settings.

To create an Access profile, follow these steps:

1. Under Views, select one of these options: **Logical View**, **Location View**, **Device View** or **Custom Group View**.

TIP: Do not select **Dashboard View**, **Datacenter View**, or **Topology View**.

2. Click  in the Network Director banner.
3. Under Select View, select either **Logical View**, **Location View**, **Device View** or **Custom Group View**.

TIP: Do not select **Dashboard View**, **Datacenter View** or **Topology View**.

4. From the Tasks pane, select the type of network (Wired), the appropriate functional area (System or AAA), and select the name of the profile that you want to create. For example, to create a port profile for a wired device, click **Wired > Profiles > Port**. The Manage Profile page opens.

5. Click **Add** to add a new profile.

If you chose to create a profile for the wired network, Network Director opens the Device Family Chooser window.

- a. From the Device Family Chooser, select the device family for which you want to create a profile. The available device families are **Switching (EX)**, **Campus Switching ELS** (Enhanced Layer 2 Software) and **Data Center Switching ELS**.

- b. Click **OK**.

The Create Access Profile page for the selected device family is displayed.

6. Click **Add**.

The Device Family Chooser window opens.

7. From the Device Family Chooser, select the device family for which you want to create a profile. The available device families are **Switching EX** and **Campus Switching ELS**.

8. Click **OK**.

The Create Access Profile wizard for the selected device family opens—it consists of two sections: Basic Settings and RADIUS and LDAP configuration.

9. Specify the access settings for the Access profile by doing one of the following:

- For EX Series switches, specify the access settings described in online help, or in [“Specifying Basic Settings for an EX Series Switching Access Profile” on page 222](#) and [“Specifying RADIUS Accounting Settings for an EX Switching Access Profile” on page 224](#).
- For Campus Switching ELS, specify the access settings as described in [“Specifying Basic Settings for a Campus Switching ELS Access Profile” on page 227](#) and [“Specifying RADIUS and LDAP Settings for Campus Switching ELS” on page 227](#).

10. Click either **Next** or **Review**. The Review page appears.

You can either save your profile or make changes to your profile from the Review page. For directions, see [“Reviewing and Modifying the Access Profile Settings” on page 234](#).

11. Click **Done** to save the Access profile.

The system saves the Access profile and then displays the Manage Access Profiles page. Your new or modified Access profile is listed in the table of Access profiles.

Specifying Basic Settings for an EX Series Switching Access Profile

Basic settings for EX Series switching Access profile include the profile name, authentication server order, and the RADIUS authentication details.

To configure the basic settings for an EX Series switch Access profile, enter the settings described in [Table 59 on page 222](#). Required settings are indicated in the user interface by a red asterisk (*) that appears next to the field label.

Table 59: Access Profile Basic Settings for EX Series Switches

Field	Action
Access Profile Details	
Profile Name	Type a unique name that identifies the profile. You can use up to 64 characters for profiles created for wired devices. Profile name must not contain special characters or spaces. Note that profiles that are automatically created by Network Director as part of device discovery or out-of-band changes may contain the underscore (_) character.
Description	Type the description of the profile.
Revert Interval	Specify the number of seconds the switch waits after an authentication server becomes unreachable. The switch rechecks the connection to the server when the specified interval expires. Default is 3 seconds.
RADIUS Servers: Authentication	
View	Select a server entry from the list and then click View to see the details of that entry.

Table 59: Access Profile Basic Settings for EX Series Switches *(continued)*

Field	Action
Task: Create and add a new RADIUS server configuration	<p>To both create and add a RADIUS server configuration to this Access profile for authentication:</p> <ol style="list-style-type: none"> Click Add > Create RADIUS. The Create RADIUS Server window opens. Complete these fields: <ul style="list-style-type: none"> Server Name—Type the name of the RADIUS server that you want to create. Server Address—Type the IP address of the RADIUS server. Authentication Port—The default RADIUS authentication port is 1812. You can change the port number by using the up and down arrows. Secret—Provide a password. If the password contains spaces, enclose it in quotation marks. The secret password used by the switch must match the one used by the server. Expand the RADIUS server and change any of these configurations: <ul style="list-style-type: none"> Accounting Port—You can change the default port number (1813) by using the up and down arrows. Retry Count—Specify the number of times that a device attempts to contact the LDAP authentication server. The default retry count is 3. You can change this value by using the up and down arrows to 1 through 10 times. Timeout (seconds)—Specify the number of seconds the switch waits to receive a response from a RADIUS server. The default timeout is 5 seconds. You can change this value, using the up and down arrows, to 1 through 90 seconds. Click OK. The RADIUS server is automatically added to the list of authentication servers assigned to this Access profile. If you have more than one RADIUS server listed, you can use the arrows to reorder the list priority so that the most preferred RADIUS server is listed first.

Table 59: Access Profile Basic Settings for EX Series Switches *(continued)*

Field	Action
Task: Add a previously configured RADIUS server for authentication	<p>The RADIUS tab is selected by default for server configuration and configured RADIUS servers are listed on this Server Settings page. To add a previously configured RADIUS server to this Access profile for authentication:</p> <ol style="list-style-type: none"> 1. Click Add > Select RADIUS. A list of available configured RADIUS servers is displayed. Servers in this list were either automatically discovered or created by using the directions in “Creating and Managing RADIUS Profiles” on page 206. 2. Select one or more RADIUS servers from the list of Available servers and use the arrows to move the server to the Selected list. 3. Click OK. The RADIUS server is added to the list of authentication servers to be used with this Access profile. 4. If you have more than one RADIUS server listed, you can use the arrows to reorder the login priority so that the most preferred RADIUS server is listed first.
Task: Delete a server	<p>To delete a RADIUS server from this Access profile:</p> <ol style="list-style-type: none"> 1. Select a RADIUS server from the list. 2. Click Delete. The RADIUS server is removed from the list of authentication servers to be used with this Access profile.

Proceed to the RADIUS Accounting settings for EX Switching Access profiles by clicking either **Accounting Settings** or **Next**. These settings are described in [“Specifying RADIUS Accounting Settings for an EX Switching Access Profile” on page 224](#).

Specifying RADIUS Accounting Settings for an EX Switching Access Profile

Configure the settings listed in [Table 60 on page 225](#) for the Access profile Accounting Settings page. Accounting settings are optional in an Access profile. You can also specify accounting settings later by modifying an existing Access profile.

Table 60: Accounting Settings for an EX Switching Access Profile

Task	Description
View	Select a RADIUS server entry from the list and then click View to see the details of that entry.
Create a new RADIUS server for both authentication and accounting	<p>To both create and add a RADIUS server configuration to this Access profile for both authentication and accounting:</p> <p>NOTE: A RADIUS profile must be configured for authentication in addition to accounting.</p> <ol style="list-style-type: none"> Click Add > Create RADIUS. The Add RADIUS Server window opens. Complete these settings: <ul style="list-style-type: none"> Server Name—Type the name of the RADIUS server that you want to create. Server Address—Type the IP address of the RADIUS server. Authentication Port—The default RADIUS authentication port is 1812. You can change the port number by using the up and down arrows. Secret—Provide a password. If the password contains spaces, enclose it in quotation marks. The secret password used by the switch must match that used by the server. Expand the Advanced Settings section and change any default settings, including the accounting port: NOTE: If you do not change the accounting configuration, default values are used. <ul style="list-style-type: none"> Accounting Port—The default RADIUS accounting port is 1813. You can change the port number by using the up and down arrows. Retry Count—Specify the number of times that a device attempts to contact the RADIUS server. The default retry count is 3. You can change this value by using the up and down arrows to 1 through 10 times. Timeout (seconds)—Specify the number of seconds the switch waits to receive a response from the RADIUS server. The default timeout is 5 seconds. You can change this value, using the up and down arrows, to 1 through 90 seconds. Click OK. The RADIUS server is automatically added to the list of RADIUS accounting servers assigned to this Access profile. If you have more than one RADIUS server listed, you can use the arrows to reorder the list priority so that the most preferred RADIUS server is listed first.

Table 60: Accounting Settings for an EX Switching Access Profile (*continued*)

Task	Description
Add a previously configured RADIUS server for accounting	<p>A RADIUS server must already be configured before you can add that server for accounting. If the server was previously configured only for authentication, default accounting settings are applied. To add a RADIUS server for accounting:</p> <ol style="list-style-type: none"> 1. Expand the Accounting Settings section of the Server Settings page. This is where RADIUS accounting is configured. A list of configured RADIUS servers is displayed. 2. Click Add > Select RADIUS. A list of eligible RADIUS servers is displayed. Servers on this list were either automatically discovered, created following the directions “Creating and Managing RADIUS Profiles” on page 206, or created on this page following the directions <i>Create and add a new RADIUS server configuration</i>. If the server was configured only for authentication, default accounting settings were applied—you can use those default settings. 3. Select a RADIUS server from the list of Available servers and then use the arrows to move it to the list of Selected servers. 4. Click OK. The RADIUS server is added to the list of accounting servers to be used with this Access profile. If the RADIUS server was previously configured only for authentication, default accounting settings are applied. 5. If you have more than one RADIUS server listed, you can use the arrows to reorder the login priority so that the most preferred RADIUS server is listed first for accounting.
Delete a server	<p>To delete a server from this Access profile:</p> <ol style="list-style-type: none"> 1. Select a server from the list. 2. Click Delete. The server is removed from the list of servers to be used with this Access profile.

Proceed to the Access profile review by clicking either **Review** or **Next**.

Specifying Basic Settings for a Campus Switching ELS Access Profile

To configure the basic settings for a Campus Switching ELS Access profile:

1. Complete the basic settings and authentication order on the Create Access Profile for Campus Switching ELS page, as described in both the online help and in [Table 61 on page 227](#). Required settings are indicated by a red asterisk (*) that appears next to the field label in the user interface.

Table 61: Access Profile Basic Settings for Campus Switching ELS

Field	Action
Access Profile Details	
Profile Name	Type a unique name that identifies the profile. Profile name must not contain special characters or spaces. Note that profiles that are automatically created by Network Director as part of device discovery or out-of-band changes might contain the underscore (_) character.
Description	Type the description of the profile.
Authentication Order	
Server settings depend on which authentication is done first, RADIUS or LDAP.	
Authentication Order	Indicate whether to authenticate first with configured RADIUS servers or with configured LDAP servers by selecting the method from <i>Based On</i> . By default, RADIUS authentication using no password is selected for initial authentication. You can change this to RADIUS authentication with a password by selecting Password . Select LDAP to authenticate first with configured LDAP servers. TIP: LDAP is not supported for EX Switching devices.

Proceed to the Server Settings for Campus Switching ELS Access profiles by clicking either **Server Settings** or **Next**. The settings are described in [“Specifying RADIUS and LDAP Settings for Campus Switching ELS” on page 227](#).

Specifying RADIUS and LDAP Settings for Campus Switching ELS

Configure either a RADIUS server, an LDAP server, or both, on the Server Settings page. A RADIUS server can provide both user accounting services and user authentication but you must be using the RADIUS server for authentication in order to use it for accounting. An LDAP server provides only user authentication. The server settings in this section determine the options used for the access servers in this Access profile.

Configure the Server settings for a Campus Switching ELS Access profile by following the directions in [Table 62 on page 228](#).

Table 62: Authentication and Accounting Server Settings for ELS Campus Switching

Task	Action
AAA: Authentication Server	
RADIUS servers are selected for configuration by default. RADIUS servers can do both authentication and accounting.	
View configured servers in this profile	Select a server entry from the list and then click View to see the details of that entry.

Table 62: Authentication and Accounting Server Settings for ELS Campus Switching (*continued*)

Task	Action
Create and add a new RADIUS server for authentication	<p>The RADIUS tab is selected by default for AAA Authentication Server configuration. To configure a RADIUS accounting server and add it to this Access profile:</p> <ol style="list-style-type: none"> 1. Click Add > Create RADIUS on the RADIUS tab. The Create RADIUS Server window opens. 2. Provide the following RADIUS authentication server information: <ul style="list-style-type: none"> • Server Name • Server Address • Authentication Port—The default RADIUS authentication port is 1812. You can change the port number by using the up and down arrows. • Secret—Provide the authentication secret password. If the password contains spaces, enclose it in quotation marks. The secret password used by the local router must match the one used by the server. 3. Optionally, expand the Advanced Settings for a RADIUS server and change any of these configurations: <ul style="list-style-type: none"> • Accounting Port—You can change the default accounting port number (1813) by using the up and down arrows. • Retry Count—Specify the number of times that a device attempts to contact the LDAP authentication server. The default retry count is 3. You can change this value by using the up and down arrows to 1 through 10 times. • Timeout (seconds)—Specify the number of seconds the switch waits to receive a response from a RADIUS server. The default timeout is 5 seconds. You can change this value, using the up and down arrows, to 1 through 90 seconds. 4. Click OK. The Create RADIUS Server window closes and the RADIUS server is automatically added to the list of RADIUS servers assigned to this Access profile. 5. If you have more than one RADIUS server listed, you can use the arrows to reorder the list priority so that the most preferred RADIUS server is listed first.

Table 62: Authentication and Accounting Server Settings for ELS Campus Switching (continued)

Task	Action
Add a previously configured RADIUS server for authentication	<p>The RADIUS tab is selected by default for server configuration and configured RADIUS servers are listed on this Server Settings page. To add a previously configured RADIUS server to this Access profile:</p> <ol style="list-style-type: none"> 1. Click Add > Select RADIUS on the RADIUS tab. <p>The Select RADIUS Server window opens, displaying a list of available RADIUS servers is displayed. Servers on this list were either automatically discovered, created following the directions “Creating and Managing RADIUS Profiles” on page 206, or created on this page following the directions in <i>Create and add a new RADIUS server configuration</i>.</p> 2. Select one or more RADIUS servers from the list of previously configured RADIUS servers. 3. Click OK. <p>The Select RADIUS Server window closes and the RADIUS server is added to the list of RADIUS authentication servers to be used with this Access profile.</p> 4. Optionally, if you have more than one RADIUS server listed, use the arrows to reorder the login priority so that the most preferred RADIUS server is listed first.

Table 62: Authentication and Accounting Server Settings for ELS Campus Switching (*continued*)

Task	Action
Add a previously configured RADIUS server for accounting	<p>A RADIUS server can provide both authentication and accounting. To configure accounting settings for a RADIUS server:</p> <p>TIP: In order to provide accounting, authentication must also be configured.</p> <ol style="list-style-type: none"> 1. Expand the RADIUS Accounting Servers section of the Server Settings. A list of RADIUS servers configured for accounting is displayed. 2. Click Add > Select RADIUS. The Select RADIUS Server window opens, displaying a list of eligible RADIUS servers is displayed. Servers on this list were either automatically discovered, created following the directions “Creating and Managing RADIUS Profiles” on page 206, or created on this page following the directions <i>Create and add a new RADIUS server configuration</i>. 3. Select one or more RADIUS servers from the list of previously configured RADIUS servers. 4. Click OK. The Select RADIUS Server window closes and the RADIUS server is added to the list of RADIUS Accounting Servers to be used with this Access profile. 5. Optionally, if you have more than one RADIUS server listed, use the arrows to reorder the login priority so that the most preferred RADIUS server is listed first.

Table 62: Authentication and Accounting Server Settings for ELS Campus Switching (*continued*)

Task	Action
Create and add a new RADIUS server for both authentication and accounting	<p>RADIUS is the only server selection available for accounting. To configure a RADIUS server for both authentication and accounting, and add it to this Access profile:</p> <ol style="list-style-type: none"> Under RADIUS Accounting Server, click Add > Create RADIUS. The Create RADIUS Server window opens. Provide the following RADIUS authentication server information: <ul style="list-style-type: none"> Server Name Server Address Authentication Port—The default RADIUS authentication port is 1812. You can change the port number by using the up and down arrows. Secret—Provide the authentication secret password. If the password contains spaces, enclose it in quotation marks. The secret password used by the local router must match that used by the server. Expand the Advanced Settings and change any of these configurations: <ul style="list-style-type: none"> Accounting Port—You can change the default port number (1813) by using the up and down arrows. Retry Count—Specify the number of times that a device attempts to contact the LDAP authentication server. The default retry count is 3. You can change this value by using the up and down arrows to 1 through 10 times. Timeout (seconds)—Specify the number of seconds the switch waits to receive a response from a RADIUS server. The default timeout is 5 seconds. You can change this value, using the up and down arrows, to 1 through 90 seconds. Click OK. The Create RADIUS Server window closes and the RADIUS server is automatically added to the list of RADIUS Accounting Servers assigned to this Access profile. If you have more than one RADIUS accounting server listed, you can use the arrows to reorder the list priority so that the most preferred RADIUS server is listed first.

Table 62: Authentication and Accounting Server Settings for ELS Campus Switching (*continued*)

Task	Action
Create and add a new LDAP authentication server	<p>TIP: LDAP servers can be configured for Campus Switching ELS.</p> <p>To configure a new LDAP authentication server and add it to this Access profile:</p> <ol style="list-style-type: none"> 1. Click the LDAP tab to display the LDAP settings. 2. Provide a Base Distinguished Name for the LDAP server. LDAP APIs reference an LDAP object by its distinguished name (DN), which is a sequence of relative distinguished names (RDN) connected by commas—for example, DC=eng, DC=Juniper Networks, DC=com. You can do an LDAP query to determine the DN for the LDAP server. 3. Click Add > Create LDAP. The Create LDAP Server window opens. 4. Provide the following LDAP server information: <ul style="list-style-type: none"> • Server Name • Server Address • Server Port—The default LDAP server port is 389. You can change the port number by using the up and down arrows. 5. Optionally provide the following Advanced LDAP server information after expanding the Advanced Settings section: <ul style="list-style-type: none"> • Timeout (seconds)—Specify the number of seconds the switch waits to receive a response from a RADIUS server. The default timeout is 5 seconds. You can change this value, using the up and down arrows, to 1 through 90 seconds. • Retry—Specify the number of times that a device attempts to contact the LDAP authentication server. The default retry count is 5. You can change this value by using the up and down arrows to 1 through 10 times. 6. Click OK. The Create LDAP Server window closes and the LDAP server is added to the list of LDAP servers.

Table 62: Authentication and Accounting Server Settings for ELS Campus Switching (*continued*)

Task	Action
Add a previously configured LDAP server for authentication	<p>TIP: LDAP servers can be configured for Campus Switching ELS.</p> <p>To add a previously configured LDAP authentication server to this Access profile:</p> <ol style="list-style-type: none"> 1. Click the LDAP tab to display the LDAP settings. 2. Provide a Base Distinguished name for the LDAP server. LDAP APIs reference an LDAP object by its distinguished name (DN), which is a sequence of relative distinguished names (RDN) connected by commas. You can do an LDAP query to determine the DN for the LDAP server. 3. Click Add > Select LDAP. The Select LDAP Server window opens, displaying a list of configured LDAP servers displayed. Servers on this list were either automatically discovered, or created following the directions “Creating and Managing LDAP Profiles” on page 212, or created by clicking Add > Create LDAP on this page. 4. Select one or more LDAP servers from the list. 5. Click OK. The Select LDAP Server window closes and selected LDAP servers are added to the list of LDAP authentication servers to be used with this Access profile. 6. Optionally, use the arrows to reorder the LDAP servers so that the most preferred LDAP server is listed first. <p>TIP: LDAP is not supported for EX Switching devices.</p>
Delete a server	<p>To delete any server from this Access profile:</p> <ol style="list-style-type: none"> 1. Select a server from the list. 2. Click Delete. The server is removed from the list of servers to be used with this Access profile.

Reviewing and Modifying the Access Profile Settings

From this page, you can save or make changes to a Access profile:

- To make changes to the profile, click **Edit** associated with the configuration to be changed.

Alternatively, you can click the appropriate sections in the profile workflow at the top of the page that corresponds to the configuration to be changed.

When you are finished with your modifications, click **Review** to return to this page.

- To save a new profile or to save modified settings to an existing profile, click **Finish**.

You will be returned to the Manage Access Profiles page. Your new or modified Access profile is listed in the table of Access profiles.

What To Do Next

After you create an Access profile, you can do one of the following:

- For switching devices, configure Access profile as a attribute while assigning Port profiles to interfaces. For more information see [“Creating and Managing Port Profiles” on page 251](#).

RELATED DOCUMENTATION

[Understanding Access Profiles | 217](#)

[Creating and Managing LDAP Profiles | 212](#)

[Creating and Managing RADIUS Profiles | 206](#)

[Creating and Managing Authentication Profiles | 237](#)

[Creating and Managing Port Profiles | 251](#)

[Network Director Documentation home page](#)

Understanding Authentication Profiles

IN THIS SECTION

- [802.1X Authentication | 236](#)
- [MAC RADIUS Authentication | 236](#)
- [Captive Portal Authentication | 237](#)

Authentication profiles include the authentication method and authentication parameters to be used for client authentication. Available authentication methods are 802.1X (dot1x), MAC-RADIUS, captive portal, and last-resort. 802.1X is the default authentication method for all device types but you can change this or add additional authentication types. If you configure multiple authentication methods on a single interface, the system tries the first method listed and then falls back to another method if the first method is unsuccessful.

You can create one or more Authentication profiles to specify different authentication methods based on client devices or sessions.

Each Authentication profile is specific to a device family. After you create an Authentication profile, you can include it in a Port profile. The Authentication profile specified in a Port profile is used to authenticate all the users and devices that connect to the port.

802.1X Authentication

Newer equipment supports the IEE standard called 802.1X. 802.1X is basically an Enterprise, per-user (username and password) authentication mechanism – it is both the newest and strongest authentication you can use. Since 802.1X authentication is the most secure authentication option, it is preferable to the older PSK authentication, Web Portals, MAC authentication, or open authentication, which really means no authentication.

802.1X authentication involves three entities, a supplicant, an authenticator, and an authentication server. The supplicant is a client device, such as a laptop, that wishes to attach to a network. The authenticator would be a switch. The authentication server is usually a RADIUS server, which can interpret 802.1X EAP modes.

- *Single supplicant mode* authenticates only the first end device that connects to an authenticator port. All other end devices connecting to the authenticator port after the first has connected successfully, whether they are 802.1X-enabled or not, are permitted free access to the port without further authentication. If the first authenticated end device logs out, all other end devices are locked out until an end device authenticates.
- *Single-secure supplicant mode* authenticates only one end device to connect to an authenticator port. No other end device can connect to the authenticator port until the first logs out.
- *Multiple supplicant mode* authenticates multiple end devices individually on one authenticator port. If you configure a maximum number of devices that can be connected to a port through port security, the lesser of the configured values is used to determine the maximum number of end devices allowed per port.

MAC RADIUS Authentication

A media access control address (MAC address) is a unique identifier assigned to network interfaces for communications on the physical network segment. MAC addresses are used as a network address for most

IEEE 802 network technologies, including Ethernet. A client's MAC address can be used for authentication by mapping a password to the client's entry in the MAC address table. MAC authentication can be done either locally or with a RADIUS server.

Captive Portal Authentication

Captive Portals are frequently used to authenticate hotspots, forcing all users to use the configured logon web page. Many companies use captive portals to authenticate guest users for temporary use of the company network. The Captive Portal has one password for all users, which should be changed frequently.

RELATED DOCUMENTATION

[Creating and Managing Authentication Profiles | 237](#)

[Network Director Documentation home page](#)

Creating and Managing Authentication Profiles

IN THIS SECTION

- [Managing Authentication Profiles | 238](#)
- [Creating an Authentication Profile | 239](#)
- [Specifying Authentication Settings for Switches | 240](#)
- [What To Do Next | 245](#)

Use the Manage Authentication Profiles page to create new Authentication profiles and manage existing Authentication profiles.

To display the Manage Authentication Profiles page: In Build mode, select Authentication from Profile and Configuration Management in the Tasks pane. The Manage Authentication Profiles page appears.

This topic describes:

Managing Authentication Profiles

From the Manage Authentication Profiles page, you can:

- Create a new Authentication profile by clicking **Add**. For directions, see [“Creating an Authentication Profile” on page 239](#).
- Modify an existing profile by selecting it and clicking **Edit**.
- View information about a profile, including the interfaces it is associated with, by clicking the profile name or by selecting the profile and clicking **Details**.
- Delete an Authentication profile by selecting a profile and clicking **Delete**.

TIP: You cannot delete profiles that are in use—that is, assigned to objects or used by other profiles. To see the current assignments for a profile, select the profile and click **Details**.

- Clone a profile by selecting a profile and clicking **Clone**.

[Table 63 on page 238](#) describes the information provided about Authentication profiles on the Manage Authentication Profiles page. This page lists all Authentication profiles defined for your network, regardless of the scope you selected in the network view.

Table 63: Manage Authentication Profile Fields

Field	Description
Profile Name	Name given to the profile when the profile was created.
Family Type	The device family on which the profile was created.
Description	Description of the profile that was entered when the profile was created. TIP: To display the entire description, you might need to resize the Description column by clicking the column border in the heading and dragging it.
Creation Time	Date and time when this profile was created.
Update Time	Date and time when this profile was last modified.
User Name	The username of the user who created or modified the profile.

TIP: All columns might not be displayed. To show or hide fields in the Manage Authentication Profiles table, click the down arrow on the field header, select **Columns**, and select or clear the check box adjacent to the field that you want to show or hide.

Creating an Authentication Profile


In Network Director, you can create an Authentication profile to configure methods to be used to authenticate users. You can also specify details about the accounting servers to be used for accounting purposes.

For an Authentication profile, you must specify the following:

- A profile name
- At least one access rule

After you create an Authentication profile, you can include it in a Port profile. The Authentication profile specified in a Port profile acts as the default profile for all the users and devices that connect to the port.

To create an Authentication profile:

1. Click  in the Network Director banner.
2. Under Select View, select either **Logical View**, **Location View**, **Device View** or **Custom Group View**.

TIP: Do not select **Dashboard View** or **Topology View**.

3. From the Tasks pane, select the type of network (Wired), the appropriate functional area (System or AAA), and select the name of the profile that you want to create. For example, to create a port profile for a wired device, click **Wired > Profiles > Port**. The Manage Profile page opens.
4. Click **Add** to add a new profile.

If you chose to create a profile for the wired network, Network Director opens the Device Family Chooser window.

- a. From the Device Family Chooser, select the device family for which you want to create a profile. The available device families are **Switching (EX)**, **Campus Switching ELS** (Enhanced Layer 2 Software), and **Data Center Switching ELS**.

b. Click **OK**.

The Create Authentication Profile page for the selected device family is displayed.

5. Specify authentication settings by doing one of the following:
- For EX Series switches, Campus Switching Enhanced Layer 2 Software, specify the settings as described in [“Specifying Authentication Settings for Switches” on page 240](#).
6. Click **Done** to save the Authentication profile.

The system saves the Authentication profile and displays the Manage Authentication Profiles page. Your new or modified Authentication profile is listed in the table of Authentication profiles.

Specifying Authentication Settings for Switches

To configure an Authentication profile for switching devices, enter the Create Authentication Profile page settings described in [Table 64 on page 240](#) for creating Authentication profiles on switches. Required settings are indicated by a red asterisk (*) that appears next to the field label in the user interface.

Table 64: Authentication Profile Settings for Switches

Field	Action
Profile Name	Type the name of the profile. You can use up to 64 characters for profiles created for wired devices. Profile name must not contain special characters or spaces. Note that profiles that are automatically created by Network Director as part of device discovery or out-of-band changes may contain the underscore (_) character.
Description	Type a short description for the profile.
802.1X Authenticator	
Enable 802.1X	802.1X authentication is enabled by default for a switching profile. 802.1X authentication works by using an Authenticator Port Access Entity (the switch) to block all traffic to and from a supplicant (end device) at the port until the supplicant's credentials are presented and matched on the Authentication server (a RADIUS server). When authenticated, the switch stops blocking traffic and opens the port to the supplicant. Network access can be further defined using VLANs. NOTE: If you disable 802.1X authentication, several related settings become unavailable.

Table 64: Authentication Profile Settings for Switches (*continued*)

Field	Action
Enable MAC-RADIUS	<p>Select to enable MAC-RADIUS based authentication for this profile. MAC RADIUS authentication enables LAN access to permitted MAC addresses. When a new MAC address appears on an interface, the switch consults the RADIUS server to check whether the MAC address is a permitted address. If the MAC address is configured on the RADIUS server, the device is allowed access to the LAN.</p> <p>TIP: You can combine 802.1X and MAC-RADIUS authentication.</p>
Supplicant Mode	<p>Specify the mode authentication supplicants use, either Single, Multiple, or Single-Secure.</p> <ul style="list-style-type: none"> • Single—Allows only one host for authentication. • Single-Secure—Allows only one end device to connect to the port. No other end device is enabled to connect until the first logs out. • Multiple—Allows multiple hosts for authentication. Each host is checked before being admitted to the network.
Guest VLAN	<p>Click Select and then select the VLAN to which an interface is moved when no 802.1X supplicants are connected on the interface. The VLAN specified must already exist on the switch.</p>
Reject VLAN	<p>Click Select and then select the VLAN to which an interface is moved when the switch receives an Extensible Authentication Protocol Over LAN (EAPoL) Access-Reject message during the authentication process between the switch and the RADIUS authentication server.</p>
Server Fail Type	<p>Specify the server fail fallback action the switch takes when all RADIUS authentication servers are unreachable, either None, Deny, Permit, Use cache, or VLAN Name.</p> <ul style="list-style-type: none"> • Deny—Force fail the supplicant authentication. No traffic will flow through the interface. • Permit—Force succeed the supplicant authentication. Traffic will flow through the interface as if it were successfully authenticated by the RADIUS server. • Use cache—Force succeed the supplicant authentication only if it was previously authenticated successfully. This action ensures that already authenticated supplicants are not affected. • VLAN Name—Move supplicant on the interface to the VLAN specified by this name. This action is allowed only if it is the first supplicant connecting to an interface. If an authenticated supplicant is already connected, then the supplicant is not moved to the VLAN and is not authenticated. If you select this option, also provide a Fail VLAN name.

Table 64: Authentication Profile Settings for Switches (*continued*)

Field	Action
Captive Portal	
A Captive Portal is a special web page used for authentication by turning a web browser into an authentication mechanism.	
Enable Captive-Portal	Enable this option to display the captive portal setting for supplicant mode. When this option is enabled, additional captive portal settings are also available under Advanced Settings.
Supplicant Mode	<p>Specify the mode to be used for Captive Portal supplicants, either Single, Multiple, or Single-Secure.</p> <ul style="list-style-type: none"> • Single—Allows only one host for authentication. • Multiple—Allows multiple hosts for authentication. Each host is checked before being admitted to the network. • Single-Secure —Allows only one end device to connect to the port. No other end device is allowed to connect until the first logs out.

To skip configuring the advanced settings and accept the default settings, click **Done**. You can now link the Authentication profile to a Port profile. For directions, see [“Creating and Managing Port Profiles” on page 251](#).

To configure advanced switch settings, click **Advanced Settings** and enter the Advanced Settings described in [Table 65 on page 242](#).

Table 65: Authentication Profile Advanced Settings for Switches

Field	Action
802.1X Settings	
These settings are available only when 802.1X authentication is enabled for this Authentication profile. You can use the default settings or you can change them.	
Transmit Period (default is 30 seconds)	Specify how long, in seconds, the interface waits before retransmitting the initial EAPOL PDUs to the supplicant. The default is 30 seconds.
Maximum Requests (default is 2 requests)	Specify the maximum number of times an EAPOL request packet is transmitted to the supplicant before the authentication session times out. The default is 2 requests.

Table 65: Authentication Profile Advanced Settings for Switches (*continued*)

Field	Action
Retries (default is 3 retries)	Specify the number of times you want the switch to attempt to authenticate the port after an initial failure. The port remains in a wait state during the quiet period after the authentication attempt. The default is 3 retries.
Quiet Period (default is 60 seconds)	Specify the number of seconds the interface remains in the wait state following a failed authentication attempt by a supplicant before reattempting authentication. The default is 60 seconds.
No Reauthentication (default is unselected)	Select this check box if you do not want the switch to reauthenticate the supplicant after the Quiet Period elapses.
Reauthentication Interval (default is 3600 seconds)	If the No Reauthentication option is not checked, specify the number of seconds after which the authentication session times out. The default is 3600 seconds.
Supplicant Timeout (default is 30 seconds)	Specify how long the port waits for a response when relaying a request from the authentication server to the supplicant before resending the request. The default is 30 seconds.
RADIUS Server Timeout (default is 30 seconds)	Specify the length of time that the switch waits for a response from the RADIUS server. The default is 30 seconds.
MAC Restrict (Switches using MAC RADIUS only)	<p>When MAC-RADIUS is enabled in this Authentication profile, select this option to restrict authentication to MAC RADIUS only. When MAC-RADIUS restrict is configured, the switch drops all 802.1X packets. This option is useful when no other 802.1X authentication methods, such as guest VLAN, are needed on the interface, and eliminates the delay that occurs while the switch determines that a connected device is a non-802.1X-enabled host.</p> <p>Optionally enable Flap-On-Disconnect. When the RADIUS server sends a disconnect message to a supplicant, the switch resets the interface on which the supplicant is authenticated. If the interface is configured for multiple supplicant mode, the switch resets all the supplicants on the specified interface. This option takes effect only when the MAC Restrict option is also set.</p>

Captive Portal

If Captive Portal is enabled in this Authentication profile in the basic settings, you can either use the default advanced Captive Portal settings or change them as indicated.

Table 65: Authentication Profile Advanced Settings for Switches (*continued*)

Field	Action
Quiet Period (default is 60 seconds)	<p>Configure the time, in seconds, between when a user exceeds the maximum number of retries and when they can again attempt to authenticate.</p> <p>Range: 1 through 65,535</p> <p>Default: 60</p>
Retries (default is 3 retries)	<p>Configure the number of times the user can attempt to submit authentication information.</p> <p>Range: 1 through 65,535</p> <p>Default: 3</p>
Session Expiry (default is 3600 seconds)	<p>Configure the maximum duration in seconds of a session.</p> <p>Range: 1 through 65,535</p> <p>Default: 3600</p>
Server Time Out (default is 30 seconds)	<p>Configure the time in seconds an interface will wait for a reply when relaying a response from the client to the authentication server before timing out and invoking the server-fail action.</p> <p>Range: 1 through 65,535</p> <p>Default: 30</p>

Click **OK**.

The Advanced Settings window closes and you once again see the Create Authentication Profile for Switching page.

Click **Done**.

The Manage Authentication Profiles page reappears with your new Authentication profile listed.

You can now link the Authentication profile to a Port profile. For more details, see [“Creating and Managing Port Profiles” on page 251](#).

What To Do Next

After you create an Authentication profile, you can do the following:

- For switching devices, link the Authentication profile to a Port profile. For more details, see [“Creating and Managing Port Profiles” on page 251](#).

RELATED DOCUMENTATION

[Understanding Authentication Profiles | 235](#)

[Creating and Managing Port Profiles | 251](#)

[Network Director Documentation home page](#)

Configuring Interfaces and VLANs

IN THIS CHAPTER

- Understanding Port Profiles | 246
- Creating and Managing Port Profiles | 251
- Assigning and Unassigning Port Profiles from Interfaces | 304
- Managing Auto Assignment Policies | 312
- Creating Auto Assignments | 314
- Configuring Easy Config Setup | 317
- Understanding Port Groups | 323
- Creating and Managing Port Groups | 323
- Understanding VLAN Profiles | 327
- Creating and Managing VLAN Profiles | 329
- Assigning a VLAN Profile to Devices or Ports | 345

Understanding Port Profiles

IN THIS SECTION

- Interface Settings Configured in the Port Profile | 247
- Interface Settings Configured by Referencing Other Profiles | 248
- Default Port Profiles | 248

Port profiles provide a convenient way of provisioning interfaces on switches. You can either use predefined port profiles, or you can define your own custom port profile.

After you create a Port profile, you can assign it to interfaces on one or more switches, including aggregated interfaces. For the configuration created by the profile to take effect on the devices, you must use Deploy mode to deploy the configuration on the devices.

This topic describes:

Interface Settings Configured in the Port Profile

You can configure the following interfaces settings in a Port profile:

- Interface protocol family—You can configure an interface to be either an Ethernet switching interface, an IPv4 routing interface, or an IPv6 routing interface.
- Port mode—You can configure a switching interface port mode to be an access, trunk, or tagged-access interface for EX Series switches. Campus Switching ELS supports access mode and trunk mode. For more information about port modes, see [Ethernet Switching](#).
- PoE settings—The factory-default configuration of switches enables PoE on all interfaces that support PoE. For many implementations, no further configuration is necessary. You can, however, override the default settings for PoE interfaces in the Port profile. Most switch models have interfaces that support Power over Ethernet (PoE), but the EX9200 does not support PoE. For more information about PoE, see [Power over Ethernet \(PoE\)](#).

If you do not explicitly configure PoE in the Port profile, the existing PoE interface settings on the switch remain in effect. Device-wide PoE settings are configured in the Device Common Settings profile.

- Physical link settings—On switches, the autonegotiation of port speed and duplex mode is enabled by default. You can disable autonegotiation and set port speed and duplex mode in the Port profile. Other link settings you can configure include flow control, which is disabled by default, and maximum transmission unit (MTU).
- Storm control settings—You can optionally enable storm control settings on switches. Storm control monitors traffic levels drops broadcast, multicast, and unknown unicast packets when a specified traffic level—called the storm control level.
- RSTP settings—You can optionally enable RSTP settings on switches. RSTP this feature enables you to fine-tune STP by setting interfaces into edge, disable, or no-root-port states.
- Port security settings—You can optionally enable port security on switched access ports. Port security features help protect the access ports on your switch against address spoofing (forging) and Layer 2 denial-of-service (DoS) attacks. For more information about port security on switches, see [Port Security](#).

NOTE: For campus switching ELS devices, disabling port security settings option is not available.

Interface Settings Configured by Referencing Other Profiles

You can optionally configure other interface-related settings in the Port profile by referencing other profiles. These profiles are:

- CoS profile—Configures class-of-service settings on the interface. You can either select or create an in-line CoS profile while creating a port profile.
- Filter profile—Configures firewall filters (often called ACLs) on the interface.
- Authentication profile—(Switching interfaces only) Configures 802.1X authentication on an interface and configures related settings, such as captive portal authentication. You can either select or create an in-line authentication profile while creating a port profile.
- Access profile—Configures the access server settings used by all 802.1X authenticator interfaces on a switch. This profile is not included in the Port profile. Instead, you assign it to a device as part of the process of assigning a Port profile to the interfaces on the device.
- VLAN profile for Campus Switching ELS is mandatory. You can either select or create an in-line VLAN profile while creating a port profile.

If you want to use one or more of these profiles with the Port profile, be sure to create them before you create and assign the Port profile.

Default Port Profiles

To help with the rapid provisioning of interfaces on switches, Network Director provides default Port profiles that contain settings for common uses of switch interfaces. You can modify or assign these default profiles to interfaces using the same method used for user-created profiles. [Table 66 on page 249](#) describes the default Port profiles.

Table 66: Default Port Profiles

Profile Name	Description	Summary of Settings
Desktop_Port	Configures an untagged port that connects to desktop computer.	<ul style="list-style-type: none"> • Family Type—switching • Port Mode—access • Auto Negotiation—disabled • Flow Control—disabled • Maximum Bytes—disabled • Speed—no default provided • Link Mode—no default provided • Trust DHCP—no • MAC Limit—1 • MAC Limit Action—drop • CoS Profile—no default provided
Desktop_Phone_Port	Configures an untagged port that connects to a combined desktop and phone port.	<ul style="list-style-type: none"> • Family Type—switching • Port Mode—access • Auto Negotiation—disabled • Flow Control—disabled • Maximum Bytes—disabled • Speed—no default provided • Link Mode—no default provided • Trust DHCP—no • MAC Limit—2 • MAC Limit Action—drop • CoS Profile—juniper_CoS_template
Server_Port	Configures a tagged port that connects to a server.	<ul style="list-style-type: none"> • References the default CoS profile, juniper_CoS_template • Sets protocol family to Ethernet switching • Sets port mode to trunk • Enables port security with trust DHCP enabled

Table 66: Default Port Profiles (*continued*)

Profile Name	Description	Summary of Settings
Switched_Downlink	Configures a tagged port that connects to endpoint devices in a branch environment.	<ul style="list-style-type: none"> • Family Type—switching • Port Mode—trunk • Auto Negotiation—disabled • Flow Control—disabled • Maximum Bytes—no default provided • Speed—no default provided • Link Mode—no default provided • Trust DHCP—yes • MAC Limit—no default provided • MAC Limit Action—no default provided • CoS Profile—juniper_CoS_template
Switched_Uplink	Configures a tagged port that connects a switch to another switch or larger network. For example, a port that connects an access switch to an aggregation switch.	<ul style="list-style-type: none"> • Family Type—switching • Port Mode—trunk • Auto Negotiation—disabled • Flow Control—disabled • Maximum Bytes—no default provided • Speed—no default provided • Link Mode—no default provided • Trust DHCP—yes • MAC Limit—no default provided • MAC Limit Action—no default provided • CoS Profile—juniper_CoS_template

RELATED DOCUMENTATION

[Creating and Managing Port Profiles | 251](#)
[Understanding Access Profiles | 217](#)
[Understanding Authentication Profiles | 235](#)
[Understanding Class of Service \(CoS\) Profiles | 396](#)
[Understanding Filter Profiles | 349](#)
[Network Director Documentation home page](#)

Creating and Managing Port Profiles

IN THIS SECTION

- [Managing Port Profiles | 252](#)
- [Creating Port Profiles | 254](#)
- [Specifying Settings for an EX Switching Port Profile | 255](#)
- [Specifying Settings for a Campus Switching ELS Port Profile | 269](#)
- [Specifying Settings for a Data Center Switching ELS Port Profile | 285](#)
- [What to Do Next | 303](#)

Port profiles provide a way to provision multiple switch interfaces, including Ethernet interfaces on EX Series switches and Campus Switching ELS. In a Port profile, you can define a set of attributes to be shared by multiple interfaces. For example, you can create a Port profile for all access interfaces that connect to VoIP desk phones, configuring the appropriate class-of-service (CoS), authentication, and port security settings for these interfaces in the Port profile. You then assign the Port profile to those interfaces and deploy the resulting configuration on the interfaces.

Port profiles define only shared attributes. To enable you to configure specific attributes for an interface or a switch during the process of assigning a Port profile to an interface, the Create Port profile wizard provides two setup options: Quick Setup and Custom Setup. The Quick Setup option enables you to create initial configuration settings for a Port profile including selecting or create inline VLAN profile. The Custom Setup option enables you to configure all the advanced settings and create any inline sub-profiles. In Custom Setup option, apart from selecting the existing VLAN, CoS, and authentication sub-profiles, you can also create these sub-profiles.



CAUTION: Ports that are involved in EVPN-VXLAN is not configured through port profile. Else, it renders EVPN-VXLAN defunct.

NOTE: If you switch from Quick Setup to Custom Setup, all the configuration settings are saved. However, if you switch from Custom Setup to Quick Setup, all the advanced settings done in the Custom Setup are lost.

To manage or create Port profiles: In Build mode, select **Port** from Profiles in the Tasks pane. The Manage Port Profile page appears.

This topic describes:

Managing Port Profiles

Use the Manage Port Profiles page to manage existing Port profiles and to create new ones. Port profiles enable the definition and application of a common set of attributes to interfaces.

From the Manage Port Profiles page, you can:

- Create a new profile by clicking **Add**. For details, see [“Creating Port Profiles” on page 254](#).
- Modify an existing profile by selecting it and clicking **Edit**.
- Associate a Port profile to specific interfaces by selecting it and clicking **Assign**.

During the assignment process, you can choose to configure interface-specific settings, such as IP address.

- Change a Port profile's current interface assignments by selecting it and clicking **Edit Assignments**. This opens the Edit assignments for profile-name page, which displays the assignment state and other details of the interfaces in a grid layout. After editing an assignment, and click **Apply**. The Edit Profile Assignment Job Details window opens, which reports the status of the interface assignment that you edited.
- View information about a profile, including the interfaces it is associated with, by selecting the profile and clicking **Details** or by clicking the profile name, which opens the Profiles Details page. This page displays the profile details and the interface associations in a grid layout. It also has an option using which you can search profiles associated with a device and filter devices. Click **Show Filters** to filter an interface based on its IP address, serial number, type, or location or custom group.
- Perform the search for the following:
 - A Port profile for a specific device by specifying the device details in the search field.
 - A port profile that is assigned to a specific port on a device. In this case, you must first enter the device details and then specify the port details in the search field to view the port profile.
 - Port profiles that are assigned to interfaces that are part of the same VLAN. When you specify the VLAN name in the search field, all the Port profiles that are part of the same VLAN are listed in the table.
- Delete profiles by selecting the profiles and clicking **Delete**.

TIP: You cannot delete profiles that are in use—that is, assigned to objects or used by other profiles. To see the current assignments for a profile, click the profile name.

- Clone a profile by selecting a profile and clicking **Clone**.

Network Director provides a set of default Port profiles: Desktop Port, Desktop and Phone Port, Server Port, Switched Downlink, Switched Uplink, and Custom Port. These profiles contain configuration appropriate for the named port role. You can manage these profiles the same way that you manage a user-created profile. For more information about these profiles, see [“Understanding Port Profiles” on page 246](#).

[Table 67 on page 253](#) describes the information provided about Port profiles on the Manage Port Profiles page. This page lists all Port profiles defined for your network, regardless of your current selected scope in the network view.

Table 67: Manage Port Profiles Table

Column	Description
Profile Name	<p>Name given to the profile when the profile was created.</p> <p>Click the profile name to view profile details.</p> <p>A ★ next to the profile name indicates that the profile is assigned to a port using an auto assignment policy. For more details on auto assignment policies, see “Managing Auto Assignment Policies” on page 312.</p>
Family Type	<p>One of the following:</p> <ul style="list-style-type: none"> • EX—for EX Series switches • ELS—for Campus Switching ELS • Data Center Switching ELS—for Data Center Switching ELS devices
Description	Description of the Port profile that was entered when the profile was created.
Port Family	<p>One of the following:</p> <ul style="list-style-type: none"> • Switching—for Port profiles that configure Layer 2 interfaces • Routing—for Port profiles that configure Layer 3 interfaces • FIBRE—for Port profiles that configure Fibre Channel (FC) interfaces.
VLANs	Name of the member VLANs configured or referenced for that Port profile.
Assignment State	<p>One of the following states:</p> <ul style="list-style-type: none"> • Deployed—The profile has been assigned to interfaces and the configuration has been deployed on the devices. • Pending Deployment—The profile has been assigned to interfaces or its previous assignments have been changed, but the new or modified configuration has not yet been deployed on the devices. • Unassigned—The profile has not yet been assigned to interfaces.


Table 67: Manage Port Profiles Table (*continued*)

Column	Description
Creation Time	Date and time when the profile was created.
Last Updated Time	Date and time when the profile was last modified.
Assigned to Devices	Number of devices to which the Port profile is assigned. Click on the link to view the profile details.
Assigned to Port	Number of ports to which the Port profile is assigned. Click on the link to view the profile details.
Assigned to	Number of port assignments and device associations for a profile.
User Name	The username of the user who created or modified the profile.

TIP: All columns might not be currently displayed. To show or hide fields in the table, click the down arrow on the field header, select **Columns**, and select or clear the check box adjacent to the field that you want to show or hide.

Creating Port Profiles

To create a Port profile for EX Series switches, Campus Switching ELS, or Data Center Switching ELS:

1. Click  in the Network Director banner.
2. Under Views, select one of the following views: **Logical View**, **Location View**, **Device View** or **Custom Group**.

TIP: Do not select **Datacenter View** or **Topology View**.

3. Click **Port** under Wired > Profiles in the Tasks pane.

The Manage Port Profile page appears.

4. Click **Add**.

The Select a Device Family page opens.

5. Select **Switching (EX)**, **Campus Switching ELS** or **Data Center Switching ELS**.

The Create Port Profile page appears showing the Quick Setup and Custom Setup tabs for the selected family with the appropriate fields for configuring that family.

6. Select initial settings in the Quick Setup option and advanced settings in the Custom Setup option for the Port profile. For information about the Port profile settings, select the section for the type of port you are configuring:

- [Specifying Settings for an EX Switching Port Profile on page 255](#)
- [Specifying Settings for a Campus Switching ELS Port Profile on page 269](#)
- [Specifying Settings for a Data Center Switching ELS Port Profile on page 285](#)

NOTE: No registered multicast and no unregistered multicast are not applicable for EX Series switches that belong to Switching (EX) device family. For example, EX2200 and EX4200.

Specifying Settings for an EX Switching Port Profile

Use the Create Port Profile page to define a common set of port attributes, which you can then apply to a group of interfaces. These directions address creating a Port profile for EX Series switches.

TIP: You can reference a VLAN profile, CoS profile, Ingress Filter profile, Egress Filter profile, and an Authentication profile in a Port profile. You can either create these profiles in their respective profile pages before you create Port profiles or you can create these profiles as in-line sub-profiles while configuring Port profiles. You can also enable power over Ethernet (PoE).

After you create a Port profile, you assign it to individual interfaces or to members of a port group. During this process, you can also configure interface-specific attributes, such as IP address, and certain device-specific attributes, such as the Access profile to use for all ports on the device. You can assign only one Port profile to an interface.

[Table 68 on page 256](#) describes the Quick Setup settings available in a Port profile. [Table 69 on page 262](#) describes the Custom Setup settings. The defaults for these options depend on the Service Type you select.

Table 68: Port Profile Quick Setup Settings for an EX Switching Port Profile

Field	Action
Profile Name	A default name that corresponds to the Service Type is displayed—when you change the Service Type, this default profile name changes. You can also change the name of profile, using up to 64 alphanumeric characters and no special characters other than the underscore. The name must be unique among Port profiles.
Description	A default description of the preconfigured service types appears by default. You can change the description of the Port profile, which appears on the Manage Port Profiles page. You can use up to 256 characters.

Table 68: Port Profile Quick Setup Settings for an EX Switching Port Profile (*continued*)

Field	Action
Service Type	<p>Select one the preconfigured switching options, Desktop Port, Desktop Phone Port, Printer Port, Switched Uplink, Switched Downlink, or Server Port. To create your own switching or routing service type, select Custom.</p> <p>TIP: No preconfigured routing Service Types are provided. You must create them using the Custom option.</p> <hr/> <p>Desktop Port default service type has the following default settings:</p> <ul style="list-style-type: none"> • CoS Profile—no default provided • Family Type—switching • Port Mode—access • Power over Ethernet—disabled • Auto Negotiation—disabled • Flow Control—disabled • Maximum Size—disabled • Speed—no default provided • Link Mode—no default provided • Port Security—enabled • Trust DHCP—disabled • MAC Limit—1 • MAC Limit Action—drop • Allowed MAC List—no default provided

Table 68: Port Profile Quick Setup Settings for an EX Switching Port Profile (*continued*)

Field	Action
	<p>Desktop Phone Port preconfigured service type has the following default settings:</p> <ul style="list-style-type: none"> • CoS Profile—juniper_CoS_template • Family Type—switching • Port Mode—access • Power over Ethernet—disabled • Auto Negotiation—disabled • Flow Control—disabled • Maximum Size—disabled • Speed—no default provided • Link Mode—no default provided • Port Security—enabled • Trust DHCP—disabled • MAC Limit—2 • MAC Limit Action—drop • Allowed MAC List—no default provided
	<p>Printer Port preconfigured service type has the following default settings:</p> <ul style="list-style-type: none"> • Family Type—switching • Port Mode—access • Power over Ethernet—no default provided • Auto Negotiation—enabled • Flow Control—enabled • Maximum Size—no default provided • Speed—no default provided • Link Mode—no default provided • Port Security—no default provided • Trust DHCP—no default provided • MAC Limit—no default provided • MAC Limit Action—no default provided • Allowed MAC List—no default provided

Table 68: Port Profile Quick Setup Settings for an EX Switching Port Profile (*continued*)

Field	Action
	<p>Switched Uplink preconfigured service type has the following default settings:</p> <ul style="list-style-type: none"> • CoS Profile—juniper_CoS_template • Family Type—switching • Port Mode—trunk • Power over Ethernet—disabled • Auto Negotiation—disabled • Flow Control—disabled • Maximum Size—no default provided • Speed—no default provided • Link Mode—no default provided • Port Security—enabled • MAC Limit—no default provided • Trust DHCP—disabled • MAC Limit Action—no default provided • Allowed MAC List—no default provided
	<p>Switched Downlink preconfigured service type has the following default settings:</p> <ul style="list-style-type: none"> • CoS Profile—juniper_CoS_template • Family Type—switching • Port Mode—trunk • Power over Ethernet—disabled • Auto Negotiation—disabled • Flow Control—disabled • Maximum Size—no default provided • Speed—no default provided • Link Mode—no default provided • MAC Limit—no default provided • Port Security—enabled • Trust DHCP—enabled • MAC Limit Action—no default provided • Allowed MAC list—no default provided

Table 68: Port Profile Quick Setup Settings for an EX Switching Port Profile (*continued*)

Field	Action
	<p>Server Port preconfigured service type has the following default settings:</p> <ul style="list-style-type: none"> • CoS Profile—juniper_CoS_template • Family Type—switching • Port Mode—trunk • Power over Ethernet—disabled • Auto Negotiation—disabled • Flow Control—disabled • Maximum Size—no default provided • Speed—no default provided • Link Mode—no default provided • Port Security—enabled • MAC Limit—no default provided • MAC Limit Action—no default provided • Allowed MAC list—no default provided
Family Type	<p>This setting cannot be changed if any preconfigured Service Type was selected. If you selected the Custom Service Type, indicate whether the interface operates as a Layer 2 (Switching) or a Layer 3 (Routing) interface.</p> <p>TIP: All preconfigured Service Types are for switching.</p> <p>If you select Routing, you configure an IP address on a per-interface basis when you assign the profile to individual interfaces.</p> <p>TIP: Service Type must be set to Custom to configure a routing interface.</p>

Table 68: Port Profile Quick Setup Settings for an EX Switching Port Profile (*continued*)

Field	Action
Port Mode	<p>This setting cannot be changed if any preconfigured Service Type was selected. If you selected the Custom Service Type, select the port mode for the EX Series switching interface, either Access, Trunk, or Tagged Access.</p> <ul style="list-style-type: none"> • Access—Use for interfaces that connect to an end device, such as a desktop computer, an IP telephone, a printer, or a security camera. The interface must belong to a single VLAN. Frames sent and received over the over the interface are untagged Ethernet frames. This is the default for a Desktop Port and Desktop Phone Port . • Trunk—Use for interfaces that connect to a switch or router. Trunk interfaces can belong to more than one VLAN, enabling VLAN traffic to be multiplexed on a single physical interface. The Ethernet frames sent and received over the interface are tagged frames, in which IEEE 802.1Q tagging is used to segregate the traffic from each VLAN. This is the default for Switched Uplink, Switched Downlink, and Server Port. • Tagged Access—Use for access interfaces where VLAN tagging is required, typically when the interface connects to a server running virtual machines using virtual Ethernet port aggregator (VEPA) technology. The traffic generated by the server can contain an aggregation of VLAN packets from different virtual machines on that server, requiring that packets be tagged.

VLAN Options

Available VLAN options depend on the Service Type selected.

Member VLAN (available for Switched Uplink, Switched Downlink, Server Port)	<p>Click All if you want to assign an interface to all the VLANs.</p> <p>This option is enabled when Port Mode is Trunk or TaggedAccess.</p>
Member VLANs (available for Desktop Port, Desktop Phone Port, Switched Uplink, Switched Downlink, Server Port, , Custom Port)	<p>Select a VLAN for the interface by clicking Select, selecting one of the listed filters, and then clicking OK. The VLAN is added to the Member VLANs list.</p> <p>Or</p> <p>Configure a VLAN by clicking Configure VLAN Settings and clicking Create. Enter the VLAN name and ID and click OK.</p>

Table 68: Port Profile Quick Setup Settings for an EX Switching Port Profile (*continued*)

Field	Action
Voice VLAN (available for Desktop Phone Port, Custom Port)	<p>Select a voice VLAN for the interface by clicking Select, selecting one of the listed filters, and then clicking OK. The VLAN is added to the Member VLANs list.</p> <p>Or</p> <p>Configure a VLAN by clicking Configure VLAN Settings and clicking Create. Enter the VLAN Name and ID and click OK.</p>
Native VLAN (available for Switched Uplink, Switched Downlink)	<p>Select a native VLAN for the interface by clicking Select, selecting one of the listed VLANs, and then clicking OK. The VLAN is added to the Member VLANs list.</p> <p>Or</p> <p>Configure a VLAN by clicking Configure VLAN Settings and clicking Create. Enter the VLAN name and ID and click OK.</p>

After providing the information in the fields listed in the preceding, click **Done**.

To use default Port Profile Custom Setup settings, click **Done**. To configure Custom Setup settings, click **Custom Setup** and then provide the information in [Table 69 on page 262](#) and then click **Done**.

Clicking **Done** in either case displays the dialog Do you want to assign Port Profile to Ports. Click **Yes** to create a profile assignment; else click **No** to navigate to the Manage Port Profile page and to create the Port assignment later on.

Table 69: Port Profile Custom Setup Settings

Field	Action
-------	--------

Advanced Settings

Expand Advanced Settings to configure link settings and port security. The Link Setting in Port profile is disabled by default. On enabling Link Settings, autonegotiation and flow control are enabled by default.

Table 69: Port Profile Custom Setup Settings (*continued*)

Field	Action
Enable Auto Negotiation	<p>Autonegotiation of link speed and duplex mode is enabled by default; clear to disable autonegotiation.</p> <p>If you disable autonegotiation, you must set link speed and link mode.</p> <p>You cannot disable autonegotiation if a link speed of 1 Gbps is configured. This configuration might be accepted, but autonegotiation is not disabled.</p> <p>NOTE: This setting is ignored when you assign a Port profile to an Aggregated Ethernet interface.</p>
Enable Flow Control	<p>Select to enable flow control on the interface, which permits the switch suspend packet transmission for a set period of time in response to a PAUSE frame sent by a congested switch.</p> <p>Flow control applies only to links operating at 1 Gbps, full-duplex mode.</p>
MTU	<p>Using the arrows, indicate the maximum transmission unit (MTU), which is the maximum size of Ethernet frames sent by the interface. To calculate the MTU, add 14 bytes overhead to the maximum payload you want sent.</p> <p>Range: 256 through 9216 bytes</p>
Speed	<p>Select the link speed.</p> <p>If you select a link speed when autonegotiation is enabled, autonegotiation remains enabled and the interface advertises the link speed that you specify as its maximum link speed.</p> <p>NOTE: This setting is ignored when you assign a Port profile to an Aggregated Ethernet interface.</p>
Link Mode	<p>Select the duplex mode, either Automatic, Full Duplex, or Half Duplex. Select Automatic to enable autonegotiation when autonegotiation is disabled.</p> <p>NOTE: This setting is ignored when you assign a Port profile to an Aggregated Ethernet interface.</p> <p>You cannot select Half Duplex with link speed set to Autonegotiation or 1 Gbps.</p>

Table 69: Port Profile Custom Setup Settings (*continued*)

Field	Action
<p>Storm Control Settings</p> <p>Enabling storm control on a switching device monitors traffic levels and drops broadcast, multicast, and unknown unicast packets when a specified traffic level—called the storm control level or storm control bandwidth—is exceeded, thus preventing packets from proliferating and degrading the LAN.</p> <p>You can customize the storm control level for a specific interface by explicitly configuring either bandwidth or level.</p> <p>NOTE: You cannot configure both bandwidth and level for the same interface.</p>	<p>Unit</p> <ul style="list-style-type: none"> • Percentage—Configures the storm control level as a percentage of the available bandwidth used by the combined applicable traffic streams that are subject to storm control on that interface. The level can be set from 0% to 100%, where 0% indicates that the entire traffic is being suppressed and 100% indicates no traffic is being suppressed, in other words there is no storm control. The default level is 80%. • Kbps—Configures the storm control level as the bandwidth in kilobits per second (Kbps) of the applicable traffic streams on that interface. Set the bandwidth from 100 through 10,000,000 in Kbps. When you configure storm control bandwidth, the value you configure is rounded off internally to the closest multiple of 64 Kbps, and the rounded-off value represents the bandwidth that is actually used. For example, if you configure a bandwidth limit of 150 Kbps, storm control uses a bandwidth limit of 128 Kbps. <p>Value</p> <p>Configures the traffic storm control threshold level value as a percentage of bandwidth or bandwidth in kilobits per second depending upon the specified unit.</p> <p>No broadcast</p> <p>Select this option to enable storm control for no broadcast traffic on a specific interface or on all interfaces.</p> <p>No unknown broadcast</p> <p>Select this option to enable storm control for no unknown broadcast traffic on a specific interface or on all interfaces.</p> <p>No multicast</p> <p>Select this option to enable storm control for no multicast traffic on a specific interface or on all interfaces.</p>

Power over Ethernet (PoE)

You can enable PoE and display the configuration options by enabling **Configure Power over Ethernet**.

Table 69: Port Profile Custom Setup Settings (*continued*)

Field	Action
Configure Power over Ethernet	<p>Enable to configure PoE settings.</p> <p>If you do not enable this option, Network Director does not send any PoE configuration commands to the device when the profile is deployed on the device. For example, if PoE is enabled on an interface, it remains enabled. On EX Series switches, the factory-default configuration enables PoE on all interfaces that support PoE.</p> <p>If you enable this option, the PoE settings in this profile is deployed on the interfaces that support PoE. If you assign this Port profile to interfaces that do not support PoE, the profile is deployed successfully on those interfaces, but the PoE settings do not take effect.</p>
Maximum Power (W)	<p>Use the arrows to adjust the maximum PoE power in watts allocated to a PoE port. The powered device cannot draw more power than the wattage specified. If it does so, PoE power to the port is shut down. Maximum power for PoE is 15.4W, Extended PoE is 18.6W and PoE+ is 30W.</p> <p>The Maximum Power setting has no effect when the PoE management mode for a switch or line card is class mode, which is the default mode. In class mode, the power allocated to a PoE port is determined either by LLDP negotiation with the powered device or by the PoE class of the powered device if LLDP is not supported.</p> <p>You must set the PoE power management mode for the switch or line card to static mode for the Maximum Power setting to take effect. Do this in the Device Common Settings profile.</p> <p>If you specify a maximum wattage that is greater than the maximum wattage that can be supplied by the port, your configuration is accepted when the Port profile is deployed on the port. However, the maximum wattage is set to the port's maximum supported wattage. The maximum supported wattages for PoE ports are:</p> <ul style="list-style-type: none"> • 15.4W for ports that support IEEE 802.3af only • 18.6W for IEEE 802.3af ports on switches that support enhanced PoE • 30W for ports that support IEEE 802.3at <p>Default: 15.4W</p>

Table 69: Port Profile Custom Setup Settings (*continued*)

Field	Action
Priority	<p>Select a power priority for the PoE port—either Low or High. If there is a shortage of PoE power on the switch, power to low priority ports is shut down before power to high priority ports. Within ports with the same assigned priority, power priority is determined by the port number—ports with a lower port number have a higher power priority.</p> <p>Default: low priority</p>
Disable PoE	Select to disable PoE on the interfaces using this Port profile.

Port Security (Switching Interfaces Only)

Select to enable port security (default); clear to disable port security.

When port security is enabled, you can configure port security options such as learned MAC address limits on an interface. When port security is disabled, no port security is applied to the interface, including the default port security options.

Trust DHCP	<p>Select to permit messages from a DHCP server to be received on the interface—this is the default. Clear to block all messages from a DHCP server from being received on the interface.</p> <p>TIP: For this port security feature to work, DHCP snooping must be enabled on the VLAN the interface belongs to. You can enable DHCP snooping on the VLAN in the VLAN profile. For directions, see “Creating and Managing VLAN Profiles” on page 329.</p>
MAC Limit	<p>Type the number of MAC address that can be dynamically learned on the interface.</p> <p>Range: 1 through 163,839</p> <p>Default: For Desktop Ports, 1. For Desktop Phone Ports, 2. For all others, none.</p>

Table 69: Port Profile Custom Setup Settings (*continued*)

Field	Action
MAC Limit Action	<p>Select the action to be taken if the MAC address limit is exceeded:</p> <ul style="list-style-type: none"> • Drop—Drop any packet with a previously unlearned MAC address and generate a system log entry, and SNMP trap, or an alarm. This is the default for a Desktop Port and Desktop Phone Ports. • Log—Accept packets with new MAC addresses and learn the addresses, but generate a system log entry, and SNMP trap, or an alarm. • Shutdown—Shut down the interface and generate a system log message, SNMP trap, or an alarm. If an interface is shut down because the MAC address limit has been exceeded, you must use the CLI command clear ethernet-switching port-error interface <i>name</i> to clear the error and bring the interface back into service. TIP: You can use the CLI to configure autorecovery on an interface that has been shut down by a MAC limit error. • None—No action. This selection effectively disables MAC address limiting on the interface. This is the default for Switched Uplink Ports, Switched Downlink Ports, and Server Ports.
Allowed MAC List	<p>Indicate the MAC addresses of devices that are allowed access to the interface in the Allowed MAC List. Any device whose MAC address does not match an address in the list is not allowed access to the interface. A list with no entries means that a client with any MAC address is permitted to access the interface.</p> <p>To enter a MAC address, click Add and then type the MAC addresses in the field provided. Enter MAC addresses as two-character hexadecimal numbers separated by colons. Click Save to save the entry.</p> <p>NOTE: Configuring an allowed MAC address list does not block the switch from receiving Layer 2 control packets—such as Link Layer Discovery Protocol (LLDP) packets—transmitted from MAC addresses that are not specified in the address list. Control packets do not undergo the MAC address check. However, the switch does not forward them to another destination.</p> <p>Default: No entries</p>

Table 69: Port Profile Custom Setup Settings (*continued*)

Field	Action
<p>RSTP Settings</p> <p>In addition to enabling or disabling the Spanning Tree Protocol (STP) as part of device profiles, this feature enables you to fine-tune STP by setting interfaces into edge, disable, or no-root-port states.</p>	<p>Edge</p> <p>RSTP defines the concept of an edge port, which is a designated port that connects to non-STP-capable devices, such as PCs, servers, routers, or hubs that are not connected to other switches. Because edge ports connect directly to end stations, they cannot create network loops and can transition to the forwarding state immediately. You can manually configure edge ports, and a switch can also detect edge ports by noting the absence of communication from the end stations.</p> <p>Disable</p> <p>Disables the RSTP on interface.</p> <p>NOTE: Configuring interfaces to one of these states is not mandatory for ELS switches. Hence, the option Disable is not applicable for ELS switches and therefore not supported.</p> <p>No Root Port</p> <p>Configures an interface to be a spanning-tree designated port. If the bridge receives superior STP bridge protocol data units (BPDUs) on a root-protected interface, that interface transitions to a root-prevented STP state (inconsistency state) and the interface is blocked. This blocking prevents a bridge from being elected the root bridge. When the bridge stops receiving superior STP BPDUs on the root-protected interface, interface traffic is no longer blocked.</p>
CoS Settings	<p>Click Select Cos Profile to choose from existing CoS profiles. The CoS configuration contained in the CoS profile is applied to the interfaces that the Port profile is assigned to when you deploy the configuration. Click OK. Some preconfigured Service Types have a default CoS profile—see the description for Service Types field for details.</p> <p>Or</p> <p>Click Configure CoS settings to configure CoS profile. See “Creating and Managing Wired CoS Profiles” on page 400 for steps to configure a CoS profile.</p>

Table 69: Port Profile Custom Setup Settings (*continued*)

Field	Action
Authentication Settings (Desktop Port, Desktop Phone Port, Custom Port)	<p>Select the Authentication profile for the interface from a list of existing profiles by clicking Select, selecting one of the listed profiles, and then clicking OK. By assigning an Authentication profile to the Port profile, you can enable 802.1x and captive portal authentication on interfaces.</p> <p>If you do not specify an Authentication profile, the interface is an open port and no authentication is required to connect.</p> <p>NOTE: You cannot configure 802.1x authentication on aggregated Ethernet interfaces. If you attempt to deploy a Port profile that contains an Authentication profile on an aggregated Ethernet interface, the deployment fails.</p> <p>Or</p> <p>Click Configure Authentication Settings to configure 802.1x and captive portal authentications. See “Creating and Managing Authentication Profiles” on page 237 for steps to configure the authentication profile.</p>
Filter Settings (available for all Service Types, including Custom for routing)	<ul style="list-style-type: none"> • Ingress Filter <p>Select an Ingress Filter for the interface by clicking Select, selecting one of the listed filters, and then clicking OK.</p> • Egress Filters <p>Select an Egress Filter for the interface by clicking Select, selecting one of the listed filters, and then clicking OK.</p>
VRRP Settings (available when Service Type is Custom and Family Type is Routing)	<p>Select the VRRP profile for the interface from a list of existing profiles by clicking Select. Select one of the listed profiles, and then click OK.</p>

If you configured Custom Setup settings, click **Done**. Upon clicking **Done** displays the dialog **Do you want to assign Port Profile to Ports?**. Click **Yes** to create a profile assignment else click **No** to create the profile and navigate to the Manage Port Profile page to create the Port assignment later.

Specifying Settings for a Campus Switching ELS Port Profile

Use the Create Port Profile page to define a common set of port attributes in a Port profile. You can then apply the Port profile to interfaces on a group of Campus Switching ELS devices.

TIP: You can reference a VLAN profile, CoS profile, Ingress Filter profile, Egress Filter profile, and an Authentication profile in a Port profile. You can either create these profiles in their respective profile pages before you create Port profiles or you can create these profiles as in-line sub-profiles while configuring Port profiles. You can also enable power over Ethernet (PoE).

After you create a Port profile, you can assign it to individual interfaces or to members of a Port group. During this assignment process, you can also configure interface-specific attributes, such as IP address, and certain device-specific attributes, such as the Access profile to use for all ports on the device. You can assign only one Port profile to an interface.

[Table 70 on page 270](#) describes the Quick Setup settings available in a Port profile. [Table 71 on page 278](#) describes the Custom Setup settings. The defaults for these options depend on the Service Type you select.

Table 70: Port Profile Quick Setup Settings for Campus Switching ELS

Field	Action
Profile Name	Type the name of profile by using up to 64 alphanumeric characters and no special characters other than the underscore. The name must be unique among Port profiles.
Description	Type a description of the Port profile, which appears on the Manage Port Profiles page. You can use up to 256 characters.

Table 70: Port Profile Quick Setup Settings for Campus Switching ELS (*continued*)

Field	Action
Service Type	<p>Select one the preconfigured options Desktop Port, Desktop Phone Port, Printer Port, Switched Uplink, Switched Downlink, or Server Port. To create your own service type, select Custom.</p> <hr/> <p>Desktop Port service type has the following default settings:</p> <ul style="list-style-type: none"> • CoS Profile—juniper_CS_Hier_Ethernet_CoS for Hierarchical port scheduling; no default profile for Non-Hierarchical port scheduling • Family Type—switching • Port Mode—access • Power over Ethernet—disabled • Auto Negotiation—disabled • Flow Control—disabled • Maximum Size—no default provided • Speed—no default provided • Link Mode—no default provided • Port Security—enabled • Trust DHCP—disabled • MAC Limit—1 • MAC Limit Action—drop • Allowed MAC List—no default provided

Table 70: Port Profile Quick Setup Settings for Campus Switching ELS (*continued*)

Field	Action
	<p>Desktop Phone Port service type has the following default settings:</p> <ul style="list-style-type: none"> • CoS Profile—juniper_CoS_template for Non-Hierarchical port scheduling; juniper_CS_Hier_Ethernet_CoS for Hierarchical port scheduling • Family Type—switching • Port Mode—access • Power over Ethernet—disabled • Auto Negotiation—disabled • Flow Control—disabled • Maximum Size—no default provided • Speed—no default provided • Link Mode—no default provided • Port Security—enabled • Trust DHCP—disabled • MAC Limit—2 • MAC Limit Action—drop • Allowed MAC List—no default provided
	<p>Printer Port preconfigured service type has the following default settings:</p> <ul style="list-style-type: none"> • Family Type—switching • Port Mode—access • Power over Ethernet—no default provided • Auto Negotiation—enabled • Flow Control—enabled • Maximum Size—no default provided • Speed—no default provided • Link Mode—no default provided • Port Security—no default provided • Trust DHCP—no default provided • MAC Limit—no default provided • MAC Limit Action—no default provided • Allowed MAC List—no default provided

Table 70: Port Profile Quick Setup Settings for Campus Switching ELS (*continued*)

Field	Action
	<p>Switched Uplink service type has the following default settings:</p> <ul style="list-style-type: none"> • CoS Profile—juniper_CoS_template for Non-Hierarchical port scheduling; juniper_CS_Hier_Ethernet_CoS for Hierarchical port scheduling • Family Type—switching • Port Mode—trunk • Power over Ethernet—disabled • Auto Negotiation—disabled • Flow Control—disabled • Maximum Size—no default provided • Speed—no default provided • Link Mode—no default provided • Port Security—enabled • MAC Limit—no default provided • MAC Limit Action—no default provided • Allowed MAC List—no default provided
	<p>Switched Downlink service type has the following default settings:</p> <ul style="list-style-type: none"> • CoS Profile—juniper_CoS_template for Non-Hierarchical port scheduling; juniper_CS_Hier_Ethernet_CoS for Hierarchical port scheduling • Family Type—switching • Port Mode—trunk • Power over Ethernet—disabled • Auto Negotiation—disabled • Flow Control—disabled • Maximum Size—no default provided • Speed—no default provided • Link Mode—no default provided • MAC Limit—no default provided • MAC Limit Action—no default provided • Allowed MAC list—no default provided

Table 70: Port Profile Quick Setup Settings for Campus Switching ELS (*continued*)

Field	Action
	<p>Server Port service type has the following default settings:</p> <ul style="list-style-type: none"> • CoS Profile—juniper_CoS_template for Non-Hierarchical port scheduling; juniper_CS_Hier_Ethernet_CoS for Hierarchical port scheduling • Family Type—switching • Port Mode—trunk • Power over Ethernet—disabled • Auto Negotiation—disabled • Flow Control—disabled • Maximum Size—no default provided • Speed—no default provided • Link Mode—no default provided • Port Security—enabled • MAC Limit—no default provided • MAC Limit Action—no default provided • Allowed MAC list—no default provided
Port Family Options <p>The available settings and defaults for these options depend on the Service Type you selected.</p>	
Family Type: Switching or Routing	<p>This setting cannot be changed if any preconfigured Service Type was selected. If you selected the Custom Service Type, indicate whether the interface operates as a Layer 2 (Switching) or a Layer 3 (Routing) interface.</p> <p>TIP: Service Type must be set to Custom to configure a routing interface.</p> <p>If you select routing, you configure an IP address on a per-interface basis when you assign the profile to individual interfaces.</p>

Table 70: Port Profile Quick Setup Settings for Campus Switching ELS (continued)

Field	Action
Port Mode for switching interfaces only	<p>This setting cannot be changed if any preconfigured Service Type was selected. If you selected the Custom Service Type, select the port mode for the interface, either Access, Trunk, or Tagged Access.</p> <ul style="list-style-type: none"> • Access—Use for interfaces that connect to an end device, such as a desktop computer, an IP telephone, a printer, or a security camera. The interface must belong to a single VLAN. Frames sent and received over the over the interface are untagged Ethernet frames. • Trunk—Use for interfaces that connect to a switch or router. Trunk interfaces can belong to more than one VLAN, enabling VLAN traffic to be multiplexed on a single physical interface. The Ethernet frames sent and received over the interface are tagged frames, in which IEEE 802.1Q tagging is used to segregate the traffic from each VLAN. • Tagged Access—Use for access interfaces where VLAN tagging is required, typically when the interface connects to a server running virtual machines using virtual Ethernet port aggregator (VEPA) technology. The traffic generated by the server can contain an aggregation of VLAN packets from different virtual machines on that server, requiring that packets be tagged.

VLAN Options

Available VLAN options depend on the Service Type selected. VLAN association is required for Campus Switching ELS.

Member VLAN (Switched Uplink, Switched Downlink, Server Port)	<p>Click All if you want to assign an interface to all the VLANs.</p> <p>This option is enabled when Port Mode is Trunk or TaggedAccess.</p>
Member VLAN (all Service Types)	<p>This configuration is for one VLAN. Select a VLAN for the interface by clicking Select, selecting one of the listed filters, and then clicking OK.</p> <p>Or</p> <p>Configure a VLAN by clicking Configure VLAN Settings and clicking Create. Enter the VLAN name and ID and click OK.</p>
Voice VLAN (Desktop Phone Port, Custom Port)	<p>This configuration is for one VLAN. Select a voice VLAN for the interface by clicking Select, selecting one of the listed filters, and then clicking OK.</p> <p>Or</p> <p>Configure a VLAN by clicking Configure VLAN Settings and clicking Create. Enter the VLAN name and ID and click OK.</p>

Table 70: Port Profile Quick Setup Settings for Campus Switching ELS (*continued*)

Field	Action
Native VLAN (Switched Uplink, Switched Downlink)	<p>Select a native VLAN for the interface by clicking Select, selecting one of the listed VLANs, and then clicking OK. The VLAN is added to the Member VLANs list.</p> <p>Or</p> <p>Configure a VLAN by clicking Configure VLAN Settings and clicking Create. Enter the VLAN name and ID and click OK.</p>
Power over Ethernet (PoE)	
Configure Power over Ethernet	<p>Enable to configure PoE settings.</p> <p>If you do not enable this option, Network Director does not send any PoE configuration commands to the device when the profile is deployed on the device. For example, if PoE is enabled on an interface, it remains enabled.</p> <p>If you enable this option, the PoE settings in this profile is deployed on the interfaces that support PoE. If you assign this Port profile to interfaces that do not support PoE, the profile can be deployed successfully on those interfaces, but the PoE settings do not take effect.</p> <p>TIP: EX9200 switches do not support PoE.</p>

Table 70: Port Profile Quick Setup Settings for Campus Switching ELS (*continued*)

Field	Action
Maximum Power (W)	<p>Use the arrows to adjust the maximum PoE power allocated to a PoE port in watts. The powered device cannot draw more power than the wattage specified. If it does so, PoE power to the port is shut down.</p> <p>The Maximum Power setting has no effect when the PoE management mode for the switch or line card is class mode, which is the default mode. In class mode, the power allocated to a PoE port is determined either by LLDP negotiation with the powered device or by the PoE class of the powered device if LLDP is not supported.</p> <p>You must set the PoE power management mode for the switch or line card to static mode for the Maximum Power setting to take effect. You can do this in the Device Common Settings profile.</p> <p>If you specify a maximum wattage that is greater than the maximum wattage that can be supplied by the port, your configuration is accepted when the Port profile is deployed on the port. However, the maximum wattage is set to the port's maximum supported wattage. The maximum supported wattages for PoE ports are:</p> <ul style="list-style-type: none"> • 15.4W for ports that support IEEE 802.3af only • 18.6W for IEEE 802.3af ports on switches that support enhanced PoE • 30W for ports that support IEEE 802.3at <p>Default: 15.4W</p>
Priority	<p>Select a power priority for the PoE port—either Low or High. If there is a shortage of PoE power on the switch, power to low priority ports is shut down before power to high priority ports. Within ports with the same assigned priority, power priority is determined by port number—ports with a lower port number have a higher power priority.</p> <p>Default: low priority</p>
Disable PoE	Select to disable PoE on the interface.

After providing the information in the fields listed in [Table 70 on page 270](#), click **Done**.

To use default Port Profile Custom Setup settings, click **Done**. To configure Custom Setup settings, click **Custom Setup** and then provide the information in [Table 71 on page 278](#) and then click **Done**.

Clicking **Done** in either case displays the dialog Do you want to assign Port Profile to Ports. Click **Yes** to create a profile assignment; else click **No** to navigate to the Manage Port Profile page and to create the Port assignment later on.

Table 71: Port Profile Custom Setup Settings for Campus Switching ELS

Field	Action
Advanced Settings Expand Advanced Settings to configure link settings and port security. The Link Setting in Port profile is disabled by default. On enabling Link Settings, autonegotiation and flow control are enabled by default.	
Enable Auto Negotiation	<p>Autonegotiation of link speed and duplex mode is enabled by default; clear to disable autonegotiation.</p> <p>If you disable autonegotiation, you must set link speed and link mode.</p> <p>You cannot disable autonegotiation if a link speed of 1 Gbps is configured. This configuration might be accepted, but autonegotiation will not be disabled.</p> <p>NOTE: This setting is ignored when you assign a Port profile to an Aggregated Ethernet interface.</p>
Enable Flow Control	<p>Select to enable flow control on the interface, which permits the switch suspend packet transmission for a set period of time in response to a PAUSE frame sent by a congested switch.</p> <p>Flow control applies only to links operating at 1 Gbps, full-duplex mode.</p>
MTU	<p>Using the arrows, indicate the maximum transmission unit (MTU), which is the maximum size of Ethernet frames sent by the interface. To calculate the MTU, add 14 bytes overhead to the maximum payload you want sent.</p> <p>Range: 256 through 9216 bytes</p>
Speed	<p>Select the link speed.</p> <p>If you select a link speed when autonegotiation is enabled, autonegotiation remains enabled and the interface will advertise the link speed that you specify as its maximum link speed.</p> <p>NOTE: This setting is ignored when you assign a Port profile to an Aggregated Ethernet interface.</p>
Link Mode	<p>Select the duplex mode, either Automatic, Full Duplex, or Half Duplex. Select Automatic to enable autonegotiation when autonegotiation is disabled.</p> <p>NOTE: This setting is ignored when you assign a Port profile to an Aggregated Ethernet interface.</p> <p>You cannot select Half Duplex with link speed set to Autonegotiation or 1 Gbps.</p>

Table 71: Port Profile Custom Setup Settings for Campus Switching ELS (continued)

Field	Action
<p>Storm Control Settings</p> <p>Enabling storm control on a switching device monitors traffic levels and drops broadcast, multicast, and unknown unicast packets when a specified traffic level—called the storm control level or storm control bandwidth—is exceeded, thus preventing packets from proliferating and degrading the LAN.</p> <p>You can customize the storm control level for a specific interface by explicitly configuring either bandwidth or level.</p> <p>NOTE: You cannot configure both bandwidth and level for the same interface.</p>	

Table 71: Port Profile Custom Setup Settings for Campus Switching ELS (*continued*)

Field	Action
	<p>Unit</p> <ul style="list-style-type: none"> • Percentage—Configures the storm control level as a percentage of the available bandwidth used by the combined applicable traffic streams that are subject to storm control on that interface. The level can be set from 0% to 100%, where 0% indicates that the entire traffic is being suppressed and 100% indicates no traffic is being suppressed, in other words there is no storm control. The default level is 80%. • Kbps—Configures the storm control level as the bandwidth in kilobits per second (Kbps) of the applicable traffic streams on that interface. Set the bandwidth from 100 through 10,000,000 in Kbps. When you configure storm control bandwidth, the value you configure is rounded off internally to the closest multiple of 64 Kbps, and the rounded-off value represents the bandwidth that is actually used. For example, if you configure a bandwidth limit of 150 Kbps, storm control uses a bandwidth limit of 128 Kbps. <p>Value</p> <p>Configures the traffic storm control threshold level value as a percentage of bandwidth or bandwidth in kilobits per second depending upon the specified unit.</p> <p>No broadcast</p> <p>Select this option to enable storm control for no broadcast traffic on a specific interface or on all interfaces.</p> <p>No unknown broadcast</p> <p>Select this option to enable storm control for no unknown broadcast traffic on a specific interface or on all interfaces.</p> <p>No multicast</p> <p>Select this option to enable storm control for no multicast traffic on a specific interface or on all interfaces.</p> <p>No registered multicast</p> <p>Select this option to enable storm control for no registered multicast traffic on a specific interface or on all interfaces.</p> <p>No unregistered multicast</p> <p>Select this option to enable storm control for no unregistered multicast traffic</p>

Table 71: Port Profile Custom Setup Settings for Campus Switching ELS (*continued*)

Field	Action
	on a specific interface or on all interfaces.
Power over Ethernet (PoE)	
You can enable PoE and display the configuration options by enabling Configure Power over Ethernet .	
Configure Power over Ethernet	<p>Enable to configure PoE settings.</p> <p>If you do not enable this option, Network Director does not send any PoE configuration commands to the device when the profile is deployed on the device. For example, if PoE is enabled on an interface, it remains enabled. On EX Series switches, the factory-default configuration enables PoE on all interfaces that support PoE.</p> <p>If you enable this option, the PoE settings in this profile is deployed on the interfaces that support PoE. If you assign this Port profile to interfaces that do not support PoE, the profile is deployed successfully on those interfaces, but the PoE settings will not take effect.</p>
Maximum Power (W)	<p>Use the arrows to adjust the maximum PoE power in watts allocated to a PoE port. The powered device cannot draw more power than the wattage specified. If it does so, PoE power to the port is shut down. Maximum power for PoE is 15.4W, Extended PoE is 18.6W and PoE+ is 30W.</p> <p>The Maximum Power setting has no effect when the PoE management mode for a switch or line card is class mode, which is the default mode. In class mode, the power allocated to a PoE port is determined either by LLDP negotiation with the powered device or by the PoE class of the powered device if LLDP is not supported.</p> <p>You must set the PoE power management mode for the switch or line card to static mode for the Maximum Power setting to take effect. Do this in the Device Common Settings profile.</p> <p>If you specify a maximum wattage that is greater than the maximum wattage that can be supplied by the port, your configuration is accepted when the Port profile is deployed on the port. However, the maximum wattage is set to the port's maximum supported wattage. The maximum supported wattages for PoE ports are:</p> <ul style="list-style-type: none"> • 15.4W for ports that support IEEE 802.3af only • 18.6W for IEEE 802.3af ports on switches that support enhanced PoE • 30W for ports that support IEEE 802.3at <p>Default: 15.4W</p>

Table 71: Port Profile Custom Setup Settings for Campus Switching ELS (*continued*)

Field	Action
Priority	<p>Select a power priority for the PoE port—either Low or High. If there is a shortage of PoE power on the switch, power to low priority ports is shut down before power to high priority ports. Within ports with the same assigned priority, power priority is determined by the port number—ports with a lower port number have a higher power priority.</p> <p>Default: low priority</p>
Disable PoE	Select to disable PoE on the interfaces using this Port profile.

Port Security (Switching Interfaces Only)

Select to enable port security (default); clear to disable port security.

When port security is enabled, you can configure port security options such as learned MAC address limits on an interface. When port security is disabled, no port security is applied to the interface, including the default port security options.

Trust DHCP	<p>Select to permit messages from a DHCP server to be received on the interface—this is the default. Clear to block all messages from a DHCP server from being received on the interface.</p> <p>TIP: For this port security feature to work, DHCP snooping must be enabled on the VLAN the interface belongs to. You can enable DHCP snooping on the VLAN in the VLAN profile. For directions, see “Creating and Managing VLAN Profiles” on page 329.</p>
MAC Limit	<p>Type the number of MAC address that can be dynamically learned on the interface.</p> <p>Range: 1 through 163839</p> <p>Default: For Desktop Ports, 1. For Desktop Phone Ports, 2. For all others, none.</p>

Table 71: Port Profile Custom Setup Settings for Campus Switching ELS (*continued*)

Field	Action
MAC Limit Action	<p>Select the action to be taken if the MAC address limit is exceeded:</p> <ul style="list-style-type: none"> • Drop—Drop any packet with a previously unlearned MAC address and generate a system log entry, and SNMP trap, or an alarm. This is the default for a Desktop Port and Desktop Phone Ports. • Log—Accept packets with new MAC addresses and learn the addresses, but generate a system log entry, and SNMP trap, or an alarm. • Shutdown—Shut down the interface and generate a system log message, SNMP trap, or an alarm. If an interface is shut down because the MAC address limit has been exceeded, you must use the CLI command clear ethernet-switching port-error interface <i>name</i> to clear the error and bring the interface back into service. TIP: You can use the CLI to configure auto-recovery on an interface that has been shut down by a MAC limit error. • None—No action. This selection effectively disables MAC address limiting on the interface. This is the default for Switched Uplink Ports, Switched Downlink Ports, and Server Ports.
Allowed MAC List	<p>Indicate the MAC addresses of devices that are allowed access to the interface in the Allowed MAC List. Any device whose MAC address does not match an address in the list will not be allowed access to the interface. A list with no entries means that a client with any MAC address is permitted to access the interface.</p> <p>To enter a MAC address, click Add and then type the MAC addresses in the field provided. Enter MAC addresses as two-character hexadecimal numbers separated by colons. Click Save to save the entry.</p> <p>NOTE: Configuring an allowed MAC address list does not block the switch from receiving Layer 2 control packets—such as Link Layer Discovery Protocol (LLDP) packets—transmitted from MAC addresses that are not specified in the address list. Control packets do not undergo the MAC address check. However, the switch does not forward them to another destination.</p> <p>Default: No entries</p>

Table 71: Port Profile Custom Setup Settings for Campus Switching ELS (*continued*)

Field	Action
<p>RSTP Settings</p> <p>In addition to enabling or disabling the Spanning Tree Protocol (STP) as part of device profiles, this feature enables you to fine-tune STP by setting interfaces into edge, disable, or no-root-port states.</p>	<p>Edge</p> <p>RSTP defines the concept of an edge port, which is a designated port that connects to non-STP-capable devices, such as PCs, servers, routers, or hubs that are not connected to other switches. Because edge ports connect directly to end stations, they cannot create network loops and can transition to the forwarding state immediately. You can manually configure edge ports, and a switch can also detect edge ports by noting the absence of communication from the end stations.</p> <p>Disable</p> <p>Disables the RSTP on interface.</p> <p>NOTE: Configuring interfaces to one of these states is not mandatory for ELS switches. Hence, the option Disable is not applicable for ELS switches and therefore not supported.</p> <p>No Root Port</p> <p>Configures an interface to be a spanning-tree designated port. If the bridge receives superior STP bridge protocol data units (BPDUs) on a root-protected interface, that interface transitions to a root-prevented STP state (inconsistency state) and the interface is blocked. This blocking prevents a bridge from being elected the root bridge. When the bridge stops receiving superior STP BPDUs on the root-protected interface, interface traffic is no longer blocked.</p>
CoS Settings	<p>Click Select Cos Profile to choose from existing CoS profiles. The CoS configuration contained in the CoS profile is applied to the interfaces that the Port profile is assigned to when you deploy the configuration. Click OK. Some preconfigured Service Types have a default CoS profile—see Service Types for details.</p> <p>Or</p> <p>Click Configure CoS settings to configure CoS profile. See “Creating and Managing Wired CoS Profiles” on page 400 for steps to configure a CoS profile.</p>

Table 71: Port Profile Custom Setup Settings for Campus Switching ELS (*continued*)

Field	Action
Authentication Settings (Desktop Port, Desktop Phone Port, Custom Port)	<p>Select the Authentication profile for the interface from a list of existing profiles by clicking Select, selecting one of the listed profiles, and then clicking OK. By assigning an Authentication profile to the Port profile, you can enable 802.1x and captive portal authentication on interfaces.</p> <p>If you do not specify an Authentication profile, the interface is an open port and no authentication is required to connect.</p> <p>NOTE: You cannot configure 802.1x authentication on aggregated Ethernet interfaces. If you attempt to deploy a Port profile that contains an Authentication profile on an aggregated Ethernet interface, the deployment fails.</p> <p>Or</p> <p>Click Configure Authentication Settings to configure 802.1x and captive portal authentications. See “Creating and Managing Authentication Profiles” on page 237 for steps to configure the Authentication profile.</p>
Filter Settings (available for all Service Types, including Custom for routing)	<ul style="list-style-type: none"> ● Ingress Filter <p>Select an Ingress Filter for the interface by clicking Select, selecting one of the listed filters, and then clicking OK.</p> ● Egress Filters <p>Select an Egress Filter for the interface by clicking Select, selecting one of the listed filters, and then clicking OK.</p>
VRRP Settings (available when Service Type is Custom and Family Type is Routing)	Select the VRRP profile for the interface from a list of existing profiles by clicking Select . Select one of the listed profiles, and then click OK .

Clicking **Done** displays the dialog Do you want to assign Port Profile to Ports. Click **Yes** to create a profile assignment else click **No** to navigate to the Manage Port Profile page and to create the Port assignment later.

Click **Done** to save the Port profile for Campus Switching ELS.

Specifying Settings for a Data Center Switching ELS Port Profile

Use the Create Port Profile page to define a common set of port attributes in a Port profile. You can create a new Port profile from scratch, or select an appropriate Service Type and use the default settings that Network Director has defined for that service type to create a Port profile. You can then apply the Port profile to interfaces on a group of Data Center Switching ELS devices.

TIP: You can reference a VLAN profile, CoS profile, Ingress Filter profile, Egress Filter profile, and an Authentication profile in a Port profile. You can either create these profiles in their respective profile pages before you create Port profiles or you can create these profiles as in-line sub-profiles while configuring Port profiles.

After you create a Port profile, you can assign it to individual interfaces or to members of a Port group. During this assignment process, you can also configure interface-specific attributes, such as IP address, and certain device-specific attributes, such as the Access profile to use for all ports on the device. You can assign only one Port profile to an interface.

[Table 72 on page 286](#) describes the Quick Setup settings available in a Port profile. [Table 73 on page 295](#) describes the Custom Setup settings. The defaults for these options depend on the Service Type you select.

Table 72: Port Profile Quick Setup Settings for Data Center Switching ELS

Field	Action
Profile Name	Type the name of profile by using up to 64 alphanumeric characters and no special characters other than the underscore. The name must be unique among Port profiles.
Description	Type a description of the Port profile, which will appear on Manage Port Profiles page. You can use up to 256 characters.

Table 72: Port Profile Quick Setup Settings for Data Center Switching ELS (*continued*)

Field	Action
Service Type	<p>Select one the preconfigured options Desktop Port, Switched Uplink, Switched Downlink, Server Port, or FCoE Transit Port. To create your own service type, select Custom.</p> <hr/> <p>Desktop Port service type has the following default settings:</p> <ul style="list-style-type: none"> • CoS Profile—no default provided • Family Type—switching • Port Mode—access • Auto Negotiation—disabled • Flow Control—disabled • Maximum Size—no default provided • Speed—no default provided • Link Mode—no default provided • Port Security—enabled • Trust DHCP—disabled • MAC Limit—1 • MAC Limit Action—drop • Allowed MAC List—no default provided

Table 72: Port Profile Quick Setup Settings for Data Center Switching ELS (continued)

Field	Action
	<p>Switched Uplink service type has the following default settings:</p> <ul style="list-style-type: none"> • CoS Profile—juniper_DC_Hier_Ethernet_CoS (for Hierarchical port scheduling), juniper_DC_NonHier_CoS_Fusion (for Non-Hierarchical (Fusion) port scheduling), and juniper_DC_Hier_Fusion_CoS (for Hierarchical (Fusion) port scheduling) • Family Type—switching • Port Mode—trunk • Auto Negotiation—disabled • Flow Control—disabled • Maximum Size—no default provided • Speed—no default provided • Link Mode—no default provided • Port Security—enabled • MAC Limit—no default provided • MAC Limit Action—no default provided • Allowed MAC List—no default provided

Table 72: Port Profile Quick Setup Settings for Data Center Switching ELS (continued)

Field	Action
	<p>Switched Downlink service type has the following default settings:</p> <ul style="list-style-type: none"> • CoS Profile—juniper_DC_Hier_Ethernet_CoS (for Hierarchical port scheduling), juniper_DC_NonHier_CoS_Fusion (for Non-Hierarchical (Fusion) port scheduling), and juniper_DC_Hier_Fusion_CoS (for Hierarchical (Fusion) port scheduling) • Family Type—switching • Port Mode—trunk • Auto Negotiation—disabled • Flow Control—disabled • Maximum Size—no default provided • Speed—no default provided • Link Mode—no default provided • Port Security—enabled • MAC Limit—no default provided • MAC Limit Action—no default provided • Allowed MAC list—no default provided

Table 72: Port Profile Quick Setup Settings for Data Center Switching ELS (continued)

Field	Action
	<p>Server Port service type has the following default settings:</p> <ul style="list-style-type: none">• CoS Profile—juniper_DC_Hier_Ethernet_CoS (for Hierarchical port scheduling), juniper_DC_NonHier_CoS_Fusion (for Non-Hierarchical (Fusion) port scheduling), and juniper_DC_Hier_Fusion_CoS (for Hierarchical (Fusion) port scheduling)• Family Type—switching• Port Mode—trunk• Auto Negotiation—disabled• Flow Control—disabled• Maximum Size—no default provided• Speed—no default provided• Link Mode—no default provided• Port Security—enabled• MAC Limit—no default provided• MAC Limit Action—no default provided• Allowed MAC list—no default provided

Table 72: Port Profile Quick Setup Settings for Data Center Switching ELS (continued)

Field	Action
	<p>FCoE Transit Port service type has the following default settings:</p> <ul style="list-style-type: none"> • Port Type—Ethernet Port • CoS Profile—juniper_DC_Hier_CoS (for Hierarchical port scheduling), juniper_DC_NonHier_CoS_ELS (for Non-Hierarchical port scheduling), juniper_DC_NonHier_CoS_Fusion (for Non-Hierarchical (Fusion) port scheduling), and juniper_DC_Hier_Fusion_CoS (for Hierarchical (Fusion) port scheduling) • Family Type—switching • Port Mode—trunk • Filters—no default provided • VLAN Options—no default provided • DCBX Version—Auto • Disable DCBX—disabled • Disable Priority Flow Control—disabled • ETS No Auto Negotiation—disabled • Recommendation TVL—no default provided • Auto Negotiation—disabled • Flow Control—disabled • Maximum Size—2500 • Speed—no default provided • Link Mode—no default provided • Port Security—enabled • FCoE Trusted—enabled • MAC Limit—no default provided • MAC Limit Action—no default provided • Allowed MAC List—no default provided

Table 72: Port Profile Quick Setup Settings for Data Center Switching ELS (continued)

Field	Action
<p>Family Type: Switching or Routing</p> <p>The available settings and defaults for these options depend on the Service Type you selected.</p>	<p>This setting cannot be changed if any preconfigured Service Type was selected. If you selected the Custom Service Type, indicate whether the interface operates as a Layer 2 (Switching) or a Layer 3 (Routing) interface.</p> <p>TIP: Service Type must be set to Custom to configure a routing interface.</p> <p>If you select routing, you configure an IP address on a per-interface basis when you assign the profile to individual interfaces.</p>
Port Mode for switching interfaces only	<p>This setting cannot be changed if any preconfigured Service Type was selected. If you selected the Custom Service Type, select the port mode for the interface, either Access or Trunk.</p> <ul style="list-style-type: none"> • Access—Use for interfaces that connect to an end device, such as a desktop computer, an IP telephone, a printer, or a security camera. The interface must belong to a single VLAN. Frames sent and received over the over the interface are untagged Ethernet frames. • Trunk—Use for interfaces that connect to a switch or router. Trunk interfaces can belong to more than one VLAN, enabling VLAN traffic to be multiplexed on a single physical interface. The Ethernet frames sent and received over the interface are tagged frames, in which IEEE 802.1Q tagging is used to segregate the traffic from each VLAN.
Port Type	For Data Center ELS profiles, the port type is always Ethernet Port.
<p>VLAN Options</p> <p>Available VLAN options depend on the Service Type selected.</p>	
Member VAN (available for Switched Uplink, Switched Downlink, Server Port, FCoE Transit Port, Custom)	<p>Click All if you want to assign an interface to all the VLANs.</p> <p>This option is enabled when Port Mode is Trunk.</p>

Table 72: Port Profile Quick Setup Settings for Data Center Switching ELS (continued)

Field	Action
Member VLANs (available for Desktop Port, Desktop Phone Port, Switched Uplink, Switched Downlink, Server Port, Custom Port)	<p>Select a VLAN for the interface by clicking Select, selecting one of the listed filters, and then clicking OK. The VLAN is added to the Member VLANs list.</p> <p>Or</p> <p>Configure a VLAN by clicking Configure VLAN Settings and clicking Create. Enter the VLAN name and ID and click OK.</p>
Voice VLAN (available for Desktop Phone Port, Custom Port)	<p>Select a voice VLAN for the interface by clicking Select, selecting one of the listed filters, and then clicking OK. The VLAN is added to the Member VLANs list.</p> <p>Or</p> <p>Configure a VLAN by clicking Configure VLAN Settings and clicking Create. Enter the VLAN name and ID and click OK.</p>
Native VLAN (available for Switched Uplink, Switched Downlink)	<p>Select a native VLAN for the interface by clicking Select, selecting one of the listed VLANs, and then clicking OK. The VLAN is added to the Member VLANs list.</p> <p>Or</p> <p>Configure a VLAN by clicking Configure VLAN Settings and clicking Create. Enter the VLAN name and ID and click OK.</p>
Member VLAN	(Access ports only) Select a VLAN profile for the interface from a list of existing profiles by clicking Select .
Member VLANs	(Trunk ports only) Select a set of VLAN profiles for the interface from a list of existing profiles by using the Add and Remove functions.
Native VLAN	(Trunk ports only) Select a native VLAN profile for the interface from a list of existing profiles by clicking Select .
DCBX Settings Data Center Bridging Capability Exchange protocol is a discovery and exchange protocol for conveying configuration and capabilities among network neighbors to ensure consistent configuration across the network. It is an extension of the Link Layer Data Protocol (LLDP, described in IEEE 802.1AB). The defaults for these settings depend on the Service Type you selected.	

Table 72: Port Profile Quick Setup Settings for Data Center Switching ELS (*continued*)

Field	Action
DCBX Version	<p>Select one of the following versions of the Data Center Bridging Capability Exchange protocol:</p> <ul style="list-style-type: none"> • Auto—automatic configuration • DCBX v1.01—The Converged Enhanced Ethernet (CEE) version of DCBX. It has a subtype of 2 and an IEEE DCBX Organizationally Unique Identifier (OUI) of 0x001b21. • IEEE DCBX—The newest DCBX version. Different TLVs have different subtypes (for example, the subtype for the ETS configuration TLV is 9); the OUI is 0x0080c2.
Disable DCBX	Select this option to turn off Data Center Bridging Capability Exchange protocol.
Disable Priority Flow Control	<p>Select this option to turn off priority flow control.</p> <p>Priority-based flow control (PFC) is a link-level flow control mechanism defined by IEEE 802.1Qbb that enables independent flow control for each class of service (as defined in the 3-bit CoS field of the Ethernet header by IEEE 802.1Q tags) to ensure that no frame loss from congestion occurs in DCB networks.</p>
ETS No Auto Negotiation	<p>Select this option to turn off ETS autonegotiation.</p> <p>Enhanced transmission selection (ETS) is a mechanism that provides finer granularity of bandwidth management within a link.</p>
Recommendation TLV	<p>Select either Enable TLV or Disable TLV.</p> <p>The enhanced transmission selection (ETS) Recommendation TLV communicates the ETS settings that the switch wants the connected peer interface to use. If the peer interface is willing, the peer interface changes its configuration to match the configuration in the ETS Recommendation TLV. By default, the switch interfaces send the ETS Recommendation TLV to the peer. The settings communicated are the egress ETS settings defined by configuring hierarchical scheduling on the interface.</p>

After providing the information in the fields listed in [Table 71 on page 278](#), click **Done**.

To use default Port profile Custom Setup settings, click **Done**. To configure Custom Setup settings, click **Custom Setup** and then provide the information in [Table 73 on page 295](#) and then click **Done**.

Clicking **Done** in either case displays the dialog Do you want to assign Port Profile to Ports. Click **Yes** to create a profile assignment; else click **No** to navigate to the Manage Port Profile page and to create the Port assignment later on.

Table 73: Port Profile Custom Setup Settings for Data Center Switching ELS

Field	Action
Advanced Settings	
Expand Advanced Settings to configure link settings and port security. The Link Setting in Port profile is disabled by default. On enabling Link Settings, autonegotiation and flow control are enabled by default.	
Enable Auto Negotiation	<p>Autonegotiation of link speed and duplex mode is enabled by default; clear to disable autonegotiation.</p> <p>If you disable autonegotiation, you must set link speed and link mode.</p> <p>You cannot disable autonegotiation if a link speed of 1 Gbps is configured. This configuration might be accepted, but autonegotiation will not be disabled.</p> <p>NOTE: This setting is ignored when you assign a Port profile to an Aggregated Ethernet interface.</p>
Enable Flow Control	<p>Select to enable flow control on the interface, which permits the switch suspend packet transmission for a set period of time in response to a PAUSE frame sent by a congested switch.</p> <p>Flow control applies only to links operating at 1 Gbps, full-duplex mode.</p>
MTU	<p>Using the arrows, indicate the maximum transmission unit (MTU), which is the maximum size of Ethernet frames sent by the interface. To calculate the MTU, add 14 bytes overhead to the maximum payload you want sent.</p> <p>Range: 256 through 9216 bytes</p>
Speed	<p>Select the link speed.</p> <p>If you select a link speed when autonegotiation is enabled, autonegotiation remains enabled and the interface will advertise the link speed that you specify as its maximum link speed.</p> <p>NOTE: This setting is ignored when you assign a Port profile to an Aggregated Ethernet interface.</p>

Table 73: Port Profile Custom Setup Settings for Data Center Switching ELS (continued)

Field	Action
Link Mode	<p>Select the duplex mode, either Automatic, Full Duplex, or Half Duplex. Select Automatic to enable autonegotiation when autonegotiation is disabled.</p> <p>NOTE: This setting is ignored when you assign a Port profile to an Aggregated Ethernet interface.</p> <p>You cannot select Half Duplex with link speed set to Autonegotiation or 1 Gbps.</p>

Table 73: Port Profile Custom Setup Settings for Data Center Switching ELS (continued)

Field	Action
<p>Storm Control Settings</p> <p>Enabling storm control on a switching device monitors traffic levels and drops broadcast, multicast, and unknown unicast packets when a specified traffic level—called the storm control level or storm control bandwidth—is exceeded, thus preventing packets from proliferating and degrading the LAN.</p> <p>You can customize the storm control level for a specific interface by explicitly configuring either bandwidth or level.</p> <p>NOTE: You cannot configure both bandwidth and level for the same interface.</p>	

Table 73: Port Profile Custom Setup Settings for Data Center Switching ELS (*continued*)

Field	Action
	<p>Unit</p> <ul style="list-style-type: none"> • Percentage—Configures the storm control level as a percentage of the available bandwidth used by the combined applicable traffic streams that are subject to storm control on that interface. The level can be set from 0% to 100%, where 0% indicates that the entire traffic is being suppressed and 100% indicates no traffic is being suppressed, in other words there is no storm control. The default level is 80%. • Kbps—Configures the storm control level as the bandwidth in kilobits per second (Kbps) of the applicable traffic streams on that interface. Set the bandwidth from 100 through 10,000,000 in Kbps. When you configure storm control bandwidth, the value you configure is rounded off internally to the closest multiple of 64 Kbps, and the rounded-off value represents the bandwidth that is actually used. For example, if you configure a bandwidth limit of 150 Kbps, storm control uses a bandwidth limit of 128 Kbps. <p>Value</p> <p>Configures the traffic storm control threshold level value as a percentage of bandwidth or bandwidth in kilobits per second depending upon the specified unit.</p> <p>No broadcast</p> <p>Select this option to enable storm control for no broadcast traffic on a specific interface or on all interfaces.</p> <p>No unknown broadcast</p> <p>Select this option to enable storm control for no unknown broadcast traffic on a specific interface or on all interfaces.</p> <p>No multicast</p> <p>Select this option to enable storm control for no multicast traffic on a specific interface or on all interfaces.</p> <p>No registered multicast</p> <p>Select this option to enable storm control for no registered multicast traffic on a specific interface or on all interfaces.</p> <p>No unregistered multicast</p> <p>Select this option to enable storm control for no unregistered multicast traffic</p>

Table 73: Port Profile Custom Setup Settings for Data Center Switching ELS (*continued*)

Field	Action
	on a specific interface or on all interfaces.
Power over Ethernet (PoE)	
You can enable PoE and display the configuration options by enabling Configure Power over Ethernet .	
Configure Power over Ethernet	<p>Enable to configure PoE settings.</p> <p>If you do not enable this option, Network Director does not send any PoE configuration commands to the device when the profile is deployed on the device. For example, if PoE is enabled on an interface, it remains enabled. On EX Series switches, the factory-default configuration enables PoE on all interfaces that support PoE.</p> <p>If you enable this option, the PoE settings in this profile is deployed on the interfaces that support PoE. If you assign this Port profile to interfaces that do not support PoE, the profile is deployed successfully on those interfaces, but the PoE settings will not take effect.</p>
Maximum Power (W)	<p>Use the arrows to adjust the maximum PoE power in watts allocated to a PoE port. The powered device cannot draw more power than the wattage specified. If it does so, PoE power to the port is shut down. Maximum power for PoE is 15.4W, Extended PoE is 18.6W and PoE+ is 30W.</p> <p>The Maximum Power setting has no effect when the PoE management mode for a switch or line card is class mode, which is the default mode. In class mode, the power allocated to a PoE port is determined either by LLDP negotiation with the powered device or by the PoE class of the powered device if LLDP is not supported.</p> <p>You must set the PoE power management mode for the switch or line card to static mode for the Maximum Power setting to take effect. Do this in the Device Common Settings profile.</p> <p>If you specify a maximum wattage that is greater than the maximum wattage that can be supplied by the port, your configuration is accepted when the Port profile is deployed on the port. However, the maximum wattage is set to the port's maximum supported wattage. The maximum supported wattages for PoE ports are:</p> <ul style="list-style-type: none"> • 15.4W for ports that support IEEE 802.3af only • 18.6W for IEEE 802.3af ports on switches that support enhanced PoE • 30W for ports that support IEEE 802.3at <p>Default: 15.4W</p>

Table 73: Port Profile Custom Setup Settings for Data Center Switching ELS (*continued*)

Field	Action
Priority	<p>Select a power priority for the PoE port—either Low or High. If there is a shortage of PoE power on the switch, power to low priority ports is shut down before power to high priority ports. Within ports with the same assigned priority, power priority is determined by port number—ports with a lower port number have a higher power priority.</p> <p>Default: low priority</p>
Disable PoE	Select to disable PoE on the interfaces using this Port profile.

Port Security (Switching Interfaces Only)

Select to enable port security (default); clear to disable port security.

When port security is enabled, you can configure port security options such as learned MAC address limits on an interface. When port security is disabled, no port security is applied to the interface, including the default port security options.

Trust DHCP	<p>Select to permit messages from a DHCP server to be received on the interface—this is the default. Clear to block all messages from a DHCP server from being received on the interface.</p> <p>TIP: For this port security feature to work, DHCP snooping must be enabled on the VLAN the interface belongs to. You can enable DHCP snooping on the VLAN in the VLAN profile. For directions, see “Creating and Managing VLAN Profiles” on page 329.</p>
FCoE Trusted	Select to configure the interface to trust Fibre Channel over Ethernet (FCoE) traffic. If an interface is connected to another switch such as an FCoE forwarder (FCF) or a transit switch, you can configure the interface as trusted so that the interface forwards FCoE traffic from the switch to the FCoE devices without installing FIP snooping filters.
MAC Limit	<p>Type the number of MAC address that can be dynamically learned on the interface.</p> <p>Range: 1 through 163839</p> <p>Default: For Desktop Ports, 1. For Desktop Phone Ports, 2. For all others, none.</p>

Table 73: Port Profile Custom Setup Settings for Data Center Switching ELS (continued)

Field	Action
MAC Limit Action	<p>Select the action to be taken if the MAC address limit is exceeded:</p> <ul style="list-style-type: none"> • Drop—Drop any packet with a previously unlearned MAC address and generate a system log entry, and SNMP trap, or an alarm. This is the default for a Desktop Port and Desktop Phone Ports. • Log—Accept packets with new MAC addresses and learn the addresses, but generate a system log entry, and SNMP trap, or an alarm. • Shutdown—Shut down the interface and generate a system log message, SNMP trap, or an alarm. If an interface is shut down because the MAC address limit has been exceeded, you must use the CLI command clear ethernet-switching port-error interface <i>name</i> to clear the error and bring the interface back into service. TIP: You can use the CLI to configure auto-recovery on an interface that has been shut down by a MAC limit error. • None—No action. This selection effectively disables MAC address limiting on the interface. This is the default for Switched Uplink Ports, Switched Downlink Ports, and Server Ports.
Allowed MAC List	<p>Indicate the MAC addresses of devices that are allowed access to the interface in the Allowed MAC List. Any device whose MAC address does not match an address in the list will not be allowed access to the interface. A list with no entries means that a client with any MAC address is permitted to access the interface.</p> <p>To enter a MAC address, click Add and then type the MAC addresses in the field provided. Enter MAC addresses as two-character hexadecimal numbers separated by colons. Click Save to save the entry.</p> <p>NOTE: Configuring an allowed MAC address list does not block the switch from receiving Layer 2 control packets—such as Link Layer Discovery Protocol (LLDP) packets—transmitted from MAC addresses that are not specified in the address list. Control packets do not undergo the MAC address check. However, the switch does not forward them to another destination.</p> <p>Default: No entries</p>

Table 73: Port Profile Custom Setup Settings for Data Center Switching ELS (*continued*)

Field	Action
RSTP Settings In addition to enabling or disabling the Spanning Tree Protocol (STP) as part of device profiles, this feature enables you to fine-tune STP by setting interfaces into edge, disable, or no-root-port states.	Edge RSTP defines the concept of an edge port, which is a designated port that connects to non-STP-capable devices, such as PCs, servers, routers, or hubs that are not connected to other switches. Because edge ports connect directly to end stations, they cannot create network loops and can transition to the forwarding state immediately. You can manually configure edge ports, and a switch can also detect edge ports by noting the absence of communication from the end stations. Disable Disables the RSTP on interface. NOTE: Configuring interfaces to one of these states is not mandatory for ELS switches. Hence, the option Disable is not applicable for ELS switches and therefore not supported. No Root Port Configures an interface to be a spanning-tree designated port. If the bridge receives superior STP bridge protocol data units (BPDUs) on a root-protected interface, that interface transitions to a root-prevented STP state (inconsistency state) and the interface is blocked. This blocking prevents a bridge from being elected the root bridge. When the bridge stops receiving superior STP BPDUs on the root-protected interface, interface traffic is no longer blocked.
CoS Settings (All except Fibre Channel Type)	Click Select Cos Profile to choose from existing CoS profiles. The CoS configuration contained in the CoS profile is applied to the interfaces that the Port profile is assigned to when you deploy the configuration. Select the type of port scheduling for the CoS profile. Port scheduling depends on the device model. When you select a port scheduling type, Network Director displays the devices that support the selected port scheduling type. Click OK . Some preconfigured Service Types have a default CoS profile—see Service Types for details. Or Click Configure CoS settings to configure CoS profile. Select the type of port scheduling for the CoS profile. Port scheduling depends on the device model. When you select a port scheduling type, Network Director displays the devices that support the selected port scheduling type. See “Creating and Managing Wired CoS Profiles” on page 400 for steps to configure a CoS profile.

Table 73: Port Profile Custom Setup Settings for Data Center Switching ELS (*continued*)

Field	Action
Authentication Settings (Desktop Port, Desktop Phone Port, Custom Port)	<p>Select the Authentication profile for the interface from a list of existing profiles by clicking Select, selecting one of the listed profiles, and then clicking OK. By assigning an Authentication profile to the Port profile, you can enable 802.1x and captive portal authentication on interfaces.</p> <p>If you do not specify an Authentication profile, the interface is an open port and no authentication is required to connect.</p> <p>NOTE: You cannot configure 802.1x authentication on aggregated Ethernet interfaces. If you attempt to deploy a Port profile that contains an Authentication profile on an aggregated Ethernet interface, the deployment fails.</p> <p>Or</p> <p>Click Configure Authentication Settings to configure 802.1x and captive portal authentications. See “Creating and Managing Authentication Profiles” on page 237 for steps to configure the authentication profile.</p>
Filter Settings (available for all Service Types, including Custom for routing)	<ul style="list-style-type: none"> ● Ingress Filter <p>Select an Ingress Filter for the interface by clicking Select, selecting one of the listed filters, and then clicking OK.</p> ● Egress Filters <p>Select an Egress Filter for the interface by clicking Select, selecting one of the listed filters, and then clicking OK.</p>
VRRP Settings (available when Service Type is Custom and Family Type is Routing)	Select the VRRP profile for the interface from a list of existing profiles by clicking Select . Select one of the listed profiles, and then click OK .

Clicking **Done** displays the dialog Do you want to assign Port Profile to Ports. click **Yes** to create a profile assignment else click **No** to navigate to the Manage Port Profile page and to create the Port assignment later.

What to Do Next

After you create a Port profile, you can assign it to interfaces or members of port groups. During this process, you can also configure interface-specific attributes, such as IP address, and certain device-specific attributes, such as which Access profile to use for all ports on the device. For more information, see [“Assigning and Unassigning Port Profiles from Interfaces” on page 304](#).

RELATED DOCUMENTATION

Understanding Port Profiles 246
Assigning and Unassigning Port Profiles from Interfaces 304
Creating and Managing Port Groups 323
Creating and Managing VLAN Profiles 329
Creating and Managing Authentication Profiles 237
Assigning and Unassigning Port Profiles from Interfaces 304
Understanding VRRP Profiles 544
Creating and Managing VRRP Profiles 545
Network Director Documentation home page

Assigning and Unassigning Port Profiles from Interfaces

IN THIS SECTION

- [Selecting Devices for Assignment | 305](#)
- [Selecting Interfaces for Assignment | 306](#)
- [Reviewing and Accepting the Assignments | 309](#)
- [Editing Profile Assignments | 310](#)
- [Unassigning a Port Profile from an Interface | 311](#)

You can assign an existing user-created or system-created Port profile to network interfaces (including aggregated Ethernet interfaces), or Port Group member interfaces on one or more devices.


During the process of assigning a Port profile to interfaces, you can also:

- Configure IPv4 or IPv6 addresses on interfaces to which you have assigned a routing Port profile.

TIP: IPv4 filters are separate from IPv6 filters.

- Configure certain authentication attributes—such as the RADIUS server or servers to use—for all 802.1X interfaces on the device. Because configuring these attributes involves assigning an Access profile to the device, you must have previously created an Access profile.

To assign a Port profile to interfaces:

1. Click  in the Network Director banner.
2. Under Select View, select one of the following views: **Logical View**, **Location View**, **Device View** or **Custom Group**.

TIP: Do not select **Datacenter View** or **Topology View**.

3. In the Tasks pane, select **Wired > Profiles > Port**.
The Manage Port Profile page is displayed.
4. Select the Port profile you want to assign and then click **Assign**.
The Assign Port Profile wizard appears. It has three parts—Device Selection, Profile Assignment, and Review.
5. Complete device selection for assignment by following the directions [“Selecting Devices for Assignment” on page 305](#).
6. Assign the port profile to one or more objects by following the directions [“Selecting Interfaces for Assignment” on page 306](#).
7. Review your configuration by following the directions [“Reviewing and Accepting the Assignments” on page 309](#).
8. Click **Finish**.

After you assign a Port profile to ports, you can modify your assignments by selecting the Port profile from the Manage Port Profiles page and clicking **Edit Assignments**.

The following sections describe how to use the Assign Port Profile wizard and the Edit Assignments page.

Selecting Devices for Assignment

Use the Device Selection page in the Assign Port Profile wizard to select one or more devices that have ports. You can select container nodes, individual devices, or port groups. For more information about Port Groups, see [“Creating and Managing Port Groups” on page 323](#).

To select devices for Port profile assignment:

1. Enable either **Select Devices** or **Select Port Groups**.
2. If you enabled **Select Devices**, expand the list of objects and select the objects that contain the devices and interfaces you want to assign by clicking the check box next to the them. If you select a container node, all devices under that node are selected.

TIP: The list of objects is filtered to include only devices that match the profile's family type. If you do not see a device that you expected to see, verify that the device matches the profile's family type.

3. If you enabled **Select Port Groups**, select one or more port groups from the Select Port Group list.
4. Click either **Next** or **Profile Assignment** to proceed to the next step in the wizard, Profile Assignment.

For directions to complete Port Profile Assignment, see ["Selecting Interfaces for Assignment" on page 306](#).

Selecting Interfaces for Assignment

Use Profile Assignment in the Assign Port Profile wizard to select the interfaces to which you want to assign the Port profile. After you have selected the interfaces, you can configure specific attributes on the interfaces or on the devices to which the interfaces belong.

TIP: Before you start the procedure below, you might want to select a device and click **View Assignments** to view what profiles and attributes are already configured on the device. Any profile assignments or attributes you define during this procedure replace the existing ones.

One optional attribute you can configure for switching interfaces is the Access profile that defines RADIUS server authentication for 802.1X ports. If you will be configuring this optional attribute, make sure that an Access profile has been created.

NOTE: During the assignment of Port profiles, Network Director excludes the details of already assigned ports that enable you to assign Port Groups to unassigned Port Profiles.

If you enabled **Select Port Groups** during Object Selection, you can assign the Port profile to any or all existing port groups.

If you enabled **Select Devices** during Object Selection, assign the Port profile to interfaces and configure the port-specific or device-specific attributes:

1. Select one or more container nodes or devices from the Assignments list:

- To assign the profile to nonconsecutive interfaces or to aggregated Ethernet interfaces, select a single device.
- To assign the profile to interfaces in the same consecutive interface range (for example, ge-0/0/0 through ge-0/0/15) on one or more devices, select one or more devices. To make multiple selections, press Shift or Ctrl while making the selections.
- To assign a profile to ports within a Virtual Chassis Fabric (VCF), you can select any container nodes or member devices within the VCF, including the VCF container node.
- To assign a profile to aggregated Ethernet ports within a Virtual Chassis or VCF, select the Virtual Chassis or VCF container node. To assign a profile to physical device ports within a Virtual Chassis or VCF, select one or more member devices.
- Channelized ports are only applicable for Data Center Switching ELS devices and only XE interfaces can be used as channelized ports.

NOTE: If Network Director fails to read the configuration of one or more devices after the device discovery, such devices are not displayed in the Assignments list. You will not be able to assign profiles to such devices. The Manage Jobs page in System mode displays details of the device discovery jobs. Use the information displayed on this page to take appropriate corrective steps to enable Network Director to reread the configuration of the failed device. For more information, see [“Discovering Devices in a Physical Network” on page 94](#).

2. Click **Assign to Port**.

The Assign Profile to Ports window opens.

3. Select either **Ports** (default) or **Port Range**. If you selected multiple devices in the previous step, you cannot choose the Port option.

4. If you selected the Port option, select the ports from the list of ports.

By default, aggregated Ethernet interfaces are listed after the ge- and xe- interfaces in the list of ports. Members of aggregated Ethernet interfaces are not included in the port list.

5. If you selected the Port Range option, enter the port range:
 - a. In the Normal Ports section, enter a first and last port name in the text boxes, then click **Add**. The port range appears in the Selected Port Range section.
 - b. Repeat the add process to add any additional port ranges.
 - c. To delete a port range, select its check box, then click **Delete**.

At least one port within the port range must be available on each selected device for the port range to succeed. Channelized ports are supported in a port range. Assignments are created for the ports within the port range that are available. You can assign the profile to the same interface on multiple devices by entering the interface name in both fields of the port range.

6. Click **Assign** to complete the port assignments and close the window.

The port assignment appears in the list of Assignments, with the Device, Type, Assigned To, and Attributes columns completed. In the Attributes column, you see a triangle and the link **Define**.

7. Configure the following port-specific or device-specific attributes:

- If the Port profile is a switching profile that contains an Authentication profile—in other words, the profile is enabling 802.1X authentication on ports—click the **Define** link in the Attributes column for a device to define additional authentication attributes.

The Configure attributes window opens. Fill in the fields described in [Table 74 on page 308](#).

Table 74: Configure Device Attributes for Port Profile Assignments

Field	Action
Access Profile	<p>Select an Access profile.</p> <p>The RADIUS server attributes defined in the Access profile is configured on the device when you deploy the configuration.</p>
Radius Server Source IP Address	Type an IP address to be used as the source IP address for RADIUS server requests sent by the switch. The source address must a valid IPv4 or IPv6 (either format) address configured on one of the switch interfaces.
Post authentication URL	Type a URL to be used for the captive portal post-authentication website.

TIP: If you see the message *Port profile does not have an associated Authentication profile. Please configure the Authentication profile.*; then click **OK**, and edit the Port profile by selecting **Port** under Profile and Configuration Management, selecting the Port profile from the list and clicking **Edit**. The Authentication profile association is located in the Port Family Options section.

NOTE: The attributes you define for the device apply to all 802.1X authenticator interfaces on the switch. Different sets of interfaces on the switch cannot have different attributes.

- If the Port profile is a routing profile, click the **Define** link in the port's Attributes field to configure an IPv4 or IPv6 address on the interface.

Repeat this step for all the ports on which you want to configure IPv4 or IPv6 interfaces.

8. Repeat the previous steps as needed to complete the port assignments and then click either **Next** or **Review**.

For review directions, see [“Reviewing and Accepting the Assignments” on page 309](#).

Reviewing and Accepting the Assignments

Use the Review step of the Assign Port Profile wizard to review and accept your assignments:

- Click **Edit** to return to the Profile Assignment step and make changes to your assignments.
- Click **Finish** to accept the assignments.

After you click Finish, the Create Profile Assignments Job Details window opens, which reports on the status of the profile assignment job. If you have assigned the profile to a large number of objects, the profile assignment job can take some time to complete. Instead of waiting for the Job Details window to report job completion status, you can close it and check the details of the profile assignment job at a later time by using the Manage Job task in System mode.

NOTE: If any assignment fails, the profile assignment job fails and none of the assignments are created. Check the details for the profile assignment job for information about why the assignment failed.

After the profile assignment job completes, you can deploy the configuration defined in the Port profile and in the port-specific and device-specific attributes on the affected devices. See [“Deploying Configuration to Devices” on page 580](#).

Editing Profile Assignments

Use the Edit Assignments page to change Port profile assignments. You can:

- Delete a port from the profile assignments.

If the profile has been already deployed on the port, then the configuration is removed from the port when you next deploy the configuration on the device. The configuration removed includes any port configuration that was defined in associated profiles, such as the CoS, Authentication, and IPv4 or IPv6 Filter profiles.
- Change the IPv4 or IPv6 address for ports associated with a routing Port profile.
- Change the device-specific authentication attributes, such as the Access profile associated with the device. For more information about these attributes, see [Table 74 on page 308](#).

NOTE: You cannot assign the Port profile to additional ports by using the Edit Assignment page. To add port assignments, use the Assign Port Profile wizard.

[Table 75 on page 310](#) describes the fields in the Edit Assignments page and how to use them to change the profile assignments. When you are finished with your modifications, click **Apply**. You can then deploy your modifications on the affected devices.

Table 75: Edit Assignment for Port Profile Fields

Field	Description
Objects	Expand the device nodes to see the ports or port group the profile is assigned to.
Assigned State	Indicates the current state of profile assignment on the port: <ul style="list-style-type: none">• Deployed—The profile configuration has already been deployed on the port.• Pending—The profile configuration has not yet been deployed on the port.• Pending Removal—The profile configuration was deployed on this port, but will be removed from the port the next time the device configuration is deployed.
Attributes	If the attributes for a port or device are currently undefined, you can click the Define link to define them. If attributes have been defined and you want to view them or change them, click the Change link.
Operation	Click the Delete link to delete the profile assignment from the port.

Table 75: Edit Assignment for Port Profile Fields (*continued*)

Field	Description
Record Status	<p>Shows the current assignment status:</p> <ul style="list-style-type: none"> • An X indicates that you have marked a port for deletion. • A pencil indicates that you have modified the associated attribute. <p>After you apply your assignment changes, these indicators disappear.</p>

Unassigning a Port Profile from an Interface

Starting Network Director Release 3.5, you can unassign multiple port profiles that are assigned to multiple ports, at the same time.

To unassign port profiles:

1. On the Network Director banner, under **Views**, select one of the following views—Logical View, Location View, Device View, or Custom Group.

2. On the Tasks pane, click **Wired > Profiles > Port**.

The Manage Port Profile page appears.

3. Select one or more port profiles that you want to unassign from the ports and click **Unassign**.

A confirmation message indicating the profiles were successfully unassigned appears and the status of the profiles change to Pending Deployment.

RELATED DOCUMENTATION

[Creating and Managing Port Profiles | 251](#)

[Creating and Managing Access Profiles | 218](#)

[Creating and Managing Port Groups | 323](#)

[Network Director Documentation home page](#)

Managing Auto Assignment Policies

To support rapid network deployment, Network Director enables you to define your network configuration in a set of profiles that you can apply to multiple objects in your network. Auto assignment policies go one step ahead and further automate profile assignment. When Network Director detects the devices included in a policy, in a campus network, the Port profiles that you have created in Network Director are automatically assigned to various switch ports on supported devices.

Network Director uses LLDP to detect the type of network device. When the devices such as Desktop, Desktop Phone, Server Port, and Printer are LLDP enabled, Network Director triggers auto assignment of Port profiles for these devices.


NOTE: Network Director detects most of the printers by using organizationally unique identifier (OUI). For more information, see [“Adding and Managing OUI Data in Network Director” on page 150](#).

To create an auto assignment policy, you specify one or more Port profiles, devices, ports or port ranges on the devices to which the Port profile is to be deployed, and a few additional parameters.

After you create an auto assignment policy, when any of the device ports that you specified in the auto assignment policy are connected to a Desktop, Desktop Phone, Server Port, or Printer, Network Director performs the following tasks:

- a. Initiates a job to update the Port profile configuration on the connected ports. If you enable a policy, Network Director overwrites the configuration that is already deployed on the ports and deploys the configuration from the profile that you specified in the auto assignment policy. You can view details of this job in the Job Management page and also from the Policy Assignment Log window.
- b. Updates the port associations and displays the results in the Manage Port Profiles page. In the Manage Port Profiles page, the Port profiles that are assigned through an auto assignment policy are highlighted with a ★ next to the profile name. You can view more details about the auto assignment in the Port Profile Details window.

To manage auto assignment policies:

1. Click  in the Network Director banner.
2. Under Views, select one of the following views: **Logical View**, **Location View**, **Device View**, or **Custom Group**.

TIP: Auto assignment is not available in **Datacenter View** or **Topology View**.

3. Click **Wired > Tasks > Manage Auto Assignments** under the Tasks pane.

The Manage Auto Assignment Policy page opens.

4. From the Manage Auto Assignment Policy page, you can:

- Create a new policy by clicking **Create**. For details, see [“Creating Auto Assignments” on page 314](#).
- Modify an existing policy by selecting the policy and clicking **Edit**.
- View information about a policy by selecting the policy and clicking **Details**. Network Director opens the Auto Assignment Policy Details page.
- Delete a policy by selecting the policy and clicking **Delete**.
- Deploy a policy by selecting the policy and clicking **Run Now**.

[Table 76 on page 313](#) describes the information provided about the policies on the Manage Auto Assignment Policy page. This page lists all auto assignment policies defined for your network, regardless of your current selected scope in the network view.

Table 76: Manage Auto Assignment Policy Page

Column	Description
Name	Unique name given to the auto assignment policy when the policy was created.
Description	Description of the policy that was entered when the policy was created.
Device Family Mode	Displays one of the following: <ul style="list-style-type: none"> • EX—for EX Series switches • ELS—for Campus Switching ELS
Log	Displays link details corresponding to each policy. Click Details corresponding to each policy. The Policy Assignment Log window opens listing the internal log of a policy.
Creation Time	Date and time when the policy was created.
Update Time	Date and time when the policy was last modified.

TIP: All columns might not be currently displayed. To show or hide fields in the table, click the down arrow on the field header, select **Columns**, and select or clear the check box adjacent to the field that you want to show or hide.

RELATED DOCUMENTATION

| [Creating Auto Assignments](#) | 314

Creating Auto Assignments

IN THIS SECTION

- [Adding Port Profiles using the Select Port Profiles Page](#) | 314
- [Adding Devices and Ports for Auto Assignment](#) | 315
- [Viewing the Auto Assignment Policy Summary](#) | 316

Auto assignments help in automating Port profile assignments. Auto assignments automatically assign Port profiles that you have defined, to various switch ports based on the endpoints detected on the port on supported devices when Network Director detects the devices included in the auto-profile in the network. Auto assignments are supported for desktop ports, desktop phone ports, server ports, and printer ports.

You can create auto assignment policies using the Create Auto Assignment wizard. The Create Auto Assignment wizard consists of three pages—Select Port Profile, Select Devices, and Summary.

To open the Create Auto Assignment wizard, click **Create** in the Manage Auto Assignment Policy page. Perform the following tasks to create an auto assignment:

Adding Port Profiles using the Select Port Profiles Page

To create an auto assignment policy and add Port profiles:

1. Enter a name and a description for the auto assignment.
2. Select the **Enable Policy** check box to enable the policy for the auto assignment.

3. Select a device and specify a port range for this auto-assignment in the next wizard page.
4. Select a device family. Auto assignment supports devices that belong to *Switching (EX)* or *Campus Switching ELS*.
5. Click **Select** in the Port Profiles table to add Port profiles for auto assignment. Network Director opens the Select Profiles window and displays only those Port profiles that are created for the device family that you selected. Select one or more Port profiles that you want to add to the auto assignment and click **Add**.

NOTE: If there are multiple profiles of the same service type for a device you can select only one of the profiles for auto assignment. This ensures that you can assign only one service type profile across all auto assignment policies. For example, if the Select Profiles window lists two Port profiles with Server as the service type, then you can select only one of these two profiles for auto assignment.

Network Director adds the selected Port profiles to the Port Profiles table. To remove a Port profile from the list, select a profile and click **Remove**.

6. Click **Next** to add devices for auto assignment.

The Select Device page opens.

Adding Devices and Ports for Auto Assignment

To add devices and specify ports for auto assignment:

1. Click **Select Devices**. The Select Devices window opens.
2. Expand the list of objects and select the objects that contain the devices you want to add by clicking the check box next to them. If you select a container node, all devices under that node are selected.

NOTE: The list of objects is filtered to include only devices that match the Port profile family type. If you do not see a device that you expected to see, verify that the device matches the profile's family type. For example, a Port profile created for Switching (EX) family type cannot be assigned to an auto assignment policy that you are creating for Campus Switching ELS.

3. Click **OK** to add the selected devices to the Select Device(s) table.

4. Select a device from the Select Device(s) table and click **Configure Range** to specify the ports or port ranges to which you want to auto assign the Port profiles.

The Configure Port Inclusion window opens.

NOTE: Network Director deploys the Port profiles only if you specify a port range for the auto assignment policy and when the end points match.

5. Select the starting and ending port numbers and click **Add** to add the port range to the Selected Port Range table. You can add multiple port ranges that do not overlap with each other, for each device across policies.
6. When you have added all the required port ranges, click **OK**.
7. Repeat steps 4 through 6 to add port ranges for all the devices that you have added to the Select Devices table.
8. To remove a device from the Select Devices table, select the device and click **Remove**.
9. Click **Next** to view a summary of the auto assignment policy.

Viewing the Auto Assignment Policy Summary

The Summary page displays the details of the auto assignment policy. You can review and make modifications to the auto assignment policy. To modify the configuration details, click the appropriate buttons in the Auto Assignment Policy workflow at the top of the page that corresponds to the configuration you want to change.

When you are finished with your modifications, click **Summary** to return to this page.

To complete the auto assignment policy creation, click **Finish**.

Network Director displays the policy that you created in the Manage Auto Assignment Policy page.

When any of the device ports that you specified in the auto assignment policy are connected to a Desktop, Desktop Phone, Server Port, and Printer, Network Director automatically assigns Port profiles to these device ports. You can view the profile deployed for a port in the Policy Assignment Log window for each policy.

RELATED DOCUMENTATION

Configuring Easy Config Setup

IN THIS SECTION

- [Configuring Interface Settings | 317](#)

In addition to the Port profile configuration, Network Director enables users to quickly configure interfaces on devices by using the Easy Config Setup task. You can perform configurations by directly selecting the device, instead of creating a new profile and assigning a profile to the device port. You can also deploy the configuration changes without creating additional profiles which results in growing number of profiles in Port profile configurations. Easy Config Setup is supported only for the configurations that are automatically approved; for configurations that require manual approvals, this task is disabled.


Configuring Interface Settings

This section describes the steps to configure the interface settings by using Easy Config Setup.

You can configure the following interface settings in the EX Switching, Campus Switching ELS (MX series devices are supported in L2NG mode only and are not supported in native MX series mode), and Data Center ELS devices.

- Port Settings
- VLAN Settings
- PoE Settings
- 802.1x Settings
- Access Settings

To configure an interface in a device by using the Easy Config Setup:

1. Select a switching device from the left navigation pane.
2. Click  in the Network Director banner.
3. Under Select View, select either **Logical View**, **Location View**, or **Device View**.

4. In the Tasks pane, click **Wired > Tasks > Easy Config Setup**.

NOTE: This task is not visible for a fabric device.

5. Enter the settings for the interface described in [Table 77 on page 318](#).

Table 77: Easy Config Setup Settings

Field	Action
Port Settings	
Ports/Interfaces	Select an Ethernet switching interface, an IPv4 routing interface, or an IPv6 routing interface. All the ports associated with the device (except Layer 3 interfaces) are available in the list.
Description(optional)	Provide a description of the device configuration or port details of the device. You can use up to 256 characters.
Port Mode	Configure a switching interface port to be an access, trunk, or tagged-access port for EX Series switches. Campus Switching ELS and Data Center Switching ELS series devices supports access mode and trunk mode. For more information about port modes, see Creating and Managing Port Profiles .
Disable Port	Disables the port. You can still configure and deploy all the settings but these settings become active only when you enable the port by clearing this selection.

Table 77: Easy Config Setup Settings (*continued*)

Field	Action
<p>Member VLAN Settings</p> <p>You can enable VLAN settings and display the configuration options by enabling Member VLAN Settings.</p>	<p>You can either select an existing VLAN profile or create a new VLAN profile that you want to assign to the port.</p> <p>To select an existing profile:</p> <ol style="list-style-type: none"> Select the option Select VLAN Profile. Select the option Select. The Choose VLAN profile window opens. Select the VLAN profile name and click OK. <p>To create a new profile:</p> <ol style="list-style-type: none"> Select Configure VLAN Settings. Click Create. The Create VLAN Profile window opens. Enter the VLAN name. Under VLAN ID, select Single and enter a VLAN ID from 1 to 4094 if you want to configure a single VLAN. or Under VLAN ID, select Range and enter a range of VLAN IDs that you want to assign to the VLAN profile. TIP: Single VLAN IDs can be configured for all products. VLAN lists or VLAN ID ranges are available for some products, depending on the technology used for implementation. Click OK.
PoE Settings	
You can enable PoE and display the configuration options by enabling PoE Settings .	
Maximum Power(W)	Use the arrows to adjust the maximum PoE power in watts allocated to a PoE port. The powered device cannot draw more power than the wattage specified. If it does, PoE power to the port is shut down.

Table 77: Easy Config Setup Settings (*continued*)

Field	Action
Priority	<p>Select a power priority for the PoE port—either Low or High. If there is a shortage of PoE power on the switch, power to low priority ports is shut down before power to high priority ports. Within ports with the same assigned priority, power priority is determined by the port number—ports with a lower port number have a higher power priority.</p> <p>Default: low priority</p>
Disable PoE	Select to disable PoE on the interfaces that use this Port profile.

802.1x Settings (Authentication)

You can configure 802.1x and display the configuration options by enabling **802.1x Settings (Authentication)**.

Enable 802.1x	<p>802.1x authentication is enabled by default for a switching profile. 802.1x authentication works by using an Authenticator Port Access Entity (the switch) to block all traffic to and from a supplicant (end device) at the port until the supplicant's credentials are presented and matched on the Authentication server (a RADIUS server). When authenticated, the switch stops blocking traffic and opens the port to the supplicant. Network access can be further defined using VLANs.</p>
Enable MAC-RADIUS	<p>Select to enable MAC-RADIUS based authentication for this profile. MAC RADIUS authentication enables LAN access to permitted MAC addresses. When a new MAC address appears on an interface, the switch consults the RADIUS server to check whether the MAC address is a permitted address. If the MAC address is configured on the RADIUS server, the device is allowed access to the LAN.</p>
Supplicant Mode	<p>Specify the mode authentication supplicants use, either Single, Multiple, or Single-Secure.</p> <ul style="list-style-type: none"> • Single—Allows only one host for authentication. This is the default mode. • Single-Secure—Allows only one end device to connect to the port. No other end device is enabled to connect until the first logs out. • Multiple—Allows multiple hosts for authentication. Each host is checked before being admitted to the network.
Guest VLAN	<p>Click Select and then select the VLAN to which an interface is moved when no 802.1x supplicants are connected on the interface. The VLAN specified must already exist on the switch.</p>

Table 77: Easy Config Setup Settings (*continued*)

Field	Action
Reject VLAN	Click Select and then select the VLAN to which an interface is moved when the switch receives an Extensible Authentication Protocol over LAN (EAPoL) Access-Reject message during the authentication process between the switch and the RADIUS authentication server.
Server Fail Type	<p>Specify the server fail fallback action the switch takes when all RADIUS authentication servers are unreachable; one of None, Deny, Permit, Use cache, or VLAN Name.</p> <ul style="list-style-type: none"> • None—No server fallback action is used. This option is selected by default. • Deny—Force fails supplicant authentication. No traffic will flow through the interface. • Permit—Force succeeds the supplicant authentication. Traffic will flow through the interface as if it were successfully authenticated by the RADIUS server. • Use cache—Force succeeds the supplicant authentication only if it was previously authenticated successfully. This action ensures that already authenticated supplicants are not affected. • VLAN Name—Move supplicant on the interface to the VLAN specified by this name. This action is allowed only for the first supplicant connecting to an interface. If an authenticated supplicant is already connected, then the supplicant is not moved to the VLAN and is not authenticated. If you select this option, you must provide a Fail VLAN name.

Access Settings

You can configure authentication parameters and accounting parameters on the network and display the configuration options by enabling **Access Settings**.

Server Address	Enter the IP address of the RADIUS server.
Authentication Port	The default RADIUS authentication port is 1812. You can change the port number by using the up and down arrows.
Secret	Provide a password. If the password contains spaces, enclose it in quotation marks. The secret password used by the switch must match the one used by the server.
Retry Count	Specify the number of times that a device attempts to contact the LDAP authentication server. The default retry count is 3. You can change this value by using the up and down arrows to 1 through 100 times.

Table 77: Easy Config Setup Settings (*continued*)

Field	Action
Timeout (seconds)	Specify the number of seconds the switch waits to receive a response from a RADIUS server. The default timeout is 5 seconds. You can change this value, using the up and down arrows, to 1 through 65535 seconds.

6. Click **Preview** to view the configuration changes that will be deployed to a device when a job runs. Use the Configuration window to see changes that were deployed to a device when a completed job ran. The configuration changes are shown in these formats:

- Select the CLI View tab to view the configuration changes in CLI format. This view shows the Junos configuration statements that will be deployed to the device.
- Select the XML View tab to view the configuration changes in XML format. This view shows the XML-formatted configuration that will be deployed to the device by using the Device Management Interface (DMI), which is used to remotely manage devices.

7. Click **Deploy** to deploy configuration to a device.

After you deploy the configuration, the device goes out of sync and Network Director triggers auto-resynchronization of the device. If there is any conflicting configuration, the new port profile (created during easy config setup) is prompted during the assignment.

The Deploy EasyPortal Configuration window opens. If you chose to deploy the changes immediately, the Deployment Status column shows the status as **INPROGRESS** and changes to **SUCCESS** after the deployment is successfully completed.

Click **Cancel** if you want to cancel your changes.

8. Click **Close** to close the deployment page.

NOTE: Clicking either **Close** or **Cancel** takes you to the Device Inventory My Network page, which displays the details of the device you selected in the first step.

Understanding Port Groups

Ports on virtual and physical devices regulate data-packet traffic to both ensure security and a guaranteed rate of packet flow, and prevent unsolicited traffic. Since configuring each port individually would be tedious, especially if the ports are configured with the same settings, port groups enable configuration of multiple ports simultaneously. First you create the port group (see [“Creating and Managing Port Groups” on page 323](#), and then you can assign a Port profile to the Port Group—see [“Assigning and Unassigning Port Profiles from Interfaces” on page 304](#).

The ports in a port group can be located on any of your switches and can include ports from different devices and from different series. Group port configuration has precedence over any individual port configuration.

RELATED DOCUMENTATION

[Creating and Managing Port Groups | 323](#)

[Assigning and Unassigning Port Profiles from Interfaces | 304](#)

[Assigning a VLAN Profile to Devices or Ports | 345](#)

[Network Director Documentation home page](#)

Creating and Managing Port Groups

IN THIS SECTION

- [Managing Port Groups | 324](#)
- [Creating Port Groups | 325](#)
- [Specifying Settings for a Port Group | 326](#)
- [What to Do Next | 326](#)

From Network Director, you can group ports and then name that port group. The ports can be located on any of your switches and can include ports from different devices and from different series. Creating Port groups enables simultaneous multiple port configuration. For example, when a Port profile (see [“Creating and Managing Port Profiles” on page 251](#)) is assigned to the members of a Port group (see [“Assigning and Unassigning Port Profiles from Interfaces” on page 304](#)), all ports in the group are configured with the Port profile.

NOTE: Configuration applied to members of a port group has precedence over any individual port configuration.

This topic describes:

Managing Port Groups

Use the Manage Port Groups page to manage existing Port groups and to create new ones. Port groups enable simultaneous multiple port configuration.

From the Manage Port Groups page, you can:

- Create a new Port Group by clicking **Add**. For directions, see [“Creating Port Groups” on page 325](#).
- Modify an existing Port Group by selecting it and clicking **Edit**.
- View information about a Port Group by selecting the group and clicking **Details** or by clicking the group name.
- Delete a Port Group by selecting a group and then clicking **Delete**.

TIP: To see the current assignments for a group, click the group name.

[Table 78 on page 324](#) describes the information provided about Port Groups on the Manage Port Groups page. This page lists all Port Groups defined for your network, regardless of your current selected scope in the network view.


Table 78: Manage Port Groups Information

Field	Description
Name	Name given to the group when the group was created. Click the group name to view group details.
Last Updated	Date the port group was last altered.
User	User name of the person who last altered the Port Group.

Creating Port Groups

TIP: The required configurations for a Port Group are a Port Group name and the configuration of at least one device port.

To create a Port group for switches:

1. Select one of the following device views in the Network Director banner:
 - **Logical View**—Displays devices in hierarchal groupings based on logical relationships.
 - **Location View**—Displays devices in hierarchal groupings based on physical locations.
 - **Device View**—Displays devices in hierarchal groupings based on device type.
 - **Custom Group View**—Displays devices in hierarchal groupings based on custom group.
2. Click  in the Network Director banner.
3. Under Select View, select either **Logical View**, **Location View**, **Device View** or **Custom Group View**.

TIP: Do not select **Datacenter View** or **Topology View**.

4. Click **Manage Port Groups** under Device Management in the Tasks pane.
The Manage Port Groups page appears.
5. Click **Add**.
The Create Port Group page appears.
6. Enter settings for the Port Group as described in [“Specifying Settings for a Port Group” on page 326](#).
7. Click **Done**.
The message *Port group successfully created* is displayed.
8. Click **OK**.
The new port group appears on the list of managed port groups.

Specifying Settings for a Port Group

Use the Create Port Group page to define the members of a Port Group.

[Table 79 on page 326](#) describes the settings available in the Port Group.

Table 79: Port Group Settings

Field	Action
Port Group Name	Type the name of the Port Group, using up to 64 alphanumeric characters and no special characters other than the underscore. The name must be unique among Port groups.
Port Group Description	Type a description of the Port Group, which will appear on Manage Port Groups page. You can type up to 256 characters.

Select Devices and Ports

Add Ports to the Port Group (task)	<p>To add ports to this Port Group:</p> <ol style="list-style-type: none"> 1. Select a device with ports from the tree displayed in the column labelled Select Device and Ports. 2. For the current selection (highlighted device), click either All Ports or Selected Ports. If you click All Ports, the ports or port range is immediately listed under Selected Ports and Port Ranges. If you click Selected Ports, you must then select individual ports from the Port Selector list, and click Done. 3. Click Done. An information window displays the message <i>Port group created successfully</i>. 4. Click OK to close the information window. The new Port Group is now listed under Manage Port Groups. <p>TIP: To edit a port group, select it from the Manage Port Groups list and then click Edit. To remove a port group, select it and then click Delete.</p>
------------------------------------	--

What to Do Next

After you create a Port Group, you can treat it like an individual port—see [“Assigning and Unassigning Port Profiles from Interfaces” on page 304](#).

RELATED DOCUMENTATION

[Creating and Managing Port Profiles | 251](#)

[Assigning and Unassigning Port Profiles from Interfaces | 304](#)

[Assigning a VLAN Profile to Devices or Ports | 345](#)

[Understanding Port Groups | 323](#)

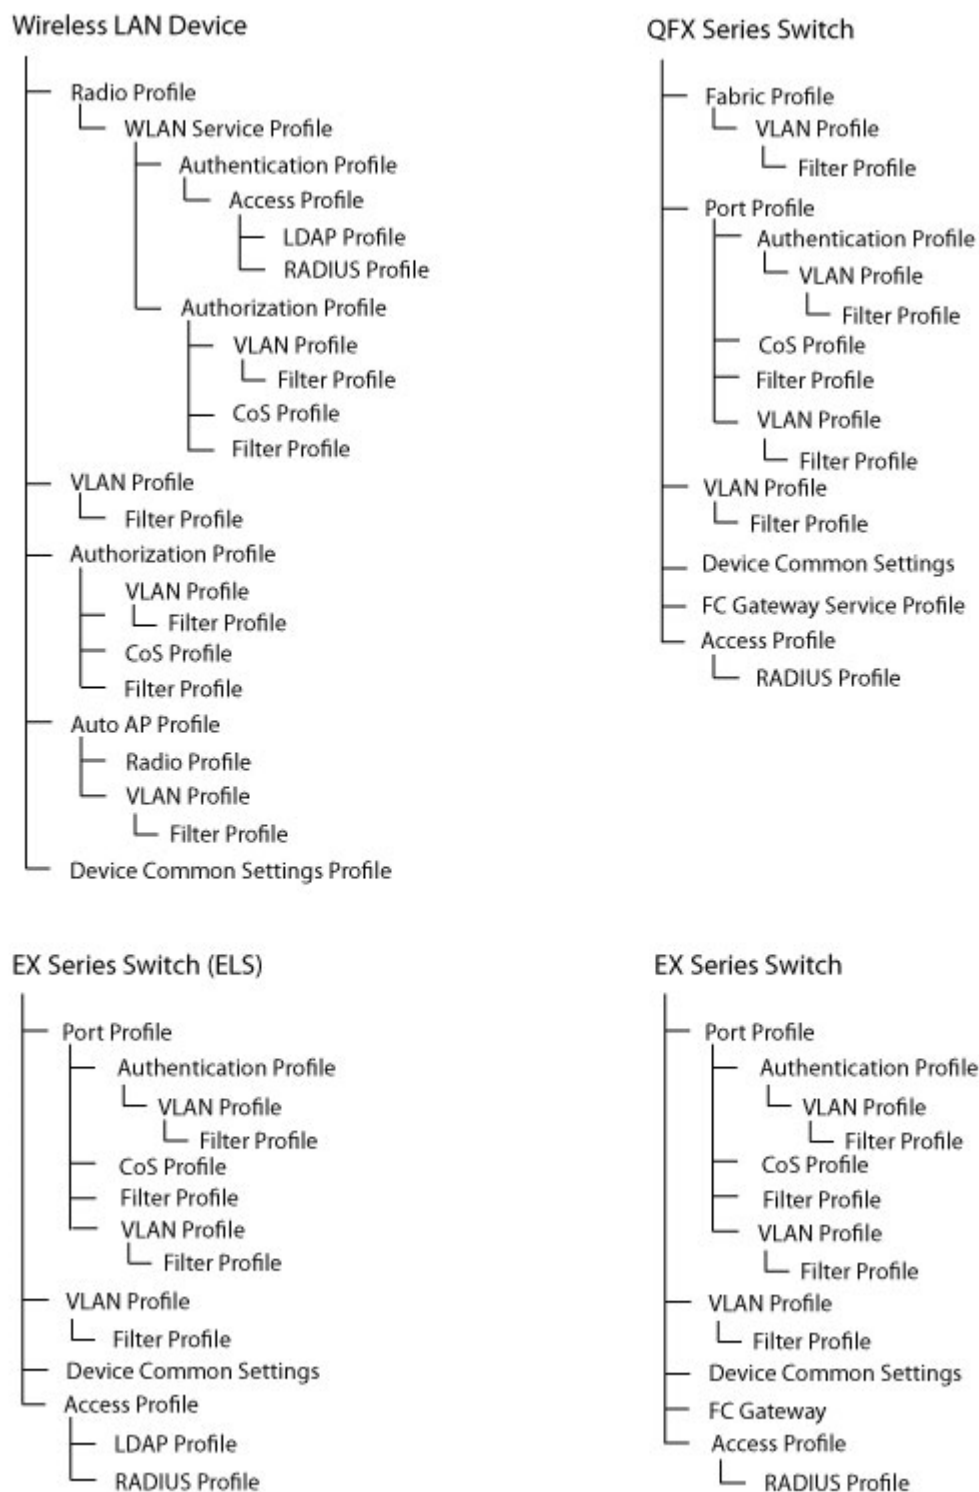
[Network Director Documentation home page](#)

Understanding VLAN Profiles

A virtual LAN (VLAN) is a Layer 2 broadcast domain that can span multiple wired segments. Each VLAN is a separate logical network, grouping hosts with common requirements, regardless of their physical location.

VLAN profiles in Network Director enable multiple VLAN configuration from a single profile. Each VLAN profile is specific to a device family: EX Switching, Campus Switching ELS. In addition, the VLANs are created for different purposes at different levels, as shown in [Figure 19 on page 328](#).

Figure 19: VLANs Are Specific to Device Families and Function Levels



Note: An Access profile is assigned to a switch when you assign a Port profile to the switch interfaces.

For EX Series Switches and Campus Switching ELS, apart from the basic settings, you can specify:

- MAC parameters
- Switching and routing parameters
- L2 and L3 Filters
- VLAN security DHCP, ARP inspection, and MAC movement.

TIP: Single VLAN IDs can be configured for all products and circumstances. VLAN lists or VLAN ranges of IDs are available for some products, depending on the technology used for implementation. For example, EX Switching does not support a VLAN list. Campus Switching ELS supports a VLAN ID range only as part of a VLAN ID list. You will only see the available configurations for the selected device family.

RELATED DOCUMENTATION

[Creating and Managing VLAN Profiles | 329](#)

[Assigning a VLAN Profile to Devices or Ports | 345](#)

[Understanding Network Configuration Profiles | 88](#)

[Network Director Documentation home page](#)

Creating and Managing VLAN Profiles

IN THIS SECTION

- [Managing VLAN Profiles | 330](#)
- [Creating a VLAN Profile | 331](#)
- [Specifying Basic EX Switching VLAN Settings | 333](#)
- [Specifying Basic Campus Switching ELS VLAN Settings | 334](#)
- [Specifying Basic VLAN Settings for Data Center Switching ELS | 336](#)
- [Specifying Advanced VLAN Profile Settings for EX Series Switches | 337](#)
- [Specifying Advanced VLAN Settings for Campus Switching ELS | 339](#)
- [Specifying Advanced VLAN Settings for Data Center Switching ELS | 341](#)
- [Reviewing and Saving the VLAN Profile Configuration | 344](#)
- [What to Do Next | 344](#)

You can create and manage VLAN profiles on switches and QFX Series devices by using the Manage VLAN Profiles window. Each VLAN profile is specific to a device family. After you create a VLAN profile, you can assign the profile at port level, or switch level.

Use the Manage VLAN Profiles page to create new VLAN profiles and to manage existing VLAN profiles.

This topic describes:

Managing VLAN Profiles

From the Manage VLAN Profiles page, you can:

- Create a new profile by clicking **Add**. For directions, see [“Creating a VLAN Profile” on page 331](#).
- Modify an existing profile by selecting the profile and clicking **Edit**.
- Assign a profile to a port or a switch, by selecting the profile and clicking **Assign**. For directions, see [“Assigning a VLAN Profile to Devices or Ports” on page 345](#).
- Modify an existing assignment of a profile by selecting the profile and clicking **Edit Assignment**.
- View information about a VLAN profile, including the interfaces it is associated with, by either clicking the profile name or by selecting the profile and clicking **Details**.
- Delete profiles by selecting the profiles and clicking **Delete**.

TIP: You cannot delete profiles that are in use—that is, assigned to objects or being used by other profiles. To see the current assignments for a profile, select the profile, click **Details**, and then click the Assigned Objects Tab in the Details window.

- Clone a VLAN profile by selecting the profile and clicking **Clone**.

[Table 80 on page 330](#) describes the fields in the Manage VLAN Profiles page. This page lists all VLAN profiles defined for your network.

Table 80: Manage VLAN Profile Fields

Field	Description
Profile Name	Name given to the profile when the profile was created.
VLAN Name	Name given to the VLAN when the VLAN profile was created.
Family Type	The device family; an EX Series switch or Campus Switching ELS.
VLAN ID	VLAN ID assigned when the profile was created.

Table 80: Manage VLAN Profile Fields (*continued*)

Field	Description
VLAN Range	<p>Range of VLAN IDs assigned when the profile was created.</p> <p>TIP: If a VLAN ID is displayed, VLAN range will be null. Also, Campus Switching ELS supports a VLAN ID range only as part of a VLAN ID list.</p>
VLAN ID List	<p>VLAN IDs can be either individually listed (with a space to separate each ID), an inclusive list separating the starting VLAN ID and ending VLAN ID with a hyphen, or a combination of both.</p> <p>TIP: If a VLAN ID is displayed, VLAN range will be null. Also, this column will never have a value for EX Switching because it is not available.</p>
Description	Description of the VLAN profile entered when the profile was created.
Assignment State	<p>Displays the assignment state of the profile. A profile can be:</p> <ul style="list-style-type: none"> • Unassigned—When the profile is not assigned to any object. • Deployed—When the profile is assigned and is deployed from Deploy mode. • Pending Deployment—When the profile is assigned, but not yet deployed in the network.
Creation Time	Date and time when the profile was created.
Last Updated Time	Date and time when the profile was last modified.
User Name	The username of the person who created or modified the profile.

TIP: All columns might not be displayed. To show or hide fields listed in the Manage Authorization Profiles table, click the down arrow on the field header, select **Columns**, and select or clear the check box adjacent to the field that you want to show or hide.

Creating a VLAN Profile

To create a VLAN profile, at minimum, you must specify the VLAN name and the IEEE 802.1Q VLAN tag for the profile. You also must indicate a device family for the VLAN: EX Series Switches, Campus Switching ELS.


In the VLAN, you can specify additional VLAN profile configuration such as:

- Ingress or egress filters to be used on the VLAN
- Parameters for handling the MAC forwarding table

To create a VLAN profile:

1. Under Views, select one of these options: **Logical View**, **Location View**, **Device View** or **Custom Group View**.

TIP: Do not select **Dashboard View**, **Datacenter View**, or **Topology View**.

2. Click  in the Network Director banner.
3. From the Tasks pane, select the type of network (Wired), the appropriate functional area, and then select the name of the profile that you want to create. For example, to create a PORT profile for a wired device, click **Wired** > **Profiles** > **PORT**. The appropriate Manage Profile page opens.
4. Click **Add** to add a new profile.
If you chose to create a profile for the wired network, Network Director opens the Device Family Chooser window.
 - a. From the Device Family Chooser, select the device family for which you want to create a profile. The available device families are **Switching (EX)**, **Campus Switching ELS** (Enhanced Layer 2 Software), and **Data Center Switching ELS**.
 - b. Click **OK**.
The Create VLAN Profile page for the selected device family is displayed. It consists of three sections, Basic Settings, Advanced Settings, and Review.
5. Specify the basic VLAN settings by using the appropriate directions:
 - [Specifying Basic EX Switching VLAN Settings on page 333](#)
 - [Specifying Basic Campus Switching ELS VLAN Settings on page 334](#)
6. When you have completed the basic settings, click **Next** or click **Advanced Settings** at the top of the wizard window.
7. Specify the advanced settings. Complete the Advanced Settings options as described in the online help:
 - [“Specifying Advanced VLAN Profile Settings for EX Series Switches” on page 337](#) for EX Series switches.

- [“Specifying Advanced VLAN Settings for Campus Switching ELS” on page 339](#) for Campus Switching ELS.
- When you have completed the advanced settings, click **Next** or click **Review** at the top of the wizard window.
 - You can make changes to your profile from the **Review** page. Click **Save > Finish** to save the profile. For directions, see [“Reviewing and Saving the VLAN Profile Configuration” on page 344](#).
 - Click **Finish**.

The system saves the VLAN profile and displays the Manage VLAN Profiles page. Your new or modified VLAN profile is listed in the table of VLAN profiles.

Specifying Basic EX Switching VLAN Settings

To configure the basic settings for an EX Switching VLAN profile, enter the settings described in [Table 81 on page 333](#). Required settings are indicated by a red asterisk (*) that appears next to the field label.

Table 81: VLAN Profile Basic Settings for EX Switching

Field	Action
Profile Name	Type a name for the profile. Profile name must not contain special characters or spaces. Note that profiles that are automatically created by Network Director as part of device discovery or out-of-band changes might contain the underscore (_) character.
VLAN Name	Type the name of VLAN. The profile name and the VLAN name can be the same or different.
Description	Type a description to identify the group or function the VLAN will be part of. The character limit is 256 characters.
VLAN ID	
You can indicate a single VLAN ID or a VLAN Range for EX Switching.	
Single VLAN ID	To specify a single VLAN ID, type the single unique IEEE 802.1Q identifier for the VLAN (VLAN tag). The range for VLAN IDs is 1 through 4094.

Table 81: VLAN Profile Basic Settings for EX Switching (*continued*)

Field	Action
Range of VLAN IDs	<p>To indicate a range of VLAN IDs for EX Series switches, follow these steps:</p> <ol style="list-style-type: none"> 1. Select Range instead of Single in the VLAN ID section. 2. Provide the first and last VLAN IDs in the range. <p>TIP: For example, if you enter 10 and 12, when you deploy the profile on a device, three VLANs are created with VLAN IDs 10, 11, and 12. The names of the VLANs are created from the name you specified by adding the VLAN ID as a suffix to the name, for example vlanname_10.</p>

Click **Next** or click **Advanced Settings** at the top of the wizard window to configure advanced VLAN EX Switching profile settings. Advanced Settings are described in [“Specifying Advanced VLAN Profile Settings for EX Series Switches” on page 337](#).

Specifying Basic Campus Switching ELS VLAN Settings

To configure the basic settings for a Campus Switching ELS VLAN profile, enter the settings described in [Table 82 on page 334](#). Required settings are indicated by a red asterisk (*) that appears next to the field label.

Table 82: VLAN Profile Basic Settings for Campus Switching ELS

Field	Action
Profile Name	<p>Type a unique name that identifies the profile.</p> <p>Profile name must not contain special characters or spaces. Note that profiles that are automatically created by Network Director as part of device discovery or out-of-band changes might contain the underscore (_) character.</p>
VLAN Name	Type the name of VLAN. The profile name and the VLAN name can be the same or be different.
Description	Type a description to identify the group or function of the VLAN. The character limit is 256 characters.

VLAN ID

NOTE: Campus Switching ELS supports a VLAN ID range only as part of a VLAN ID list. Follow the directions for adding a list of VLAN IDs if you are adding a VLAN range.

Table 82: VLAN Profile Basic Settings for Campus Switching ELS (*continued*)

Field	Action
Single VLAN ID	To specify a single VLAN ID (default), type the single unique IEEE 802.1Q identifier for the VLAN—the VLAN tag. The range for VLAN IDs is 1 through 4094.
List of VLAN IDs	<p>To create a list of VLAN IDs for switches, follow these steps:</p> <ol style="list-style-type: none"> 1. Select List instead of Single in the VLAN ID section. 2. Click Add under VLAN IDs. The Add VLAN Details window opens. 3. To add a single VLAN ID to the list, type the VLAN ID and then click either Add which closes this window or Add More which allows you to continue adding to the list. 4. To add a range of VLAN IDs to this list: <ol style="list-style-type: none"> a. In the Add VLAN Details window, select Range to add VLAN IDs in the range format 1 - 3. b. In the Add VLAN Details window, provide the first and last VLAN IDs in the range. TIP: For example, if you enter 10 and 12, when you deploy the profile on a device, three VLANs are created with VLAN IDs 10, 11, and 12. The names of the VLANs are created from the name you specified by adding the VLAN ID as a suffix to the name, for example vlanname_10. c. Click either Add to close this window, or Add More to allow you to continue adding to the list. 5. When you are finished creating the list, close the window (if it is still open). All VLAN IDs you added appear in the VLAN IDs list.

Click **Next** or click **Advanced Settings** at the top of the wizard window to configure advanced Campus Switching ELS VLAN profile settings. Advanced settings are described in [“Specifying Advanced VLAN Settings for Campus Switching ELS” on page 339](#).

Specifying Basic VLAN Settings for Data Center Switching ELS

To configure the basic settings for a Data Center Switching ELS VLAN profile, specify the parameters described in [Table 83 on page 336](#) for an Ethernet VLAN profile. All settings are optional.

Table 83: VLAN Profile Basic Settings for Data Center Switching ELS

Field	Action
Data Center Switching ELS MAC Parameters	
Profile Name	Type a unique name that identifies the profile. Profile name must not contain special characters or spaces. Note that profiles that are automatically created by Network Director as part of device discovery or out-of-band changes might contain the underscore (_) character.
VLAN Name	Type the name of VLAN. The profile name and the VLAN name can be the same or be different.
Description	Type a description to identify the group or function of the VLAN. The character limit is 256 characters.
VLAN ID	
NOTE: Data Center Switching ELS supports a VLAN ID range only as part of a VLAN ID list. Follow the directions for adding a list of VLAN IDs if you are adding a VLAN range.	
Single VLAN ID	To specify a single VLAN ID (default), type the single unique IEEE 802.1Q identifier for the VLAN—the VLAN tag. The range for VLAN IDs is 1 through 4094.

Table 83: VLAN Profile Basic Settings for Data Center Switching ELS (*continued*)

Field	Action
List of VLAN IDs	<p>To create a list of VLAN IDs for switches, follow these steps:</p> <ol style="list-style-type: none"> 1. Select List instead of Single in the VLAN ID section. 2. Click Add under VLAN IDs. The Add VLAN Details window opens. 3. To add a single VLAN ID to the list, type the VLAN ID and then click either Add which closes this window or Add More which allows you to continue adding to the list. 4. To add a range of VLAN IDs to this list: <ol style="list-style-type: none"> a. In the Add VLAN Details window, select Range to add VLAN IDs in the range format 1 - 3. b. In the Add VLAN Details window, provide the first and last VLAN IDs in the range. TIP: For example, if you enter 10 and 12, when you deploy the profile on a device, three VLANs are created with VLAN IDs 10, 11, and 12. The names of the VLANs are created from the name you specified by adding the VLAN ID as a suffix to the name, for example vlnname_10. c. Click either Add to close this window, or Add More to allow you to continue adding to the list. 5. When you are finished creating the list, close the window (if it is still open). All VLAN IDs you added appear in the VLAN IDs list.

Click **Next** or click **Advanced Settings** at the top of the wizard window to configure advanced Data Center Switching ELS VLAN profile settings. Advanced Settings are described in [“Specifying Advanced VLAN Settings for Campus Switching ELS” on page 339](#).

Specifying Advanced VLAN Profile Settings for EX Series Switches

To configure the EX Switching advanced settings for the VLAN profile, enter the MAC parameters and Layer 2 filters described in [Table 84 on page 337](#) for EX Series switching. All settings are optional.

Table 84: VLAN Profile Advanced Settings for an EX Series Switch

EX Switching MAC Parameters

Table 84: VLAN Profile Advanced Settings for an EX Series Switch (*continued*)

MAC Limit	<p>Type the number of dynamic MAC addresses that can be learned on the VLAN. If this number is exceeded, packets containing new MAC addresses are dropped and an alarm is raised.</p> <p>Setting a limit on the number of dynamic MAC addresses protects against an Ethernet switching table overflow attack.</p>
MAC Aging Time (ms)	<p>Indicate the number of milliseconds that unused dynamic MAC addresses remain in the MAC forwarding table before being deleted. If you specify the time as unlimited, entries are never removed from the table. Generally, use this setting only if the switch or the VLAN has a fairly static number of end devices—otherwise the table will eventually fill up. You can use this setting to minimize traffic loss and flooding that might occur when traffic arrives for MAC addresses that have been removed from the table.</p> <p>The range is from 60 through 1,000,000.</p>
EX Switching L2 Filters	
L2 Ingress Filter	<p>Click Select to choose from existing Filter profiles. Select a profile from the Choose Filter Profile window and click OK. The filter configuration contained in the profile is applied to ingress traffic on the VLAN.</p> <p>To remove the selected Filter profile, click Clear.</p>
L2 Egress Filter	<p>Click Select to choose from existing Filter profiles. Select a profile from the Choose Filter profile window and click OK. The filter configuration contained in the profile is applied to egress traffic on the VLAN.</p> <p>To remove the selected Filter profile, click Clear.</p>
EX Switching L3 Routing Filters	
<p>If you indicated a single VLAN ID under the Basic Settings, you can specify one or more routing parameters (Layer 3 filters) for the profile.</p>	
L3 Ingress Filter L3 IPv6 Ingress Filter	<p>Click Select to choose from existing Filter profiles. Select a profile from the Choose Filter profile window and click OK. The filter configuration contained in the profile is applied to ingress traffic on the VLAN.</p> <p>To remove the selected Filter profile, click Clear.</p>

Table 84: VLAN Profile Advanced Settings for an EX Series Switch (*continued*)

L3 Egress Filter	Click Select to choose from existing Filter profiles. Select a profile from the Choose Filter profile window and click OK . The filter configuration contained in the profile is applied to egress traffic on the VLAN. To remove the selected Filter profile, click Clear .
L3 IPv6 Egress Filter	

VLAN Security Settings

Optionally, select VLAN Security Settings to display the security options DHCP, ARP inspection, and MAC movement limit for VLAN profiles for EX switching.

Enable DHCP Snooping	Check to apply a series of security techniques to the DHCP infrastructure.
Enable ARP Inspection	The Address Resolution Protocol (ARP), which provides IP communication within a Layer 2 broadcast domain by mapping an IP address to a MAC address, has security issues. Select this option to apply inspection to untrusted interfaces.
MAC Movement Limit	Indicate the number of times a MAC address entry can be moved in the MAC address table without consequences.
MAC Movement Action	When a MAC Movement Limit is specified, select an action to be applied to MAC addresses that exceed the MAC Movement Limit: None , Log , Drop , Shut Down , or Drop and Log .
VRRP Settings	Select the VRRP profile for the interface from a list of existing profiles by clicking Select . Select one of the listed profiles, and then click OK .

Click **Next** or click **Review** to see the Review page of the wizard. For review directions, see [“Reviewing and Saving the VLAN Profile Configuration” on page 344](#).

Specifying Advanced VLAN Settings for Campus Switching ELS

To configure the advanced settings for a Campus Switching ELS VLAN profile, specify the parameters described in [Table 85 on page 339](#) for Campus Switching ELS. All settings are optional.

Table 85: VLAN Profile Advanced Settings for Campus Switching ELS

Field	Action
Campus Switching ELS MAC Parameters	

Table 85: VLAN Profile Advanced Settings for Campus Switching ELS (*continued*)

Field	Action
Interface MAC Limit	<p>Indicate the number of dynamic MAC addresses that can be learned on the VLAN. If this number is exceeded, packets containing new MAC addresses are dropped and an alarm is raised.</p> <p>Setting a limit on the number of dynamic MAC addresses protects against an Ethernet switching table overflow attack.</p>
Packet Action	Indicate the packet action for MAC addresses that exceed the Interface MAC Limit, by selecting None , Log , Drop , Shut Down , or Drop and Log .
MAC Table Size	If you indicated an Interface MAC limit, provide a table size here by using the up and down arrows. The MAC table must allow for at least 16 entries—you can increase this limit with the arrow.

L2 Filters

Ingress Filter	<p>Click Select to choose from existing Filter profiles. Select a profile from the Choose Filter Profile window and then click OK. The filter configuration contained in the profile is applied to ingress traffic on the VLAN.</p> <p>To remove a selected Filter profile, click Clear.</p>
Egress Filter	<p>Click Select to choose from existing Filter profiles. Select a profile from the Choose Filter profile window and then click OK. The filter configuration contained in the profile is applied to egress traffic on the VLAN.</p> <p>To remove a selected Filter profile, click Clear.</p>

Routing

If you selected a single VLAN ID under Basic Settings, you can specify Layer 3 filter routing parameters for the VLAN profile.

NOTE: If an IP address is configured for a VLAN on some devices, then the configured IP address will be retained and a DHCP client will not be enabled on those devices. Also, if you indicated a VLAN range for basic ELS switching configuration, this option is not available.

Routing L3 Filters

Ingress Filter IPv6 Ingress Filter	<p>Click Select to choose from existing Filter profiles. Select a profile from the Choose Filter profile window and then click OK. The filter configuration contained in the profile is applied to ingress traffic on the VLAN.</p> <p>To remove a selected Filter profile, click Clear.</p>
---	---

Table 85: VLAN Profile Advanced Settings for Campus Switching ELS (*continued*)

Field	Action
Egress Filter	Click Select to choose from existing Filter profiles. Select a profile from the Choose Filter profile window and then click OK . The filter configuration contained in the profile is applied to egress traffic on the VLAN. To remove a selected Filter profile, click Clear .
IPv6 Egress Filter	

VLAN Security Settings

Optionally, enable VLAN Security Settings to display the security options DHCP, ARP inspection, and MAC movement limit for VLAN profiles for ELS switching.

Enable DHCP Snooping	When checked (default), this option applies a series of security techniques to the DHCP infrastructure.
Enable ARP Inspection	The Address Resolution Protocol (ARP), which provides IP communication within a Layer 2 broadcast domain by mapping an IP address to a MAC address, has security issues. Select this option to apply inspection to untrusted interfaces.
MAC Movement Limit	Indicate the number of times a MAC address entry can be moved in the MAC address table without consequences.
MAC Movement Action	When a MAC Movement Limit is specified, select an action to be applied to MAC addresses that exceed the MAC Movement Limit: None , Log , Drop , Shut Down , or Drop and Log .

Click **Next** or click **Review** to see the Review page of the wizard. For review directions, see [“Reviewing and Saving the VLAN Profile Configuration” on page 344](#).

Specifying Advanced VLAN Settings for Data Center Switching ELS

To configure the advanced settings for a Data Center Switching ELS VLAN profile, specify the parameters described in [Table 86 on page 341](#) for an Ethernet VLAN profile. All settings are optional.

Table 86: VLAN Profile Advanced Settings for Data Center Switching ELS Ethernet VLAN

Field	Action
Data Center Switching ELS MAC Parameters	

Table 86: VLAN Profile Advanced Settings for Data Center Switching ELS Ethernet VLAN (continued)

Field	Action
Interface MAC Limit	<p>Indicate the number of dynamic MAC addresses that can be learned on the VLAN. If this number is exceeded, packets containing new MAC addresses are dropped and an alarm is raised.</p> <p>Setting a limit on the number of dynamic MAC addresses protects against an Ethernet switching table overflow attack.</p>
Packet Action	Indicate the packet action for MAC addresses that exceed the Interface MAC Limit. The options are: None , Log , Drop , Shut Down , and Drop and Log .
MAC Table Size	If you indicated an Interface MAC limit, provide a table size here by using the up and down arrows. The MAC table must allow for at least 16 entries—you can increase this limit by using the arrow.

L2 Filters

Ingress Filter	<p>Click Select to choose from existing Filter profiles. Select a profile from the Choose Filter Profile window and then click OK. The filter configuration contained in the profile is applied to ingress traffic on the VLAN.</p> <p>To remove a selected Filter profile, click Clear.</p>
Egress Filter	<p>Click Select to choose from existing Filter profiles. Select a profile from the Choose Filter profile window and then click OK. The filter configuration contained in the profile is applied to egress traffic on the VLAN.</p> <p>To remove a selected Filter profile, click Clear.</p>

Routing

If you selected a single VLAN ID under Basic Settings, you can specify Layer 3 filter routing parameters for the VLAN profile.

NOTE: If an IP address is configured for a VLAN on some devices, then the configured IP address will be retained and a DHCP client will not be enabled on those devices. Also, if you indicated a VLAN range for basic ELS switching configuration, this option is not available.

Routing L3 Filters

Ingress Filter IPv6 Ingress Filter	<p>Click Select to choose from existing Filter profiles. Select a profile from the Choose Filter profile window and then click OK. The filter configuration contained in the profile is applied to ingress traffic on the VLAN.</p> <p>To remove a selected Filter profile, click Clear.</p>
---------------------------------------	---

Table 86: VLAN Profile Advanced Settings for Data Center Switching ELS Ethernet VLAN (*continued*)

Field	Action
Egress Filter	Click Select to choose from existing Filter profiles. Select a profile from the Choose Filter profile window and then click OK . The filter configuration contained in the profile is applied to egress traffic on the VLAN. To remove a selected Filter profile, click Clear .
IPv6 Egress Filter	

VLAN Security Settings

Optionally, enable VLAN Security Settings to display the security options DHCP, ARP inspection, and MAC movement limit for VLAN profiles for ELS switching.

Enable DHCP Snooping	When checked (default), this option applies a series of security techniques to the DHCP infrastructure.
Enable ARP Inspection	The Address Resolution Protocol (ARP), which provides IP communication within a Layer 2 broadcast domain by mapping an IP address to a MAC address, has security issues. Select this option to apply inspection to untrusted interfaces.
MAC Movement Limit	Indicate the number of times a MAC address entry can be moved in the MAC address table without consequences.
MAC Movement Action	When a MAC Movement Limit is specified, select an action to be applied to MAC addresses that exceed the MAC Movement Limit. The options are: None , Log , Drop , Shut Down , and Drop and Log .

FIP Snooping Settings

Enable VN2VN Snooping	Select to enable VN_Port to VN_Port (VN2VN) FIP snooping on the VLAN.
Beacon Period (ms)	Set the interval between periodic beacons, in milliseconds. Beacons perform virtual link maintenance for VN_Ports in a way that is similar to FIP keepalive advertisements. Range: 250 through 90000 milliseconds. Default: 8000 milliseconds.
FC Map	Set the FCoE mapped address prefix (FC-MAP) value for the FCoE VLAN to match the FC switch (or FCoE forwarder) FC-MAP value for the FC fabric. The FC-MAP value is a unique MAC address prefix an FC switch uses to identify FCoE traffic for a given FC fabric (traffic on a particular FCoE VLAN). Range: 0x0EFC00 through 0x0EFCFF. Default: 0xEFC00.

Click **Next** or click **Review** to see the Review page of the wizard. For review directions, see [“Reviewing and Saving the VLAN Profile Configuration” on page 344](#).

Reviewing and Saving the VLAN Profile Configuration

From this page, you can either save the VLAN profile or make changes to the VLAN profile:

- To make changes to the profile, click the **Edit** associated with the configuration you want to change.

Alternatively, you can click **Basic Settings** or **Advanced Settings** from the wizard workflow at the top of the page and make changes there.

When you are finished with your modifications, click **Review** to return to this page.

- To save a new profile or to save modified settings to an existing profile, click **Finish**.

The Manage VLAN Profiles page is displayed and your new or modified VLAN profile is listed in the table of VLAN profiles.

What to Do Next

Once the VLAN profile is created, you must assign the VLAN profile from the Assign VLAN Profile page to the required ports or switches, . To assign a VLAN profile, see [“Assigning a VLAN Profile to Devices or Ports” on page 345](#). After you assign a VLAN profile to a port or switch, you must deploy the profile configuration from the Deploy mode. For directions on deploying your configurations, see [“Deploying Configuration to Devices” on page 580](#).

FCoE VLANs are assigned to Fabric profiles, where they define the FCoE VLAN for a gateway FC fabric.

RELATED DOCUMENTATION

[Assigning a VLAN Profile to Devices or Ports | 345](#)

[Deploying Configuration to Devices | 580](#)

[Understanding VLAN Profiles | 327](#)

[Understanding VRRP Profiles | 544](#)

[Creating and Managing VRRP Profiles | 545](#)

[Network Director Documentation home page](#)

Assigning a VLAN Profile to Devices or Ports

IN THIS SECTION


- [Assigning a VLAN Profile | 345](#)
- [Editing Profile Assignments | 347](#)

After a VLAN profile is created, assign it to switches, aggregation devices in a Junos Fusion fabric, Virtual Chassis Fabric, members of Layer 3 Fabric, or members of custom groups.

You must have one or more existing VLAN profiles, either user-configured or system-created, before you can assign a VLAN profile to a switch, or member of a custom group or port group. For further directions, see [“Creating and Managing VLAN Profiles” on page 329](#), [“Creating Custom Device Groups” on page 155](#), and [“Creating and Managing Port Groups” on page 323](#).

Assigning a VLAN Profile

To assign a VLAN profile:

1. Click  in the Network Director banner.
The Manage VLAN profiles page is displayed. The page displays all user-configured and system-created VLAN profiles for discovered devices.
2. Select **VLAN** from the Profile and Configuration Management menu in the Tasks pane.
The Assign VLAN Profile page for the selected VLAN appears.
3. Select a VLAN profile from the list of VLAN profiles and then click **Assign**.
The Assign VLAN Profile page for the selected VLAN appears.

NOTE: If Network Director fails to read the configuration of one or more devices after the device discovery, those devices are not displayed in the Assign Profile page. You will not be able to assign profiles to those devices. The Manage Jobs page in System mode displays details of the device discovery jobs. Use the information displayed on this page to take appropriate corrective steps to enable Network Director to reread the configuration of the failed device. For more information, see [“Discovering Devices in a Physical Network” on page 94](#).

4. If you are assigning a VLAN profile to a device, a tree is displayed. Choose at least one device for VLAN assignment. Be sure that a check mark appears in front of the device - just highlighting the name does not select it.

If you are assigning a VLAN profile to members of a Custom Group or Port Group, selecting the group selects all members of that group.

5. Click either **Next** or **Profile Assignment**.

The Profile Assignment page displays the list of existing assignments in the Assignments table.

6. Select a device from the Assignments table and click **Assign to Device**.
7. If you are assigning the profile to an EX9200 and a routing instance has been created on this network, you have a choice between assigning the profile to the EX9200 device or assigning the profile to the routing instance. If a routing instance has been created with the Junos CLI, select either the EX9200 device or the routing instance in the window that opens. If no routing instance exists, the profile is automatically assigned to the device.

NOTE: 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.

8. Click **Define** from the **Attributes** column in the Assignments table to modify the attributes. The Configure Attributes page is displayed:

For switches, specify the following details:

- Enter the Layer 3 interface IP address and port number (between 0 through 32).
- Enter the unit number. By default, the VLAN ID is populated in this field. You can optionally change the value.

9. You can view the assignment details for the selected device or delete any assignments.

- To view assignment details, select a device and click **View Assignments**.

The Profile Details page for selected device appears. Expand the **Device** name to see the details of the assignment. The assignment status displays the status whether the device is deployed or is pending device update, etc.

- To delete a VLAN profile assignment for a device, select the device from the Assignments table and click **Remove**.

10. Click **Next** or click **Review** from the top wizard workflow to review the assignments. Alternately, click **Edit** to edit the profile assignment.

11. Click **Finish** once you are done reviewing the profile assignment.

After you click Finish, the Create Profile Assignments Job Details dialog box appears, which reports on the status of the profile assignment job. If you have assigned the profile to a large number of objects, the profile assignment job can take some time to complete. Instead of waiting for the Job Details dialog box to report job completion status, you can close it and check the details of the profile assignment job at a later time using the Manage Job task in System mode.

NOTE: If any assignment fails, the profile assignment job fails and none of the assignments are created. Check the details for the profile assignment job for information about why the assignment failed.

After you assign a VLAN profile to a device or port, you can deploy the VLAN profile from the **Deploy** mode. For details, see [“Deploying Configuration to Devices” on page 580](#).

Editing Profile Assignments

You can edit VLAN assignments from the Manage VLAN Profiles page. To edit an existing assignment:

1. Select a profile from the Manage VLAN profiles page and then click **Edit Assignment**.

The Edit Assignments page for the selected device appears.

2. Expand the **Devices** cabinet select a device.

3. Make the required changes in the **Operation** column of the table.

To change the attributes, see step 8, from the Assign VLAN profiles tasks.

4. Click **Apply**.

The Manage VLAN Profiles page reappears.

RELATED DOCUMENTATION

[Creating and Managing VLAN Profiles](#) | 329

[Deploying Configuration to Devices | 580](#)

[Creating Custom Device Groups | 155](#)

[Creating and Managing Port Groups | 323](#)

[Understanding VLAN Profiles | 327](#)

[Network Director Documentation home page](#)

Configuring Firewall Filters (ACLs)

IN THIS CHAPTER

- [Understanding Filter Profiles | 349](#)
- [Creating and Managing Wired Filter Profiles | 350](#)

Understanding Filter Profiles

Filter profiles are a set of rules that define whether to accept or discard packets that are transiting on an interface on a Juniper Networks EX Series Ethernet Switch. You configure Filter profiles to determine whether to accept or decline traffic before it enters or exits a port to which the Filter profile is applied to.

A Filter profile must contain at least one term. Each term consists of the following components:

- **Match conditions**—Specify the values or fields that the packet must contain. You can define various match conditions, depending on the device for which you are defining these conditions. For example, for EX Series switches, you can specify a match condition based on the IP source address field, IP destination address field, Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) source port field, IP protocol field, Internet Control Message Protocol (ICMP) packet type, TCP flags, interfaces, and so on.
- **Action**—Specifies what to do if a packet matches the match conditions. Possible actions are to accept or discard the packet or to send the packet to a specific virtual routing interface. In addition, packets can be counted to collect statistical information. If no action is specified for a term, the default action is to accept the packet.
- **Action modifier**—Specifies one or more actions for the switch if a packet matches the match conditions. You can specify action modifiers such as the loss priority, policer details, and forwarding class, depending on the type of device on which you are creating the Filter profile.

The maximum number of terms allowed per Filter profile for EX Series switches is:

- 512 for EX2300 switches

NOTE: Firewall filters are categorized into two different pools. Port and VLAN filters are pooled together (the memory threshold for this pool is 22K) while router firewall filters are pooled separately (the threshold for this pool is 32K). The assignment happens based on the filter pool type. You can share free space blocks only among the firewall filters belonging to the same filter pool type. An error message is generated if you attempt to configure a firewall filter beyond the TCAM threshold.

The Manage Filter Profiles page enables you create, modify, view, and delete Filter profiles.

RELATED DOCUMENTATION

[Creating and Managing Wired Filter Profiles | 350](#)

[Network Director Documentation home page](#)

Creating and Managing Wired Filter Profiles

IN THIS SECTION

- [Managing Wired Filter Profiles | 351](#)
- [Creating a Wired Filter Profile | 352](#)
- [Specifying Settings for an EX Series Switch Filter Profile | 353](#)
- [Specifying Settings for a Campus Switching ELS Switch Filter Profile | 365](#)
- [Specifying Settings for a Data Center Switching ELS Filter Profile | 379](#)
- [What to Do Next | 394](#)

Filter profiles are sets of rules that determine whether to accept or discard packets transiting on switch.

Use the Manage Filter Profiles page to create new wired Filter profiles and manage existing Filter profiles.

This topic describes:

Managing Wired Filter Profiles

From the Manage Filter Profiles page, you can:

- Create a new wired Filter profile by clicking **Add**. For directions, see [“Creating a Wired Filter Profile” on page 352](#).
- Modify an existing wired Filter profile by selecting it and clicking **Edit**.
- View information about a wired Filter profile, including the associated interfaces, by either clicking the profile name or by selecting the profile and clicking **Details**.
- Delete a wired Filter profile by selecting the profile and clicking **Delete**.

TIP: You cannot delete profiles that are in use—that is, profiles assigned to objects or used by other profiles. To see the current assignments for a profile, select the profile and click **Details**.

- Clone a wired Filter profile by selecting a profile and clicking **Clone**.

[Table 87 on page 351](#) describes the information provided about wired Filter profiles on the Manage Filter Profiles page. This page lists all Filter profiles defined for your network, regardless of the scope you selected in the network view.

Table 87: Manage Wired Filter Profile Fields


Field	Description
Profile Name	Name given to the profile when the profile was created.
Family Type	The device family on which the profile was created: Switching (EX) , or Campus Switching ELS .
Description	Description of the profile entered when the profile was created. TIP: To display the entire description, you might need to resize the Description column by clicking the column border in the heading and dragging it.
Creation Time	Date and time when the profile was created.
Last Updated Time	Date and time when the profile was last modified.
User Name	The username of the user who created or modified the profile.

TIP: All columns might not be displayed. To show or hide fields in the table, click the down arrow on the field header, select **Columns**, and select or clear the check box adjacent to the field that you want to show or hide.

Creating a Wired Filter Profile

To create a wired Filter profile, you must provide a filter name and configure at least one term. A term is a collection of one or more match conditions, and actions that the system takes when match conditions are met. A term must have at least one match condition.

To create a wired Filter profile:

1. Click  **Build** in the Network Director banner.
2. Under Views, select one of these options: **Logical View**, **Location View**, **Device View** or **Custom Group View**.

TIP: Do not select **Dashboard View**, **Datacenter View**, or **Topology View**.

3. From the Tasks pane, expand **Wired**, expand **System**, and then select **Filter**.
4. Click **Add** to add a new profile.
Network Director opens the Device Family Chooser window.
5. From the Device Family Chooser, select the wired device family for which you want to create a profile. The available device families are **Switching (EX)**, **Campus Switching ELS** (Enhanced Layer 2 Software), and **Data Center Switching**.
6. Click **OK**.
The Create Filter Profile wizard for the selected device family is displayed.
7. Specify the filter settings by following these directions:
 - For EX Series switches, specify the settings as described in both the online help and in [“Specifying Settings for an EX Series Switch Filter Profile” on page 353](#).

- For Campus Switching ELS, specify the settings as described in both the online help and in [“Specifying Settings for a Campus Switching ELS Switch Filter Profile” on page 365](#).

8. Click **Done** to save the Filter profile.

The system saves the Filter profile and displays the Manage Filter Profiles page. Your new or modified Filter profile is listed in the table of Filter profiles.

Specifying Settings for an EX Series Switch Filter Profile

A Filter profile must have at least one term in it. Each term has one filtering function. For example, if a term is evaluating the source of packets, then that term cannot also evaluate the protocols used by the packets. Some switch models do accommodate multiple terms in one filter. When you have more than one term in a filter, the ordering of the terms is important. The system evaluates multiple filter terms as follows:

- The packet is evaluated against the first term's conditions. If the packet matches all of the conditions in that term, the action specified for that condition is taken and evaluation ends. Subsequent terms in the filter are not evaluated.
- If a packet does not match all conditions in the first term, the packet is then evaluated against the conditions in the second term. This process continues until either the packet matches all conditions in a term or there are no more terms in the filter. Whenever a match occurs, the term's corresponding action is taken and evaluation ends—subsequent terms in the filter are not evaluated.
- If a packet passes through all the terms in the filter without a match, the packet is discarded.

To configure a Filter profile for EX Series switches:

1. Specify a filter name and description for the Filter profile.
2. Select the switch filter family for which you want to create the profile:
 - If you want to create a Layer 2 based filter, select **Ethernet switching**.
 - If you want to create a Layer 3 based filter for IPv4, select **INET**.
 - If you want to create a Layer 3 based filter for IPv6, select **INET6**.
3. Under Terms, click **Add** to add one or more terms with match condition(s) for this filter.

The Create Term window opens, displaying a section for each type of term you can create, Source and Destination Parameters, Protocols, DSCP Settings, TCP Settings, and ICMP Settings. The Action section applies to all of those types.

NOTE: The order of the terms within a Filter profile configuration is important. Packets are tested against each term in the order in which terms are listed.

- 4. Enter a name for the filter term.
- 5. Specify the match condition(s) for the filter term as described in [Table 88 on page 354](#). Required settings are indicated by a red asterisk (*) that appears next to the field label in the user interface.

Table 88: Create Term Fields for EX Switching

Task	Description
Source and Destination Parameters	
You can specify match conditions for either packets' origin (source) or packets' destination, or both. You are indicating the location of the filtering here—either specifying that packets that originate at a specific place (source) will be filtered or packets destined for a specific location (destination) will be filtered. You can have multiple sources and destinations for one filter term.	
Add Source Parameters and Destination Parameters	<p>To add source and destination parameters to the named filter term:</p> <ul style="list-style-type: none">a. Click Add to the right of the Destination Parameters list. <p>The Add Source/Destination Parameter window appears.</p> <ul style="list-style-type: none">b. Select either Source (default) or Destination from the Add Source/Destination Parameter page.c. Select one of following available Parameter Types from the Add Source/Destination Parameter page and provide the corresponding information: <p>TIP: Available parameter types vary.</p> <ul style="list-style-type: none">● IP Address—also provide the IP address of the source or destination device● MAC Address—also provide the MAC address of the source or destination device● Port—also provide the port type of the source or destination port. Select either AFS (Andrew File System), BGP (Border Gateway Protocol), BIFF (UNIX mail notification), Bootpc (bootstrap protocol client), Bootps, Cmd, CVS pserver, DHCP, Domain, EK login, EK shell, EXEC, Finger (protocol), or FTP. <p>NOTE: If you selected Port as the parameter and do not find the type of port that you want to add from the Port list, then select Other and enter a port number.</p> <ul style="list-style-type: none">d. Click OK <p>The parameter term is added to the appropriate list, either Source Parameters or Destination Parameters.</p>

Table 88: Create Term Fields for EX Switching (*continued*)

Task	Description
Protocols and EtherTypes	
For either INET family, you can apply a filter term based on protocols being used by packets. For the Ethernet-switching family, you can apply a filter term based on either the protocols being used by packets or on the EtherTypes being used by packets. EtherType indicates a protocol that is encapsulated in the payload of an Ethernet Frame. Expand the Protocols section to see the configuration.	
Add a Protocol Match Condition (Ethernet-switching family or INET family)	<p>To add a protocol match condition to the named filter term:</p> <ol style="list-style-type: none"> Expand the Protocols and EtherTypes section. Click Add under Protocols. The Select Protocols window opens, displaying a list of protocols. From the list of protocols, select one or more. The options are AH, DSTOPTS, EGP, ESP, Fragment, GRE, Hop-by-hop, ICMP, ICMP6, IPIP, IPv6, No-text-header, OSPF, PIM, Routing, RSVP, SCTP, TCP, UDP, and VRRP. Click OK. The protocols are added to the Protocols list.
Add an EtherType Match Condition (Ethernet-switching family)	<p>To add an EtherType match condition to the named filter Ethernet-switching family term:</p> <ol style="list-style-type: none"> Expand the Protocols and EtherTypes section. Click Add under EtherTypes. The Select EtherTypes window opens, displaying a list of protocols. From the list of EtherTypes, select one or more. The options are AARP, AppleTalk, ARP, IPv4, IPv6, MPLS multicast, MPLS unicast, OAM, PPP, PPPOE discovery, PPPOE session, and SNA. Click OK. The EtherTypes are added to the EtherTypes list.

Table 88: Create Term Fields for EX Switching *(continued)*

Task	Description
------	-------------

DSCP Settings

Expand this section to see the DSCP term settings. DiffServ is a simple mechanism for classifying and managing network traffic and providing quality-of-service (QoS) on IP networks. DiffServ can, for example, be used to apply low-latency to critical network traffic such as voice or streaming media while providing simple best-effort service to non-critical services such as Web traffic. Here, you can apply a filter term based on the Differentiated Services code point (DSCP) which is a field in IPv4 and IPv6 headers.

NOTE: With IPv6 packets, the DS field and ECN field replace the IPv4 TOS field.

Table 88: Create Term Fields for EX Switching (*continued*)

Task	Description
Add a DSCP Match Condition	<p>To add a DSCP match condition to the named filter term:</p> <p>NOTE: A DSCP IP match condition and a precedence match condition cannot be both specified for the same term.</p> <p>a. Click Add in the DSCP section to see a list of match conditions.</p> <p>The Select DSCP list appears.</p> <p>b. Select one or more of the following DSCP types from the list:</p> <ul style="list-style-type: none"> • AF11—Assured forwarding class 1, low drop precedence • AF12—Assured forwarding class 1, medium drop precedence • AF21—Assured forwarding class 2, low drop precedence • AF22—Assured forwarding class 2, medium drop precedence • AF23—Assured forwarding class 2, high drop precedence • AF31—Assured forwarding class 3, low drop precedence • AF32—Assured forwarding class 3, medium drop precedence • AF33—Assured forwarding class 3, high drop precedence • AF41—Assured forwarding class 4, low drop precedence • AF42—Assured forwarding class 4, medium drop precedence • AF43—Assured forwarding class 4, high drop precedence • BE—Best effort (default) • EF—Expedited forwarding • CS0—Class selector 0 • CS1—Class selector 1 • CS2—Class selector 2 • CS3—Class selector 3 • CS4—Class selector 4 • CS5—Class selector 5 • CS6—Class selector 6 • CS7—Class selector 7 <p>c. Click OK.</p> <p>The DSCP code term for the named filter is added to the DSCP list.</p>

Table 88: Create Term Fields for EX Switching (*continued*)

Task	Description
Add a Precedence match condition	<p>You can apply an IP precedence match condition to the named term. With IP precedence, a device prioritizes traffic by class first. Then it differentiates and prioritizes same-class traffic.</p> <p>NOTE: The two match conditions IP Precedence and DSCP cannot be simultaneously applied to a term.</p> <p>To apply an IP precedence value match condition to the named term:</p> <ol style="list-style-type: none"> Click Add in the Precedence section. <p>The Select Precedence list appears.</p> <ol style="list-style-type: none"> Select one of the following precedence settings from the list: Routine (0 or lowest, also called Best Effort), Priority (1), Immediate (2), Flash (3, mainly used for voice signaling or for video), Flash-override (4), Critical-ECP (5, mainly used for voice RTP), Internet-control (6, used for IP routing protocols), or Net-control (7 or highest, used for link layer and routing protocol keep alive). Click OK. <p>The precedence match condition is added to the named term, and the condition is listed in the Precedence list.</p>

TCP Settings

Expand this section to see the TCP term settings. The Transmission Control Protocol (TCP) is the most common core protocol of the Internet protocol suite (IP). TCP provides reliable, ordered, error-checked delivery of a stream of octets between programs running on computers connected to the Internet or an intranet. You can use the TCP initial flag for a match condition.

Enable TCP Initial flag match condition	Select to use the TCP initial flag for a match condition. The TCP flags option becomes unavailable as a result.
---	---

Table 88: Create Term Fields for EX Switching (*continued*)

Task	Description
Enable other TCP flag match conditions	<p>If you are not using the TCP initial flag for a match condition, select one of the TCP flags from the list—RST, ACK, SYN, Urgent, Push, FIN, None. These flags have the following meaning:</p> <ul style="list-style-type: none"> • RST—Reset flag indicates that the TCP connection will be reset. • ACK—Third step in TCP three-way handshake for connection. In response to a server's SYN-ACK, the client replies with an ACK. • SYN—First step in TCP three-way handshake for connection. The active open is performed by the client sending a SYN to the server. • Urgent—If the URG flag is set, then the 16-bit field is an offset from the sequence number indicating the last urgent data byte. • Push—Push flags request that buffered data to the receiving application be sent now. • FIN—The final flag indicates that no more data will be sent.

ICMP Settings

Expand the ICMP Settings section to select an ICMP code value for the filter item's match condition. The Internet Control Message Protocol (ICMP) is one of the core IP protocols used by operating systems of networked computers to send error messages. ICMP can also be used to relay query messages.

Add an ICMP Code match condition	<p>To apply an ICMP code match condition to the named term:</p> <ol style="list-style-type: none"> Click Add in the ICMP Codes section. The Select ICMP Code list appears. Select one or more ICMP codes from the list. These codes vary, depending on the Filter Family you selected. Click OK. The ICMP code match condition is listed in the ICMP Code list and added to the named term. <p>NOTE: ICMP code specifies more specific information than ICMP type. Because the value's meaning depends upon the associated ICMP type, you must specify an ICMP type along with an ICMP code. The keywords are grouped by the ICMP type with which they are associated.</p>
----------------------------------	--

Table 88: Create Term Fields for EX Switching (*continued*)

Task	Description
Add an ICMP Type match condition	<p>NOTE: ICMP type specifies the ICMP packet type field. Typically, you specify this match condition in conjunction with the protocol match condition to determine which protocol is being used on the port.</p> <p>ICMP code specifies more specific information than ICMP type. Because the value's meaning depends upon the associated ICMP type, you must specify an ICMP type along with the ICMP code. The keywords are grouped by the ICMP type with which they are associated.</p> <p>To apply an ICMP type match condition to the named term:</p> <ol style="list-style-type: none"> Click Add in the ICMP Type section. The Select ICMP Type list appears. Select one or more ICMP types from the list. Options vary, depending on which Filter Family you selected. Click OK. The ICMP type match condition is listed in the ICMP Type list and is added to the named term.
Action	
Select the action that the system performs on an IP packet if all match conditions that you specified above are met. Possible actions are Discard and Accept. The default action is to discard a packet that matches the filter term's conditions.	
Action	<p>Select either Discard or Accept to indicate what the filter term does with a packet when a match is made.</p> <p>NOTE: The remaining fields in this section are enabled only if you select Accept as the action.</p>
Counter Name	When Accept is the action, specify a counter name.
Loss Priority	<p>When Accept is the action, specify the packet loss priority, Low, High, or None.</p> <p>NOTE: Forwarding class and loss priority must be specified together for the same term.</p>

Table 88: Create Term Fields for EX Switching (continued)

Task	Description
Policer	<p>When you create a Filter profile, you can specify a policer action for any term or terms within the filter. Policing, or rate limiting, enables you to limit the amount of traffic that passes into or out of an interface. All traffic that matches a term that contains a policer action goes through the policer that the term references.</p> <p>You have two options with a policer. You can specify that an existing policer be used for the packet that matches the match condition. Or, you can create a new policer for the packet that matches the match condition.</p> <p>To select a policer from an existing list of policers, click Select. The Select Policer page appears. Select the policer that you want to use for the term and click OK. The system displays the selected policer in the Policer field in the Create Term page.</p>

Table 88: Create Term Fields for EX Switching (*continued*)

Task	Description
	<p>To create a new policer:</p> <ol style="list-style-type: none"> Click Create. The Create Policer page appears. Type a name for the policer—you can use this policer again in the future. Select a policer type from the list, either a single-rate-two-color policer, or a three-color-policer. The type of policer that you select here affects the rest of the configurations available for the policer. If you selected a three-color-policer, then also select a rate for it, either single-rate or two-rate. <ul style="list-style-type: none"> Single-rate two-color—A two-color policer (sometimes called simply <i>policer</i>) meters the traffic stream and classifies packets into two categories of packet loss priority (PLP) according to a configured bandwidth and burst-size limit. You can mark packets that exceed the bandwidth and burst-size limit or simply discard them. A two-color policer is most useful for metering traffic at the port (physical interface) level. Single-Rate Three-color—This type of policer is defined in RFC 2697, A Single Rate Three Color Marker, as part of an assured forwarding (AF) per-hop-behavior (PHB) classification system for a Differentiated Services (DiffServ) environment. This type of policer meters traffic based on the configured committed information rate (CIR), committed burst size (CBS), and the excess burst size (EBS). Traffic is marked as belonging to one of three categories (green, yellow, or red) based on whether the packets are arriving at rates that are below the CBS (green), exceed the CBS but not the EBS (yellow), or exceed the EBS (red). A single-rate three-color policer is most useful when a service is structured according to packet size and not according to peak arrival rate. Two-rate three-color—This type of policer is defined in RFC 2698, A Two Rate Three Color Marker, as part of an assured forwarding (AF) per-hop-behavior (PHB) classification system for a Differentiated Services (DiffServ) environment. This type of policer meters traffic based on the configured CIR and the peak information rate (PIR), along with their associated burst sizes; the CBS, and the peak burst size (PBS). Traffic is marked as belonging to one of three categories (green, yellow, or red) based on packets are arriving at rates that are below the CIR (green), exceed the CIR but not the PIR (yellow), or exceed the PIR (red). A two-rate three-color policer is most useful when a service is structured according to arrival rates and not to packet size. <p>NOTE: The system displays and hides various fields in the Create Policer page depending on the type of policer that you want to create.</p>

Table 88: Create Term Fields for EX Switching (*continued*)

Task	Description
	<p>d. Configure these fields for a single-rate-two-color policer:</p> <ul style="list-style-type: none"> ● Bandwidth Limit—Specify the traffic rate in bits per second, 1000 through 102,300,000,000 (102.3g) bps. ● Burst Size Limit—Specify the maximum number of bytes allowed for incoming packets to burst above the peak information rate (PIR) and still be marked with medium-high packet loss priority (yellow). Packets that exceed the peak burst size (PBS) are marked with high packet loss priority (red). The range is 1500 through 100,000,000,000 bytes. ● Action—Select either Discard or None. ● Loss Priority—Select either High or None. <p>e. Configure these fields for a single-rate-three-color policer:</p> <ul style="list-style-type: none"> ● Committed Information Rate—Specify the guaranteed bandwidth (in bits per second) under normal line conditions and the average rate up to which packets are marked with low packet loss priority (green). The range is 32,000 through 40,000,000,000 bps. ● Committed Burst Size—Specify the maximum number of bytes allowed for incoming packets to burst above the committed information rate (CIR) and still be marked with low packet loss priority (green). The range is 1500 through 100,000,000,000 bytes. ● Excess Burst Size—Specify the maximum number of bytes allowed for incoming packets to burst above the committed information rate and still be marked with medium-high packet loss priority (yellow). Packets that exceed the excess burst size (EBS) are marked with high packet loss priority (red). The range is 1500 through 100,000,000,000 bytes. ● Color Mode—Select the way the preclassified packets are to be metered: <ul style="list-style-type: none"> ● Color-aware—The local switch can assign a higher packet loss priority but cannot assign a lower packet loss priority. ● Color-blind—The local switch ignores the preclassification of packets and can assign a higher or lower packet loss priority. ● None—The preclassified packets are not metered. ● Action—Options are Discard and None. ● Loss Priority—Options are High and None.

Table 88: Create Term Fields for EX Switching (*continued*)

Task	Description
	<p>f. Configure these fields for a three-color two-rate policer:</p> <ul style="list-style-type: none"> ● Committed Information Rate—Specify the guaranteed bandwidth (in bits per second) under normal line conditions and the average rate up to which packets are marked with low packet loss priority (green). The range is 32,000 through 40,000,000,000 bps. ● Committed Burst Size—Specify the maximum number of bytes allowed for incoming packets to burst above the committed information rate (CIR) and still be marked with low packet loss priority (green). The range is 1500 through 100,000,000,000 bytes. ● Peak Burst Size—Specify the maximum number of bytes allowed for incoming packets to burst above the peak information rate (PIR) and still be marked with medium-high packet loss priority (yellow). Packets that exceed the peak burst size (PBS) are marked with high packet loss priority (red). The range is 1500 through 100,000,000,000 bytes. ● Peak Information Rate—Specify the maximum achievable rate in bits per second. Packets that exceed the peak information rate (PIR) are marked with high packet loss priority (red). You can configure a discard action for packets that exceed the PIR. The range is 32,000 through 40,000,000,000 bps. ● Color Mode—Select the way the preclassified packets are to be metered: <ul style="list-style-type: none"> ● Color-aware—The local switch can assign a higher packet loss priority but cannot assign a lower packet loss priority. ● Color-blind—The local switch ignores the preclassification of packets and can assign a higher or lower packet loss priority. ● None—The preclassified packets are not metered. ● Action—Options are Discard and None. ● Loss Priority—Options are High and None. <p>g. Click OK.</p> <p>The policer is added to the list of applied policers and the list of available policers.</p>

Table 88: Create Term Fields for EX Switching (*continued*)

Task	Description
Forwarding Class	<p>When Accept is the action, specify the forwarding class (or output queue) that is to be used for the packet that matches the match condition. You can create a new forwarding class or select from a list of available forwarding classes.</p> <p>To select a forwarding class from an existing list of classes, click Select. The Select Forwarding Class page appears. Select the forwarding class that you want to use for the packet and click OK. The system displays the selected forwarding class in the Forwarding Class field in the Create Term page.</p> <hr/> <p>To create a new forwarding class:</p> <ol style="list-style-type: none"> Click Create. The Create Forwarding Class page appears. Type a name for the forwarding class—you can use this forwarding class again in the future. Select a queue number from the list, and then click OK. The system creates a new forwarding class and displays it in the Forwarding Class field in the Create Term page.

Click **OK** to save the term and return to the Create Filter Profile page.

Specifying Settings for a Campus Switching ELS Switch Filter Profile

A Filter profile must have at least one term in it. Each term has one filtering function. For example, if a term is evaluating the source of packets, then that term cannot also evaluate the protocols used by the packets. Some switch models accommodate multiple terms in one filter. When you have more than one term in a filter, the ordering of the terms is important. The system evaluates multiple filter terms as follows:

- The packet is evaluated against the first term's conditions. If the packet matches all of the conditions in that term, the corresponding action for that condition is taken and evaluation ends. Subsequent terms in the filter are not evaluated.
- If the packet does not match all conditions in the first term, the packet is evaluated against the conditions in the second term. This process continues until either the packet matches all the conditions in one of the subsequent terms or there are no more terms in the filter. If a match is found, the action specified in the Action section of the matched term is taken and the evaluation ends. Subsequent terms in the filter are not evaluated.

- The term conditions for protocol, EtherType, DSCP, precedence, ICMP code and ICMP type must all be either match conditions or except conditions.
- If a packet passes through all the terms in the filter without a match, the packet is discarded.

To configure a Filter profile for Campus switching ELS:

1. Specify a filter name and description for the Filter profile.
2. Select the switch filter family for which you want to create the profile:
 - If you want to create a Layer 2 based filter, select **Ethernet switching**.
 - If you want to create a Layer 3 based filter for IPv4, select **INET**.
 - If you want to create a Layer 3 based filter for IPv6, select **INET6**.
3. Under Terms, click **Add** to add one or more terms with match condition(s) to the named filter. You need at least one term for this filter.

The Create Term window opens.

NOTE: The order of the terms within a Filter profile configuration is important. Packets are tested against each term in the order in which the terms are listed.

4. Enter a name for the filter term.
5. Specify the match condition(s) for the filter term as described in [Table 89 on page 366](#). Required settings are indicated by a red asterisk (*) that appears next to the field label in the user interface.

Table 89: Create Term Fields for Campus Switching ELS

Field	Description
-------	-------------

Source and Destination Parameters

You can specify match conditions based on the packets' origin (source) or the packets' destination, or both. You are indicating the location of the filtering here—either specifying that packets that originate at a specific place (source) will be filtered or packets destined for a specific location (destination) will be filtered. You can have multiple sources and destinations for one filter.

Table 89: Create Term Fields for Campus Switching ELS (continued)

Field	Description
Source Parameters and Destination Parameters	<p>To add source and destination parameters to the named filter term:</p> <ol style="list-style-type: none"> Click Add to the right of the Destination Parameters lists. The Add Source/Destination Parameter window opens. Select either Source (default) or Destination from the Add Source/Destination Parameter window. Select one of following available Parameter Types from the Add Source/Destination Parameter page and provide the corresponding information: <ul style="list-style-type: none"> • IP Address—Provide the IP address of the source or destination device. • MAC Address—Provide a MAC address. • Port—Provide the port type of the source or destination port. Select either AFS (Andrew File System), BGP (Border Gateway Protocol), BIFF (UNIX mail notification), Bootpc (bootstrap protocol client), Bootps, Cmd, CVS pserver, DHCP, Domain, EK login, EK shell, EXEC, Finger protocol, FTP, FTP data, HTTP, HTTPS, Ident protocol, IMAP (Internet Message Access protocol), Kerberos-sec (Kerberos security), Klogin forwarding, Kpasswd command, KRB-prop (Kerberos database propagation), Krbupdate (Kerberos database update), Kshell (Kerberos rsh), LDAP, Login (UNIX rlogin), Mobilip-agent (Mobile IP agent), Mobilip-mn (Mobile IP MN), MSDP (Multicast Source Discovery Protocol), NetBIOS dgm, NetBIOS-ns (NetBIOS name service), NetBIOS-ssn (NetBIOS session service), NFSD, NNTP (Network News Transport Protocol), Ntalk, NTP (Network Time Protocol), POP3 (Post Office Protocol3), PPTP, Printer, RADacct (RADIUS accounting), RADIUS, RIP, RKINIT (Kerberos remote kinit), SMTP, SNMP trap, SNPP, SUNRPC, Syslog, TACACS, TACACS-ds, Talk (UNIX Talk), Telnet, TFTP, Timed (UNIX time daemon), Who (UNIX rwho), XDMCP (X Display Manager Control Protocol), Zephyr-clt (Zephyr serv-hm connection), Zephyr-hm (Zephyr hostmanager), Zephyr-srv (Zephyr server), or Other. <p>NOTE: If you selected Port as the parameter and do not find the type of port that you want to add from the Port list, then select other and enter a port number.</p> To select any other source/destination than the one indicated, enable Except. TIP: You cannot indicate both matching and except for a parameter. Click OK The parameter term is added to the appropriate list, either Source Parameters or Destination Parameters.

Table 89: Create Term Fields for Campus Switching ELS (continued)

Field	Description
Protocols and EtherTypes <p>Depending on the Filter Family you selected, you can sometimes apply a filter term based on either protocols being used by packets or on EtherTypes being used by packets. Recognized protocols are listed where applicable. Recognized EtherTypes, which indicate the protocol that is encapsulated in the payload of an Ethernet Frame, are also listed where applicable.</p>	
Protocols (apply to Ethernet and INET filter families)	<p>To add a protocol match condition to the named filter term:</p> <ol style="list-style-type: none"> Expand the Protocols and EtherTypes section. Click Add under Protocols. The Select Protocols window opens, displaying a list of protocols. From the list of protocols, select one or more. The options are AH, DSTOPTS, EGP, ESP, Fragment, GRE, Hop-by-hop, ICMP, IPIP, IPv6, No-text-header, OSPF, PIM, Routing, RSVP, SCTP, TCP, UDP, and VRRP. To make the filter exclude the specified protocol, select Except. NOTE: The term conditions for protocol, EtherType, DSCP, precedence, ICMP code and ICMP type must all be either match conditions or except conditions. Click OK. The protocols are added to the Protocols list.

Table 89: Create Term Fields for Campus Switching ELS (*continued*)

Field	Description
EtherTypes (apply to Ethernet filter family)	<p>To add an EtherTypes match condition to the named filter term:</p> <ol style="list-style-type: none"> Expand the Protocols and EtherTypes section. Click Add under EtherTypes. The Select EtherTypes window opens, displaying a list of protocols. From the list of EtherTypes, select one or more. The options are AARP, AppleTalk, ARP, IPV4, MPLS multicast, MPLS unicast, OAM, PPP, PPPOE discovery, PPPOE session, and SNA. To make the filter exclude the specified EtherType, select Except. NOTE: Term values must all be either match conditions or all except conditions. Click OK. The EtherTypes are added to the EtherTypes list.

DSCP Settings

Expand the DSCP section to see the DSCP match settings. DiffServ is a simple mechanism for classifying and managing network traffic and providing quality-of-service (QoS) on IP networks. DiffServ can, for example, be used to apply low-latency to critical network traffic such as voice or streaming media while providing simple best-effort service to non-critical services such as Web traffic. Here, you can apply a filter term based on the Differentiated Services code point (DSCP) which is a field in IPv4 and IPv6 headers.

NOTE: With IPv6 packets, the DS field and ECN field replace the IPv4 TOS field.

Table 89: Create Term Fields for Campus Switching ELS (*continued*)

Field	Description
DSCP (Ethernet and INET filter families)	<p>To add a DSCP match condition to the named filter term:</p> <p>NOTE: A DSCP IP match condition and a precedence match condition cannot be both specified for the same term.</p> <p>a. Click Add in the DSCP section.</p> <p>The Select DSCP Match Condition list appears.</p> <p>b. Select one of the following DSCP types from the list:</p> <ul style="list-style-type: none"> • AF11—Assured forwarding class 1, low drop precedence • AF12—Assured forwarding class 1, medium drop precedence • AF21—Assured forwarding class 2, low drop precedence • AF22—Assured forwarding class 2, medium drop precedence • AF23—Assured forwarding class 2, high drop precedence • AF31—Assured forwarding class 3, low drop precedence • AF32—Assured forwarding class 3, medium drop precedence • AF33—Assured forwarding class 3, high drop precedence • AF41—Assured forwarding class 4, low drop precedence • AF42—Assured forwarding class 4, medium drop precedence • AF43—Assured forwarding class 4, high drop precedence • BE—Best effort (default) • EF—Expedited forwarding • CS0—Class selector 0 • CS1—Class selector 1 • CS2—Class selector 2 • CS3—Class selector 3 • CS4—Class selector 4 • CS5—Class selector 5 • CS6—Class selector 6 • CS7—Class selector 7 <p>c. To make the filter exclude a specified DSCP type, select Except.</p> <p>NOTE: Term values must all be either match conditions or all of them need to be except conditions.</p> <p>d. Click OK.</p> <p>The DSCP code term for the named filter is added to the DSCP list.</p>

Table 89: Create Term Fields for Campus Switching ELS (*continued*)

Field	Description
Precedence for DSCP (Ethernet and INET filter families)	<p>You can apply an IP precedence match condition to the named term. With IP precedence, a device prioritizes traffic by class first. Then it differentiates and prioritizes same-class traffic.</p> <p>NOTE: The match conditions IP Precedence and DSCP cannot be simultaneously applied to a term.</p> <p>To apply an IP precedence value match condition to the named term:</p> <ol style="list-style-type: none"> Click Add in the Precedence section. <p>The Select Precedence list appears.</p> <ol style="list-style-type: none"> Select one of the following precedence settings from the list: Routine (0 or lowest, also called Best Effort), Priority (1), Immediate (2), Flash (3, mainly used for voice signaling or for video), Flash-override (4), Critical-ECP (5, mainly used for voice RTP), Internet-control (6, used for IP routing protocols), or Net-control (7 or highest, used for link layer and routing protocol keep alive). To make the filter exclude the specified IP precedence value, select Except. <p>NOTE: Term values must all be either match conditions or all of them need to be except conditions.</p> <ol style="list-style-type: none"> Click OK. <p>The precedence match condition is added to the named term, and the condition is listed in the Precedence list.</p>

TCP Settings

Expand this section to access the TCP settings. The Transmission Control Protocol (TCP) is the most common core protocol of the Internet protocol suite (IP). TCP provides reliable, ordered, error-checked delivery of a stream of octets between programs running on computers connected to the Internet or an intranet. You can use the TCP initial flag for a match condition.

Enable TCP Initial (all families)	<p>Select to use the TCP initial flag for an Ethernet, INET, or INET6 match condition.</p> <p>TIP: If you use the TCP initial flag for filtering, you cannot use any other TCP flag.</p>
-----------------------------------	---

Table 89: Create Term Fields for Campus Switching ELS (*continued*)

Field	Description
TCP Flags	<p>If you are not using the TCP initial flag for a match condition, you can select one of the TCP flags from the list for a match condition—RST, ACK, SYN, Urgent, Push, FIN, or None. These flags have the following meaning:</p> <ul style="list-style-type: none"> • RST—Reset flag indicates that the TCP connection will be reset. • ACK—Third step in TCP three-way handshake for connection. In response to a server's SYN-ACK, the client replies with an ACK. • SYN—First step in TCP three-way handshake for connection. The active open is performed by the client sending a SYN to the server. • Urgent—If the URG flag is set, then the 16-bit field is an offset from the sequence number indicating the last urgent data byte. • Push—Push flags request that buffered data to the receiving application be sent now. • FIN—The final flag indicates that no more data will be sent.

ICMP Settings

You can select the ICMP code value for the filter item's match condition—expand this section to access the ICMP settings. The Internet Control Message Protocol (ICMP) is one of the core IP protocols used by operating systems of networked computers to send error messages. ICMP can also be used to relay query messages.

ICMP Code	<p>To apply an ICMP code match condition to the named term:</p> <ol style="list-style-type: none"> Click Add in the ICMP Code section. The Select ICMP Code window appears. Select one or more ICMP codes from the list. These codes vary, depending on the Filter Family you selected. To make the filter exclude the specified ICMP code, select Except. NOTE: Term values must all be either match conditions or all of them need to be except conditions. Click OK. The ICMP code match condition is added to the named term, and the condition is listed in the ICMP Code list. You can now enable Except. <p>NOTE: An ICMP code specifies more specific information than an ICMP type. Because the value's meaning depends upon the associated ICMP type, you must specify an ICMP type along with ICMP code. The keywords are grouped by the ICMP type with which they are associated.</p>
-----------	---

Table 89: Create Term Fields for Campus Switching ELS (*continued*)

Field	Description
ICMP Type	<p>NOTE: ICMP type specifies the ICMP packet type field. Typically, you specify this match condition in conjunction with the protocol match condition to determine which protocol is being used on the port.</p> <p>ICMP code specifies more specific information than ICMP type. Because the value's meaning depends upon the associated ICMP type, you must specify ICMP type along with ICMP code. The keywords are grouped by the ICMP type with which they are associated.</p> <p>To apply an ICMP type match condition to the named term:</p> <ol style="list-style-type: none"> Click Add in the ICMP Type section. The Select ICMP Types window appears. Select one or more ICMP types from the list. These types vary, depending on the Filter Family selected. To make the filter exclude the specified ICMP type, select Except. NOTE: Term values must all be either match conditions or all of them need to be except conditions. Click OK. The ICMP type match condition is added to the named term, and the condition is listed in the ICMP Type list. You can now enable Except.
Action	
Select the action that the system performs on an IP packet if all match conditions that you specified above are met. Possible actions are Discard and Accept. The default action is to discard packet that matches the filter term conditions.	
Action	<p>Select either Discard or Accept to indicate what the filter term does with a packet when a match is made.</p> <p>NOTE: All other fields in this section are enabled only if you select Accept as the action.</p>
Counter Name	When the action selected is accept, specify the maximum packet count for this filter, term, or policer.

Table 89: Create Term Fields for Campus Switching ELS (continued)

Field	Description
Loss Priority	<p>When the action selected is accept, specify the packet loss priority, Low, High, Medium-low, Medium-high, or None.</p> <p>NOTE: Forwarding class and loss priority must be specified together for the same term.</p>

Table 89: Create Term Fields for Campus Switching ELS (continued)

Field	Description
Policer	<p>When you create a Filter profile with the action accept, you can specify a policer action for any term or terms within the filter. Policing, or rate limiting, enables you to limit the amount of traffic that passes into or out of an interface. All traffic that matches a term that contains a policer action goes through the policer that the term references.</p> <p>You have two options with a policer. You can specify that an existing policer be used for the packet that matches the match condition. Or, you can create a new policer for the packet that matches the match condition.</p> <hr/> <p>To select an existing policer:</p> <ol style="list-style-type: none"> Click Select. The Select Policer page appears. Click OK. The policer is added to the list of applied policers. <hr/>

Table 89: Create Term Fields for Campus Switching ELS (*continued*)

Field	Description
	<p>To create a new policer:</p> <ol style="list-style-type: none"> Click Create. The Create Policer page appears. Type a name for the policer—you can use this policer again in the future. Select a policer type from the list, either a single-rate-two-color policer, or a three-color-policer. The type of policer that you select here affects the rest of the configurations available for the policer. If you selected a three-color-policer, then also select a rate for it, either single-rate or two-rate. <ul style="list-style-type: none"> Single-rate two-color—A two-color policer (sometimes called simply <i>policer</i>) meters the traffic stream and classifies packets into two categories of packet loss priority (PLP) according to a configured bandwidth and burst-size limit. You can mark packets that exceed the bandwidth and burst-size limit or simply discard them. A two-color policer is most useful for metering traffic at the port (physical interface) level. Single-Rate Three-color—This type of policer is defined in RFC 2697, A Single Rate Three Color Marker, as part of an assured forwarding (AF) per-hop-behavior (PHB) classification system for a Differentiated Services (DiffServ) environment. This type of policer meters traffic based on the configured committed information rate (CIR), committed burst size (CBS), and the excess burst size (EBS). Traffic is marked as belonging to one of three categories (green, yellow, or red) based on whether the packets are arriving at rates that are below the CBS (green), exceed the CBS but not the EBS (yellow), or exceed the EBS (red). A single-rate three-color policer is most useful when a service is structured according to packet size and not according to peak arrival rate. Two-rate three-color—This type of policer is defined in RFC 2698, A Two Rate Three Color Marker, as part of an assured forwarding (AF) per-hop-behavior (PHB) classification system for a Differentiated Services (DiffServ) environment. This type of policer meters traffic based on the configured CIR and the peak information rate (PIR), along with their associated burst sizes; the CBS, and the peak burst size (PBS). Traffic is marked as belonging to one of three categories (green, yellow, or red) based on packets are arriving at rates that are below the CIR (green), exceed the CIR but not the PIR (yellow), or exceed the PIR (red). A two-rate three-color policer is most useful when a service is structured according to arrival rates and not to packet size. <p>NOTE: The system displays and hides various fields in the Create Policer page depending on the type of policer that you want to create.</p>

Table 89: Create Term Fields for Campus Switching ELS (*continued*)

Field	Description
	<p>d. Configure these fields for a single-rate-two-color policer:</p> <ul style="list-style-type: none"> ● Bandwidth Limit—Specify the traffic rate in bits per second, 1000 through 102,300,000,000 (102.3g) bps. ● Burst Size Limit—Specify the maximum number of bytes allowed for incoming packets to burst above the peak information rate (PIR) and still be marked with medium-high packet loss priority (yellow). Packets that exceed the peak burst size (PBS) are marked with high packet loss priority (red). The range is 1500 through 100,000,000,000 bytes. ● Action—Select either Discard or None. ● Loss Priority—Select either High or None. <p>e. Configure these fields for a single-rate-three-color policer:</p> <ul style="list-style-type: none"> ● Committed Information Rate—Specify the guaranteed bandwidth (in bits per second) under normal line conditions and the average rate up to which packets are marked with low packet loss priority (green). The range is 32,000 through 40,000,000,000 bps. ● Committed Burst Size—Specify the maximum number of bytes allowed for incoming packets to burst above the committed information rate (CIR) and still be marked with low packet loss priority (green). The range is 1500 through 100,000,000,000 bytes. ● Excess Burst Size—Specify the maximum number of bytes allowed for incoming packets to burst above the committed information rate and still be marked with medium-high packet loss priority (yellow). Packets that exceed the excess burst size (EBS) are marked with high packet loss priority (red). The range is 1500 through 100,000,000,000 bytes. ● Color Mode—Select the way the preclassified packets are to be metered: <ul style="list-style-type: none"> ● Color-aware—The local switch can assign a higher packet loss priority but cannot assign a lower packet loss priority. ● Color-blind—The local switch ignores the preclassification of packets and can assign a higher or lower packet loss priority. ● None—The preclassified packets are not metered. ● Action—Options are Discard and None. ● Loss Priority—Options are High and None.

Table 89: Create Term Fields for Campus Switching ELS (*continued*)

Field	Description
	<p>f. Configure these fields for a three-color two-rate policer:</p> <ul style="list-style-type: none"> ● Committed Information Rate—Specify the guaranteed bandwidth (in bits per second) under normal line conditions and the average rate up to which packets are marked with low packet loss priority (green). The range is 32,000 through 40,000,000,000 bps. ● Committed Burst Size—Specify the maximum number of bytes allowed for incoming packets to burst above the committed information rate (CIR) and still be marked with low packet loss priority (green). The range is 1500 through 100,000,000,000 bytes. ● Peak Burst Size—Specify the maximum number of bytes allowed for incoming packets to burst above the peak information rate (PIR) and still be marked with medium-high packet loss priority (yellow). Packets that exceed the peak burst size (PBS) are marked with high packet loss priority (red). The range is 1500 through 100,000,000,000 bytes. ● Peak Information Rate—Specify the maximum achievable rate in bits per second. Packets that exceed the peak information rate (PIR) are marked with high packet loss priority (red). You can configure a discard action for packets that exceed the PIR. The range is 32,000 through 40,000,000,000 bps. ● Color Mode—Select the way the preclassified packets are to be metered: <ul style="list-style-type: none"> ● Color-aware—The local switch can assign a higher packet loss priority but cannot assign a lower packet loss priority. ● Color-blind—The local switch ignores the preclassification of packets and can assign a higher or lower packet loss priority. ● None—The preclassified packets are not metered. ● Action—Options are Discard and None. ● Loss Priority—Options are High and None. <p>g. Click OK.</p> <p>The policer is added to the list of applied policers and the list of available policers.</p>

Table 89: Create Term Fields for Campus Switching ELS (continued)

Field	Description
Forwarding Class	<p>Specify the forwarding class (or output queue) that is to be used for the packet that matches the match condition. You can either select from a list of available forwarding classes or create a new forwarding class.</p> <p>To select a forwarding class from an existing list of classes, click Select. The Select Forwarding Class page appears. Select the forwarding class that you want to use for the packet and click OK. The system displays the selected forwarding class in the Forwarding Class field in the Create Term page.</p> <hr/> <p>To create a new forwarding class:</p> <ol style="list-style-type: none"> Click Create. The Create Forwarding Class page appears. Type a name for the forwarding class—you can use this forwarding class again in the future. Select a queue number from the list, and then click OK. The system creates a new forwarding class and displays it in the Forwarding Class field in the Create Term page.

6. Click **OK** to save the term and return to the Create Filter Profile page.

7. Click **Done**.

The new filter is added to the Manage Filter Profile list.

Specifying Settings for a Data Center Switching ELS Filter Profile

A Filter profile must have at least one term in it. Each term has one filtering function. For example, if a term is evaluating the source of packets, then that term cannot also evaluate the protocols used by the packets. Some switch models accommodate multiple terms in one filter. When you have more than one term in a filter, the ordering of the terms is important. The system evaluates multiple filter terms as follows:

- The packet is evaluated against the first term's conditions. If the packet matches all of the conditions in that term, the corresponding action for that condition is taken and evaluation ends. Subsequent terms in the filter are not evaluated.
- If the packet does not match all conditions in the first term, the packet is evaluated against the conditions in the second term. This process continues until either the packet matches all the conditions in one of the subsequent terms or there are no more terms in the filter. If a match is found, the action specified in the Action section of the matched term is taken and the evaluation ends. Subsequent terms in the filter are not evaluated.
- The term conditions for protocol, EtherType, DSCP, precedence, ICMP code and ICMP type must all be either match conditions or except conditions.
- If a packet passes through all the terms in the filter without a match, the packet is discarded.

To configure a Filter profile for Data Center switching ELS:

1. Specify a filter name and description for the Filter profile.
2. Select the switch filter family for which you want to create the profile:
 - If you want to create a Layer 2 based filter, select **Switching**.
 - If you want to create a Layer 3 based filter for IPv4, select **INET**.
 - If you want to create a Layer 3 based filter for IPv6, select **INET6**.
3. Under Terms, click **Add** to add one or more terms with match condition(s) to the named filter. You need at least one term for this filter.

The Create Term window opens.

NOTE: The order of the terms within a Filter profile configuration is important. Packets are tested against each term in the order in which the terms are listed.

4. Enter a name for the filter term.
5. Specify the match condition(s) for the filter term as described in [Table 90 on page 381](#). Required settings are indicated by a red asterisk (*) that appears next to the field label in the user interface.

Table 90: Create Term Fields for Data Center Switching ELS

Field	Description
-------	-------------

Source and Destination Parameters

You can specify match conditions based on the packets' origin (source) or the packets' destination, or both. You are indicating the location of the filtering here—either specifying that packets that originate at a specific place (source) will be filtered or packets destined for a specific location (destination) will be filtered. You can have multiple sources and destinations for one filter.

Table 90: Create Term Fields for Data Center Switching ELS (*continued*)

Field	Description
Source Parameters and Destination Parameters	<p>To add source and destination parameters to the named filter term:</p> <ol style="list-style-type: none"> Click Add to the right of the Destination Parameters lists. The Add Source/Destination Parameter window opens. Select either Source (default) or Destination from the Add Source/Destination Parameter window. Select one of following available Parameter Types from the Add Source/Destination Parameter page and provide the corresponding information: <ul style="list-style-type: none"> • IP Address—Provide the IP address of the source or destination device. • MAC Address—Provide a MAC address. • Port—Provide the port type of the source or destination port. Select either AFS (Andrew File System), BGP (Border Gateway Protocol), BIFF (UNIX mail notification), Bootpc (bootstrap protocol client), Bootps, Cmd, CVS pserver, DHCP, Domain, EK login, EK shell, EXEC, Finger protocol, FTP, FTP data, HTTP, HTTPS, Ident protocol, IMAP (Internet Message Access protocol), Kerberos-sec (Kerberos security), Klogin forwarding, Kpasswd command, KRB-prop (Kerberos database propagation), Krbupdate (Kerberos database update), Kshell (Kerberos rsh), LDAP, Login (UNIX rlogin), Mobilip-agent (Mobile IP agent), Mobilip-mn (Mobile IP MN), MSDP (Multicast Source Discovery Protocol), NetBIOS dgm, NetBIOS-ns (NetBIOS name service), NetBIOS-ssn (NetBIOS session service), NFSD, NNTP (Network News Transport Protocol), Ntalk, NTP (Network Time Protocol), POP3 (Post Office Protocol3), PPTP, Printer, RADacct (RADIUS accounting), RADIUS, RIP, RKINIT (Kerberos remote kinit), SMTP, SNMP trap, SNPP, SUNRPC, Syslog, TACACS, TACACS-ds, Talk (UNIX Talk), Telnet, TFTP, Timed (UNIX time daemon), Who (UNIX rwho), XDMCP (X Display Manager Control Protocol), Zephyr-clt (Zephyr serv-hm connection), Zephyr-hm (Zephyr hostmanager), Zephyr-srv (Zephyr server), or Other. <p>NOTE: If you selected Port as the parameter and do not find the type of port that you want to add from the Port list, then select other and enter a port number.</p> To select any other source/destination than the one indicated, enable Except. TIP: You cannot indicate both matching and except for a parameter. Click OK The parameter term is added to the appropriate list, either Source Parameters or Destination Parameters.

Table 90: Create Term Fields for Data Center Switching ELS (continued)

Field	Description
Protocols and EtherTypes	
<p>Depending on the Filter Family you selected, you can sometimes apply a filter term based on either protocols being used by packets or on EtherTypes being used by packets. Recognized protocols are listed where applicable. Recognized EtherTypes, which indicate the protocol that is encapsulated in the payload of an Ethernet Frame, are also listed where applicable.</p>	
Protocols (apply to Ethernet and INET filter families)	<p>To add a protocol match condition to the named filter term:</p> <ol style="list-style-type: none"> Expand the Protocols and EtherTypes section. Click Add under Protocols. The Select Protocols window opens, displaying a list of protocols. From the list of protocols, select one or more. The options are AH, DSTOPTS, EGP, ESP, Fragment, GRE, Hop-by-hop, ICMP, IPIP, IPv6, No-text-header, OSPF, PIM, Routing, RSVP, SCTP, TCP, UDP, and VRRP. To make the filter exclude the specified protocol, select Except. NOTE: The term conditions for protocol, EtherType, DSCP, precedence, ICMP code and ICMP type must all be either match conditions or except conditions. Click OK. The protocols are added to the Protocols list.

Table 90: Create Term Fields for Data Center Switching ELS (*continued*)

Field	Description
EtherTypes (apply to Ethernet filter family)	<p>To add an EtherTypes match condition to the named filter term:</p> <ol style="list-style-type: none"> Expand the Protocols and EtherTypes section. Click Add under EtherTypes. The Select EtherTypes window opens, displaying a list of protocols. From the list of EtherTypes, select one or more. The options are AARP, AppleTalk, ARP, FCoE, FIP, IPv4, MPLS multicast, MPLS unicast, OAM, PPP, PPPOE discovery, PPPOE session, and SNA. To make the filter exclude the specified EtherType, select Except. NOTE: Term values must all be either match conditions or all except conditions. Click OK. The EtherTypes are added to the EtherTypes list.

DSCP Settings

Expand the DSCP section to see the DSCP match settings. DiffServ is a simple mechanism for classifying and managing network traffic and providing quality-of-service (QoS) on IP networks. DiffServ can, for example, be used to apply low-latency to critical network traffic such as voice or streaming media while providing simple best-effort service to non-critical services such as Web traffic. Here, you can apply a filter term based on the Differentiated Services code point (DSCP) which is a field in IPv4 and IPv6 headers.

NOTE: With IPv6 packets, the DS field and ECN field replace the IPv4 TOS field.

Table 90: Create Term Fields for Data Center Switching ELS (*continued*)

Field	Description
DSCP (Ethernet and INET filter families)	<p>To add a DSCP match condition to the named filter term:</p> <p>NOTE: A DSCP IP match condition and a precedence match condition cannot be both specified for the same term.</p> <p>a. Click Add in the DSCP section.</p> <p>The Select DSCP Match Condition list appears.</p> <p>b. Select one of the following DSCP types from the list:</p> <ul style="list-style-type: none"> • AF11—Assured forwarding class 1, low drop precedence • AF12—Assured forwarding class 1, medium drop precedence • AF21—Assured forwarding class 2, low drop precedence • AF22—Assured forwarding class 2, medium drop precedence • AF23—Assured forwarding class 2, high drop precedence • AF31—Assured forwarding class 3, low drop precedence • AF32—Assured forwarding class 3, medium drop precedence • AF33—Assured forwarding class 3, high drop precedence • AF41—Assured forwarding class 4, low drop precedence • AF42—Assured forwarding class 4, medium drop precedence • AF43—Assured forwarding class 4, high drop precedence • BE—Best effort (default) • EF—Expedited forwarding • CS0—Class selector 0 • CS1—Class selector 1 • CS2—Class selector 2 • CS3—Class selector 3 • CS4—Class selector 4 • CS5—Class selector 5 • CS6—Class selector 6 • CS7—Class selector 7 <p>c. To make the filter exclude a specified DSCP type, select Except.</p> <p>NOTE: Term values must all be either match conditions or all of them need to be except conditions.</p> <p>d. Click OK.</p> <p>The DSCP code term for the named filter is added to the DSCP list.</p>

Table 90: Create Term Fields for Data Center Switching ELS (*continued*)

Field	Description
Precedence for DSCP (Ethernet and INET filter families)	<p>You can apply an IP precedence match condition to the named term. With IP precedence, a device prioritizes traffic by class first. Then it differentiates and prioritizes same-class traffic.</p> <p>NOTE: The match conditions IP Precedence and DSCP cannot be simultaneously applied to a term.</p> <p>To apply an IP precedence value match condition to the named term:</p> <ol style="list-style-type: none"> Click Add in the Precedence section. <p>The Select Precedence list appears.</p> <ol style="list-style-type: none"> Select one of the following precedence settings from the list: Routine (0 or lowest, also called Best Effort), Priority (1), Immediate (2), Flash (3, mainly used for voice signaling or for video), Flash-override (4), Critical-ECP (5, mainly used for voice RTP), Internet-control (6, used for IP routing protocols), or Net-control (7 or highest, used for link layer and routing protocol keep alive). To make the filter exclude the specified IP precedence value, select Except. <p>NOTE: Term values must all be either match conditions or all of them need to be except conditions.</p> <ol style="list-style-type: none"> Click OK. <p>The precedence match condition is added to the named term, and the condition is listed in the Precedence list.</p>

TCP Settings

Expand this section to access the TCP settings. The Transmission Control Protocol (TCP) is the most common core protocol of the Internet protocol suite (IP). TCP provides reliable, ordered, error-checked delivery of a stream of octets between programs running on computers connected to the Internet or an intranet. You can use the TCP initial flag for a match condition.

Enable TCP Initial (all families)	<p>Select to use the TCP initial flag for an Ethernet, INET, or INET6 match condition.</p> <p>TIP: If you use the TCP initial flag for filtering, you cannot use any other TCP flag.</p>
-----------------------------------	---

Table 90: Create Term Fields for Data Center Switching ELS (*continued*)

Field	Description
TCP Flags	<p>If you are not using the TCP initial flag for a match condition, you can select one of the TCP flags from the list for a match condition—RST, ACK, SYN, Urgent, Push, FIN, or None. These flags have the following meaning:</p> <ul style="list-style-type: none"> • RST—Reset flag indicates that the TCP connection will be reset. • ACK—Third step in TCP three-way handshake for connection. In response to a server's SYN-ACK, the client replies with an ACK. • SYN—First step in TCP three-way handshake for connection. The active open is performed by the client sending a SYN to the server. • Urgent—If the URG flag is set, then the 16-bit field is an offset from the sequence number indicating the last urgent data byte. • Push—Push flags request that buffered data to the receiving application be sent now. • FIN—The final flag indicates that no more data will be sent.

ICMP Settings

You can select the ICMP code value for the filter item's match condition—expand this section to access the ICMP settings. The Internet Control Message Protocol (ICMP) is one of the core IP protocols used by operating systems of networked computers to send error messages. ICMP can also be used to relay query messages.

ICMP Code	<p>To apply an ICMP code match condition to the named term:</p> <ol style="list-style-type: none"> Click Add in the ICMP Code section. The Select ICMP Code window appears. Select one or more ICMP codes from the list. These codes vary, depending on the Filter Family you selected. To make the filter exclude the specified ICMP code, select Except. NOTE: Term values must all be either match conditions or all of them need to be except conditions. Click OK. The ICMP code match condition is added to the named term, and the condition is listed in the ICMP Code list. <p>NOTE: An ICMP code specifies more specific information than an ICMP type. Because the value's meaning depends upon the associated ICMP type, you must specify an ICMP type along with ICMP code. The keywords are grouped by the ICMP type with which they are associated.</p>
-----------	---

Table 90: Create Term Fields for Data Center Switching ELS (continued)

Field	Description
ICMP Type	<p>NOTE: ICMP type specifies the ICMP packet type field. Typically, you specify this match condition in conjunction with the protocol match condition to determine which protocol is being used on the port.</p> <p>ICMP code specifies more specific information than ICMP type. Because the value's meaning depends upon the associated ICMP type, you must specify ICMP type along with ICMP code. The keywords are grouped by the ICMP type with which they are associated.</p> <p>To apply an ICMP type match condition to the named term:</p> <ol style="list-style-type: none"> Click Add in the ICMP Type section. The Select ICMP Types window appears. Select one or more ICMP types from the list. These types vary, depending on the Filter Family selected. To make the filter exclude the specified ICMP type, select Except. NOTE: Term values must all be either match conditions or all of them need to be except conditions. Click OK. The ICMP type match condition is added to the named term, and the condition is listed in the ICMP Type list.
Action	
Select the action that the system performs on an IP packet if all match conditions that you specified above are met. Possible actions are Discard and Accept. The default action is to discard packet that matches the filter term conditions.	
Action	<p>Select either Discard or Accept to indicate what the filter term does with a packet when a match is made.</p> <p>NOTE: All other fields in this section are enabled only if you select Accept as the action.</p>
Counter Name	When the action selected is accept, specify the maximum packet count for this filter, term, or policer.

Table 90: Create Term Fields for Data Center Switching ELS (*continued*)

Field	Description
Loss Priority	<p>When the action selected is accept, specify the packet loss priority, Low, High, Medium-low, Medium-high, or None.</p> <p>NOTE: Forwarding class and loss priority must be specified together for the same term.</p>

Table 90: Create Term Fields for Data Center Switching ELS (continued)

Field	Description
Policer	<p>When you create a Filter profile with the action accept, you can specify a policer action for any term or terms within the filter. Policing, or rate limiting, enables you to limit the amount of traffic that passes into or out of an interface. All traffic that matches a term that contains a policer action goes through the policer that the term references.</p> <p>You have two options with a policer. You can specify that an existing policer be used for the packet that matches the match condition. Or, you can create a new policer for the packet that matches the match condition.</p> <hr/> <p>To select an existing policer:</p> <ol style="list-style-type: none">Click Select.The Select Policer page appears.Click OK.The policer is added to the list of applied policers. <hr/>

Table 90: Create Term Fields for Data Center Switching ELS (*continued*)

Field	Description
	<p>To create a new policer:</p> <ol style="list-style-type: none"> Click Create. The Create Policer page appears. Type a name for the policer—you can use this policer again in the future. Select a policer type from the list, either a single-rate-two-color policer, or a three-color-policer. The type of policer that you select here affects the rest of the configurations available for the policer. If you selected a three-color-policer, then also select a rate for it, either single-rate or two-rate. <ul style="list-style-type: none"> Single-rate two-color—A two-color policer (sometimes called simply <i>policer</i>) meters the traffic stream and classifies packets into two categories of packet loss priority (PLP) according to a configured bandwidth and burst-size limit. You can mark packets that exceed the bandwidth and burst-size limit or simply discard them. A two-color policer is most useful for metering traffic at the port (physical interface) level. Single-Rate Three-color—This type of policer is defined in RFC 2697, A Single Rate Three Color Marker, as part of an assured forwarding (AF) per-hop-behavior (PHB) classification system for a Differentiated Services (DiffServ) environment. This type of policer meters traffic based on the configured committed information rate (CIR), committed burst size (CBS), and the excess burst size (EBS). Traffic is marked as belonging to one of three categories (green, yellow, or red) based on whether the packets are arriving at rates that are below the CBS (green), exceed the CBS but not the EBS (yellow), or exceed the EBS (red). A single-rate three-color policer is most useful when a service is structured according to packet size and not according to peak arrival rate. Two-rate three-color—This type of policer is defined in RFC 2698, A Two Rate Three Color Marker, as part of an assured forwarding (AF) per-hop-behavior (PHB) classification system for a Differentiated Services (DiffServ) environment. This type of policer meters traffic based on the configured CIR and the peak information rate (PIR), along with their associated burst sizes; the CBS, and the peak burst size (PBS). Traffic is marked as belonging to one of three categories (green, yellow, or red) based on packets are arriving at rates that are below the CIR (green), exceed the CIR but not the PIR (yellow), or exceed the PIR (red). A two-rate three-color policer is most useful when a service is structured according to arrival rates and not to packet size. <p>NOTE: The system displays and hides various fields in the Create Policer page depending on the type of policer that you want to create.</p>

Table 90: Create Term Fields for Data Center Switching ELS (*continued*)

Field	Description
	<p>d. Configure these fields for a single-rate-two-color policer:</p> <ul style="list-style-type: none"> ● Bandwidth Limit—Specify the traffic rate in bits per second, 8000 through 50,000,000,000 bps. ● Burst Size Limit—Specify the maximum number of bytes allowed for incoming packets to burst above the peak information rate (PIR) and still be marked with medium-high packet loss priority (yellow). Packets that exceed the peak burst size (PBS) are marked with high packet loss priority (red). The range is 1 through 2,147,450,880 bytes. ● Action—The default action is Discard. ● Loss Priority—Not available. <p>e. Configure these fields for a single-rate-three-color policer:</p> <ul style="list-style-type: none"> ● Committed Information Rate—Specify the guaranteed bandwidth (in bits per second) under normal line conditions and the average rate up to which packets are marked with low packet loss priority (green). The range is 1500 through 100,000,000,000 bps. ● Committed Burst Size—Specify the maximum number of bytes allowed for incoming packets to burst above the committed information rate (CIR) and still be marked with low packet loss priority (green). The range is 1500 through 100,000,000,000 bytes. ● Excess Burst Size—Specify the maximum number of bytes allowed for incoming packets to burst above the committed information rate and still be marked with medium-high packet loss priority (yellow). Packets that exceed the excess burst size (EBS) are marked with high packet loss priority (red). The range is 1500 through 100,000,000,000 bytes. ● Color Mode—Select the way the preclassified packets are to be metered: <ul style="list-style-type: none"> ● Color-aware—The local switch can assign a higher packet loss priority but cannot assign a lower packet loss priority. ● Color-blind—The local switch ignores the preclassification of packets and can assign a higher or lower packet loss priority. ● None—The preclassified packets are not metered. ● Action—Options are Discard and None. ● Loss Priority—Options are High and None.

Table 90: Create Term Fields for Data Center Switching ELS (*continued*)

Field	Description
	<p>f. Configure these fields for a three-color two-rate policer:</p> <ul style="list-style-type: none"> ● Committed Information Rate—Specify the guaranteed bandwidth (in bits per second) under normal line conditions and the average rate up to which packets are marked with low packet loss priority (green). The range is 1500 through 100,000,000,000 bps. ● Committed Burst Size—Specify the maximum number of bytes allowed for incoming packets to burst above the committed information rate (CIR) and still be marked with low packet loss priority (green). The range is 1500 through 100,000,000,000 bytes. ● Peak Burst Size—Specify the maximum number of bytes allowed for incoming packets to burst above the peak information rate (PIR) and still be marked with medium-high packet loss priority (yellow). Packets that exceed the peak burst size (PBS) are marked with high packet loss priority (red). The range is 1500 through 100,000,000,000 bytes. ● Peak Information Rate—Specify the maximum achievable rate in bits per second. Packets that exceed the peak information rate (PIR) are marked with high packet loss priority (red). You can configure a discard action for packets that exceed the PIR. The range is 1500 through 100,000,000,000 bps. ● Color Mode—Select the way the preclassified packets are to be metered: <ul style="list-style-type: none"> ● Color-aware—The local switch can assign a higher packet loss priority but cannot assign a lower packet loss priority. ● Color-blind—The local switch ignores the preclassification of packets and can assign a higher or lower packet loss priority. ● None—The preclassified packets are not metered. ● Action—Options are Discard and None. ● Loss Priority—Options are High and None. <p>g. Click OK.</p> <p>The policer is added to the list of applied policers and the list of available policers.</p>

Table 90: Create Term Fields for Data Center Switching ELS (continued)

Field	Description
Forwarding Class	<p>Specify the forwarding class (or output queue) that is to be used for the packet that matches the match condition. You can either select from a list of available forwarding classes or create a new forwarding class.</p> <p>To select a forwarding class from an existing list of classes, click Select. The Select Forwarding Class page appears. Select the forwarding class that you want to use for the packet and click OK. The system displays the selected forwarding class in the Forwarding Class field in the Create Term page.</p> <hr/> <p>To create a new forwarding class:</p> <ol style="list-style-type: none"> Click Create. The Create Forwarding Class page appears. Type a name for the forwarding class—you can use this forwarding class again in the future. Select a queue number from the list, and then click OK. The system creates a new forwarding class and displays it in the Forwarding Class field in the Create Term page.

6. Click **OK** to save the term and return to the Create Filter Profile page.

7. Click **Done**.

The new filter is added to the Manage Filter Profile list.

What to Do Next

After you create a Filter profile, you can do one of the following:

- Link the Filter profile as ingress and egress filters to a Port profile. For more information, see [“Creating and Managing Port Profiles” on page 251](#).
- Link the Filter profile as ingress and egress filters to a VLAN profile. For more information, see [“Creating and Managing VLAN Profiles” on page 329](#). You can then assign the VLAN profile to a device or port in case of switching devices.

RELATED DOCUMENTATION

[Understanding Filter Profiles | 349](#)

[Network Director Documentation home page](#)

Configuring Class of Service (CoS)

IN THIS CHAPTER

- Understanding Class of Service (CoS) Profiles | 396
- Creating and Managing Wired CoS Profiles | 400

Understanding Class of Service (CoS) Profiles

IN THIS SECTION

- How Would I Use CoS (also known as QoS)? | 397
- What CoS Parameters Can I Control? | 397
- What Are the Default CoS Traffic Types? | 398
- Data Center Switching CoS Configuration | 399
- How Do I Implement Class of Service? | 399
- Editing Discovered CoS Profiles | 399

When a network experiences congestion and delay, some packets must be prioritized to avoid random loss of data. Class of service (CoS) (also known as QoS) accomplishes this prioritization by dividing similar types of traffic, such as e-mail, streaming video, voice, large document file transfer, into classes. You then apply different levels of priority, such as those for throughput and packet loss, to each group, and thereby control traffic behavior. For example, when packets must be dropped, you can ensure that packet loss takes place according to your configured rules. CoS also enables you to rewrite the Differentiated Services code point (DSCP), IP precedence, or 802.1p CoS bits of packets exiting a specific interface, thus enabling you to tailor outgoing packets to meet the network requirements of remote peers.

How Would I Use CoS (also known as QoS)?

On an Ethernet trunk, you can mark frames with a class-of-service (CoS) value. CoS is used to define trunk connections as full-duplex, incoming only, or outgoing only.

Network devices such as routers and switches can be configured to use existing CoS values on incoming packets from other devices (trust mode), or can rewrite the CoS values to something completely different. Layer 2 markings also can extend to the WAN; for example, with a frame relay network. CoS is usually limited to use within an organization's intranet.

With legacy telephone systems, CoS can be used to define the permissions an extension will have on a private branch exchange (PBX) or Centrex. Some users might need extended voicemail message retention or the ability to forward calls to a cell phone, while others have no need to make calls outside the office. Permissions for a group of extensions can be changed by modifying a CoS variable applied to the entire group.

NOTE: CoS configurations can be complicated, so unless it is required, we recommend that you do not alter the default class names or queue number associations.

How Do I Create CoS Groups?

Use 802.1Q tagged VLANs to group users and enable CoS to set priorities supported by downstream devices.

How Is CoS Different From QoS?

CoS operates only on 802.1Q VLAN Ethernet at the data link layer (layer 2), while quality-of-service (QoS) mechanisms operate at the IP network layer (layer 3). 802.1p Layer 2 tagging can be used by QoS to differentiate and shape network traffic.

What CoS Parameters Can I Control?

You can use CoS profiles to group a set of class of service (CoS) parameters and apply it to one or more interfaces. You can configure the following parameters within a CoS profile:

- **Classifiers**—Packet classification refers to the examination of an incoming packet. This function associates the packet with a particular CoS servicing level.
- **Scheduler maps**—Schedulers define the properties of output queues. These properties include the amount of interface bandwidth assigned to the queue, the size of the memory buffer allocated for storing packets, the priority of the queue, and the drop profiles associated with the queue. You associate the schedulers with forwarding classes by means of scheduler maps. You can then associate each scheduler map with an interface, thereby configuring the queues, packet schedulers, and tail drop processes that operate according to this mapping.
- **Rewrite values**—A rewrite rule modifies the appropriate CoS bits in an outgoing packet. Modification of CoS bits enables the next downstream device to classify the packet into the appropriate service group. Rewriting or marking outbound packets is useful when the device is at the border of a network and must alter the CoS values to meet the policies of the targeted peer.
- **Traffic-control profile**—Traffic-control profiles enable traffic limitation of a certain class to a specified bandwidth and burst size. Packets exceeding the limits can be discarded, or can be assigned to a different forwarding class, a different loss priority, or both.

What Are the Default CoS Traffic Types?

On EX Series switches, the system provides you with these four predefined traffic types—Data, Voice, Video, and Network Control—with these default traffic configuration and shaping details:

- **Data**—Forwarding queue 0 (nd_best-effort), Buffer size 50%, Bandwidth reserved 30%
- **Voice**—Forwarding queue 5 (nd_expedited-forwarding), Buffer size 20%, Bandwidth reserved 0%
- **Video**—Forwarding queue 4 (nd_video-forwarding), Buffer size 20%, Bandwidth reserved 70%
- **Network Control**—Forwarding queue 7 (nd_network-control), Buffer size 10%, Bandwidth reserved 0%

For Campus Switching ELS, the system provides you with these four predefined traffic types—Data, Voice, Video, and Network Control—with these default traffic configuration and shaping details:

- **Data**—Forwarding queue 0 (nd_best-effort), Buffer size 50%, Bandwidth reserved 30%
- **Voice**—Forwarding queue 1 (nd_expedited-forwarding), Buffer size 20%, Bandwidth reserved 0%
- **Video**—Forwarding queue 2 (nd_video-forwarding), Buffer size 20%, Bandwidth reserved 70%
- **Network Control**—Forwarding queue 3 (nd_network-control), Buffer size 10%, Bandwidth reserved 0%

For Campus Switching ELS with *Hierarchical Port Scheduling* (Juniper Networks EX2300 Ethernet switches), Network Director provides you with predefined forwarding classes—nd_cs_best-effort, nd_cs_video-forwarding, nd_cs_expedited-forwarding, and nd_cs_network-control. These forwarding classes are grouped under two priority groups—data_video_pg and voice_control_pg.

For Campus Switching ELS with *Hierarchical Port Scheduling*, you can modify and customize each of these priority groups and forwarding classes. For more details, see [“Creating and Managing Wired CoS Profiles” on page 400](#).

Data Center Switching CoS Configuration

For data center switching devices, these additional CoS features are available:

- Hierarchical Port Scheduling (ETS)—Hierarchical port scheduling (Enhanced Transmission Selection, or ETS) is a two-tier process that provides better port bandwidth utilization and greater flexibility to allocate resources to queues and to groups of queues.
- Priority-based flow control (PFC)—A link-level flow control mechanism.

How Do I Implement Class of Service?

CoS can be implemented from the MSS CLI, from Network Director. RingMaster configures unicast traffic but does not configure multicast traffic. For directions to implement CoS from Network Director, see [“Creating and Managing Wired CoS Profiles” on page 400](#).

Editing Discovered CoS Profiles

Duplicate scheduler configuration is deployed to the device when you edit a CoS profile that are automatically created by Network Director as part of device discovery or out-of-band changes. In CoS configuration, a single classifier can be associated to multiple ports regardless of the other CoS configuration. When Network Director discovers a device with such configuration it will create multiple profiles, based on the difference in other CoS configurations, and mapped to same classifier configuration. If you modify classifier settings in such a CoS profile that is created automatically by Network Director, Network Director cannot modify the configuration because it is mapped to multiple profiles. Whenever you modify such a CoS profile that is created automatically, Network Director will create new classifier settings configuration on the device and map the same to it, without affecting the existing classifier settings. Newly created classifier settings will have a name generated based on the profile name. Even if only one profile is mapped to the classifier settings, Network Director creates new classifier settings and the old settings are orphaned.

NOTE: This behavior is applicable to both hierarchical and non hierarchical profiles, and is applicable for congestion notification profile name, traffic control profile name, scheduler map name, classifier name and rewrite rule settings.

RELATED DOCUMENTATION

Creating and Managing Wired CoS Profiles

IN THIS SECTION

- [Managing Wired CoS Profiles | 400](#)
- [Using the Default CoS Profiles for Switches | 401](#)
- [Using the Default CoS Profiles for Data Center Switching | 402](#)
- [Creating a Wired CoS Profile | 402](#)
- [Specifying Settings for a Switching and Campus Switching ELS CoS Profile | 403](#)
- [Specifying Settings for a Data Center Switching CoS Profile | 407](#)
- [What to Do Next | 416](#)

CoS profiles enable the grouping of class-of-service (CoS) parameters and apply them to one or more interfaces. Network Director provides you with predefined traffic types for each CoS profile that you create. These traffic types represent the most common types of traffic for the device type. Each of these templates has preconfigured values for all CoS parameters based on the typical application requirements. You can change the preconfigured values of these parameters to suit your requirements.

This topic describes:

Managing Wired CoS Profiles

From the Manage CoS Profiles page, you can:

- Create a new CoS profile by clicking **Add**. For details, see [“Creating a Wired CoS Profile” on page 402](#).
- Modify an existing CoS profile by selecting it and clicking **Edit**.
- View information about a profile by selecting the profile and clicking **Details**.
- Delete a CoS profile by selecting a profile and clicking **Delete**.

TIP: You cannot delete profiles that are in use—that is, assigned to objects or used by other profiles. To see the current assignments for a profile, select the profile and click **Details**.

- Clone an existing CoS profile by selecting it and clicking **Clone**.

Table 91 on page 401 describes the information provided about wired CoS profiles on the Manage CoS Profiles page. This page lists all CoS profiles defined for your network, regardless of the scope you selected in the network view.

Table 91: Managing Wired CoS Profile Fields

Field	Description
Profile Name	Name given to the profile when the profile was created.
Family Type	The device family on which the profile was created: EX Series Switches or Campus Switching ELS.
Description	<p>Description of the profile that was entered when the profile was created. If the profile was created by using the CLI and then discovered by Network Director, the description is <i>Profile created as part of device discovery</i>.</p> <p>TIP: To display the entire description, you might need to resize the Description column by clicking the column border in the heading and dragging it.</p>
Creation Time	Date and time when the profile was created.
Update Time	Date and time when the profile was last modified.
User Name	The username of the user who created or modified the profile.

TIP: All columns might not be displayed. To show or hide fields listed in the Manage Authorization Profiles table, click the down arrow on the field header, select **Columns**, and select or clear the check box adjacent to the field that you want to show or hide.

Using the Default CoS Profiles for Switches

When you install Network Director, a default CoS profile (juniper_CoS_template) is added to the Manage CoS Profiles page for EX Series switches and another with the same name is added for Campus Switching ELS. Default CoS profiles have most basic settings preconfigured. For example, the forwarding classes in

the default CoS profile have already been assigned with default scheduler values. However, you can use the Edit CoS Profile page to optimize your communication with the network by customizing the bandwidth and buffer size assigned to each of the forwarding classes in the default CoS profile.

Using the Default CoS Profiles for Data Center Switching

When you install Network Director, the following default CoS profiles are installed for Data Center Switching:

- juniper_DC_NonHier_Ethernet_CoS
- juniper_DC_Hier_Ethernet_CoS
- juniper_DC_NonHier_CoS
- juniper_DC_Hier_CoS
- juniper_DC_Hier_FCoE_CoS
- Juniper_DC_Hier_CoS_Fusion


To see the settings configured for a default profile, select it on the Manage CoS Profiles page, then click **Details**.

Creating a Wired CoS Profile

In Network Director, you can create a CoS profile to group a set of Class of Service parameters and apply it to one or more network sessions.

For a CoS profile, you must specify the profile name. You can use defaults for the other values.

To create a wired CoS profile:

1. Click  in the Network Director banner.
2. Under Select View, select one of the following: **Logical View**, **Location View**, **Device View** or **Custom Group View**.

TIP: Do not select **Dashboard View**, **Virtual View** or **Topology View**.

3. From the Tasks pane, expand **Wired**, expand **Profiles**, and then select **CoS**.
4. Click **Add** to add a new profile.

Network Director opens the Device Family Chooser window.

5. From the Device Family Chooser, select the wired device family for which you want to create a profile. The available device families are **Switching (EX)** or **Campus Switching ELS**.
6. Click OK.
7. Complete the appropriate settings using the steps mentioned in [“Specifying Settings for a Switching and Campus Switching ELS CoS Profile” on page 403](#), or [“Specifying Settings for a Data Center Switching CoS Profile” on page 407](#).

Specifying Settings for a Switching and Campus Switching ELS CoS Profile

Create a CoS profile for switching by providing a profile name and, optionally, changing any default settings for Traffic Configuration and Shaping.

1. Enter the CoS switching settings described in [Table 92 on page 403](#). Required settings are indicated by a red asterisk (*) that appears next to the field label in the user interface.

Table 92: CoS Profile Settings for EX and Campus Switching ELS

Field	Action
Profile Name	<p>Type the name of the profile.</p> <p>You can use up to 64 characters for profiles created for wired devices. Profile name must not contain special characters or spaces. Note that profiles that are automatically created by Network Director as part of device discovery or out-of-band changes may contain the underscore (_) character.</p>
Description	Type a description of the profile.

2. Network Director includes four predefined traffic types, Data, Voice, Video, and Network Control. You can either modify those traffic types or you can create your own traffic type. Modify and customize any listed traffic type by selecting the traffic type from the list and clicking **Edit**, then changing any of the settings described in [Table 93 on page 404](#).
3. To create your own traffic type, click **Add** and then configure the settings described in [Table 93 on page 404](#).

Table 93: Traffic Configuration and Shaping for EX Switching and Campus Switching ELS

Field	Description
Traffic Type	If you are editing a Network Director default traffic type, this field cannot be changed. If you are adding a traffic type, indicate the type of traffic—this can be any value, such as a server name or something to do with your business.
Forwarding Name	<p>If you are editing a Network Director default traffic type, this field cannot be changed. If you are adding a traffic type, you can use one of the predefined forwarding classes for your switch or you can create your own forwarding class. These forwarding classes are always provided: nd_best-effort, nd_network-control, nd_video-forwarding, and nd_expedited-forwarding. To create your own forwarding class, type a name instead of selecting an option.</p> <p>Most switches support the four predefined forwarding classes listed above. The exception is the EX4300 switch, which has eight default forwarding classes, including the standard four classes, plus multicast-network-connect, multicast-assured-forwarding, multicast-expedited-forwarding, and multicast-network-connect.</p>
Forwarding Queue	<p>Existing forwarding classes already have associated queues that cannot be altered. If you defined a new forwarding class by specifying your own Forwarding Name, then select an internal queue number to which forwarding classes are assigned. Most switches support queues 0 - 10. The exception is the EX4300 switch, which supports queues 0 - 11.</p> <p>By default, if a packet is not classified, it is assigned to the class associated with queue 0. You can assign more than one forwarding class to a queue number.</p>

Scheduler Map

A note in the Scheduler Map section indicates how much buffer size and bandwidth you have available to configure. For example, the message “You have been left with 0 percent buffer size and 0 percent bandwidth.” means that you have no available buffer or bandwidth, and you must reconfigure existing traffic types to free some bandwidth before configuring additional traffic types.

Low Priority	Enable Low Priority if you want the queue to receive low priority.
---------------------	---

Table 93: Traffic Configuration and Shaping for EX Switching and Campus Switching ELS (continued)

Field	Description
Strict High Priority	<p>Enable Strict High Priority if you want the queue to receive preferential treatment over a low-priority queue. Unlimited bandwidth is assigned to a strict-high priority queue.</p> <p>A strict-high priority queue receives preferential treatment over a low-priority queue. Unlimited bandwidth is assigned to a strict-high priority queue. Queues are scheduled according to the queue number, starting with the highest queue, 7, with decreasing priority down through queue 0. Traffic in higher-numbered queues is always scheduled prior to traffic in lower-numbered queues. In other words, in case of two high-priority queues, the queue with the higher queue number is processed first.</p> <p>NOTE: You can modify this field in the Traffic Configuration and Shaping table or from the Traffic Configuration and Shaping window.</p>
Buffer Size (%)	<p>Buffer Size (%) is the size of the memory buffer allocated for storing packets. Use the slider to specify the scheduler Buffer Size percentage.</p> <p>NOTE: You can modify this value by double-clicking this field in the Traffic Configuration and Shaping table or by sliding the bar in the Traffic Configuration and Shaping window.</p>
Bandwidth Reserved (%)	<p>Bandwidth Reserved (%) is the amount of interface bandwidth assigned to the queue. Move the slider to specify the Bandwidth Reserved percentage. Defaults are:</p> <ul style="list-style-type: none"> • Data: 30% • Voice: Strict High • Video: 70% • Network control: 0% <p>If Strict-High is enabled for this traffic type, you cannot reserve bandwidth.</p> <p>NOTE: This field displays the value based on either your input or on the transmit-rate parameter from the switch, if that parameter is configured. While specifying transmit-rate on the EX Series switch, if you choose to specify the value as an exact rate, Network Director converts this value and displays it as a percentage in the Bandwidth Reserved (%) field. You can modify this percentage value from the CoS Profile page.</p>
Shaping Rate	<p>Move the Shaping Rate slider to throttle the rate of packet transmission by setting a maximum bandwidth (rate in bits per second) or a maximum percentage of bandwidth for a queue or a forwarding class.</p>

Table 93: Traffic Configuration and Shaping for EX Switching and Campus Switching ELS (continued)

Field	Description
Traffic Classification Behavior aggregate classification classifies packets. The DSCP or DSCP IPv6 precedence bits of the IP header convey the behavior aggregate class information. The information might also be found in the IEEE 802.1ad, or IEEE 802.1p CoS bits.	
Classifier Type	Select a classifier type— DSCP , DSCP-IPv6 , INET-precedence , or IEEE-802.1 —and associate the corresponding code-point aliases to loss priorities. NOTE: You can specify code-point—loss priority associations for one or more classifier types. <ul style="list-style-type: none"> • DSCP—Differentiated services code point, a field in IPv4 headers, is used to classify traffic. • DSCP-IPv6—Differentiated services code point, a field in IPv6 headers, is used to classify traffic. • INET precedence—Field that indicates class of service rewrite rules are used to classify traffic. • IEEE-802.1—IEEE 802.1ad, or IEEE 802.1p CoS bits are used to classify traffic.
Classifier Code Points	
Code Points	The code points list includes all available and unselected code points for the selected classifier type. Specify one or more code-point aliases or bit sets to associate with a forwarding class by moving the value to one of the two lists, Loss Priority Low or Loss Priority High.
Loss Priority Low	Indicate that packets have low loss priority by selecting code-point aliases from the Code Points table and using the LEFT and RIGHT arrows to move them into the appropriate loss priority table.
Loss Priority Medium-Low	Indicate that packets have medium-low loss priority by selecting code-point aliases from the Code Points table and using the LEFT and RIGHT arrows to move them into the appropriate loss priority table.
Loss Priority Medium-High	Indicate that packets have medium-high loss priority by selecting code-point aliases from the Code Points table and using the LEFT and RIGHT arrows to move them into the appropriate loss priority table.

Table 93: Traffic Configuration and Shaping for EX Switching and Campus Switching ELS (*continued*)

Field	Description
Loss Priority High	Indicate that packets have high loss priority by selecting code-point aliases from the Code Points table and using the LEFT and RIGHT arrows to move them into the appropriate loss priority table.

- Click **OK** to close the Add Traffic and Classification window and save your configuration.

Your changes are added to this CoS profile.

NOTE: If all bandwidth has already been reserved, your changes are not made. Reduce the bandwidth reserved from another Traffic Type, then repeat the configuration.

- To configure rewrite rules for a forwarding queue, click **Configure Rewrite Rules** at the bottom of the screen. The Configure Rewrite Rules window appears. Specify rewrite rule settings as described below to alter CoS values in outgoing packets on the outbound interfaces of an edge switch:
 - Select the forwarding class for which you want to create or modify rewrite rules. Network Director lists all the forwarding classes that you have used for configuring traffic in the Traffic Configuration and Shaping section.
 - For each classifier's loss priority, select a code-point alias for each loss-priority type—Low, Medium-Low, Medium-High, and High.
- Click **OK** to save the rewrite rules and close the Configure Rewrite Rules window.

The system saves the rewrite rules and returns to the **Create CoS Profile** page.

- Click **Done**.

After you create a CoS profile for switching devices, associate the CoS profile with a Port profile. For directions, see [“Creating and Managing Port Profiles” on page 251](#).

Specifying Settings for a Data Center Switching CoS Profile

You can create a CoS profile by specifying the profile settings and the traffic configuration and shaping details.

To specify the settings for the CoS profile:

1. Enter the settings described in [Table 94 on page 408](#). Required settings are indicated by a red asterisk (*) that appears next to the field label in the user interface.

Table 94: CoS Profile Basic Settings for Data Center Switching

Field	Action
Profile Name	Type the name of the profile. You can use up to 64 characters for profiles created for wired devices. Profile name must not contain special characters or spaces. Note that profiles that are automatically created by Network Director as part of device discovery or out-of-band changes may contain the underscore (_) character.
Description	Type the description of the profile.

2. In the Traffic Classification and Shaping Settings section, select one of these options:
 - **Hierarchical Port Scheduling (ETS)**—Hierarchical port scheduling (Enhanced Transmission Selection, or ETS) is a two-tier process that provides better port bandwidth utilization and greater flexibility to allocate resources to queues and to groups of queues (for QFX devices).
 - **Non Hierarchical Port Scheduling**—Non-hierarchical scheduling is a one-tier process that provides port bandwidth utilization and allocates resources to queues (for EX4600 and EX4650 transit switches).
 - **Hierarchical (Fusion)**—Select this scheduling type if you plan to assign the CoS profile to QFX10002 and QFX10008 switches
3. If you selected Hierarchical Port Scheduling (ETS), specify settings in the Priority Group and Traffic Settings section.

The table lists priority groups and the forwarding classes they contain in an expandable list. Priority groups refer to forwarding class sets in the device. You can perform these tasks on priority groups and forwarding classes:

- To add a new priority group, click **Add Priority Group**. The Add Priority Group and Traffic Control Profile Window opens. Enter the settings as described in [Table 95 on page 409](#).

Table 95: Add Priority Group and Traffic Control Profile Window

Field	Description
Priority Group Name	Enter a name for the priority group.
Traffic Control Profile Settings	
Transmit Rate (%)	Select a transmit rate percentage for the priority group.
Shaping Rate (%)	Select a shaping rate percentage for the priority group.

- To edit a priority group or forwarding class's properties, click the field that you want to edit in the table.
- To edit a forwarding class's properties, click its name. The Edit Traffic Classification and Shaping for priority group window opens. Enter the settings as described in [Table 96 on page 409](#).

Table 96: Edit and Add Traffic Classification and Shaping for Priority Group Window

Field	Description
Forwarding Class Name	Select or specify a name for the forwarding class.
Forwarding Class Queue	Specify the internal queue numbers to which forwarding classes are assigned.
No Loss	Select to make the forwarding class lossless.
Scheduler Map	
Strict High	Select if you want the queue to receive preferential treatment over a low-priority queue. Unlimited bandwidth is assigned to a strict-high priority queue.
Transmit Rate	Select the percentage of interface bandwidth assigned to the forwarding class. If you have enabled Strict-High , you cannot reserve bandwidth for this traffic type.
Shaping Rate	Select a shaping rate percentage for the forwarding class.
Buffer Size (%)	Select the percentage of the memory buffer allocated for storing packets for the forwarding class.
Traffic Classification	
Classifier Type	Select the classifier type that maps packets to a forwarding class and a loss priority.

Table 96: Edit and Add Traffic Classification and Shaping for Priority Group Window (continued)

Field	Description
Code Points	Specify one or more code-points for associating with a forwarding class.
Loss Priority Low	Indicates that packets have low loss priority. Select code points from the Code Points table and use the LEFT and RIGHT arrows to move them into the appropriate loss priority table.
Loss Priority Medium High	Indicates that packets have medium high loss priority. Select code points from the Code Points table and use the LEFT and RIGHT arrows to move them into the appropriate loss priority table.
Loss Priority High	Indicates that packets have high loss priority. Select code points from the Code Points table and use the LEFT and RIGHT arrows to move them into the appropriate loss priority table.

- To add a forwarding class to a priority group, click the **Add Forwarding Class** link at the end of the priority group's list of forwarding classes. The Add Traffic Classification and Shaping for priority group window opens. Enter the settings as described in [Table 96 on page 409](#).
 - To remove a priority group or forwarding class, click the **X** at the end of its table row.
4. If you selected Non Hierarchical Port Scheduling, specify settings in the Traffic Configuration and Shaping table.

The table lists forwarding classes. You can perform these tasks on forwarding classes:

- To add traffic configuration and shaping details for different types of traffic, click **Add** in the Traffic Configuration and Shaping box. The Add Traffic Classification and Shaping window opens.
- To modify the details of an existing traffic configuration, select the traffic configuration from the list and click **Edit**. The Edit Traffic Classification and Shaping window opens.

NOTE: You can modify some of the details in the Traffic Configuration and Shaping table without having to open the Edit Traffic Classification and Shaping window—by clicking on the field that you want to modify.

- To delete a traffic configuration entry, select the traffic configuration from the list and click **Remove**.

The system deletes the selected traffic configuration entry.

To create your own traffic type, click **Add** and then configure the settings described in [Table 97 on page 411](#).

Table 97: Traffic Configuration and Shaping for EX Switching and Campus Switching ELS

Field	Description
Traffic Type	If you are editing a Network Director default traffic type, this field cannot be changed. If you are adding a traffic type, indicate the type of traffic—this can be any value, such as a server name or something to do with your business.
Forwarding Name	<p>If you are editing a Network Director default traffic type, this field cannot be changed. If you are adding a traffic type, you can use one of the predefined forwarding classes for your switch or you can create your own forwarding class. These forwarding classes are always provided: nd_best-effort, nd_network-control, nd_video-forwarding, and nd_expedited-forwarding. To create your own forwarding class, type a name instead of selecting an option.</p> <p>Most switches support the four predefined forwarding classes listed above. The exception is the EX4300 switch, which has eight default forwarding classes, including the standard four classes, plus multicast-network-connect, multicast-assured-forwarding, multicast-expedited-forwarding, and multicast-network-connect.</p>
Forwarding Queue	<p>Existing forwarding classes already have associated queues that cannot be altered. If you defined a new forwarding class by specifying your own Forwarding Name, then select an internal queue number to which forwarding classes are assigned. Most switches support queues 0 - 10. The exception is the EX4300 switch, which supports queues 0 - 11.</p> <p>By default, if a packet is not classified, it is assigned to the class associated with queue 0. You can assign more than one forwarding class to a queue number.</p>

Scheduler Map

A note in the Scheduler Map section indicates how much buffer size and bandwidth you have available to configure. For example, the message “You have been left with 0 percent buffer size and 0 percent bandwidth.” means that you have no available buffer or bandwidth, and you must reconfigure existing traffic types to free some bandwidth before configuring additional traffic types.

Low Priority	Enable Low Priority if you want the queue to receive low priority.
---------------------	---

Table 97: Traffic Configuration and Shaping for EX Switching and Campus Switching ELS (continued)

Field	Description
Strict High Priority	<p>Enable Strict High Priority if you want the queue to receive preferential treatment over a low-priority queue. Unlimited bandwidth is assigned to a strict-high priority queue.</p> <p>A strict-high priority queue receives preferential treatment over a low-priority queue. Unlimited bandwidth is assigned to a strict-high priority queue. Queues are scheduled according to the queue number, starting with the highest queue, 7, with decreasing priority down through queue 0. Traffic in higher-numbered queues is always scheduled prior to traffic in lower-numbered queues. In other words, in case of two high-priority queues, the queue with the higher queue number is processed first.</p> <p>NOTE: You can modify this field in the Traffic Configuration and Shaping table or from the Traffic Configuration and Shaping window.</p>
Buffer Size (%)	<p>Buffer Size (%) is the size of the memory buffer allocated for storing packets. Use the slider to specify the scheduler Buffer Size percentage.</p> <p>NOTE: You can modify this value by double-clicking this field in the Traffic Configuration and Shaping table or by sliding the bar in the Traffic Configuration and Shaping window.</p>
Bandwidth Reserved (%)	<p>Bandwidth Reserved (%) is the amount of interface bandwidth assigned to the queue. Move the slider to specify the Bandwidth Reserved percentage. Defaults are:</p> <ul style="list-style-type: none"> • Data: 30% • Voice: Strict High • Video: 70% • Network control: 0% <p>If Strict-High is enabled for this traffic type, you cannot reserve bandwidth.</p> <p>NOTE: This field displays the value based on either your input or on the transmit-rate parameter from the switch, if that parameter is configured. While specifying transmit-rate on the EX Series switch, if you choose to specify the value as an exact rate, Network Director converts this value and displays it as a percentage in the Bandwidth Reserved (%) field. You can modify this percentage value from the CoS Profile page.</p>
Shaping Rate	<p>Move the Shaping Rate slider to throttle the rate of packet transmission by setting a maximum bandwidth (rate in bits per second) or a maximum percentage of bandwidth for a queue or a forwarding class.</p>

Table 97: Traffic Configuration and Shaping for EX Switching and Campus Switching ELS (*continued*)

Field	Description
Traffic Classification <p>Behavior aggregate classification classifies packets. The DSCP or DSCP IPv6 precedence bits of the IP header convey the behavior aggregate class information. The information might also be found in the IEEE 802.1ad, or IEEE 802.1p CoS bits.</p>	
Classifier Type	<p>Select a classifier type—DSCP, DSCP-IPv6, INET-precedence, or IEEE-802.1—and associate the corresponding code-point aliases to loss priorities.</p> <p>NOTE: You can specify code-point—loss priority associations for one or more classifier types.</p> <ul style="list-style-type: none"> • DSCP—Differentiated services code point, a field in IPv4 headers, is used to classify traffic. • DSCP-IPv6—Differentiated services code point, a field in IPv6 headers, is used to classify traffic. • INET precedence—Field that indicates class of service rewrite rules are used to classify traffic. • IEEE-802.1—IEEE 802.1ad, or IEEE 802.1p CoS bits are used to classify traffic.
Classifier Code Points	
Code Points	<p>The code points list includes all available and unselected code points for the selected classifier type.</p> <p>Specify one or more code-point aliases or bit sets to associate with a forwarding class by moving the value to one of the two lists, Loss Priority Low or Loss Priority High.</p>
Loss Priority Low	Indicate that packets have low loss priority by selecting code-point aliases from the Code Points table and using the LEFT and RIGHT arrows to move them into the appropriate loss priority table.
Loss Priority Medium-Low	Indicate that packets have medium-low loss priority by selecting code-point aliases from the Code Points table and using the LEFT and RIGHT arrows to move them into the appropriate loss priority table.
Loss Priority Medium-High	Indicate that packets have medium-high loss priority by selecting code-point aliases from the Code Points table and using the LEFT and RIGHT arrows to move them into the appropriate loss priority table.

Table 97: Traffic Configuration and Shaping for EX Switching and Campus Switching ELS (*continued*)

Field	Description
Loss Priority High	Indicate that packets have high loss priority by selecting code-point aliases from the Code Points table and using the LEFT and RIGHT arrows to move them into the appropriate loss priority table.

- If you selected Hierarchical Port Scheduling (ETS), specify priority-based flow control (PFC) settings in the PFC Settings section. Enter the settings as described in [Table 98 on page 414](#).

Table 98: PFC Settings for Data Center Switching Hierarchical Port Scheduling (ETS) CoS Profile

Field	Description
Input Cable Length (meter)	Enter the length of the cable attached to the input interface, in meters.
Input	
Add	Click to add an input congestion notification profile (CNP). A new entry appears in the table.
Remove	Click to remove the selected input CNP.
IEEE Code Point	Select the IEEE code point for the input CNP.
Maximum Receive Size (bytes)	Enter the maximum receive unit (MRU) on an interface for traffic that matches the PFC priority, in bytes.
Output	
Add	Click to add an output CNP. A new entry appears in the table.
Remove	Click to remove the selected output CNP.
IEEE Code Point	Select the IEEE code point for the output CNP.
Queue List	Select output queues on which to enable flow control (PFC pause).

- If you selected Non-Hierarchical Port Scheduling, specify priority-based flow control (PFC) settings in the PFC Settings section. Enter the settings as described in [Table 99 on page 415](#).

Table 99: PFC Settings for Data Center Switching Non-Hierarchical Port Scheduling CoS Profile

Field	Description
Input	
Add	Click to add an input congestion notification profile (CNP). A new entry appears in the table.
Remove	Click to remove the selected input CNP.

7. If you selected Hierarchical Port Scheduling (ETS), specify rewrite rule settings in the Rewrite Rule Settings section as described in [Table 100 on page 415](#).

Table 100: Rewrite Rule Settings for Data Center Switching CoS Profile

Field	Description
Forwarding Name	The name of the forwarding class.
Queue	The number corresponding to the forwarding queue. You cannot modify this field.
Rewrite Type	Select a rewrite-rules mapping for the traffic that passes through the various queues on the interface.
Egress Code Point - Loss Priority Low	Specify a code-point for association with a forwarding class for loss priority low.
Egress Code Point - Loss Priority Medium High	Specify a code-point for association with a forwarding class for loss priority medium high.
Egress Code Point - Loss Priority High	Specify a code-point for association with a forwarding class for loss priority high.

8. If you selected Non-Hierarchical Port Scheduling, click **Configure Rewrite Rules** at the bottom of the screen to configure rewrite rules for a forwarding queue. The Configure Rewrite Rules window appears. Specify rewrite rule settings as described below to alter CoS values in outgoing packets on the outbound interfaces of an edge switch:
- Select the forwarding class for which you want to create or modify rewrite rules. Network Director lists all the forwarding classes that you have used for configuring traffic in the Traffic Configuration and Shaping section.

- b. For each classifier's loss priority, select a code-point alias for each loss-priority type—Low, Medium-Low, Medium-High, and High.

9. Click **Done** to save the changes to the profile.

What to Do Next

After you have created a CoS profile for switching devices, you can associate the CoS profile to a Port profile.

RELATED DOCUMENTATION

| [Understanding Class of Service \(CoS\) Profiles](#) | 396

Configuring Media Access Control Security (MACsec)

IN THIS CHAPTER

- [Media Access Control Security Overview | 417](#)
- [Configuring and Managing MACsec Profiles | 418](#)
- [Assigning the MACsec Profiles | 423](#)

Media Access Control Security Overview

Media Access Control Security (MACsec) is an industry-standard security technology that provides secure communication on Ethernet links. MACsec enables you to secure Ethernet links between two MACsec-capable devices. You can enable MACsec on point-to-point Ethernet links using static secure association key (SAK) security mode or static connectivity association key (CAK) security mode.

When you enable MACsec using the static CAK security mode, a connectivity association key and a randomly generated secure association key are exchanged between the devices on each point-to-point Ethernet link. After the matching pre-shared keys are successfully exchanged, MACsec enables MKA protocol on the devices. The MKA protocol maintains MACsec on the link, and decides which switch on the point-to-point link becomes the key server. The key server then creates an SAK that is shared with the switch at the other end of the point-to-point link only, and that SAK is used to secure all data traffic traversing the link. The key server will continue to periodically create and share a randomly-created SAK over the point-to-point link for as long as MACsec is enabled. A pre-shared key includes a connectivity association name (CKN) and its own connectivity association key (CAK). The CKN and CAK are configured by the user in the connectivity association and must match on both ends of the link to initially enable MACsec.

NOTE: A pre-shared key includes a connectivity association name (CKN) and its own connectivity association key (CAK). You can configure the CKN and CAK in the connectivity association and these values must match on both ends.

When you enable MACsec using static SAK security mode, you must configure the secure channels between the point-to-point Ethernet link. The secure channels are responsible for transmitting and receiving data

on the MACsec-enabled link, and also responsible for transmitting SAKs across the link to enable and maintain MACsec. A typical connectivity association when MACsec is enabled using SAK security mode contains two secure channels—one secure channel for inbound traffic and another secure channel for outbound traffic. You must configure the SAK settings manually, there is no key server or other tool that creates SAKs. Security is maintained on the point-to-point Ethernet link by periodically rotating between the two security keys. Each security key name and value must have a corresponding matching value on the interface at the other end of the point-to-point Ethernet link to maintain MACsec on the link.

MACsec is widely used in campus deployments to secure network traffic between endpoints and access switches. You can enable MACsec on extended ports in a Junos Fusion Enterprise topology to provide secure communication between the satellite device and connected hosts. Network Director supports MACsec configuration for a Junos Fusion Enterprise setup. You can create a profile for the MACsec configuration and assign the profiles to the extended ports of the satellite devices in a Junos Fusion Enterprise setup.

For more information about MACsec, see [Understanding Media Access Control Security \(MACsec\)](#).

Configuring and Managing MACsec Profiles

IN THIS SECTION

- [Creating a MACsec Profile | 419](#)
- [Specifying Settings for a MACsec Profile | 420](#)
- [What to Do Next | 423](#)

From the MACsec Profile page of the Network Director UI you can create and manage MACsec profiles that specify MACsec settings for the extended ports in the aggregation device in a Junos Fusion Enterprise device. From the Manage MACsec Profile page, you can:

- Create a new MACsec profile by clicking **Add**.
- Modify an existing MACsec profile by selecting the profile and clicking **Edit**.
- Associate a profile to the extended ports by selecting the profile and clicking **Assign**.
- Change current assignments for a profile by selecting the profile and clicking **Edit Assignment**.
- Delete a MACsec profile by selecting the profile and clicking **Delete**.
- Clone an existing MACsec profile by selecting the profile and clicking **Clone**.
- View information about a profile by selecting the profile and clicking **Details**.

Table 101 on page 419 describes the information provided about wired MACsec profiles on the Manage MACsec Profiles page. This page lists all the MACsec profiles defined for the Junos Fusion Enterprise device, regardless of the scope you selected in the network view.

Table 101: Managing MACsec Profile Fields

Field	Description
Profile Name	Name of the profile.
Connection Association Name	Name of the MACsec connectivity association.
Description	Description of the profile.
MACsec Mode	Static secure association key (static-SAK) security mode or static connectivity association key (static-CAK) using which you enabled MACsec on the device.
Assignment State	Profile assignment state. One of the following: <ul style="list-style-type: none"> • Deployed—The profile has been assigned and the configuration has been deployed on the devices. • Pending Deployment—The profile has been assigned or its previous assignments have been changed, but the new or modified configuration has not yet been deployed on the devices. • Unassigned—The profile has not yet been assigned.
User Name	The username of the user who created or modified the profile.


This topic describes:

Creating a MACsec Profile

To create a MACsec profile:

1. Under Views, select one of these options: **Logical View**, **Location View**, **Device View**, or **Custom Group View**.

TIP: Do not select **Dashboard View**, **Datacenter View**, or **Topology View**.

2. Click  in the Network Director banner.
3. In the Tasks pane, expand **Wired**, expand **Profiles**, and then select **MACsec**.

The Manage MACsec Profile page appears, displaying the list of currently configured MACsec profiles.

4. Click **Add** to add a new profile.

The Create MACsec Profile page appears.

5. Enter the MACsec settings described in [“Specifying Settings for a MACSsec Profile” on page 420](#).

6. Click **Done**.

Specifying Settings for a MACSsec Profile

[Table 102 on page 420](#) describes the MACsec Profile settings. Required settings are indicated by a red asterisk (*) that appears next to the field label in the user interface.

Table 102: MACsec Profile Settings

Field	Action
Profile Name	Type the name of the profile.
Description	Type a description of the profile.
Connection Association Name	Type the name for the MACsec connectivity association.
MACsec Mode	Select the mode using which you can enable MACsec on the device. The available modes are static secure association key (static-SAK) security mode or static connectivity association key (static-CAK) security mode.
CAK Settings	If you want to enable MACsec by using the CAK mode, configure the CAK settings specified in Table 103 on page 420 .
SAK Settings	If you want to enable MACsec by using the SAK mode, configure the SAK settings specified in Table 104 on page 422 for the inbound and outbound secure channels.

Table 103: CAK Settings

Field	Description
Connectivity Association Key Name	Type a name for the connectivity association key that you want to use for enabling MACsec.
Connectivity Association Key	Specify the key to exchange with the other end of the link on the secure channel. You must use a hexadecimal string of 32 digits.

Table 103: CAK Settings (*continued*)

Field	Description
Confirm Connectivity Association Key	Specify the connectivity association key again. If there is a mismatch (between the connectivity association keys), an error message is shown.
Enable Include Secure Channel Identifier	Enable Include Secure Channel Identifier tagging on a device that is enabling MACsec on an Ethernet link connecting to an Junos Fusion Enterprise device.
Key Server Priority	Specify the MACsec Key Agreement (MKA) server election priority number. You can specify a value between 0 and 255. The lower the number, the higher the priority.
Transmit Interval (milli sec)	<p>Specify the transmit interval for MACsec Key Agreement (MKA) protocol data units (PDUs). The MKA transmit interval setting sets the frequency for how often the MKA PDU is sent to the directly connected device to maintain MACsec on a point-to-point Ethernet link. A lower interval increases bandwidth overhead on the link; a higher interval optimizes the MKA protocol data unit exchange process.</p> <p>The default transmit interval is 2000 milliseconds</p>
Disable Encryption	Select this option if you want to disable the MACsec encryption for a connectivity association that has MACsec already enabled on it.
Offset	<p>Specify the offset 0, 30, or 50 for all the packets traversing the link. The default offset is 0. All traffic in the connectivity association is encrypted when encryption is enabled and an offset is not set.</p> <p>When the offset is set to 30, the IPv4 header and the TCP/UDP header are unencrypted while encrypting the rest of the traffic.</p> <p>When the offset is set to 50, the IPv6 header and the TCP/UDP header are unencrypted while encrypting the rest of the traffic.</p> <p>You would typically forward traffic with the first 30 or 50 octets unencrypted if a feature needed to see the data in the octets to perform a function, but you otherwise prefer to encrypt the remaining data in the frames traversing the link. Load balancing features, in particular, typically need to see the IP and TCP/UDP headers in the first 30 or 50 octets to properly load balance traffic.</p>
Replay Window Size	<p>Specify the size of the replay protection window.</p> <p>NOTE: When this variable is set to 0, all packets that arrive out-of-order are dropped.</p>

Table 103: CAK Settings (*continued*)

Field	Description
Exclude Protocols	<p>Specify the name of the protocol that should not be MACsec-secured. Options include:</p> <ul style="list-style-type: none"> • cdp—Cisco Discovery Protocol. • lacp—Link Aggregation Control Protocol. • lldp—Link Level Discovery Protocol.

Table 104: SAK Settings

Field	Description
Secure Channel name	Type a name for the secure channel.
MAC address	Specify a MAC address on which you want to enable MACsec using static secure association key (SAK) security mode. The mac-address variables must match on the sending and receiving ends of a link to enable MACsec using static SAK security mode.
Port	<p>Specify the port ID number in a secure channel when enabling MACsec using static secure association key (SAK) security mode. The port IDs must match on a sending and receiving secure channel on each side of a link to enable MACsec.</p> <p>After the port numbers match, MACsec is enabled for all traffic on the connection.</p>
Enable Encryption	<p>Select this option if you want to Enable MACsec encryption within an outbound secure channel.</p> <p>NOTE: You can enable MACsec without enabling encryption. If a connectivity association with an outbound secure channel that has not enabled MACsec encryption is associated with an interface, traffic is forwarded across the Ethernet link in clear text. You are, therefore, able to view this unencrypted traffic when you are monitoring the link.</p>
Offset	<p>Specify the number of octets in an Ethernet frame that you want to send in unencrypted plain-text when encryption is enabled for MACsec.</p> <p>Setting the offset to 30 allows a feature to see the IPv4 header and the TCP/UDP header while encrypting the remaining traffic. Setting the offset to 50 allows a feature to see the IPv6 header and the TCP/UDP header while encrypting the remaining traffic.</p>
Secure Association	<p>Specify the secure association keys corresponding to the secure association number. The key string is a 32-digit hexadecimal number.</p> <p>Re-enter the secure association key for every secure association number. If there is a mismatch between the connectivity association key and their respective confirmation keys, an error message is shown.</p>

What to Do Next

After you create the MACsec profile, you must assign the profile to the Junos Fusion Enterprise satellite device by using the Manage MacSec Profile page and then deploy the Device profile by using the **Deploy** mode.

To assign a MACsec Settings profile to a device, see [“Assigning the MACsec Profiles” on page 423](#). For information about deploying the configurations, see [“Deploying Configuration to Devices” on page 580](#).

NOTE: You can assign the MACsec profile to the extended ports on Junos Fusion Enterprise Aggregation Device.

In the CAK mode, if you change the connection association key name of a deployed MACsec profile, you must re-configure the connectivity association key and the confirmation key for that profile. Similarly, in the SAK mode, if you change the inbound or outbound channel names of the deployed MACSec profiles, you must re-configure the key and the confirmation key for that profile.

RELATED DOCUMENTATION

| [Media Access Control Security Overview | 417](#)


Assigning the MACsec Profiles

IN THIS SECTION

- [Assigning a MACsec Profile to a Device | 423](#)
- [Editing the MACsec Profile Assignments | 424](#)

Assigning a MACsec Profile to a Device

You can assign a MACsec profile to extended ports of satellite devices only and not to the native ports on the aggregation devices. Therefore, the tree view shows only the satellite devices (standalone and that are part of the cluster) to which you can assign a profile.

1. Click  in the Network Director banner.
2. Under Select View, select one of the following views: **Logical View**, **Location View**, **Device View** or **Custom Group**.

TIP: Do not select **Datacenter View** or **Topology View**.

3. In the Tasks pane, select **Wired > Profiles > MACsec**.

The Manage MACsec Profile page is displayed.

4. Select the MACsec profile that you want to assign and then click **Assign**.

The Assign MACsec profile page appears displaying a list of Junos Fusion Enterprise devices managed by Network Director.

Editing the MACsec Profile Assignments

Use the Edit Assignments page to change MACsec profile assignments. To edit an existing assignment:

1. Select a profile from the **Manage MACsec Profile Settings** page and click **Edit Assignment**.

The Edit Assignments page for the selected device appears.

2. Expand the **Devices** cabinet and make the desired change from the **Operation** column of the table.

3. Click **Apply** once you are done with the changes.

The Manage MACsec Profile page is displayed.

RELATED DOCUMENTATION

[Media Access Control Security Overview | 417](#)

[Configuring and Managing MACsec Profiles | 418](#)

Configuring Link Aggregation Groups (LAGs)

IN THIS CHAPTER

- [Understanding Link Aggregation | 425](#)
- [Managing and Creating a Link Aggregation Group | 426](#)
- [Understanding Multichassis Link Aggregation | 431](#)
- [Creating and Managing Multichassis Link Aggregation Groups \(MC-LAGs\) | 432](#)
- [Creating and Managing ESI Link Aggregation Groups \(ESI-LAGs\) | 448](#)

Understanding Link Aggregation

You can combine multiple physical Ethernet ports to form a logical point-to-point link, known as a link aggregation group (LAG) or bundle. A LAG provides more bandwidth than a single Ethernet link can provide. Additionally, link aggregation provides network redundancy by load-balancing traffic across all available links. If one of the links fails, the system automatically load-balances traffic across all remaining links. In a Virtual Chassis, LAGs can be used to load-balance network traffic between member switches.

The maximum number of interfaces that can be grouped into a LAG and the maximum number of LAGs supported on a switch varies according to the switch model and the version of and the version of Juniper Networks Junos operating system (Junos OS) that is running on that switch. [Table 105 on page 425](#) lists the maximum number of interfaces per LAG and the maximum number of LAGs that are supported on EX Series switches running Junos OS Release 19.1. If your switch is running a different version of Junos OS, refer to the device specific documentation, [EX Series Ethernet Switches](#), before implementing LAG in your network.

Table 105: Maximum Interfaces per LAG and Maximum LAGs per Switch

Switch	Maximum Interfaces per LAG	Maximum LAGs
EX2300	8	128
EX3400	16	128

RELATED DOCUMENTATION

[Managing and Creating a Link Aggregation Group | 426](#)

[Network Director Documentation home page](#)

Managing and Creating a Link Aggregation Group

IN THIS SECTION

- [Link Aggregation Group Options | 427](#)
- [Creating a Link Aggregation Group | 428](#)
- [Managing ICCP Settings | 430](#)
- [What To Do Next | 431](#)

IEEE 802.3ad link aggregation enables you to group Ethernet interfaces to form a single link layer interface, also known as a link aggregation group (LAG) or bundle.

Aggregating multiple links between physical interfaces creates a single logical point-to-point trunk link or a LAG. Link Aggregation Control Protocol (LACP), a component of IEEE 802.3ad, provides additional functionality for LAGs.

LACP ensures that both ends of the Ethernet link are functional and are members of the aggregation group before the link is added to the LAG. If you use LACP, make sure that LACP is enabled at both the local and remote ends of the link. When LACP is configured, it detects misconfigurations on the local end or the remote end of the link. Thus, LACP can help to prevent communication failure. When LACP is not enabled, a local LAG might attempt to transmit packets to a remote single interface, which causes the communication to fail. However, when LACP is enabled, a local LAG cannot transmit packets unless a LAG with LACP is also configured on the remote end of the link.

The maximum number of interfaces that can be grouped into a LAG and the maximum number of LAGs supported on a switch varies according to the switch model and the version of Juniper Networks Junos operating system (Junos OS) that is running on that switch. Be aware of the maximum number of interfaces per LAG and the maximum number of LAGs that are supported on your switches by referring to your device specific documentation before implementing LAG in your network.

NOTE: You only see the Manage Lag option under Device Management when a qualified switch is selected in the View Pane.

When creating LAGs, follow these guidelines:

- You must configure the LAG on both sides of the link.
- You must set the interfaces on either side of the link to the same speed.
- You can configure and apply firewall filters on a LAG.

NOTE: You only see the Manage Lag option under Device Management when a qualified switch is selected in the View Pane.

NOTE: MC-LAG, or Multi-Chassis Link Aggregation Group, is a type of LAG with constituent ports that terminate on separate chassis, thereby providing node-level redundancy. Unlike link aggregation in general, MC-LAG is not covered under IEEE 802.1AX-2008. Its implementation varies by vendor. For directions to create an MC-LAG, see [“Creating and Managing Multichassis Link Aggregation Groups \(MC-LAGs\)” on page 432](#).

This topic includes:

Link Aggregation Group Options

From the Manage LAG page, you can:

- Create a new Link Aggregation by clicking **Create**. The Create Link Aggregation window opens—for directions, see [“Creating a Link Aggregation Group” on page 428](#).
- Modify an existing Link Aggregation by selecting it and clicking **Edit**. The Modify Link Aggregation window opens. You can modify all the fields in the Modify Link Aggregation window, except the Interface Name field.
- Delete a Link Aggregation Group by selecting it and clicking **Delete**.
- Manage ICCP settings for the selected device by clicking **Manage ICCP Settings**. See [“Managing ICCP Settings” on page 430](#) for more information.

[Table 106 on page 428](#) describes the information provided about the link aggregation configurations on the LACP (Link Aggregation Control Protocol) Configuration page. This page lists all link aggregation groups defined on the selected device.

Table 106: LACP (Link Aggregation Control Protocol) Configuration Fields

Field	Description
Logical Interface Name	Name given to the aggregated interface when the LAG was created.
Member Interfaces	Names of individual member interfaces.
LACP Mode	<p>Mode in which LACP packets are exchanged between the interfaces.</p> <p>The possible modes are:</p> <ul style="list-style-type: none"> • Active—Indicates that the interface initiates transmission of LACP packets • Passive—Indicates that the interface responds only to LACP packets.
Description	<p>The description for the LAG.</p> <p>TIP: If you cannot view the entire description, you can resize the Description column by clicking the column border in the heading and dragging it.</p>
Deployment State	<p>The deployment state of the link aggregation. Deployment state can be:</p> <ul style="list-style-type: none"> • Pending Deployment—Indicates that the LAG is not yet deployed on the device. • Deployed—Indicates that the LAG is deployed on the device. • Pending Removal—Indicates that the LAG is deleted.
Creation Time	Date and time when this profile was created.
Update Time	Date and time when this profile was last modified.
User Name	The username of the user who created or modified the profile.

TIP: All columns might not be displayed. To show or hide fields in the LACP (Link Aggregation Control Protocol) Configuration table, click the DOWN arrow on the field header, select **Columns**, and select or clear the check box adjacent to the field that you want to show or hide.

Creating a Link Aggregation Group


You can create one or more LAGs in Network Director. The number of interfaces that can be grouped into a LAG and the total number of LAGs supported on a switch varies according to switch model.

TIP: You can also create one or more MC-LAGs for Virtual Chassis—see [“Creating and Managing Multichassis Link Aggregation Groups \(MC-LAGs\)”](#) on page 432.

To create a link aggregation group:

1. In the View pane, select a switch for link aggregation.

NOTE: The Manage LAG task is only available when a qualified switch is selected in the View pane.

2. Click  in the Network Director banner.
3. Select **Wired > Manage LAG** in the Tasks pane.
The Manage LAG page opens.
4. Click **Create**.
The Create Link Aggregation window opens.
5. Use the up and down arrows to select an AE Name for the aggregation interface. The interface name begins with *ae* followed by an interface number.
6. Select the mode in which LACP packets are to be exchanged between interfaces, either **Active** or **Passive**.
 - **Active**—Indicates that the interface initiates transmission of LACP packets
 - **Passive**—Indicates that the interface responds only to LACP packets.
7. Enter a description for the link aggregation.
8. Configure up to eight available interfaces on the LAG. Select one or more interfaces from the Available list and then click the RIGHT arrow to move them to the Selected list.

NOTE: The Available interfaces list displays only those interfaces that are not part of any link aggregation.

9. If the device is capable of using MC-LAGs, an MC-LAGs section also appears in the Create Link Aggregation window. For information about MC-LAG configuration, see [“MC-LAG Settings” on page 432](#).

10. Click **OK** to save the link aggregation configuration.

A message confirms that the link aggregation is created successfully and ready to be deployed to a device. If the configuration contains an error, the message instead indicates the error.

11. Click **OK** to close the information message.

The LAG appears in the Manage LAG list.

Managing ICCP Settings

When a QFX Series device is the selected scope, you can use the ICCP LAG Settings window to manage ICCP on the selected device. [Table 107 on page 430](#) describes the fields in this window.

Table 107: ICCP Settings

Field	Description
Disable (Delete) ICCP Settings	Disable ICCP on the device.
AE Name	Select the aggregated Ethernet interface to use for the ICCP connection.
Local IP	Configure the local IP address to be used by all switches hosting the MC-LAG.
Peer IP	Configure the IP address of the ICCP peer.
VLAN	Enter the name of the VLAN to use for the ICCP connection.
VLAN ID	Enter the ID of the VLAN to use for the ICCP connection.
Liveness detection min receive interval	Configure the minimum interval at which the switch must receive a reply from the other switch with which it has established a Bidirectional Forwarding Detection (BFD) session.
Liveness detection min transmit interval	Configure the minimum transmit interval during which a switch must receive a reply from a switch with which it has established a BFD session.
Liveness detection backup peer IP	Configure the IP address of the liveness detection backup.

Table 107: ICCP Settings (*continued*)

Field	Description
Session establish hold time	Configure the time during which an ICCP connection must succeed between the switches hosting the MC-LAG. Configured session establishment hold time results in faster ICCP connection establishment. The recommended value is 50 seconds.

What To Do Next

The configuration changes that you make in the Build mode are not deployed to devices automatically. After you create a link aggregation group, you must manually deploy the changes to the switches in Deploy mode. For details, see [“Deploying Configuration to Devices” on page 580](#).

TIP: Even though link aggregation configuration is not contained within a profile, you can view the link aggregation groups assigned to a switch by using the View Assigned Profiles task in Build mode.

RELATED DOCUMENTATION

[Understanding Link Aggregation | 425](#)

[Viewing Profiles Assigned to a Device | 558](#)

[Creating and Managing Multichassis Link Aggregation Groups \(MC-LAGs\) | 432](#)

[Network Director Documentation home page](#)

Understanding Multichassis Link Aggregation

Layer 2 networks are increasing in scale mainly because of technologies such as virtualization. Protocol and control mechanisms that limit the disastrous effects of a topology loop in the network are necessary. Spanning Tree Protocol (STP) is the primary solution to this problem because it provides a loop-free Layer 2 environment. STP has gone through a number of enhancements and extensions, and although it scales to very large network environments, it still provides only one active path from one device to another, regardless of the number of actual connections existing in the network. Although STP is a robust and scalable solution to redundancy in a Layer 2 network, the single logical link creates two problems: At least half of the available system bandwidth is off-limits to data traffic, and network topology changes occur.

The Rapid Spanning Tree Protocol (RSTP) reduces the overhead of the rediscovery process and allows a Layer 2 network to reconverge faster, but the delay is still high.

Link aggregation (IEEE 802.3ad) solves some of these problems by enabling users to use more than one link connection between switches. All physical connections are considered one logical connection. The problem with standard link aggregation is that the connections are point to point.

Multichassis link aggregation groups (MC-LAGs) enable a client device to form a logical LAG interface between two MC-LAG peers (QFX5100 and QFX10002 devices). An MC-LAG provides redundancy and load balancing between the two MC-LAG peers, multihoming support, and a loop-free Layer 2 network without running STP.

On one end of an MC-LAG, there is an MC-LAG client device, such as a server, that has one or more physical links in a link aggregation group (LAG). This client device does not need to have an MC-LAG configured. On the other side of the MC-LAG, there are two MC-LAG peers. Each of the MC-LAG peers has one or more physical links connected to a single client device.

The MC-LAG peers use Interchassis Control Protocol (ICCP) to exchange control information and coordinate with each other to ensure that data traffic is forwarded properly.

Link Aggregation Control Protocol (LACP) is a subcomponent of the IEEE 802.3ad standard. LACP is used to discover multiple links from a client device connected to an MC-LAG peer. LACP must be configured on all member links for an MC-LAG to work correctly.

RELATED DOCUMENTATION

[Creating and Managing Multichassis Link Aggregation Groups \(MC-LAGs\) | 432](#)

[Network Director Documentation home page](#)

Creating and Managing Multichassis Link Aggregation Groups (MC-LAGs)

IN THIS SECTION

- [Accessing the MC-LAG Page | 433](#)
- [Creating an MC-LAG | 434](#)
- [MC-LAG Automation Parameters | 440](#)
- [Editing an MC-LAG | 441](#)
- [Deleting an MC-LAG | 447](#)
- [Managing an MC-LAG Created Through CLI Mode | 447](#)

Multichassis link aggregation groups (MC-LAGs) enable a device to form a logical link aggregation group (LAG) interface between two switches. An MC-LAG provides redundancy and load balancing between the two switches, multihoming support, and a loop-free Layer 2 network without running Spanning Tree Protocol (STP).

MC-LAG peer switches use the Inter-Chassis Control Protocol (ICCP) to exchange control information and interchassis link (ICL) to exchange data.

At one end of an MC-LAG are the MC-LAG client devices, such as servers or switches, that have one or more physical links in a LAG. Client devices do not need to detect the MC-LAG. At the other end of the MC-LAG are two peer devices. Each of these switches has one or more physical links connected to a single client device. The switches coordinate with each other to ensure that data traffic is forwarded properly.

You can create MC-LAGs using QFX Series and EX9200 devices. However, both the peer devices must be the same type. Network Director can manage MC-LAG devices that are created and configured through the CLI mode also. If MC-LAG devices are configured through the CLI mode, ensure that LLDP is enabled on MC-LAG, ICCP LAG, ICL LAG, and client LAG links.

Supported devices in an MC-LAG:

- Peer devices: QFX5100, QFX10002, and EX9200 switches
- Client devices: All standalone and Virtual Chassis devices managed by Network Director except Virtual Chassis Fabric (VCF) members, IP Fabric devices and MX Series devices

For detailed steps on creating MC-LAGs using Network Director, follow the procedure given below or the steps shown in this video-based tutorial:



Video: [Configuring MC-LAG using Network Director](#)

This topic includes:

Accessing the MC-LAG Page

To access the MC-LAG page:

1. Click the Build mode icon  **Build** in the Network Director banner.
2. Select **Wired > Tasks > Manage MC-LAG** in the Tasks pane.

The Manage MC-LAG page opens, which displays the existing MC-LAG peers and enables you to create, edit, or delete an MC-LAG. In the Manage MC-LAG page, the peer devices for each MC-LAG that is created using Network Director or created using the CLI mode and discovered for management by Network Director, are listed. The Manage MC-LAG page displays the device name, device model, deployment status, and local IP address of the MC-LAG peer devices. If any peer device is not managed by Network Director, the MC-LAG Peer displays as *Unknown*. Click the peer devices of any MC-LAG

to view details of the MC-LAG, such as, descriptions of the peer devices, peer-to-peer link details, and client-to-peer link details.

Creating an MC-LAG

IN THIS SECTION

- [Selecting Peer Devices and Configuring Peer-to-Peer Link Settings | 434](#)
- [Selecting Client Devices and Configuring Client-to-Peer Link Settings | 436](#)
- [Saving MC-LAG Settings | 439](#)
- [Deploying MC-LAG Configuration | 439](#)

To create an MC-LAG:

1. Click **Create MC-LAG** in the Manage MC-LAG page.

The Create MC-LAG page opens. It displays two tabs—Peer Devices and Client Devices. By default, the Peer Devices tab is selected and displays in orange color.

On the left of the Create MC-LAG page, the Peer Devices tab lists QFX Series and EX9200 devices that are managed by Network Director. These are the available devices from which you can select the peer devices for the MC-LAG you create. On the right, a schematic diagram of the two peer devices PEER1 and PEER2 and a representation of the client devices as boxes are displayed.

Creating an MC-LAG involves four tasks:

Selecting Peer Devices and Configuring Peer-to-Peer Link Settings

To select the peer devices and configure peer-to-peer link settings:

1. From the list of devices in the Peer Devices tab in the Create MC-LAG page, select a device, and drag and drop it into one of the boxes labeled PEER1 or PEER2. 2.

After you drag and drop the first peer device, the list refilters and displays only devices that qualify to be the second peer.

For example, if you select a QFX10002 switch as one of the peer devices, then only QFX10002 switches are listed for you to select as the second peer device.

2. Select the second device from the refiltered list of peer devices and drag and drop it into the second peer box.

The Peer to Peer Link Settings window opens. The Client Devices tab is automatically enabled in the background in the Create MC-LAG page.

3. In the Peer to Peer Link Settings window, select **Combine Data and Control Links** if you want to combine the data and control links, that is, if you want a single link to act as both the control link and data link between the two ports that you selected. Network Director configures this link as an ICCP link.

If you want to have ICCP (control) and ICL (data) links separately between the peer devices, do not select this option.

NOTE: By default, Combine Data and Control Links is not selected. If you select this, you must specify the **VLAN Name** and **VLAN ID** in the respective fields in the Peer to Peer Link Settings window.

The physically connected ports of the peer devices are displayed in the Data and Control Link Ports* table in the Peer to Peer Link Settings window, if you have refreshed the topologies of the peer devices in the Topology View in Network Director. If the LLDP or topology information of the peer devices are not available for Network Director, port details are not displayed.

TIP: To refresh the topology, select Topology View in Views and then select Discovery-Topology > Refresh Topology in the Tasks pane. For the topology to refresh, LLDP must be enabled on the interfaces that are connected to the peer and client devices.

4. Click **Add Port**.

A new row is added to the table under Data and Control Link Ports*, where you must enter the port details for the peer devices.

5. From the drop-down menu for the PEER₁ device you selected, select a port to assign to the MC-LAG.
6. From the drop-down menu for the PEER₂ device you selected, select a port to assign to the MC-LAG.
7. Specify the type of link between the ports on the two peer devices by selecting **Data**, **Control**, or **Data & Control** from the Port Type list.

NOTE: If you have selected Combine Data and Control Links in Step 3, Data & Control is the default port type. If you have not selected Combine Data and Control Links in Step 3, the available options are Data and Control.

For the Data & Control port type, a single link between the peer devices acts as both the control and data link. If you select Data as the port type, then you must add a new pair of ports in the next row by clicking Add Ports, and then selecting Control as the port type for the control between the peer devices. If you select Control as the port type, then you must add a new pair of ports in the next row by clicking Add Ports, and then selecting Data as the port type for the control between the peer devices.

NOTE: You must specify at least one link between the peer devices.

8. Click **Update**.

A new row is added with the port details.

9. Enter the **IPv4 Address** and mask.

This IPv4 address is configured for the control link inet address and used as the local IP address for the ICCP. Network Director configures the peer IP addresses from this local IP address internally.

10. Click **OK**.

The Peer to Peer Link Settings window closes, and the Create MC-LAG page is displayed.

The Client Devices tab is selected by default. In the schematic diagram, the links that you configured between the peer devices changes to display in green, indicating that the links are successfully configured. The color does not indicate the operational status of the link.

Selecting Client Devices and Configuring Client-to-Peer Link Settings

To select a client device and configure client-to-peer link settings:

1. In the Client Devices tab on the Create MC-LAG page, select the device type for the client you want to be part of the MC-LAG by clicking an option from the drop-down menu in the **Type** field. Options available are: Switches, Bare Metal Servers, and Hypervisors.

The list of client devices displays the devices (switches, hypervisors, or bare metal servers) depending on the type that you selected. By default, only switches are listed.

2. Select a device from the list of client devices, and drag and drop it into one of the boxes labeled *Drag & Drop Clients here to add*.

NOTE: If you select a Virtual Chassis switch, the client box shows a graphical representation of Virtual Chassis; if you select a bare metal server or hypervisor server, the client box shows the graphical representation of that respective type of server.

The Client to Peer Link Settings window opens.

3. Select **MC-AE Mode**. The modes available are Active-Active and Active-Standby.

NOTE: Only EX9200 and QFX10002 devices support both Active-Active and Active-Standby modes. The other devices support only the Active-Active mode.

- **Active-Active mode:** If the client-to-peer setting mode is set to Active-Active mode, all peer port links will be active in the MC-LAG. In this mode, MAC addresses discovered in one MC-LAG peer device is propagated to the other peer device. Traffic is load balanced, and convergence is faster.
- **Active-Standby mode:** If the client-to-peer setting mode is set to Active-Standby mode, only one of the MC-LAG peer devices is active at any given time. The other peer device is in backup, that is standby, mode.

The ports that are physically connected between the client and peer devices are displayed in **PEER?** 1 and **PEER?** 12 you have refreshed the topologies of the peer devices in the Topology View in Network Director. If the LLDP or topology information of the peer devices are not available for Network Director, the port details are not displayed.

TIP: To refresh the topology, select Topology View in Views and then select Discovery-Topology > Refresh Topology in the Tasks pane. For the topology to refresh, LLDP must be enabled on the ports that are connected to the peer and client devices.

4. Click **Add Port** to select the client and peer ports.

A new row is added to the table, where you must enter the port details for the peer and client devices.

5. Select the client port from the drop-down menu corresponding to the **Client Port**.

NOTE: If you selected Switches as the type of client device, then Client Port is a mandatory field. If you selected Bare Metal Servers or Hypervisors, then the drop-down menu does not display any client port, as Network Director does not enable you to configure VLANs or ports in the servers.

6. From the drop-down menu select the peer port you want to connect to the client port.

7. Click **Update**. A row is added that displays the client port and peer port.

NOTE: If you have selected Peer 1 Port and linked it to a client port first, then select Peer 2 Port and link it to a client port. Both Peer 1 Port and Peer 2 Port cannot be selected in one row. The client device must be connected to both peer devices.

8. In the Client to Peer VLANs* table in the Client to Peer Link Settings window, Network Director displays all the VLANs of the client. If the client has the same VLAN ID as that of a peer or the peers, Network Director automatically populates the Routed Interface Address and VRRP Attributes for those peers in the respective fields. If there are no VLANs displayed in the table, add a VLAN by clicking **Select VLAN** or **Add VLAN**. This VLAN is configured in the PEER 1 and PEER 2 devices to ensure connectivity and data flow between the peers. You can configure multiple clients.

To edit a VLAN, click the fields of the VLAN that you want to edit.

NOTE:

- If you want to select a VLAN other than the VLANs displayed in the Client to Peer Link Settings window, click **Select VLAN** and select the VLANs from the list that displays in the Choose VLAN Profile pop-up window.
- You can remove a VLAN that you have created, but not deployed, in the client device by selecting the VLAN and clicking **Remove VLAN**.
- Do not remove VLANs that are deployed in the devices.

Network Director Release 2.5 supports Layer 3 routing. To enable Layer 3 routing, configure the Routed Interface Address and VRRP Attributes by clicking the respective fields.

Select the IP type by clicking the arrow in the **IP Type** field. The available options are IPv4 and IPv6.

9. Enter the IP addresses and mask for the peer devices in the corresponding fields. The IP addresses must be the IP addresses of the integrated routing and bridging interface.
10. Enter the VRRP group ID in **Group ID** and enter the virtual IP address in **Virtual IP** to assign the virtual IP that is shared between each switch in the VRRP group.
11. Click **Update**.

To add a VLAN, click **Add VLAN**. A new row is created in the Client to Peer VLANs* table. Enter the VLAN ID and VLAN name in their corresponding fields, and perform Steps 8 through 11.

To remove a VLAN, select the VLAN, and click **Remove VLAN**.

12. Click **OK** to submit the settings that you entered in the Client to Peer Link Settings window and to close the window.

The Client to Peer Link Settings window closes.

Network Director configures different IP addresses on IRB interfaces on the MC-LAG peers and runs the VRRP on the IRB interfaces. The virtual IP address is the gateway IP address for the MC-LAG clients. To provide Layer 3 routing functions to downstream clients, the MC-LAG network peers must be configured to provide the same gateway address to the downstream clients.

Saving MC-LAG Settings

To save the MC-LAG settings that you configured:

1. Click **Save** in the Create MC-LAG page.

Network Director saves the MC-LAG settings and displays the message **MC-LAG save is successful and is ready to be deployed to the devices**.

2. Click **OK**.

The Manage MC-LAG page lists the MC-LAG that you created. By default, the Deployment State for the MC-LAG displays as Pending Deployment.

Deploying MC-LAG Configuration

To deploy a new or edited MC-LAG configuration:

1. In the **Deploy** mode, click **Configuration Deployment > Deploy Configuration Changes** in the Tasks pane.

The Devices with Pending Changes page opens, displaying devices that have pending configuration changes.

2. In the list in the Devices with Pending Changes page, select the devices that you configured as the peer and client devices of the MC-LAG.

NOTE: To view the deployment information for a device, select the device and click **View**. The Configuration window opens, which shows the CLI and XML view of the configuration that will be deployed in the device.

3. Click **Deploy Now** to deploy the configuration.

The Device Configuration window opens. The Deployment Status shows the status as **INPROGRESS** and changes to **SUCCESS** once the deployment is successfully completed.

MC-LAG Automation Parameters

Network Director configures a number of parameters internally and automates the creation or modification of MC-LAGs.

Table 108 on page 440 describes the parameters that are internally configured by Network Director.

Table 108: MC-LAG Automation Parameters

Parameter	Description
LAG	LAG is created for ICCP, ICL, and MC-AE in peer devices, and a LAG is created for client devices.
mc-ae-id	Specifies which MC-LAG the aggregated Ethernet interface belongs to.
redundancy-group (supported only in QFX10002 and EX9200 devices)	Used by ICCP to associate multiple chassis that perform similar redundancy functions. It is used to establish a communication channel so that applications running on the peer devices can exchange messages.
init-delay-time:240ms:	Specifies the delay in number of seconds to bring the MC-LAG interface back to the Up state when an MC-LAG peer is rebooted.
chassis-id 0 for Peer 1, and 1 for Peer 2	Used by LACP for calculating the port number of the MC-LAG physical member links. Each MC-LAG peer must have a unique chassis ID.
status-control Active for Peer 1, and Standby for Peer 2	Specifies whether this node becomes active or goes into standby mode when an ICL failure occurs; must be active on one node and standby on the other node.
LACP active	Configured in ICL LAG, ICCP LAG, MC-LAG and client switch LAG. LACP is used to discover multiple links from a client device connected to an MC-LAG peer. LACP must be configured on all member links for an MC-LAG to work correctly.
LACP system-id and admin-key	Configures the same LACP system ID and admin-key for the MC-LAG on each MC-LAG peer. This displays Peer1 and Peer2 as a single switch to the edge switch when negotiating LACP.

Table 108: MC-LAG Automation Parameters (*continued*)

Parameter	Description
LACP periodic fast	Configured on ICCP LAG, ICL LAG and MC-LAG. LACP fast periodic is achieved by configuring fast intervals (in seconds) for periodic transmission of LACP.
Hold time Up 100000 down 0 for interfaces used for MC LAG. Up 0 down 2000 for interfaces used for ICL LAG.	Specifies the hold-time value to use to damp interface transitions. When an interface goes down, it is not broadcast to the rest of the system till it remains down for the hold-time period. Similarly, an interface is not broadcast as being Up till it remains up for the hold-time period.
multi-chassis-protection	Specifies the peer's ICCP IP address and the ICL link used for protection if the MC-AE interface goes down.
session-establishment-hold-time 300	Establishes ICCP connection quickly.
backup-liveness-detection: management IP of peer device	Is invoked when the ICCP link goes down. With backup liveness detection enabled, the MC-LAG peers establish an out-of-band channel through the management network in addition to the ICCP channel.
liveness-detection minimum-receive-interval 500, multiplier 3, transmit-interval 500	Determines whether a peer is up or down by exchanging keepalive messages over the management link between the two ICCP peers.
RSTP	Is enabled on peer devices MC-LAG and switch client LAG in point to point mode . If client is a server, then it enables bpdu-block-on-edge and edge on MC-LAG peer devices. Bridge-priority is set to 0 on both the peer devices.
ARP, MAC, arp-l2-validate, l2-interface ICL LAG on IRB	Provides IRB-to-IRB connectivity across the ICL. Using the VRRP over IRB method to enable Layer 3 functionality, it configures static ARP entries through the ICL for the IRB interface of the remote MC-LAG peer, which enables routing protocols to run over the IRB interfaces.

Editing an MC-LAG

IN THIS SECTION

- [Managing Peer Devices and Peer-to-Peer Link Settings | 442](#)
- [Managing Client Devices and Client-to-Peer Link Settings | 443](#)

In the Manage MC-LAG page, you can add, edit or delete peer ports, edit existing peer-to-peer link settings, add client, remove client, and edit client-to-peer link settings. You cannot add or delete peer devices if both the peers are part of MC-LAG.

1. Click **Edit** corresponding to the MC-LAG peers that you want to modify, in the Manage MC-LAG page.

The Edit MC-LAG page opens. It displays two tabs—Peer Devices and Client Devices. If both the peer devices of the MC-LAG are already configured as part of the MC-LAG, the Client Devices tab is selected, and it displays in orange color. On the left of the Edit MC-LAG page, a list of client devices are displayed.

If one of the peer devices is *Unknown*, the Peer Devices tab is selected, and it displays in orange color. On the left of the Edit MC-LAG page, a list of peer devices, that are of the same type and ELS capability as of the discovered peer, are displayed.

On the right of the Edit MC-LAG page, a schematic diagram of the existing two peer devices PEER¹ 1, PEER² 2 and a representation of the client devices as boxes are displayed.

Managing Peer Devices and Peer-to-Peer Link Settings

To add, edit, or delete a peer port, or edit peer-to-peer link settings:

1. Click **Control Link** or **Data Link** that is displayed between PEER¹ 1 and PEER² 2 in the schematic diagram.

The Peer to Peer Link Settings window opens.

NOTE: The Combine Data and Control Links option is unavailable.

The peer ports that you already configured are displayed in the table Data and Control Links Ports*.

2. To add a port, click **Add Port**.

A new row is added to the table, where you must enter the port details for the peer devices.

3. From the drop-down menu for the PEER¹ 1 device, select a port to assign to the MC-LAG.
4. From the drop-down menu for the PEER² 2 device, select a port to assign to the MC-LAG.
5. Specify the type of link between the ports on the two peer devices by selecting **Data**, **Control**, or **Data & Control** from the Port Type list.

NOTE: If you have selected Combine Data and Control Links, Data & Control is the default port type. If you have not selected Combine Data and Control Links, the available options are Data and Control.

For the Data & Control port type, a single link between the peer devices act as both the control and data link. If you select Data as the port type, then you must add a new pair of ports in the next row by clicking Add Ports, and then selecting Control as the port type for the control between the peer devices. If you select Control as the port type, then you must add a new pair of ports in the next row by clicking Add Ports, and then selecting Data as the port type for the control between the peer devices.

NOTE: You must specify at least one link between the peer devices.

NOTE: To delete a peer port, select the port that you want to remove from the MC-LAG, and click **Remove Port**.

To edit a peer port, click the port and modify the port details.

6. Click **Update**.

7. If you want to edit the control link IPv4 address, edit the **IPv4 Address** and mask fields in the Control Link table.

This IPv4 address is configured for the control link inet address and is used as the local IP address for the ICCP. Network Director internally configures the peer IP addresses from this local IP address.

If the data and control links are combined, VLAN ID and VLAN name are displayed and can be edited here.

8. Click **OK**.

The Peer to Peer Link Settings window closes, and the Edit MC-LAG page is displayed.

Managing Client Devices and Client-to-Peer Link Settings

To add or remove client devices, and edit client-to-peer link settings:

1. Click Client Devices tab in the Edit MC-LAG page.

The Client Devices tab on the Create MC-LAG page lists switches that are managed by Network Director. On the right, a schematic diagram of the two peer devices **PEER1** and **PEER2** and a representation of the client devices as boxes are displayed.

2. Select the device type for the client you want to be part of the MC-LAG by clicking an option from the drop-down menu in the **Type** field. Options available are: Switches, Bare Metal Servers, and Hypervisors.

The list displays the devices (switches, hypervisors, or bare metal servers) depending on the type that you selected. By default, only switches are listed.

3. Select a device from the list of client devices, and drag and drop it into one of the boxes labeled as *Drag & Drop Clients here to add*.

NOTE: If you select a Virtual Chassis switch, the client box shows a graphical representation of Virtual Chassis; if you select a bare metal server or hypervisor server, the client box shows the graphical representation of that respective type of server.

To delete a client device from an MC-LAG configuration, click the **x** mark on the client device in the carousel. The client device is removed from the carousel and the Deployment State changes to Pending Deployment in the Manage MC-LAG page.

The Client to Peer Link Settings window opens.

4. Select the **MC-AE Mode**. The available options are: Active-Active and Active-Standby.

NOTE:

- The MC-AE mode is enabled if you are adding a client device.
- Only EX9200 and QFX10002 devices support both Active-Active and Active-Standby modes. The other devices support only the Active-Active mode.
- Active-Active mode: If the client-to-peer setting mode is set to Active-Active mode, all peer port links will be active in the MC-LAG. In this mode, MAC addresses discovered in one MC-LAG peer device is propagated to the other peer device. Traffic is load balanced, and convergence is faster.
- Active-Standby mode: If the client-to-peer setting mode is set to Active-Standby mode, only one of the MC-LAG peer devices is active at any given time. The other peer device is in backup, that is standby, mode.

The ports that are physically connected between the client and peer devices are displayed in PEER? 1 and PEER? 2 you have refreshed the topologies of the peer devices in the Topology View in Network Director. If the LLDP or topology information of the peer devices are not available for Network Director, the port details are not displayed.

TIP: To refresh the topology, select Topology View in Views and then select Discovery-Topology > Refresh Topology in the Tasks pane. For the topology to refresh, LLDP must be enabled on the ports that are connected to the peer and client devices.

5. Click **Add Port** to select the client and peer ports.

A new row is added to the table, where you must enter the port details for the peer and client devices.

6. Select the client port from the drop-down menu corresponding to the **Client Port**.

NOTE: If you selected Switches as the type of client device, then Client Port is a mandatory field. If you selected Bare Metal Servers or Hypervisors, then the drop-down menu does not display any client port, as Network Director does not enable you to configure VLANs or ports in the servers.

7. Select **Peer 1 Port** or **Peer 2 Port** from the drop-down list in the corresponding fields.

NOTE: The client port is displayed only if you have selected Switches as the device type for the client device.

8. Click **Update**.

9. To add a new peer port and link it to a client port:

Click **Add Port**.

A blank row is added in the Client to Peer Ports table.

10. Select the client port by clicking the drop-down menu corresponding to the **Client Port**.

11. Select **Peer 1 Port** or **Peer 2 Port**, from the drop-down menu in the corresponding fields.

NOTE: If you have selected Peer 1 Port and linked it to a client port earlier, then select Peer 2 Port and link it to the a client port. Both Peer 1 Port and Peer 2 Port cannot be selected in one row.

12. Click **Update**.

13. In the Client to Peer VLANs* table, in the Client to Peer Link Settings window, Network Director displays all the VLANs of the client. If the client has the same VLAN ID as that of a peer or the peers, Network Director automatically populates the Routed Interface Address and VRRP Attributes for those peers in the respective fields. If there are no VLANs displayed in the table, add a VLAN by clicking Select VLAN or Add VLAN. This VLAN is configured in the Peer 1 and Peer 2 devices to ensure connectivity and data flow between the peers. You can configure multiple clients.

NOTE:

- If you want to select a VLAN other than the VLANs displayed in the Client to Peer Link Settings window, click **Select VLAN** and select the VLANs from the list that displays in the Choose VLAN Profile pop-up window.
- You can remove a VLAN that you have created, but not deployed, in the client device by selecting the VLAN and clicking **Remove VLAN**.
- Do not remove VLANs that are deployed in the devices.

Network Director Release 2.5 supports Layer 3 routing. To enable Layer 3 routing, configure the Routed Interface Address and VRRP Attributes in the respective fields.

Select the IP type by clicking the arrow in the **IP Type** field. The available options are IPv4 and IPv6.

Enter the IP addresses and mask for the peer devices in the corresponding fields.

NOTE: While editing an existing client device link settings, you cannot edit the VLAN Name, VLAN ID, Routed interface Address, and VRRP Attributes if they are already configured. If they are not configured, you can add Routed interface Address and VRRP Attributes.

14. Enter the VRRP group ID in **Group ID** and enter the virtual IP address in **Virtual IP** to assign the virtual IP address that is shared between each switch in the VRRP group.

15. Click **Update**.

To add a VLAN, click **Add VLAN**. A new row is created. Enter the VLAN ID and VLAN name in their corresponding fields, and perform Steps 13 through 15.

To remove a VLAN, select the VLAN and click **Remove VLAN**.

16. Click **OK** to submit the settings that you entered in the Client to Peer Link Settings window.

The Client to Peer Link Settings window closes.

17. Click **Save** in the Manage MC-LAG page.

Network Director saves the MC-LAG settings and displays the message **MC-LAG save is successful and is ready to be deployed to the devices**.


18. Click **OK**.

The Manage MC-LAG page lists the newly created MC-LAG. By default, the Deployment State for the newly created MC-LAG displays as Pending Deployment.

To deploy the edited MC-LAG, see [“Deploying MC-LAG Configuration” on page 439](#)

Deleting an MC-LAG

To delete MC-LAG:

1. Click the Build mode icon  in the Network Director banner.
Click the Build mode icon in the Network Director banner.
2. Select **Wired > Tasks > Manage MC-LAG** in the Tasks pane.
The Manage MC-LAG page opens, displaying the existing MC-LAG peers and enables you to delete MC-LAGs. The MC-LAGs displayed can be MC-LAGs that are created using Network Director or through the CLI mode.
3. Click **Delete** for the corresponding MC-LAG Peers that you want to delete, in the Manage MC-LAG page.

NOTE: If you delete an MC-LAG, Network Director removes the MC-LAG configuration settings from the peer devices and also deletes the LAG configuration from the client devices. The Deployment State changes to Pending Removal if the MC-LAG is already deployed. If it is not deployed, that is, if it is Pending Deployment, then the MC-LAG is removed from the Manage MC-LAG page.

Managing an MC-LAG Created Through CLI Mode

IN THIS SECTION

- [MC-LAG Peer Pairing | 447](#)
- [Mapping Client Devices to Peer Devices | 448](#)
- [Ports Mapping Between Peer-to-Peer and Client-to-Peer Devices | 448](#)

MC-LAG Peer Pairing

Once the MC-LAG devices are discovered by Network Director and Network Director successfully retrieves the MC-LAG configuration from the peer devices, Network Director pairs the MC-LAG peers based on the ICCP local IP address and peer IP address. For example, if Peer 1 is configured with ICCP local IP address 192.0.2.1 and Peer IP address 192.0.2.2, and Peer 2 is configured with ICCP local IP address 192.0.2.2

and Peer IP address 192.0.2.1, then based on the local IP address of Peer 1, Network Director searches for devices that have the same peer IP address as the local IP address. Because Peer 2 has the same peer IP address as the Peer 1 IP address, these two devices form MC-LAG peers. If in case, the local IP address is not found, then Network Director displays one of the peer devices in the MC-LAG pair as *Unknown* in the MC-LAG Manage page.

Mapping Client Devices to Peer Devices

If LLDP is enabled in the connected ports of the peer devices and client devices, after refreshing the topology, the Edit MC-LAG page displays the Network Director managed client switches connected to peer devices. If the client device is not managed by Network Director, or if the client device is not a switch (it is a bare metal server or a hypervisor), or if the topology information is not available for the devices, then the Edit MC-LAG page displays the client device as **Client_MC-AE ID (Unknown)**, where MC-AE ID specifies which MC-LAG the aggregated Ethernet port belongs to.

Ports Mapping Between Peer-to-Peer and Client-to-Peer Devices

On refreshing the topology, peer-to-peer link settings display the port mapping between the peer devices, and client-to-peer link settings display the port mapping between the client and peer devices.

RELATED DOCUMENTATION

[Understanding Link Aggregation | 425](#)

[Managing and Creating a Link Aggregation Group | 426](#)

[Viewing Profiles Assigned to a Device | 558](#)

[Network Director Documentation home page](#)

Creating and Managing ESI Link Aggregation Groups (ESI-LAGs)

IN THIS SECTION

- [Accessing the ESI-LAG Page | 449](#)
- [Creating an ESI-LAG | 449](#)
- [Editing an ESI-LAG | 454](#)
- [Deleting an ESI-LAG | 457](#)
- [ESI-LAG Automation Parameters | 458](#)

Ethernet Switch Identifier (ESI) refers to the set of Ethernet links that connect one or more access devices (called client devices) to a pair of core devices (called as peers) in a campus environment. ESI link aggregation groups (ESI-LAGs) enable one or more client devices to form a logical link aggregation group (LAG) interface with the peers. The peer should already be connected with each other before forming an ESI-LAG between them.

You can create ESI-LAGs by using EX9200 devices as the core devices.

NOTE: Network Director supports an ESI-LAG configuration only if the ESI-LAG is created by using Network Director.

Supported devices in an ESI-LAG:

- Peer devices in a core network: EX9200
- Client devices in an access network: EX4300, EX3400, EX4600, EX2300

For creating an ESI-LAG, follow the procedure described in this topic:

This topic includes:

Accessing the ESI-LAG Page

To access the ESI-LAG page:

1. Click the Build mode icon  in the Network Director banner.
2. Select **Wired** > **Tasks** > **Manage ESI-LAG** in the Tasks pane.

The Manage ESI-LAG page opens, which displays the existing ESI-LAG peers and enables you to create, edit, or delete an ESI-LAG. The Manage ESI-LAG page also displays the device name, device model, deployment status, and local IP address of the ESI-LAG peer devices. Click the peer devices of any ESI-LAG to view details such as, descriptions of the peer devices, peer-to-peer link details, and client-to-peer link details, of the ESI-LAG.

Creating an ESI-LAG

IN THIS SECTION

- [Selecting Peer Devices and Configuring Peer-to-Peer Link Settings | 450](#)
- [Selecting Client Devices and Configuring Client-to-Peer Link Settings | 452](#)

- [Saving ESI-LAG Settings | 453](#)
- [Deploying ESI-LAG Configuration | 454](#)

To create an ESI-LAG:

1. Click **Create ESI-LAG** on the Manage ESI-LAG page.

The Create ESI-LAG page opens. It displays two tabs—Peer Devices and Client Devices. By default, the Peer Devices tab is selected and displays in orange color.

On the left of the Create ESI-LAG page, the Peer Devices tab lists EX9200 devices that are managed by Network Director. These are the available devices from which you can select the peer devices for the ESI-LAG you create. On the right, a schematic diagram of the two peer devices PEER1 and PEER2 and boxes representing the client devices is displayed.

Creating an ESI-LAG involves the following tasks:

Selecting Peer Devices and Configuring Peer-to-Peer Link Settings

To select the peer devices and configure peer-to-peer link settings:

1. From the list of devices in the Peer Devices tab on the Create ESI-LAG page, select a device, and drag and drop it into the box labeled PEER1 or PEER2.

After you drag and drop the first peer device, the list of devices refilters and displays only devices that qualify to be the second peer.

For example, if you select an EX9200 device as one of the peer devices, then other EX9200 devices that are discovered by Network Director are listed for you to select as the second peer device.

2. Select the second device from the refiltered list of peer devices and drag and drop it into the second peer box.

The Peer to Peer Link Settings window opens. The Client Devices tab is automatically enabled in the background on the Create ESI-LAG page.

3. In the Peer to Peer Link Settings window, click **Add Port** to add ports of the peer devices to be used in the LAG.

A new row is added to the Peer to Peer Ports* table, where you must enter the port details for the peer devices.

4. From the drop-down menu for the Peer 1 device, select a port to assign to the ESI-LAG.

5. From the drop-down menu for the Peer 2 device, select a port to assign to the ESI-LAG.
6. In the Loop Back section, configure the loopback IPv4 address for Peer 1 and Peer 2 in the **Peer 1 IPv4 Address** and **Peer 2 IPv4 Address** text fields, respectively.

The loopback address ensures that the peer devices are reachable to management applications and other entities that want to communicate with the devices.

7. In the Logical Interface section,
 - Provide the logical IPv4 address for Peer 1 and Peer 2 devices in the **Peer 1 IPv4 Address** and **Peer 2 IPv4 Address** text fields, respectively.

The logical interfaces are created on the interfaces on which the ESI-LAG is to be configured.

- (Optional) In the **BGP Group Name** text field, edit the name for the BGP group to which the peer devices will be assigned.

When the ESI-LAG is deployed on your network, a BGP group with the name assigned is created. Other BGP-related parameters are autogenerated and assigned to the peer devices.

- (Optional) In the **Virtual Switch Instance** text field, edit the name of the virtual switch instance assigned to the peer devices.

When the ESI-LAG is deployed on your network, a virtual switch instance with the name assigned is created in the routing instance of the peer devices. You can configure only one routing instance by using Network Director.

- (Optional) In the **VRF Instance Name** text field, edit the name of the VRF instance assigned to the peer devices.

When the ESI-LAG is deployed on your network, a VRF instance by the name assigned is created and assigned to the peer devices.

- In the **Autonomous System Number** text field, enter the autonomous system number to which the peer devices are assigned as part of ESI-LAG.

8. In the Peer To Peer VLAN table:

- Click **Add VLAN** to add VLANs. This VLAN is configured between the peers to ensure connectivity and data flow between the peers.

When you click Add VLAN, enter the following values for the VLAN:

- **VLAN ID**
- **VLAN Name**
- **VNI ID**

- **IP Type, IP Address** of the peers and Mask for the routed interface address of the VLAN
- **Anycast Gateway** for the VLAN

9. Click **Update**.

A row is added that displays the VLAN associated with the peer devices.

10. (Optional) Click a VLAN and click **Remove VLAN** to remove any VLAN that you do not want to be part of the ESI-LAG configuration.

11. Click **OK**.

The Peer to Peer Link Settings window closes, and the Create ESI-LAG page appears.

The Client Devices tab is selected by default. In the schematic diagram, the links that you configured between the peer devices changes to green, indicating that the links are successfully configured. The color does not indicate the operational status of the link.

Selecting Client Devices and Configuring Client-to-Peer Link Settings

To select a client device and configure client-to-peer link settings:

1. In the Client Devices tab on the Create ESI-LAG page, select the client device.
2. Select a device from the list of client devices, and drag and drop it into one of the boxes labeled *Drag & Drop Clients here to add*.

The Client to Peer Link Settings window opens.

3. Click **Add Port** to select the client and peer ports.

A new row is added to the table, where you must enter the port details for the peer and client devices.

4. Select the client port from the **Client Port** drop-down list.
5. From the drop-down menu, select the peer port you want to connect to the client port.

NOTE: You can configure only one interface to connect the client to a peer.

6. Click **Update**. A row is added that displays the client port and peer port.

NOTE: If you have selected Peer 1 Port and linked it to a client port first, then select Peer 2 Port and link it to a client port. Both Peer 1 Port and Peer 2 Port cannot be selected in one row. The client device must be connected to both peer devices.

7. In the Client to Peer VLANs* table in the Client to Peer Link Settings window, click **Select VLAN** to assign the client to one or more VLANs.

The Choose VLANs pop-up window appears listing the VLANs to which the Peers are assigned.

8. Select one or more VLANs from the Choose VLANs pop-up window to which you want to assign the client.

9. Click **OK**.

The selected VLANs are added to the Client to Peer VLANs* table.

10. (Optional) Remove a VLAN that you have created, in the client device by selecting the VLAN and clicking **Remove VLAN**.

11. Click **Update**.

The client is assigned to the selected VLANs.

12. Click **OK** to submit the settings that you entered in the Client to Peer Link Settings window and close the window.

The Client to Peer Link Settings window closes.

Saving ESI-LAG Settings

To save the ESI-LAG settings that you configured:

1. Click **Save** on the Create ESI-LAG page.

Network Director saves the ESI-LAG settings and displays the message **ESI-LAG save is successful and is ready to be deployed to the devices**.

2. Click **OK**.

The Manage ESI-LAG page lists the ESI-LAG that you created. By default, the Deployment State for the ESI-LAG displays as Pending Deployment.

Deploying ESI-LAG Configuration

To deploy a new or edited ESI-LAG configuration:

1. In the **Deploy** mode, click **Configuration Deployment > Deploy Configuration Changes** in the Tasks pane.

The Devices with Pending Changes page opens, displaying devices that have pending configuration changes.

2. In the list on the Devices with Pending Changes page, select the devices that you configured as the peer and client devices of the ESI-LAG.

NOTE: To view the deployment information for a device, select the device and click **View**. The Configuration window opens, which shows the CLI and XML view of the configuration that will be deployed on the device.

3. Click **Deploy Now** to deploy the configuration.

The Device Configuration window opens. The Deployment Status shows the status as **INPROGRESS** and changes to **SUCCESS** once the deployment is successfully completed.

Editing an ESI-LAG

IN THIS SECTION

- [Managing Peer Devices and Peer-to-Peer Link Settings | 455](#)
- [Managing Client Devices and Client-to-Peer Link Settings | 456](#)

On the Manage ESI-LAG page, you can add, edit, or delete peer ports, edit existing peer-to-peer link settings, add client, remove client, and edit client-to-peer link settings. However, you cannot add or delete peer devices of the ESI-LAG.

NOTE: You cannot edit an ESI-LAG after it is configured and the devices are deployed on the network (that is, the state of the devices is DEPLOYED.)

1. On the Manage ESI-LAG page, click **Edit** corresponding to the ESI-LAG peers that you want to modify.

The Edit ESI-LAG page opens. It displays two tabs—Peer Devices and Client Devices. If both the peer devices of the ESI-LAG are already configured as part of the ESI-LAG configuration, the Client Devices tab is selected, and it displays in orange color. On the left of the Edit ESI-LAG page, a list of client devices are displayed.

If one of the peer devices is *Unknown*, the Peer Devices tab is selected, and it displays in orange color. On the left of the Edit ESI-LAG page, a list of peer devices, that are of the same type and ELS capability as of the discovered peer, are displayed.

On the right of the Edit ESI-LAG page, a schematic diagram of the existing two peer devices PEER? 1, PEER? 2 and a representation of the client devices as boxes are displayed.

Managing Peer Devices and Peer-to-Peer Link Settings

To add, edit, or delete a peer port, or edit peer-to-peer link settings:

1. Click **EVPN-VXLAN** link that is displayed between PEER? 1 and PEER? 2 in the schematic diagram.

The Peer to Peer Link Settings window opens.

The peer ports that you already configured are displayed in the Peer to Peer Ports table.

2. Do one of the following:

- To add a port, click **Add Port**.

A new row is added to the table, where you must enter the port details for the peer devices.

From the drop-down menu for the PEER? 1 device or PEER 2 device, select a port to assign to the ESI-LAG.

- To edit a peer port, click the port and edit the port.
- To delete a peer port, select the port that you want to remove from the ESI-LAG, and click **Remove Port**.

NOTE: You must specify at least one link between the peer devices.

3. Click **Update**.
4. (Optional) Edit the Peer 1 or Peer 2 loop back address, Peer 1 or Peer 2 logical interfaces. BGP Group Name, Virtual Instance Switch Name, VRF Instance Name, or Autonomous System Number.
5. (Optional) Edit the Peer To Peer VLANs as follows:
 - To add a VLAN, click **Add VLAN**.
A new row is added to the Peer To Peer VLANs table. Enter the values for VLAN ID, VLAN Name, VNI ID, IP type, IP addresses and mask for the integrated routing and bridging interface of the VLAN and the anycast gateway in the corresponding fields. Click **Update** to save the new VLAN in the table.
 - To edit a VLAN, click on the VLAN and edit one or more attributes of the VLAN—VLAN ID, VLAN Name, VNI ID, IP type, IP addresses and mask for the integrated routing and bridging interface of the VLAN and the anycast gateway.
 - To remove a VLAN, select the VLAN and click **Remove VLAN**.
The VLAN is removed from the table.

6. Click **OK**.

The Peer to Peer Link Settings window closes, and the Edit ESI-LAG page is displayed.

Managing Client Devices and Client-to-Peer Link Settings

To add or remove client devices, and edit client-to-peer link settings:

1. Click Client Devices tab on the Edit ESI-LAG page.

The Client Devices tab on the Create ESI-LAG page lists switches that are managed by Network Director. On the right, a schematic diagram of the two peer devices PEER1 PEER2 and a representation of the client devices as boxes is displayed.

2. Select a device from the list of client devices, and drag and drop it into one of the boxes labeled as *Drag & Drop Clients here to add*.

NOTE: To delete a client device from an ESI-LAG configuration, click the x mark on the client device in the carousel. The client device is removed from the carousel.

You cannot delete a client device if the client device is already deployed.

The Client to Peer Link Settings window opens.

3. In the Client to Peer Ports table, you can edit or delete the configured client and peer ports.

To remove the configured client and peer ports, click on a row and click **Remove Port**. The port is removed from the Client to Peer Ports table.

To add a port, click **Add Port** and select the client and peer ports. Click **Update** to save the port.

To edit a port, click on the port and edit the port. Click **Update** to save the edited port values.

NOTE: You can add only one interface connection between a client and a peer.

4. In the Client to Peer VLANs* table, in the Client to Peer Link Settings window, Network Director displays the VLAN configured in the ESI-LAG

In this table, you can edit the configured VLAN, add a new VLAN or remove a configured VLAN.

To add a new VLAN, click **Select VLAN** and select the VLANs from the list that displays in the Choose VLAN Profile pop-up window.

To edit a VLAN, click on the VLAN and edit the VLAN ID or VLAN Name.

You can remove a VLAN that you have created, but not deployed, in the client device by selecting the VLAN and clicking **Remove VLAN**.

5. Click **Update**.
6. Click **OK** to submit the settings that you entered in the Client to Peer Link Settings window.

The Client to Peer Link Settings window closes.

7. Click **Save** in the Manage ESI-LAG page.

Network Director saves the ESI-LAG settings and displays the message **ESI-LAG save is successful and is ready to be deployed to the devices**.


8. Click **OK**.

The Manage ESI-LAG page lists the newly created ESI-LAG. By default, the Deployment State for the edited ESI-LAG displays as Pending Deployment.

To deploy the edited ESI-LAG, see [“Deploying ESI-LAG Configuration” on page 454](#).

Deleting an ESI-LAG


To delete ESI-LAG:

1. Click the Build mode icon  **Build** in the Network Director banner.

2. Select **Wired > Tasks > Manage ESI-LAG** in the Tasks pane.

The Manage ESI-LAG page opens, displaying the configured ESI-LAG peers and enables you to delete ESI-LAGs.

3. Click **Delete** for the corresponding ESI-LAG Peers that you want to delete, in the Manage ESI-LAG page.



NOTE: If you delete an ESI-LAG, Network Director removes the ESI-LAG configuration settings from the peer devices and also deletes the LAG configuration from the client devices. The Deployment State changes to Pending Removal if the ESI-LAG is already deployed. If it is not deployed, that is, if it is Pending Deployment, then the ESI-LAG is removed from the Manage ESI-LAG page.

ESI-LAG Automation Parameters

Network Director configures a number of parameters internally and automates the creation or modification of ESI-LAGs.

[Table 109 on page 458](#) describes the parameters that are internally configured by Network Director.

Table 109: ESI-LAG Automation Parameters

Parameter	Description
LAG	Used to create a LAG between peer devices, and between the client and peer devices.
LACP active	Used to configure LACP in peer devices and the client device. LACP is used to discover multiple links from a client device connected to peers. LACP must be configured on all member links to work properly.
LACP periodic fast	Used to configure LACP periodic fast in Peer switches and client switch. LACP fast periodic is achieved by configuring fast intervals (in seconds) for periodic transmission of LACP.
Loopback Address	Used to configure loopback address. The loopback address ensures that the device provides an IP address to management applications as the device must always be available to hosts attempting to route packets to the device. Setting a loopback address ensures that the device can receive packets addressed to the loopback address as long as the device is reachable through any entry (ingress) interface.

Table 109: ESI-LAG Automation Parameters (continued)

Parameter	Description
ESI ID	Used to configure an Ethernet Segment Identifier (ESI) on a per-interface basis. All interfaces configured with the same ESI, on any devices within the same EVPN domain, appear as part of the same L2 segment or LAG.
ESI mode	Used to configure ESI all-active mode to enable Active-Active Multihoming in peers.
Policy options	Used to specify routing policy evpn-pplb for EVPN.
Routing option	Used to specify forwarding table with per-packet load balancing (PPLB) export policy for EVPN, autonomous system number and router-id.
Virtual Switch Configuration	Used to create a virtual switch routing instance and auto-generate related parameters.
vtep-source-interface	Used to specify the source interface for a Virtual Extensible LAN (VXLAN) tunnel and configure a logical interface unit 0 (lo0.0) on the loopback interface.
instance-type - (virtual-switch)	Used to provide support for Layer 2 bridging. This routing instance type is used to isolate a LAN segment with its Spanning Tree Protocol (STP) instance and to separate its VLAN identifier space.
route distinguisher id	Used to specify a route distinguisher for the routing instance.
vrf-target	Used to specify a virtual routing and forwarding (VRF) target community
Protocol	Used to enable the Ethernet VPN (EVPN) protocol.
VLAN-VxLAN Mapping	Used to map the VLAN and VxLAN configuration to the virtual switch
VPN routing and forwarding (VRF) instance	Used to create a VRF routing instance and auto-generate the VRF routing parameters.
instance-type (vrf)	Used to provide support for Layer 3 VPNs, where interface routes for each instance goes only into the corresponding forwarding table
vrf-target	Used to specify a virtual routing and forwarding (VRF) target community
route distinguisher id	Used to assign a route distinguisher to the routing instance automatically.
BGP Protocol	Used to enable BGP protocol on peer devices and auto-generates the BGP parameters.

Table 109: ESI-LAG Automation Parameters (*continued*)

Parameter	Description
BGP sessions	Used to auto-generate internal group type, set the loopback address as the local address and the peer loop back address as the neighbor address.
VPN family	Used to auto-generate inet, inet-vpn, and evpn signaling for BGP.
mutipath	Used to allow load sharing among multiple eBGP and multiple iBGP paths.
OSPF Protocol	Used to enable OSPF on peer devices and configur egaArea IP as 0.0.0.0 on loopback and LAG Interface.

RELATED DOCUMENTATION

[Understanding Link Aggregation | 425](#)
[Managing and Creating a Link Aggregation Group | 426](#)
[Viewing Profiles Assigned to a Device | 558](#)
[Network Director Documentation home page](#)

Creating and Managing Fabrics

IN THIS CHAPTER

- Understanding Junos Fusion | 461
- Understanding Junos Fusion Enterprise | 463
- Software Requirements for Junos Fusion | 465
- Creating and Managing Fusion Configuration Templates | 466
- Managing Fusion Fabrics | 480
- Creating and Managing Satellite Software Upgrade Groups | 486
- Understanding Layer 3 Fabrics | 489
- User Privileges Required for the DHCP and File Server While Using Zero Touch Provisioning | 490
- Managing Layer 3 Fabrics | 491
- Creating Layer 3 Fabrics | 493
- Editing Layer 3 Fabrics | 508
- Viewing Layer 3 Fabric Connectivity | 511
- Performing Layer 3 Fabric Connectivity Checks | 512
- Setting Up Virtual Chassis Fabrics | 513
- Managing Virtual Chassis Fabrics | 521

Understanding Junos Fusion

Junos Fusion technology, based on the IEEE 802.1BR standard, is a rich, open framework that makes networks highly versatile, extensible, and responsive in multivendor environments. With Junos Fusion technology, network administrators can reduce network complexity and operational expenses by collapsing underlying network elements into a single, logical point of management using QFX Series and EX Series switches running the Junos operating system.

Junos Fusion consists of two major components—*aggregation devices* and *satellite devices*.

Aggregation devices serve as the core of a Junos Fusion fabric and are responsible for almost all management tasks, including interface configuration for every satellite device interface in the topology. An aggregation device runs Junos OS for the entire Junos Fusion. The network-facing interfaces on the satellite

devices—extended ports—are configured from the aggregation device and support features that are supported on Junos OS running on the aggregation device. A Junos Fusion fabric can have one aggregation device (single-home Junos Fusion) or two aggregation devices (multihome or dual-home Junos Fusion).

Satellite devices form the access layer of a Junos Fusion fabric. These devices, which are connected through uplink ports to the aggregation devices, need not be individually managed as the control plane resides on the aggregation device.

Network Director supports Junos Fusion technologies — Junos Fusion Enterprise.

[Table 110 on page 462](#) lists the device models that are supported as aggregation and satellite devices for each Junos Fusion technology.

Table 110: Devices Supported in Junos Fusion

Junos Fusion Technology	Aggregation Device	Satellite Device
Junos Fusion Enterprise	EX9200	EX2300
		EX3400
		EX4300

Network Director provides a single pane of glass (SPOG) solution for managing Junos Fusion Enterprise fabrics, enabling network agility and reducing costs.

Setting up a Junos Fusion Enterprise using Network Director involves the following tasks:

1. Create a configuration template for the instance (Junos Fusion Enterprise). In a configuration template, you specify details such as the fusion topology, ports to be used, software image to be used for satellite devices, and so on.
2. Apply the template to one or more Junos Fusion systems. While applying a template, you select the aggregation devices, specify the software image to be used for the aggregation devices, the DHCP and file server details that Network Director uses to bring up the aggregation devices, and the ICL or ICCP port details for multi-home Junos Fusion.

Network Director creates the Fusion solution based on the template that you applied, and displays the Fusion instance in the Manage Fusion Fabric page.

RELATED DOCUMENTATION

Understanding Junos Fusion Enterprise	463
Creating and Managing Fusion Configuration Templates	466

Understanding Junos Fusion Enterprise

For enterprise networks, Junos Fusion Enterprise provides automated network configuration and simplifies scalability for medium to large enterprise networks with the Juniper Networks EX9200 line of Ethernet switches, EX4300, EX2300, and EX3400 switches. Junos Fusion Enterprise technology can be deployed across a building, or multiple buildings, to connect large numbers of devices in a fabric that can be managed as a single device.

Figure 20 on page 463 displays a typical Junos Fusion Enterprise topology.

In Junos Fusion Enterprise deployments, satellite devices do not need to be individually connected to aggregation devices. Up to 10 satellite devices can be interconnected through standard 10-Gigabit Ethernet or 40-Gigabit Ethernet interfaces to form a satellite cluster (as shown in Figure 20 on page 463), which in turn can be connected to the aggregation devices over a pair of fiber uplink ports.

Figure 20: Junos Fusion Enterprise Topology




Figure 20 on page 463 displays a multihome Junos Fusion Enterprise topology with two aggregation

devices—*Aggregation device 1* and *Aggregation device 2*.

Connections—Aggregation device 1 is connected to *Satellite device 1* and a satellite device in the satellite cluster. Satellite device 1 is connected to hosts 1 through 4. Aggregation device 2 is connected to Satellite device 1 and a satellite cluster comprising four satellite devices connected in a ring topology. In a satellite cluster, it is not necessary that all satellite devices are directly connected to an aggregation device. In this topology, we have two satellite devices in the cluster that are directly connected to Aggregation device 2 and one satellite device that is directly connected to Aggregation device 1. The satellite devices, in turn, are connected to each other using cluster ports—cluster port 1 (ClSP1) and cluster port 2 (ClSP2) of each device connects to similar ports on the neighboring satellite devices. One of the satellite device is connected to Aggregation Device 1 through uplink port UP5. The satellite devices in the cluster are connected to hosts 5 through 11.

Port Usage—Both the aggregation devices are connected using ICL and ICCP ports on both ends. ICL ports are used to forward data traffic, whereas ICCP ports are used to exchange control information. You can also choose to use a single link for both ICL and ICCP traffic.

Aggregation devices use cascade ports (CP1, CP2, CP3, CP4, and CP5) to connect to the satellite devices. Satellite device 1 uses uplink port UP1 to connect to the Aggregation Device 1 and UP4 to connect to Aggregation device 2. The satellite cluster uses uplink ports UP2 and UP3 to connect to Aggregation device 2 and UP5 to connect to Aggregation device 1. The satellite devices in the cluster are connected to each other using cluster ports ClSP1 and ClSP2. The satellite devices use extended ports (EP1 through EP11) to connect to the hosts.

To set up a Junos Fusion Enterprise similar to the topology shown in [Figure 20 on page 463](#) using Network Director, you must perform the following tasks:

1. Create a configuration template. While creating the template, you specify that this is a Junos Fusion for a multihome enterprise fabric; plan the chassis for the aggregation device; identify the ICL, ICCP, and cascade ports on the aggregation device and the cluster ports on the satellite devices.
2. Apply the template to a Junos Fusion fabric. The device to which you apply the template should:
 - Support Junos Fusion. For the list of supported devices see, [Juniper Networks Devices Supported by Junos Space Network Management Platform](#).
 - Have pre-configured with Multichassis Link Aggregation Groups (MC-LAGs) configuration (as MAC-LAG provides multihome support) and should be managed by Network Director.
 - Have the required DMI schema uploaded. For the list of supported DMI schema, see [Uploading DMI Schemas](#)

You can apply the template to a device that is already managed by Network Director or to a new unmanaged device, to make the device the aggregation device for the fusion fabric. When you apply the configuration template to one or more unmanaged devices, you specify the software image that must be installed on the aggregation devices and the Zero Touch Provisioning details for the aggregation devices. Whereas, when you apply the template to a managed device (device that is already managed by Network Director), you select the device that you want to convert to an aggregation device in your

fusion fabric and manually refresh the device topology. To refresh the topology, navigate to Topology View and click **Refresh Topology** under the Tasks menu. Network Director converts the device to a fusion fabric aggregation device and deploys all the necessary configurations on the device.

After the template is applied successfully to a Junos Fusion fabric, there are quite a few tasks that Network Director performs internally that makes building your Junos Fusion fabric simple and error-free. Network Director converts the device that you specified in the Apply Template workflow to an aggregation device and applies the port settings on the various ports that you specified in the configuration template. When an EX4300 device is connected to one of the configured cascade ports, a *link up* event is triggered. The link up event initiates a syslog message to Network Director and Network Director initiates a *topology refresh* job. Network Director then installs the appropriate satellite software on the satellite device and performs the necessary configurations. If a second satellite device is connected to the first satellite device to form a satellite cluster, or another satellite device is connected to the aggregation device, another link up event is triggered and the same steps are repeated. This process continues for all additional satellite devices that are connected to the aggregation device.

Network Director lists the fusion fabric in the Manage Fusion Fabrics page. You can see the details and status of the fabric in the Manage Fusion Fabrics page. You can also edit the fusion fabric, download the cabling plan, and view the fabric connectivity using the Manage Fusion Fabrics page.

RELATED DOCUMENTATION

[Understanding Junos Fusion | 461](#)

[Media Access Control Security Overview | 417](#)

Software Requirements for Junos Fusion

An aggregation device in a Junos Fusion always runs Junos OS software and is responsible for almost all management tasks, including configuring all network-facing ports—the *extended ports*—on all satellite devices in the Junos Fusion. The extended ports in a Junos Fusion, therefore, support all features that are supported by the Junos OS running on the aggregation device.

An aggregation device in a Junos Fusion runs the same Junos OS regardless of whether it is or is not part of a Junos Fusion. Hence, Junos OS is acquired, installed, and managed on an aggregation device in a Junos Fusion in the same manner that it is acquired, installed, and managed on a standalone device that is not part of a Junos Fusion.

The satellite devices in a Junos Fusion run satellite software that has the built-in intelligence to extend the feature set on Junos OS to the satellite device. The satellite software is a Linux-based operating system that enables the satellite devices to communicate with the aggregation device for control plane data while also passing network traffic. Satellite software is also known as satellite network operating system software.

Make sure that the devices that are to be converted as satellite devices run the requisite Junos OS version that Junos Fusion supports. For more details on supported Junos OS version, see *Network Director Release Notes, Release 3.2*.

All satellite devices in a Junos Fusion must run the satellite software. The satellite software, notably, applies the feature set on Junos OS to the aggregation device to the satellite device. The satellite software enables the satellite device to participate in the Junos Fusion, but does not provide any other software features for the satellite device.

You can run the same version of satellite software on satellite devices that are different hardware platforms. For instance, if your Junos Fusion included EX4300 and QFX5100 switches as satellite devices, the EX4300 and QFX5100 switches acting as satellite devices can install the satellite software from the same satellite software package.

RELATED DOCUMENTATION

| [Understanding Junos Fusion](#) | 461

Creating and Managing Fusion Configuration Templates

IN THIS SECTION

- [Create a Configuration Template for Junos Fusion Enterprise](#) | 467
- [Clone a Configuration Template](#) | 473
- [Apply Configuration Template to Devices](#) | 474
- [View Details about a Configuration Template](#) | 479
- [Delete a Configuration Template](#) | 480

Large campus might have many similar network topology instances. These instances can be Virtual Chassis Fabrics (VCFs), Layer 3 Fabrics, or Junos Fusion Enterprise. Creating each of these instances afresh can be tedious, time-consuming, and error-prone.

Network Director enables you to create a configuration template for a Junos Fusion Fabric and reuse it in all similar instances. A configuration template combines the common settings that apply to an instance, such as the chassis details, ports to be used as cascade and cluster ports, software image for the satellite devices, and so on. However, you might still specify some additional configuration attributes while applying the template to a fabric or Junos Fusion setup.

Before you create a configuration template:

- Understand the software requirements for the aggregation devices and the satellite devices. For more information, see [“Software Requirements for Junos Fusion” on page 465](#).
- Ensure that the software images for the aggregation devices and the satellite devices are uploaded to Network Director using the **Image Management > Manage Image Repository** in the Deploy mode.


You can perform the following tasks from the Manage Fusion Configuration Templates page.

Create a Configuration Template for Junos Fusion Enterprise

To create a Junos Fusion Enterprise configuration template:

1. Under Views, select one of these options: **Logical View**, **Location View**, **Device View**, or **Custom Group View**.

TIP: Do not select **Dashboard View**, **Datacenter View**, or **Topology View**.

2. Click  in the Network Director banner.
3. From the Tasks menu, select **Network Builder > Manage Fusion Config. Templates**.
The Manage Fusion Configuration Templates page opens.
4. Click **Create**.
The Create Fusion Configuration Template wizard opens.
5. In the Template Type page of the Create Fusion Configuration Template wizard, select the type of deployment for which you want to create the template. Select Campus to create a template for Junos Fusion Enterprise.
6. Select one of the following depending on the type of Junos Fusion topology:
 - **Fusion Enterprise Single Aggregation Device**—Indicates that the one or more satellite devices are connected to a single aggregation device.
 - **Fusion Enterprise Multiple Aggregation Devices**—Indicates that each satellite device is connected to two aggregation devices forming an MC-LAG cluster at the aggregation layer.
7. For a Junos Fusion Enterprise, select one of the following depending on the type of Junos Fusion topology:

- **Fusion Enterprise Single Aggregation Device**—Indicates that the one or more satellite devices are connected to a single aggregation device.
 - **Fusion Enterprise Multiple Aggregation Devices**—Indicates that each satellite device is connected to two aggregation devices forming an MC-LAG cluster at the aggregation layer.
8. Click + in the Available Satellite Images box to ensure that the appropriate satellite software image is available.

Network Director lists the satellite software image only if you have uploaded the software image in Network Director Image Repository.

9. Click **Next**.

The Settings page opens.

10. Enter a name for the configuration template.

If you chose to configure with multiple aggregation devices, the Settings page displays two tabs—*Aggregation Device 1* and *Aggregation Device 2*.

11. In the Aggregation Device 1 tab, select the device model that you want to use as the aggregation device.

Network Director supports the following as aggregation devices:

- EX9204, EX9208, and EX9214 in a Junos Fusion Enterprise setup.

NOTE: In a Fusion Enterprise Multiple Aggregation Device topology, when you select a device as the first aggregation device, Network Director considers the second aggregation device also to be of the same device model.

12. When you select a device model in a Junos Fusion Enterprise setup, you need to create a new chassis for the aggregation device model that you selected.

To create a new chassis:

- a. Click **Build New Chassis**.

The Build Chassis window opens. The Build Chassis window has two panes—the *Available line cards* pane and the *Chassis: FPC slots* pane.

The Available line cards pane lists the all the EX9200 line card models and the Chassis: FPC slots pane lists the available FPC slots on the device.

NOTE: Network Director does not allow you to import chassis details. You must create a chassis for each configuration template irrespective of whether the template is for a single-home or a multihome Junos Fusion topology.

- b. Drag and drop the line cards that you want to add to the chassis from the Available line cards pane to the appropriate FPC slots in the Chassis: FPC slots pane.

From all the EX9200 line cards that are listed in the Available Line Card list, select at least one line card that is supports the cascade port. If you do not choose any line card that is supported by cascade port, Network Director returns an error message that prompts you to select a line card that is supported by cascade port.

- c. Click **Set** after you have added all the required line cards to the FPC slots.

The Build Chassis window closes.

After you setup the chassis, the line cards that support the cascade ports are listed in the Select Cascade Ports window. All the other line cards are filtered out. For the list of EX9200 line cards that support cascade ports see, [Line Cards on EX9200 Switch Cascade Port Support](#).

- d. Mouse over **Preview** to preview the chassis with the line cards that you added.

13. A cascade port is a port on an aggregation device that sends and receives control and network traffic from an attached satellite device. Click **Select Cascade Ports** and perform the following steps to select cascade ports for the Junos Fusion Enterprise:

- a. To select ports individually from the ports list, click **Identify Port > By List** or to select ports by specifying a range, click **Identify Port > By Range**. The Select Ports window opens.

NOTE: In the Select Ports window, Network Director displays only the Gigabit Ethernet interfaces, 10-Gigabit Ethernet interfaces, and the 40-Gigabit Ethernet interfaces as these are the recommended interfaces for a cascade port.

- b. If you want to select ports individually, select the ports that you want to use as cascade ports on each FPC. Use the left navigation pane to view and select the ports on each FPC slot.
- c. If you want to select ports using a port range, specify the starting and ending port ID.
- d. Click **Add**.

Network Director adds the selected ports or the ports that are available on the FPC from the range that you specified, to the list of cascade ports in the Create Template page.

- e. Select a cascade port and optionally, you can enter the corresponding FPC slot ID of the satellite device.

Specify the FPC slot ID in the range 65-254.

NOTE: When you add a satellite-enabled device in a Junos Fusion setup, Network Director checks for the FPC slot ID that you specify and uses this ID while Network Director configures the satellite devices. If you have not provided an FPC slot ID, Network Director automatically generates the FPC slot ID for satellite device from the allowed range and provisions the same on the aggregation device.

14. For a Fusion Enterprise Multiple Aggregation Device topology, you must select the port type as an **ICL Port** to configure an Inter-Chassis Link (ICL port). However, it is optional to select the **ICCP PORT** as

the port type to configure an Inter-Chassis Control Protocol (ICCP) port. Select the **ICCP PORT** as the port type, to manually configure the ICCP port. If you do not select the **ICCP PORT** as the port type, Junos platform automatically assigns one of the ports as an ICCP port and deploys all the necessary configurations to the port. At the time of applying the template to one or more aggregation devices, the ICCP configuration is pushed to the aggregation devices.

NOTE: Automatic ICCP provisioning is enabled by default and if you manually configure an ICCP parameter that is normally set by default, your configuration automatically overrides the default parameter. If you decide to configure ICCP, you must configure matching configurations on both the aggregation devices.

An ICL port is used to forward data traffic across the aggregation devices. This link provides redundancy when a link failure occurs in one of the active links.

An ICCP port is used to exchange the control information between two aggregation devices to ensure that data traffic is forwarded properly.

Click **Select ICL & ICCP Ports** and perform the following steps to select the ports:

- a. To select ports individually from the ports list, click **Identify Port > By List** or to select ports by specifying a range, click **Identify Port > By Range**. The Select Ports window opens.
- b. If you want to select ports individually, select the ports that you want to use as ICL ports on each FPC. Use the left navigation pane to view and select the ports on each FPC slot.
- c. If you want to select ports using a port range, specify the starting and ending port ID.
- d. Click **Add**. Network Director adds the selected ports or the ports that are available on the FPC from the range that you specified, to the list of ICL ports in the Create Template page.
- e. For a Junos Fusion Enterprise, specify the port type as **ICL port** or **ICCP port** depending on the port usage.

In Select ICL & ICCP window, you can select more than one ICL and ICCP port as a LAG interface. The LAG interface provides redundancy and load balancing between the two aggregation devices.

When you select an ICL port as the LAG interface, there is no additional configuration required for the LAG on the selected ICL port. Network Director generates a default configuration for the aggregation devices. This configuration is pushed to the aggregation device at the time of link up event when the connection between the aggregation devices on the ICL interfaces are made.

When you select an ICCP port as the LAG interface, you must specify the local IP address of the ICCP port. For information on specifying the ICCP local address, see [“Apply Configuration Template](#)

to Devices” on page 474. If you do not select the ICCP port, Network Director automatically configures the ICCP port to establish ICCP session between the connected aggregated devices.

NOTE: If ICL or ICCP physical connections are not made between aggregation devices when you apply the template to aggregation devices, Network Director displays the each aggregation device as separate single home Junos Fusion fabrics. When the ICL and ICCP connections between aggregation devices are established, a link up is triggered. Network Director deletes the device, rediscovers the device, and converts the Junos Fusion to a multihome fabric.

15. If the Junos Fusion Enterprise uses a satellite cluster, you must select the ports on satellite devices that will act as cluster ports. In a satellite cluster, up to 10 satellite devices can be interconnected using the standard 10-Gigabit Ethernet or 40-Gigabit Ethernet interfaces to form a cluster, which in turn can be connected to the aggregation devices over a pair of fiber uplinks. In a cluster, all the satellite devices do not need to be directly connected to the aggregation device. One or two satellite devices in a satellite cluster connects to the aggregation device through cluster ports.

Click **Select Satellite Cluster Ports** to select the cluster ports that the satellite devices use to connect to other satellite devices in a satellite cluster. Perform the following steps to add cluster ports:

NOTE: If you have multiple satellite devices that connect to one or more aggregation devices, the ports that you select in is applied to all the directly connected satellite devices that are part of the cluster.

- a. To select ports individually from the ports list, click **Identify Port > By List** or to select ports by specifying a range, click **Identify Port > By Range**. The Select Ports window opens.
- b. Select the device model of the satellite device.
- c. If you want to select ports individually, select the ports that you want to use as cluster ports on the selected satellite device.

- d. If you want to select ports using a port range, specify the starting and ending port ID.
 - e. Click **Add**. Network Director adds the selected ports that you specified or the ports that are available on the device from the range that you specified, to the list of cluster ports in the Create Template page.
16. Select the software image that you want to install on all the satellite devices from the Satellite Image list. The supported satellite image version is **satellite-3.1R1.3**. For more information about software images in Junos Fusion, see [“Software Requirements for Junos Fusion” on page 465](#).
 17. For a Fusion Enterprise Multiple Aggregation Device topology, click the Aggregation Device 2 tab.
 Network Director copies the same settings that you specified for the first aggregation device for the second aggregation device. Review the settings.
 If you want to modify the setting including the device model, click **Edit**. Select a device model and follow steps Step 7 through Step 15 to specify details for the second aggregation device.
 18. Click **Next** to review the template details.
 19. After you have reviewed the details, click **Finish**. Network Director creates the template and displays it in the Manage Fusion Configuration Templates page.

Clone a Configuration Template

You can make a copy of an existing template by cloning the template. When you clone a template, Network Director copies all the settings to the cloned template. However, you can modify the settings in the cloned template based on your requirements.

To clone a Junos Fusion Enterprise:

1. Select the template that you want to clone and click **Clone**.
 The Clone Fusion Configuration Template page opens.
2. Network Director copies all the settings in the original template to the cloned template. However, you can modify all the settings in the cloned template as per your requirement. Follow step 6 through 16 in the [“Create a Configuration Template for Junos Fusion Enterprise” on page 467](#) for instructions on specifying settings for the configuration template.

Apply Configuration Template to Devices

After you have created a configuration template, you can apply the template to aggregation devices. You can apply the template to aggregation devices that are already managed by Network Director or new aggregation devices.

If you plan to apply the template to a managed device, make sure that the device runs the Junos OS software image that is supported on Junos Fusion systems. For detailed steps on installing or upgrading software image using Network Director, see [“Managing Software Images” on page 632](#). See Network Director release notes to know the Junos OS software version that is supported for Junos Fusion systems.

To apply a configuration template to Junos Fusion Enterprise:

1. From the Manage Fusion Configuration Template page, select the template that you want to apply to an aggregation device.
2. Do one of the following:
 - If the aggregation device is not managed by Network Director, click **Apply** and select **Unmanaged Devices**. The Apply Fusion Configuration Template page opens.
 - If the aggregation device is already managed by Network Director, but is not part of a fusion fabric, click **Apply** and select **Managed Devices**. The Apply Fusion Configuration Template page opens.

Skip to Step 4 and follow the instructions to apply the template to managed devices.

3. If you selected to apply the template to Unmanaged devices, perform the following steps to specify details about the aggregation devices:

- a. In the Apply Fusion Configuration Template page, specify the DHCP server settings by following the descriptions given in [Table 111 on page 475](#).

Table 111: DHCP Server Settings

Field	Description
DHCP Server	IP address or the hostname of the DHCP server.
DHCP Server Type	<p>The type of DHCP server that provides the necessary information to the aggregation devices. You can choose to use a CentOS DHCP server, an Ubuntu DHCP server, or any other DHCP server.</p> <p>NOTE: If you select Other, you must configure the DHCP server settings manually.</p>
Manually Configure Server	<p>Select to indicate that you want to manually configure the DHCP server. You can configure the CentOS and Ubuntu DHCP servers manually or from Network Director.</p> <p>If you want to use any other type of DHCP server, do the following:</p> <ol style="list-style-type: none"> a. Select the Manually Configure Server check box. Network Director hides all the other details except the DHCP Server Type. b. Follow the instructions displayed in this box to configure the DHCP server manually.
DHCP User	Username to log in to the DHCP server.
DHCP Password	Password for the specified username.
Confirm Password	Confirm the DHCP server password.

- b. Specify the File server settings described in [Table 112 on page 475](#).

Table 112: File Server Settings

Field	Description
File Server Type	The type of file server where the software image to be installed on the aggregation device is to be stored. You can choose to use an FTP, HTTP, or an TFTP file server.
File Server	IP address or hostname of the file server.
File Server Root Directory	The root directory of the file server.

- c. Select the software image that you want to install on the aggregation device from AD Image.

NOTE: Ensure that the software image is uploaded to Network Director using the **Image Management > Manage Image Repository** in the Deploy mode. If the software image is not uploaded, Network Director does not display the software image in this field.

- d. ZTP process maps the management interface MAC address or the device chassis serial number of each aggregation device to the software image, IP address, hostname, and the configuration file stored on the file server. This mapping is stored in the DHCP server. When an aggregation device starts up, the device contacts the DHCP server to obtain the IP address and the software image location. The DHCP server looks up in its MAC address or serial number mapping database to identify the device and provide details about the file server that the device must contact to get the software image and configuration file. The device uses this information to contact the file server and obtain the software image and the configuration file for deploying on the device.

Click **Add** to add a row to the Device Details table. You specify the aggregation device details in the Device Details table. You can enter the device details manually or you can specify the details in a CSV file and import it.

- e. Do the following to import the device details from a CSV file:


NOTE: When you use the Import option, you must specify either the MAC address or the Serial number, but not both.

1. Click **Import > By Mgmt MAC addresses** to import the MAC addresses of aggregation devices in CSV format. You must enter the MAC addresses in the specified format. Click **CSV Sample** to download a sample CSV file that you can use to import MAC addresses.
 2. Click **Import > By Device Serial Numbers** to import the serial numbers of aggregation devices in CSV format. You must enter the serial numbers in the specified format. Click **CSV Sample** to download a sample CSV file that you can use to import serial numbers.
- f. If you want to specify a password for all the aggregation devices that you add to this fusion fabric, click **Set Default Password**. The Set Default Password window opens.

Enter the password, confirm the password and click **Set** to set the password as the default password for all the aggregation devices that you add to this fusion fabric.

You can skip this step if you wish to specify the password individually or if you want to use a different password for each of your aggregation device.

- g. Do the following to specify details about the aggregation device manually:

1. Network Director assigns a host name for the aggregation device. You can modify this name by clicking .
2. Enter the **IP address of the management interface** or the **MAC address and the device chassis serial number** of the aggregation device.
3. If you are applying the template to a multihome Junos Fusion that has two aggregation devices, enter the ICCP IP address that this device uses to connect to the second aggregation device. If you are applying the template to a single-home Junos Fusion, you can leave this field blank.
4. Enter the MAC address of the management interface (for example, the em0 interface) or the device chassis serial number of the aggregation device in the MAC Addresses or Serial Number field.
5. Click **Import MAC Address** to import the MAC addresses of aggregation devices in CSV format. You must enter the MAC addresses in the specified format. Click **CSV Sample** to download a sample CSV file that you can use to import MAC addresses.

NOTE: When you use the Import option, you must specify either the MAC address or the Serial number, but not both.

6. Click **Import Serial Number** to import the serial numbers of the aggregation device in CSV format. You must enter the MAC addresses in the specified format. Click **CSV Sample** to download a sample CSV file that you can use to import serial numbers.
- h. Type the local IP address for the ICCP port. The IP address is used to communicate to the peers that hosts an ICCP port as a LAG interface.

NOTE: This field is available only when you choose to configure the ICCP port manually for a Junos Fusion Enterprise device.

- i. Click the text link in the Password column to set or modify the password for the aggregation device.
 - If you have not configured a default password for the aggregation device, then click **Set Password**. The Set Custom Password window opens. Specify the password, confirm the password, and click **Set**.
 - If you want to override the default password and specify a custom password for the aggregation device, then click **Edit**. The Edit Custom Password window opens. Select **Use Default Password** if you want to use the default password.

Select **Edit Custom** if you want to specify a different password for the aggregation device. Specify the password, confirm the password, and click **Set**.

- j. To view the configuration that Network Director deploys on the aggregation device, click **View** in the Config field.
 - l. Repeat above steps, Step d through Step j to add another aggregation device.
4. If you selected to apply the template to Managed devices, perform these steps:
- a. Select the devices that you want to add as aggregation device from the Select Managed Devices table.
 - b. Click on the link in the Backup Config column.

Network Director creates a ZIP archive file containing the existing configuration of the device. You save a copy of this onto your local machine.

NOTE: Before you perform this step, make sure that the device is running a Junos OS software image that is supported on Junos Fusion systems. If not, you must upgrade the software image before you proceed. For detailed steps on installing or upgrading software image using Network Director, see [“Managing Software Images” on page 632](#).

5. Click **Apply**.

Network Director converts the device to a fusion fabric aggregation device and deploys all the necessary configurations on the device and opens the Manage Fusion Fabrics page.

Manage Fusion Fabrics page displays the status and details of all fusion fabrics that are created using Network Director.

After the template is applied successfully to a Junos Fusion fabric, there are quite a few tasks that Network Director performs internally that makes building your Junos Fusion fabric simple and error-free. Network Director converts the device that you specified in the Apply Template workflow to an aggregation device and applies the port settings on the various ports that you specified in the configuration template. When an EX4300 device is connected to one of the configured cascade ports, a *link up* event is triggered. The link up event initiates a syslog message to Network Director and Network Director initiates a *topology refresh* job. Network Director then installs the appropriate satellite software on the satellite device and performs the necessary configurations. If a second satellite device is connected to the first satellite device to form a satellite cluster, or another satellite device is connected to the aggregation device, another link up event is triggered and the same steps are repeated. This process continues for all additional satellite devices that are connected to the aggregation device.

NOTE: If ICL or ICCP physical connections are not made between aggregation devices when you apply the template to aggregation devices, Network Director displays the each aggregation device as separate single home Junos Fusion fabrics. When the ICL and ICCP connections between aggregation devices are established, a link up is triggered. Network Director deletes the device, rediscovers the device, and converts the Junos Fusion to a multihome fabric.

For Junos Fusion Enterprise, if you want to form a cluster setup and if the cascade port of the Aggregate device is connected to the base port (1 GE) of EX2300, EX3400, or EX4300 devices, you need to configure the cluster-policy manually by logging into the aggregation device. Execute the following commands to configure the cluster policy:

- If the satellite device is part of a cluster, log in to the CLI of the aggregation device and execute the following commands:

```
user@aggregation-device# set groups policy-number policy-options satellite-policies port-group-alias
port-group-alias-name pic pic-identifier port port-ID1
```

```
user@aggregation-device# set groups policy-number policy-options satellite-policies port-group-alias
port-group-alias-name pic pic-identifier port port-ID2
```

```
user@aggregation-device# set groups policy-number policy-options satellite-policies port-group-alias
port-group-alias-name pic pic-identifier port port-ID3
```

```
user@aggregation-device# set groups policy-number policy-options satellite-policies port-group-alias
port-group-alias-name pic pic-identifier port port-ID4
```

```
user@aggregation-device# set groups policy-number policy-options satellite-policies port-group-alias
port-group-alias-name set groups policy-number policy-options satellite-policies candidate-uplink-port-policy
cluster-policy-name uplink-port-group pic-identifier
```

```
user@aggregation-device#set apply-groups policy-number
```

```
user@aggregation-device# set chassis satellite-management cluster cluster-number cluster-policy
cluster-policy-name
```

View Details about a Configuration Template

To view details about a configuration template, select the template and click **Details**.

The Template Details for <template-name> window opens.

For a Junos Fusion Enterprise setup, the Template Details for <template-name> page displays the template name, deployment type, configuration type, software upgrade group details, and details of ports such as satellite cluster ports, cascade ports, ICL, and ICCP ports.

Delete a Configuration Template

To delete a Junos Fusion Enterprise, select the template and click **Delete**.

RELATED DOCUMENTATION

[Understanding Junos Fusion | 461](#)

[Managing Fusion Fabrics | 480](#)

Managing Fusion Fabrics

IN THIS SECTION

- [Modify the Fusion Fabric | 481](#)
- [View the Cabling Plan | 484](#)
- [View Fabric Connectivity | 484](#)
- [Replace Aggregation Device or Satellite Device in Junos Fusion | 485](#)

With Junos Fusion technology, network administrators can reduce network complexity and operational expenses by collapsing underlying network elements into a single, logical point of management using QFX Series and EX Series switches running the Junos operating system. You can create and provision Junos Fusion fabrics using Network Director. The Manage Fusion Fabrics page displays the status and details of all fusion fabrics that are created using Network Director.

To open the Manage Fusion Fabrics page:

1. Under Views, select one of these options: **Logical View**, **Location View**, **Device View**, or **Custom Group View**.

TIP: Do not select **Dashboard View**, **Datacenter View**, or **Topology View**.


- 2. Click  in the Network Director banner.
- 3. From the Tasks menu, select **Network Builder > Manage Fusion Fabrics**.
The Manage Fusion Fabrics page opens. [Table 113 on page 481](#) describes the fields that are displayed in the Manage Fusion Fabrics page.

Table 113: Manage Fusion Fabrics Field Descriptions

Field	Description
Fusion Name	Name of the fusion fabric.
Fusion Type	Indicates the type of fusion fabric as Fusion Enterprise Single Home, Fusion Enterprise Multi Home.
Aggregation Devices	Hostname and IP address of the aggregation devices.
Summary	Summary of the fusion fabric. This field displays the number of cascade ports, number of satellite devices, and the status of the satellite devices.
Status	Status of the fusion fabric.
Action	Additional tasks that you can perform on the fusion fabric. Available actions are—View Topology, Download Cabling Plan, and View Connectivity.
Last Updated Time	Date and time when the fusion fabric was last updated.

You can perform the following tasks from the Manage Fusion Fabrics page:

Modify the Fusion Fabric

IN THIS SECTION

- [Edit Aggregation Device Details | 482](#)
- [Edit Satellite Device Details | 483](#)
- [Enable Uplink Failure Detection | 483](#)

From the Edit Fusion Fabric page:

1. Select a fabric from the table and click **Edit**.

The Edit Fusion Fabric page opens.

2. Edit the Aggregation device details. For procedure steps, see [“Edit Aggregation Device Details” on page 482](#).
3. Edit the satellite device details. For procedure steps, see [“Edit Satellite Device Details” on page 483](#).
4. Enable uplink failure detection for a satellite device. For procedure steps, see [“Enable Uplink Failure Detection” on page 483](#).
5. Click **Save** to save the changes to the fusion fabric and close the Edit Fusion Fabric page. The Configuration Deployment window opens displaying the status of the job the Network Director initiates to deploy the changes to the fusion fabric.
6. Click **Close** in the Configuration Deployment window to close the window and return to the Manage Fusion Fabrics page.

Edit Aggregation Device Details

To edit the aggregation device details:

1. Open the **Aggregation Devices** tab in the Edit Fusion Fabric page to edit the aggregation device details. You can add cascade ports or delete cascade ports from an aggregation device. To do this, select an aggregation device from the table. Network Director displays the cascade ports that are part of the selected aggregation device in the port details table.
2. To add new cascade ports to the aggregation device, click **Add Ports**. The Select Ports window opens. Click an FPC and select the ports that you want to include. When you have added the ports, click **Add**.
3. To delete a cascade port, select a port from the port details table and click **Remove**.

NOTE: In a multihome Junos Fusion fabric, if you remove a cascade port from one aggregation device, Network Director initiates a two-step commit process on both the aggregation devices. If one of the aggregation device is out-of-sync or is not reachable, the port is not deleted.

Edit Satellite Device Details

To edit the satellite device details:

1. Open the **Satellite Devices** tab in the Edit Fusion Fabric page to edit the satellite device details.

You can perform the following operations on the satellite devices in the Edit Fusion Fabric page:

- change the alias name for the satellite device. See Step 2.
- remove a satellite device. See Step 3.

2. In a multi-home setup, the configuration for a satellite device exists on both the aggregation devices. In such a topology, the alias name of the satellite device should be identical on both the aggregation devices. If the alias names of a satellite device does not match, you need to edit the alias name to be identical on both the aggregation devices. To edit the alias name of a satellite device:

- a. Select the FPC slot number.
- b. Click the alias name in the column **Alias Name**.
- c. Type the new alias name and then click **Save**.

Network Director displays a warning message **Alias name for AD1 and AD2 will be same after the edit**.

- d. Click **OK**.

The alias name of the satellite device is now identical on both the aggregation devices.

3. To remove a satellite device from the fusion fabric, select the device and click **Remove**.

NOTE: Replacing a satellite device that has no MAC address or serial number binding to the FPC slot on the aggregation device is plug-and-play. Replace the satellite device in your Junos Fusion topology and Network Director takes care of all the remaining configurations. This is irrespective of whether the device is a standalone device or a member of a satellite cluster. However, if there is a MAC address or serial number binding between the satellite device and the FPC slot on the aggregation device, you must run some commands from the CLI of the aggregation device to replace the satellite device. For more details, see [“Replace Aggregation Device or Satellite Device in Junos Fusion” on page 485](#).

Enable Uplink Failure Detection

To enable update failure detection:

1. Click the **Settings** tab in the Edit Fusion Fabric page.

2. Select **Uplink Failure Detection** check box corresponding to the device to enable uplink failure detection on a Junos Fusion.

Enabling uplink failure detection on a satellite devices detects link failures on the uplink interfaces used to connect to aggregation devices. When uplink failure detection detects uplink failure on a satellite device, all of the device's extended ports (which connect to host devices) are shut down.

You can also view the ICL and ICCP settings for a multihome Junos Fusion Enterprise fabric.

View the Cabling Plan

To download the cabling plan for the Junos Fusion fabric, click **Download Cabling Plan** from the **Action** field corresponding to a fabric. Network Director generates and downloads the cabling plan as a PDF file.

View Fabric Connectivity

After you have set up and deployed Junos Fusion devices in Network Director, you can pictorially view the physical connectivity between the various devices in the fabric. This page displays the devices and their physical connectivity in the spine-and-leaf topology.

To view the connectivity between the Junos Fusion devices:

1. Do one of the following:
 - From the Manage Fusions page, click **View Topology** in the **Actions** field.
 - While in the Logical, Location, Device, or Custom View, select the Junos Fusion fabric for which you want to view the connectivity details from the View pane and click **Connectivity > View Fabric Connectivity** from the Tasks pane.

The Fusion Connectivity page opens, displaying the connectivity between the aggregation devices and the satellite devices in the selected fusion fabric.

NOTE: If you change the position and arrangement of Junos Fusion devices in a topology and navigate to some other page in the user interface, Network Director preserves the position of the devices in the topology when you return to the topology.

2. From the Fabric Connectivity page, you can:
 - Click **Device** and select **Provisioned SDs** to view all the satellite devices that are provisioned as part of the Junos Fusion system.
 - Select Color Code Port Utilization to view the color coded port utilization in the graphical view by clicking **Links**.

- Mouse over each entity to view a window displaying the details about that entity and View Device Connectivity link. Clicking on View Device Connectivity link displays Device Connectivity page. For more information about this page, see [“Physical Topology” on page 553](#).
- Zoom in or zoom out of the connectivity view by using the + and - buttons.

Replace Aggregation Device or Satellite Device in Junos Fusion

Network Director enables you to replace faulty or non-responsive aggregation devices in your Junos Fusion fabric by using the Replace functionality. You can replace an aggregation device in a dual home Junos Fusion by selecting the device that you want to remove and by specifying the serial number or the MAC address of the replacement device.

Replacing a satellite device that has no MAC address or serial number binding to the FPC slot on the aggregation device is plug-and-play. Replace the satellite device in your Junos Fusion topology and Network Director takes care of all the remaining configurations. This is irrespective of whether the device is a standalone device or a member of a satellite cluster. However, if there is a MAC address or serial number binding between the satellite device and the FPC slot on the aggregation device, you must run some commands from the CLI of the aggregation device to replace the satellite device.

This topic describes the steps that you must perform to replace an aggregation device or a satellite device from the a Junos Fusion fabric.

To replace aggregation device or satellite device:

1. From the Manage Fusions page, select a fusion fabric for which you want to replace an aggregation and click **Replace**. The Replace Fusion Fabric page opens.
2. To replace an aggregation device, select the aggregation device that you want to replace and click **Replace**.

Enter the MAC address or the serial number of the new aggregation device.

3. Do one of the following to replace a satellite device that has a MAC address or serial number binding to the FPC slot of the aggregation device:

- If the satellite device is a standalone device, log in to the CLI of the aggregation device and execute the following commands:

```
user@aggregation-device# set chassis satellite-management fpc fpc-id system-id
MAC-address-of-the-device | Serial-number -of-the-device
user@aggregation-device# Commit
```

- If the satellite device is part of a cluster, log in to the CLI of the aggregation device and execute the following commands:

```

user@aggregation-device# set chassis satellite-management cluster cluster-name fpc fpc-id alias alias
user@aggregation-device# set chassis satellite-management cluster cluster-name fpc fpc-id description
10.204.248.62-member0
user@aggregation-device# set chassis satellite-management cluster cluster-name fpc fpc-id member-id
member-id
user@aggregation-device# set chassis satellite-management cluster cluster-name fpc fpc-id system-id
MAC-address-of-the-device | Serial-number -of-the-device
user@aggregation-device# Commit

```

RELATED DOCUMENTATION

[Understanding Junos Fusion](#) | 461

[Creating and Managing Fusion Configuration Templates](#) | 466

Creating and Managing Satellite Software Upgrade Groups

IN THIS SECTION

- [Create a Software Upgrade Group](#) | 487
- [Edit a Software Upgrade Group](#) | 488
- [View Details of a Software Upgrade Group](#) | 488
- [Delete a Software Upgrade Group](#) | 488

A satellite software upgrade group is a group of satellite devices that are designated to upgrade to the same satellite software version using the same satellite software package. One Junos Fusion can contain multiple software upgrade groups, and multiple software upgrade groups must be configured in most Junos Fusions to avoid network downtimes during satellite software installations.

In Network Director, you select a fusion fabric and create a software upgrade group. The software upgrade group can contain one or more satellite devices. When a satellite device is added to a Junos Fusion, the aggregation device checks whether the satellite device or the FPC ID that is used by the satellite device is included in the satellite software upgrade group that is assigned to the Fusion system. If it is, the device—unless it is already running the same version of satellite software—upgrades its satellite software using the satellite software associated with the satellite software upgrade group.

When the satellite software package associated with an existing satellite software group is changed, the satellite software for all member satellite devices is upgraded using a throttled upgrade. The throttled upgrade ensures that only a few satellite devices are updated at a time to minimize the effects of a traffic disruption caused by too many satellite devices upgrading software simultaneously.

You can create and manage software upgrade groups from the Manage Software Upgrade Groups page. After you create a software image group, you can open the Deploy Images to Devices page and select a satellite software image for each software upgrade group and use the Select Options tab to set the date and time when the upgrade must be performed.

To access the Manage Software Upgrade Groups page:

1. Select the fusion fabric in the View pane for which you want to create a software upgrade group.
2. Click **Deploy** in the Network Director banner.
3. In the Tasks pane, select **Image Management > Manage Software Upgrade Group**.

The Software Upgrade Groups page opens in the main window. The table lists the software upgrade groups that exist for the selected fabric, if any.

NOTE: This task is available only if you select a fabric in the View pane.

You can perform the following tasks from the Manage Software Upgrade Groups page:

Create a Software Upgrade Group

To create a software upgrade group:

1. Click **Add**. The Add Software Upgrade Group window opens.
2. Enter a name for the new software upgrade group.
3. Select one or more satellite devices from the Select Devices tab that you want to be part of the software upgrade group. The Select Devices table lists the satellite devices and the available FPC slots in each of these that are not yet part of any other software upgrade groups.
4. Click the **FPC Range** tab and select the FPCs that you want to add to the software upgrade group. To add a single FPC number, specify the number in **From** and click **Add**. To add a range of FPC numbers, enter the starting and ending FPC numbers in **From** and **To** respectively, and click **Add**. You can also add FPCs that are currently inactive to the software upgrade group.

Network Director displays the FPC numbers that you added in the Selected FPC Number/Range table.

5. Click the **Preview** tab to review the software upgrade group settings. You can click the **Select Devices** or **FPC Range** tabs to modify the configuration settings.

6. Click **Add** to create the software upgrade group.

Network Director lists the new software upgrade group in the **Manage Software Upgrade Group** page. You can now use the **deploy image** task to assign a software image to this software upgrade group and deploy the image to the devices that are part of the upgrade group. For details on selecting and deploying software images, see [“Deploying Software Images” on page 635](#).

Edit a Software Upgrade Group

To edit a software upgrade group:

1. Select a software upgrade group and click **Edit**.

The **Edit Software Upgrade group** window opens.

2. You can edit the FPC number and the FPC range.

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

View Details of a Software Upgrade Group

To view the details of a software upgrade group:

1. Select a software upgrade group and click **Details**. The **Software Upgrade Group Summary** window opens displaying the members that are part of the upgrade group and the associated satellite image name.

2. Click **OK** to close the summary window.

Delete a Software Upgrade Group

To delete a software upgrade group, select a software upgrade group and click **Delete**.

RELATED DOCUMENTATION

[Deploying Software Images | 635](#)

[Understanding Junos Fusion | 461](#)


[Software Requirements for Junos Fusion | 465](#)

Understanding Layer 3 Fabrics

Most enterprises are looking to increase resiliency and also support new technologies such as VMware NSX that allow them to deploy applications, servers, and virtual networks within seconds. Layer 3 Fabrics allow them to support better uptime, performance, and newer cloud infrastructures such as VMware NSX. In order to maintain the large scale required to host thousands of servers, the use of a multi-stage Clos architecture is required. Such an architecture allows the physical network to scale beyond the port density of a single switch. Layer 3 Fabrics use BGP as the control plane protocol to advertise prefixes, perform traffic engineering, and tag traffic. The most common designs in a multi-stage Clos architecture are a 3-stage and 5-stage networks that use the spine-and-leaf topology.

Spine-and-leaf topology is an alternate to the traditional three-layer network architecture, which consists of an access layer, aggregation layer, and a core. In the spine-and-leaf topology, all the leaf devices are connected to the spine devices in a mesh as shown in [Figure 21 on page 489](#).

Figure 21: Layer 3 Fabric in a Spine and Leaf Topology



Typically, the spine devices are high-performance switches capable of Layer 3 switching and routing combined with high port density. Spine devices constitute the core and the leaf devices constitute the

access layer in Layer 3 Fabrics. Leaf devices enable servers to connect to the Layer 3 Fabric. They also provide uplinks to spine devices.

Network Director currently supports only the 3-stage design. The 3-stage design has two roles—the spine and the leaf. It is called a 3-stage design because the traffic must traverse three switches in the worst-case scenario.

The maximum number of spine devices that you can have in your Layer 3 Fabric depends on the number of 40-Gigabit Ethernet interfaces in your leaf devices. A Layer 3 Fabric that has 8 QFX5100-24Q spine devices and 32 QFX5100-96S leaf devices (each leaf supports 96 10-Gigabit Ethernet ports) can provide 3072 usable 10-Gigabit Ethernet ports.

RELATED DOCUMENTATION

[Managing Layer 3 Fabrics](#) | 491

[Creating Layer 3 Fabrics](#) | 493

[Network Director Documentation home page](#)

User Privileges Required for the DHCP and File Server While Using Zero Touch Provisioning

Ensure that you have the following user privileges on the DHCP server and the file server prior to configuring them for zero touch provisioning (ZTP).

- DHCP server—Ensure that the DHCP user has permissions to:
 - write to the **dhcpd.conf** file on the DHCP server.

NOTE: To fetch the **dhcp.conf** file, ensure that the DHCP server and the Layer 3 Fabric devices are in the same subnets. If you are not in the same subnet, you must specify the gateway IP address that these devices can use to reach Network Director and fetch the **dhcp.conf** file. For information about specifying the gateway IP address, see, “[Creating Layer 3 Fabrics](#)” on page 493.

- write to the **/etc/dhcp/ddns-keys** directory
- copy the file **dhcpd.conf** to the file **dhcplibbacknd.conf**
- start the **isc-dhcp-server** service

For more information about file permissions, refer DHCP server documentation.

- File server—Network Director uses the *anonymous* user to connect to the file server. You must modify certain configurations in the server configuration file to enable Network Director to access the file server. Change the configuration settings for the following file servers, depending on the file server type and the operating system that is running on the file server:

- For FTP server running CentOS (or any other FreeBSD-based servers)—Modify the configuration in the `/etc/vsftpd/vsftpd.conf` file as follows:

```
anonymous_enable=YES
anon_upload_enable=YES
anon_mkdir_write_enable=YES
file_open_mode=0644
anon_umask=033
```

- For TFTP running on a Linux server—Modify the configuration in the `/etc/xinetd.d/tftp` file as follows:

```
server_args = -c -s <dir>
disable = no
```

RELATED DOCUMENTATION

[Configuring and Monitoring Zero Touch Provisioning | 650](#)

[Creating Layer 3 Fabrics | 493](#)

Managing Layer 3 Fabrics

You can view and manage Layer 3 Fabrics in your network, using the Manage Layer 3 Fabric page.

To manage Layer 3 Fabrics:

1. Under Views, select one of these options: **Logical View**, **Location View**, **Device View**, or **Custom Group View**.

TIP: Do not select **Dashboard View**, **Datacenter View**, or **Topology View**.

2. Click  **Build** or  **Deploy** in the Network Director banner.

3. In the Tasks pane, click **Network Builder > Manage Layer 3 Fabrics**.

The Manage Layer 3 Fabrics page opens.

4. [Table 114 on page 492](#) describes the information provided about Layer 3 Fabrics on the Manage Layer 3 Fabrics page. This page lists all Layer 3 Fabrics defined for your network, regardless of the scope you selected in the network view.

Table 114: Manage Layer 3 Fabrics Field Descriptions

Field	Description
Fabric Name	Name given to the Layer 3 Fabric when the fabric was created.
Description	Description given when the fabric was created.
Summary	<p>Displays the following details about the fabric:</p> <ul style="list-style-type: none"> • Type of fabric—3-stage • Number of deployed and active spine and leaf devices. <p>Click View Topology to view the topology of the Layer 3 Fabric.</p>
Status	<p>Displays the current status of the Layer 3 Fabric. Status can be—In Design, In Progress, Deployed, or Failed.</p> <ul style="list-style-type: none"> • In Design—The fabric creation is in progress. The user might have saved and exited the Create Layer 3 Fabric wizard without specifying all the details. • In Progress—All the details that Network Director requires to create the fabric was entered by the user. Network Director might be performing background actions such as copying software images or configurations. Click this field to view the status of the jobs that are running in the background. • Deployed—The fabric is deployed and provisioned in the network, but might not be connected. • Failed—Creation or deployment of the fabric failed.
Cabling	<p>You can download the cabling plan or run a connectivity check by clicking the respective buttons in this field. You can also view details about the last run connectivity check by clicking View Connectivity Results.</p> <p>Mouse over this field to view the date and time when the last connectivity check was run.</p>
Updated Time	Time when the fabric was last updated.

TIP: All columns might not be displayed. To show or hide fields listed in the table, click the down arrow on the field header, select **Columns**, and select or clear the check box adjacent to the field that you want to show or hide.

5. You can perform the following tasks from the Manage Layer 3 Fabrics page:

- Click **Create** to create a new Layer 3 Fabric. For detailed steps, see [“Creating Layer 3 Fabrics” on page 493](#).
- Select a Layer 3 Fabric and click **Edit** to modify the fabric details. For detailed steps, see [“Editing Layer 3 Fabrics” on page 508](#).
- Select a Layer 3 Fabric and click **Delete** to delete the fabric.
- Click **Download Cabling Plan** in the Cabling column to download the cabling plan of the fabric. For more details, see [“Performing Layer 3 Fabric Connectivity Checks” on page 512](#).
- Click **Run Connectivity Check** or **Re-run Connectivity Check** in the Cabling column to check the cabling plan for the fabric. Network Director performs the cabling check. You can view the result of the cabling check by clicking **View Connectivity Results**.
- Click **View Topology** in the Summary column to view the physical topology of the Layer 3 Fabric. For more details, see [“Viewing Layer 3 Fabric Connectivity” on page 511](#).

RELATED DOCUMENTATION

[Creating Layer 3 Fabrics | 493](#)

[Editing Layer 3 Fabrics | 508](#)

[Performing Layer 3 Fabric Connectivity Checks | 512](#)

[Understanding Layer 3 Fabrics | 489](#)

Creating Layer 3 Fabrics

IN THIS SECTION

- [Specifying the Fabric Requirements | 494](#)
- [Specifying the Device Details | 499](#)
- [Specifying Configuration Details | 500](#)

- [Viewing the Cabling Plan | 501](#)
- [Specifying Zero Touch Provisioning Details | 504](#)
- [Reviewing the Layer 3 Fabric Settings | 507](#)

You can create and manage 3-stage Layer 3 Fabrics in Network Director by using the Create Layer 3 Fabrics wizard. Use the various pages of the wizard to specify the requirements and configurations for a Layer 3 Fabric. You can save the data that you have entered in one or more wizard pages, and come back later to specify the remaining details and complete the fabric creation.



CAUTION: Ensure that you always create the Layer 3 Fabric using this wizard and perform the physical connections based on the cabling plan that Network Director generates for your fabric. Not following this set order might render your Layer 3 Fabric defunct.

Before you begin, ensure that you have the necessary privileges on the FTP and the file server that Network Director uses for Zero Touch Provisioning. For more details, see [“User Privileges Required for the DHCP and File Server While Using Zero Touch Provisioning” on page 490](#).

You can do the following tasks from the Create Layer 3 Fabric wizard pages:

Specifying the Fabric Requirements

To specify the fabric requirements:

1. Enter a name for the Layer 3 Fabric. The fabric name must be unique and can contain alphanumerals, hyphens, and underscores.
2. Enter a description for the Layer 3 Fabric.

NOTE: Network Director currently enables you to create 3-stage Layer 3 Fabrics and hence this is the default selection. You cannot modify the Fabric Type.

3. Select **QFX10008**, **QFX10002-36Q**, **QFX10002-72Q**, **QFX5100-24Q-2P** or **QFX5200-32C-32Q** as the device model for spine devices. All spine devices will be of the model that you select.

NOTE: If you select QFX10008 as spine, all the line cards must be homogenous across spines. For example, if you are building an IP fabric with four QFX10008 spines containing L1, L2, and L3 line cards, all the four spines must have L1, L2, and L3 line cards only and in the same slots.

4. Enter the number of spine devices that you plan to have initially and the maximum number of devices that you plan to have in this fabric, in the Initial Capacity and Max Capacity boxes respectively. You can have a minimum of 2 and a maximum of 8 spine devices.

NOTE: Initial capacity must be less than or equal to the maximum capacity. Maximum capacity must be greater than or equal to the initial capacity and must not be more than 8.

5. If you selected **QFX10008** as the spine device you must build the device chassis using the chassis builder.

To do this:

- a. Click **Build New Chassis** to create a new chassis.

The Build Chassis window opens. The Build Chassis window has two panes— the *Available line cards* pane and the *Chassis: FPC slots* pane.

The Available line cards pane lists the line cards that are supported on the selected aggregation device and the Chassis: FPC slots pane lists the available FPC slots on the device.

- b. Drag and drop the line cards that you want to add to the chassis from the Available line cards pane to the appropriate FPC slots in the Chassis: FPS slots pane.
 - c. Click **Set** after you have added all the required line cards to the FPC slots. The Build Chassis window closes.
 - d. Mouse over **Preview** to preview the chassis with the line cards that you added.
6. In the Fabric leaves section, click a row in the table to select a leaf device model and specify the capacity of the selected model that you plan to have in the fabric. Click **Add** to add subsequent rows.

NOTE: This is an optional step, however, it is mandatory to specify the maximum number of leaf devices you plan to have in this fabric. If you do not add any leaf devices, Network Director considers these devices as unknown and creates a cabling plan accordingly. After the fabric is deployed, you can plug and play any of the supported leaf device models to the fabric. After reaching the initial capacity for the spine devices, Network Director regenerates the cabling plan. Follow this plan to connect additional spine devices.

You can add one or more of the following device models as leaf devices in your Layer 3 Fabric:

- QFX5100-48S-6Q
- QFX5100-96S-8Q
- QFX5100-48T-6Q
- QFX5200-32C-32Q
- QFX5100-24Q-2P
- EX4300-32F
- EX4300-48P
- EX4300-24P
- EX4300-48T
- EX4300-24T

NOTE: QFX5100-48T-6Q can be standalone or Virtual Chassis leaf devices. QFX10008, QFX10000-36Q, QFX10000-30C, QFX10000-60S-6Q, and QFX5200 are supported only as standalone devices.

Network Director supports a maximum of two members in a Virtual Chassis.

If you want to delete a device entry, select a row and click **Remove**.

7. If you want to include Virtual Chassis as a leaf device, select **Include Virtual Chassis (VC) as a leaf**.

Enter the number of Virtual Chassis that you want to deploy immediately in **Initial Capacity** and the total number of Virtual Chassis that will be part of the Layer 3 Fabric in **Max. Capacity**. The minimum number of devices you can specify in **Initial Capacity** is 0. The **Max. Capacity** is the maximum number of devices you can specify, which depends on the spine device that you have selected. See [Table 115 on page 497](#).

Table 115: Maximum Virtual Chassis Supported on Spine Devices

If you choose the spine device as...	then, the maximum number of virtual chassis leaf devices supported is...
QFX5100-24Q-2P	16
QFX10002-36Q	18
QFX5200-32C-32Q	16
QFX10002-72Q	36
QFX10008	~144 NOTE: Depends upon the type of line card connected.

NOTE: Plug and play is not supported for Virtual Chassis leaf members. Therefore, before you physically connect the Virtual Chassis members, make sure that you add the Virtual Chassis leaf members by using this Layer 3 Fabric wizard.

Initial capacity must be less than or equal to the maximum capacity.

For example, if your selected spine model is QFX5100-24Q-2P and if all of the leaf device members are Virtual Chassis, each containing 2 members, then the maximum number of Virtual Chassis leaf devices is restricted to 16, as there is a connection from both the primary and backup member of the Virtual Chassis to each spine device. See [Table 116 on page 498](#) for the maximum number of devices supported on various spine devices.

NOTE: You cannot modify the number of Virtual Chassis after the Layer 3 Fabric is created.

Network Director helps in creating the access link aggregation group (LAG) between Virtual Chassis members and host access devices. Network Director creates the access LAG in either of the two ways.

- **Dynamic LAG creation**—As the access devices are connected to the Virtual Chassis members, Network Director creates the LAG (if there are more than one connection between the access device and the Virtual Chassis members) dynamically. To identify the connected links for LAG creation, Network Director uses the Topology Discovery, which requires LLDP to be enabled in both the host and leaf members. For Network Director to create the LAG dynamically, ensure that LLDP is enabled in both the host and leaf (Virtual Chassis) devices.

- Preprovisioning LAG configuration on Virtual Chassis members—If LLDP is not enabled in the access or host devices, Network Director generates the LAG configuration on the Virtual Chassis member devices during the workflow creation and pushes the configuration to the Virtual Chassis members when they are connected to the network. The LAG interfaces are depicted in the cabling plan graph and in the grid view generated by Network Director. You must connect the host devices to the Virtual Chassis member devices according to the cabling plan.

You can enable Network Director to create the LAG as the physical connections are established. Select the **Dynamically create LAG when hosts are connected** check box.

8. Enter the maximum number of leaf devices, which includes standalone and Virtual Chassis devices, that the fabric can accommodate in **Max Capacity**. The minimum value you can enter is 1 and the maximum value depends on the spine device that you choose. See [Table 116 on page 498](#).

Table 116: Maximum Number of Leaves

If you choose the spine device as...	then, the maximum number of leaves is...
QFX5100-24Q-2P	32
QFX10002-36Q	36
QFX10002-72Q	72
QFX10008 <ul style="list-style-type: none"> • QFX10000-36Q • QFX10000-60S-6Q • QFX10000-30C 	<ul style="list-style-type: none"> • < 288 • < 48 • < 240
QFX5200-32C	32

9. Do one of the following:
 - Click **Next** to open the Devices page where you can view and modify details of the spine and the leaf devices.
 - Click **Save & Exit** to save the changes and exit the Create Layer 3 Fabric wizard. You can resume this task later without losing any information that you entered.

Specifying the Device Details

The Devices page displays the number of spine and leaf devices that you are provisioning as part of the initial capacity, enables you to edit the hostname for all the spine and leaf devices. Select a model for each member of Virtual Chassis if you have opted for Virtual Chassis leaves, and search for a specific device in the fabric.

Network Director prefixes the name of the fabric that you specified in the Fabric Requirements page to the name of all the spine and leaf devices. If required, you can modify this prefix in the Devices page. You can also use the search box to search for specific devices in the fabric.

To specify the device details:

1. Click **Edit Host Name Prefix** if you want to change the device name prefix to something other than the name of the fabric. The Edit Host Name window opens.
2. Enter the name that you want to use as the device name and click **OK**.

Network Director replaces the device name prefix with the name that you entered.

3. The Devices page displays the details of the hostname and the devices associated with it. See [Table 117 on page 499](#).

NOTE: The details of the device in each row, which is colored blue are to be provisioned now, and those colored orange are reserved for future allocation.

Table 117: Devices Page Description

Column	Description
Host Name	Displays the hostname with the name of the fabric, which you specified in the Fabric Requirement page.
Model	<p>Displays the model of the switch.</p> <p>If you have selected Virtual Chassis to be included in your Layer 3 Fabric in the Fabric Requirements page, the <i>Type</i> of the model will be <i>Virtual Chassis</i> and the <i>Model</i> is not displayed. You can select the switch model for the Virtual Chassis member from the drop-down list, which lists all supported Virtual Chassis members.</p> <p>NOTE: It is mandatory to select the switch model for Virtual Chassis member that you are provisioning now. For the Virtual Chassis members that are <i>Reserved for future</i> you may select the model later.</p>

Table 117: Devices Page Description (*continued*)

Column	Description
Type	Displays the type of switch—standalone, virtual chassis, virtual chassis member, or FPC.
Role	Displays the role being played by the switch model.

4. Do one of the following:

- Click **Next** to open the Configuration page, where you can specify the configuration details of the Layer 3 Fabric.
- Click **Save & Exit** to save the changes and exit the Create Layer 3 Fabric wizard. You can resume this task later without losing any information that you entered.

Specifying Configuration Details

To specify the configuration details for the Layer 3 Fabric:

1. Enter details in the Configuration page by following the descriptions given in [Table 118 on page 500](#).

Table 118: Layer 3 Fabric Configuration Details

Field	Description
Loopback Network Address	Specify the IP address block that you want to use for configuring the loopback interface in each member. Each device in the fabric is assigned one IP address from the block. This IP address can be used for troubleshooting and for checking connectivity between switches.
Interconnect Network Address	Specify the IP address block that you want to use for configuring the IP addresses for interconnect links between leaves and spines. Each interconnect links is assigned two IP addresses from this block.
VLAN Network Address	Specify the IP address block to be reserved for the virtual machines or hosts that you want to connect to the leaves. Network Director allocates each leaf device with a subnet from the given IP address block.
Start Management IP	Specify the management IP address that Network Director will use to manage each switch. NOTE: If you have provisioned for Virtual Chassis members in the Layer 3 Fabric, each Virtual Chassis member is initially treated as a standalone device and it goes through the ZTP process. The Management IP address block is sufficient to provide individual unique IP address for each of the Virtual Chassis member in the fabric.

Table 118: Layer 3 Fabric Configuration Details (*continued*)

Field	Description
Max Hosts/VMs per leaf	Specify the maximum IP addresses that are required in the subnet to be allocated from the VLAN Network Address.
Spine-BGP Autonomous System Number	<p>Specify the starting autonomous system (AS) number to be assigned to the first spine device. Subsequent spine devices are assigned incremental AS numbers starting from the number you specified.</p> <p>Network Director updates the last AS number based on the number of spine devices that you plan to have in the fabric. You cannot modify the last AS number.</p>
Leaf-BGP Autonomous System Number	<p>Specify the starting autonomous system (AS) number to be assigned to the first leaf device. Subsequent leaf devices are assigned incremental AS numbers starting from the number you specified.</p> <p>Network Director updates the last AS number based on the number of leaf devices that you plan to have in the fabric. You cannot modify the last AS number.</p>
Device Password	Specify the default password that you want to set for all the devices in the fabric.
Management Gateway	<p>If Network Director and the Layer 3 Fabric devices are in different subnets, specify the gateway IP address that these devices can use to reach Network Director.</p> <p>NOTE: This is an optional field if the Layer 3 Fabric and Network Director are in the same subnet.</p>

2. Do one of the following:

- Click **Next** to open the Cabling page where you can view the cabling plan for your Layer 3 Fabric. This might take some time depending on the fabric capacity.
- Click **Save & Exit** to save the changes and exit the Create Layer 3 Fabric wizard. You can resume this task later without losing any information that you entered.

Viewing the Cabling Plan

The Cabling Plan page displays the recommended cabling plan for the device that you select in the left pane. If you specify all the spine and leaf devices, the cabling plan displays the exact port numbers that you must use to connect your spine and leaf devices. However, if you have not specified any leaf devices and have only specified the maximum leaf count, the plan displays all the leaf devices as unknown. The leaf devices in this case are plug-and-play and you can use any of the uplink ports on your plug-and-play leaf device.

This holds good until you have reached the initial capacity of the spine devices. If you are adding an additional spine device, beyond the initial capacity, Network Director regenerates the cabling plan and you must follow the recommended cabling plan for all subsequent spine to leaf connections. Note that the connections to the existing devices need not be changed as part of this change.



Network Director regenerates the cabling plan, if one of the following occurs:

- A spine device is added
- A spine device is deleted
- A leaf device is added
- A leaf device is deleted

If the selected spine device model in the Fabric Requirements page is QFX10002-72Q, the cabling plan is represented as two chassis images. The first chassis image displays the connections for the ports in the first and second rows, and the second chassis image displays the connections for the ports in the third and fourth rows.

If the selected spine device model in the Fabric Requirements page is QFX10008, and selected line card model is QFX10000-60S-6Q in the Build New Chassis section, cabling plan is represented in two chassis images. The first image displays connections for the ports in first and third rows, and the second chassis image displays the connections for the ports in the middle row.

From the Cabling page, you can:

1. View the cable connectivity that you must follow for each device in your fabric. The device table is color coded to identify the devices that are provisioned now (identified by  color), the devices reserved for future (identified by  color), the Virtual Chassis connections (identified by green color), and access LAG ports (identified by pink color).
2. Click **Grid View** to view the cabling plan in a grid. Select a device in the left pane to view the cabling details and access LAG connections of the selected device.

NOTE: The Access LAG ports are displayed if you have not selected **Dynamically create LAG when hosts are connected** in the Fabric Requirement page. Network Director preprovisions the LAG configuration in the Virtual Chassis members.

3. Click **Graph View** to view the graphical representation of the cabling plan.
4. Do one of the following:
 - Click **Next** to open the ZTP page where you can specify the Zero Touch Provisioning (ZTP) details for the Layer 3 Fabric.

- Click **Save & Exit** to save the changes and exit the Create Layer 3 Fabric wizard. You can resume this task later without losing any information that you entered.

Specifying Zero Touch Provisioning Details

Zero Touch Provisioning (ZTP) enables you to provision devices in your network automatically, without manual intervention. When a device is physically connected, it boots up with factory-default configuration and auto installs a configuration file from the network. In Network Director, the ZTP is used to provision Layer 3 fabric and all the configurations are pushed through OpenClos. To specify the ZTP details:

NOTE: When you select QFX10008 as the spine model, only the leaf models are provisioned with ZTP configuration. For the spine model QFX10008, you must either copy the config file from Network Director or manually download it from the file server. To copy the config file from the file server, SSH or Telnet must be enabled on the device (QFX10008).

1. Specify the DHCP server settings by following the descriptions given in [Table 119 on page 505](#).

Table 119: DHCP Server Details

Field	Description
DHCP Server	IP address or the hostname of the DHCP server.
DHCP Server Type	<p>The type of DHCP server that provides the necessary information to the switch. You can choose to use a CentOS DHCP server, an Ubuntu DHCP server, or any other DHCP server.</p> <p>NOTE: If you select Other, you must configure the DHCP server settings manually.</p>
Manually Configure Server	<p>Select to indicate that you want to manually configure the DHCP server. You can configure the CentOS and Ubuntu DHCP servers manually or from Network Director.</p> <p>If you want to use any other type of DHCP server, do the following:</p> <ol style="list-style-type: none"> a. Select the Manually Configure Server check box. Network Director hides all the other details except the DHCP Server Type. b. Follow the instructions displayed in this box to configure the DHCP server manually.
DHCP User	Username to log in to the DHCP server.
DHCP Password	Password for the specified username.
Confirm Password	Confirm the DHCP server password.

NOTE: *When you are replacing a member device*—If the member that is replaced is up, Network Director obtains the latest configuration from the replaced device and maps this configuration to the corresponding MAC or serial number in the DHCP server.

However, if the member that is replaced is down, Network Director is not able to reach the device to get its latest configuration. In such case, Network Director maps the configuration that is generated from OpenClos (Stage-2 for leaf devices) for the replaced device to the MAC or serial number of the new device in the DHCP server. Note that the mapped configuration in the DHCP server does not have any configuration that is pushed from Network Director to the device.

NOTE: The DHCP server configuration file does not contain entries related to the spine device QFX10008 as the device does not go through ZTP.

2. Specify the File server settings by following the descriptions given in [Table 120 on page 506](#).

Table 120: File Server and Software Details

Field	Description
File Server Type	The type of file server where the software images are to be stored. You can choose to use an FTP, HTTP, or an TFTP file server.
File Server	IP address or hostname of the file server.
File Server Root Dir	The root directory of the file server.
Spine Image	<p>The software image file that you want to use for your spine devices.</p> <p>NOTE: Ensure that the software image is uploaded to Network Director using the Image Management > Manage Image Repository in the Deploy mode. Else Network Director does not display the software image.</p>
Leaf Image	<p>The software image file that you want to use for your leaf devices.</p> <p>NOTE: Ensure that the software image is uploaded to Network Director using the Image Management > Manage Image Repository in the Deploy mode. Else Network Director does not display the software image.</p> <p>As Network Director supports two device models—EX4300 and QFX5100—and their variants as leaf devices, you can specify a software image for each of these leaf devices irrespective of the variant that you have selected for your fabric. The same software image applies to all the variants of a device.</p>

3. ZTP process maps the management interface MAC address or the device chassis serial number of each spine device to the device-specific software image, IP address, hostname, and the configuration file stored on the file server. This mapping is stored in the DHCP server. When a spine device starts up, the device contacts the DHCP server to obtain the IP address and the software image location. The DHCP server looks up in its MAC address or serial number mapping database to identify the device and provide details about the file server that the device must contact to get the software image and configuration file. The device uses this information to contact the file server and obtain the software image and the configuration file for deploying on the device.

Do one of the following to specify the MAC address or the serial number of your spine devices:

- Enter the MAC address of the management interface (for example, the em0 interface) or the device chassis serial number of the spine devices in **MAC Addresses** or **Serial Number** in the table.
- Click Import MAC Address to import the MAC addresses of spine device in CSV format. You must enter the MAC addresses in the specified format. Click **Download CSV format** to download a sample CSV file that you can use to import MAC addresses.

NOTE: When you use the Import option, you must specify either the MAC address or the Serial number, but not both.

- Click **Import Serial Number** to import the serial numbers of spine device in CSV format. You must enter the MAC addresses in the specified format. Click **Download CSV format** to download a sample CSV file that you can use to import serial numbers.

NOTE: You can specify serial number only for spine devices running Junos OS Release 14.1X53D15 or later.

NOTE: Entering serial number or MAC address of the spine device is not applicable for the device model QFX10008 as it is not provisioned through ZTP.

4. To view the configuration that is deployed on a spine device, click **Actions > View Config**.
5. To view the configuration that is deployed initially on the leaf devices, click **View Leaf Config**.
6. Do one of the following:
 - Click **Next** to open the Review page where you can review the Layer 3 Fabric settings.
 - Click **Save & Exit** to save the changes and exit the Create Layer 3 Fabric wizard. You can resume this task later without losing any information that you entered.

Reviewing the Layer 3 Fabric Settings

From the Review page you can:

- View the DHCP configuration that will be deployed on to the DHCP server by clicking **View DHCP Config** in the **ZTP Settings** sub-tab.

The DHCP configuration opens in a new window.

- Review the Layer 3 Fabric settings in the Review page, Devices sub-tab and the Configuration sub-tab.

- Click **Deploy** to deploy the Layer 3 Fabric.
- Click **Save & Exit** to save the changes and exit the Create Layer 3 Fabric wizard. You can resume this task later without losing any information that you entered.

RELATED DOCUMENTATION

Understanding Layer 3 Fabrics 489
Managing Layer 3 Fabrics 491
Editing Layer 3 Fabrics 508
Network Director Documentation home page

Editing Layer 3 Fabrics

Network Director enables you to edit some of the settings for a Layer 3 Fabric after you deploy the fabric successfully.

To edit a Layer 3 Fabric:

1. From the Manage Layer 3 Fabric page, select the fabric that you want to edit and click **Edit**.

The Edit Layer 3 Fabric page opens. The Edit Layer 3 Fabric page has similar wizard pages as in the Create Layer 3 Fabric page. However, you will not be able to modify all the fields in these wizard pages while you are editing the settings of a fabric.

2. [Table 121 on page 508](#) lists the settings that you can edit and the wizard pages to which these settings belong.

Table 121: Layer 3 Fabric Settings that can be edited

Wizard page	Field Name	Action
Fabric Requirement	Description	<p>Modify the description to a description of your choice.</p> <p>NOTE: If you add a plug-and-play leaf device to the fabric, make sure that you modify the Description field. If you do not do this, the Cabling page might not update the cabling plan for that leaf in the graph and grid views.</p>

Table 121: Layer 3 Fabric Settings that can be edited (*continued*)

Wizard page	Field Name	Action
Devices	Add Spine	<p>To add a spine device:</p> <ol style="list-style-type: none"> Click Add Spine. The Select Devices window opens. Select a device that you want to add and click OK. Network Director adds the selected device to the device table and updates the status of the device as Added. In the ZTP wizard page, specify the MAC address or serial number of the added device.
Devices	Add Leaf (VC)	<p>To add a virtual chassis device as a leaf member:</p> <ol style="list-style-type: none"> Click Add Leaf (VC). The Select Devices window opens. Select a device that you want to add and click OK. Network Director adds the selected device to the device table and updates the status of the device as Added. You can specify the model for Virtual Chassis members. Once you have specified the model for the Virtual Chassis member, then the cabling plan changes, and Network Director displays a new cabling plan in the Cabling Plan page. Follow the cabling plan to physically connect the devices.
Devices	Remove	Select a device from the list and click Remove .

Table 121: Layer 3 Fabric Settings that can be edited (*continued*)

Wizard page	Field Name	Action
Devices	Replace	<p>To replace device:</p> <ol style="list-style-type: none"> Select the check box corresponding to device that you want to replace from the list and click Replace. Network Director updates the status of the device as Replaced. NOTE: You cannot replace an inactive leaf device. For spine devices and standalone leaf devices, in the ZTP wizard page, specify the MAC address or serial number of the replaced device. NOTE: While replacing a spine device, you can specify either the MAC address or the serial number. However, while replacing a leaf device, you can specify only the MAC address. <p>NOTE: When you are replacing a spine devices or a standalone leaf devices, if the device that is replaced is up, Network Director obtains the latest configuration from the replaced device and maps this configuration to the corresponding MAC address or serial number of the new device in the DHCP server.</p> <p>However, if the device is down, Network Director is unable to reach the device to get its latest configuration. In such case, Network Director maps the configuration that is generated from OpenClos (Stage-2 for leaf devices) for the new device to the MAC address or serial number of the device in the DHCP server. Note that the mapped configuration in the DHCP server will not have any configuration that is pushed from Network Director to the device.</p> <p>NOTE: You can replace only one member of a Virtual Chassis at a time. While replacing a Virtual Chassis member, you need not specify the MAC address or serial number of the device.</p>

Table 121: Layer 3 Fabric Settings that can be edited (*continued*)

Wizard page	Field Name	Action
ZTP	<i>Image Details and the Device Details</i>	<p>You can modify the software image and the Device details in the ZTP wizard page.</p> <p>New and replaced devices are listed in the Device Details section. You must specify the MAC address or serial number for these devices.</p> <p>NOTE: While replacing a spine device, you can specify either the MAC address or the serial number. However, while replacing a leaf device, you can specify only the MAC address.</p> <p>For more details about the fields in the ZTP wizard page, see “Configuring and Monitoring Zero Touch Provisioning” on page 650.</p>

3. Click **Deploy** to save and deploy your edits.

RELATED DOCUMENTATION

[Understanding Layer 3 Fabrics | 489](#)

[Managing Layer 3 Fabrics | 491](#)

[Network Director Documentation home page](#)

Viewing Layer 3 Fabric Connectivity

After you have set up and deployed a Layer 3 Fabric in Network Director, you can pictorially view the physical connectivity between the various devices in the fabric. This page displays the devices and their physical connectivity in the spine-and-leaf topology.


To view the connectivity between the devices:

1. Do one of the following:
 - While in the Logical, Location, Device, or Custom View, select the Layer 3 Fabric for which you want to view the connectivity details from the View pane and click **Connectivity > View Layer 3 Fabric Connectivity** from the Tasks pane.
 - Click **View Topology** in the Manage Layer 3 Fabric page.

- While in the topology view, zoom in to a rack that has a Layer 3 Fabric member device, select the device and click **View Layer 3 Fabric Connectivity** from the Task pane.

The Layer 3 Fabric Connectivity page opens.

2. You can perform the following tasks from the Layer 3 Fabric Connectivity page.

- Mouse over each entity to know more details about that entity.
- You can Zoom in or zoom out of the connectivity view by using the + and - buttons.
- View the faults and alarms on an entity by zooming in to the entity. Double-click the alarm or fault count to view more details about it in the Fault mode.
- Click  on a leaf device to expand and view the host machines that are connected to the device.

RELATED DOCUMENTATION

[Understanding Layer 3 Fabrics | 489](#)

[Network Director Documentation home page](#)

Performing Layer 3 Fabric Connectivity Checks

After you have deployed a Layer 3 Fabric, you can run connectivity checks to troubleshoot any connectivity issues in the fabric. You can initiate a connectivity check from the Manage Layer 3 Fabrics page. Network Director uses LLDP to check the connectivity of each device with the device's neighbor and compares it with the recommended cabling plan and reports the results in the Cabling Check Results page.

NOTE: Network Director performs connectivity check only for the devices that are scheduled for immediate deployment and not for devices that are reserved for the future.

The Cabling Check Results page contains three tabs—Unknown Devices, Devices with Cabling Faults, and Devices Connected Properly. Each of these tabs displays a graphical representation and a grid view of each leaf device with the corresponding details. To switch between the Graphical view and the Grid view, toggle the **Switch to Grid View** and **Switch to Graph View** buttons.

From the Cabling Check Results page, you can:

1. Click the **Unknown Devices** tab to view the devices that are part of the Layer 3 Fabric, but are not physically connected to the network and have not undergone zero touch provisioning. Click each leaf device listed in the Devices box to see a detailed view of the device with the connections as per the cabling plan.
2. Click the **Cabling Faults** tab to view the devices which did not adhere to the recommended cabling plan. This might be because of missing connections or additional connections. You can view the details section to identify the faulty connections for each device.
3. Click **Devices Connected Properly** tab to view details of devices that are connected as per the cabling plan. Click each device to see the connections.
4. Click **Export Results to PDF** to export the connectivity check results as a PDF file.
5. Click **OK** to close the Cabling Check Results page.

RELATED DOCUMENTATION

[Managing Layer 3 Fabrics | 491](#)

[Network Director Documentation home page](#)

Setting Up Virtual Chassis Fabrics

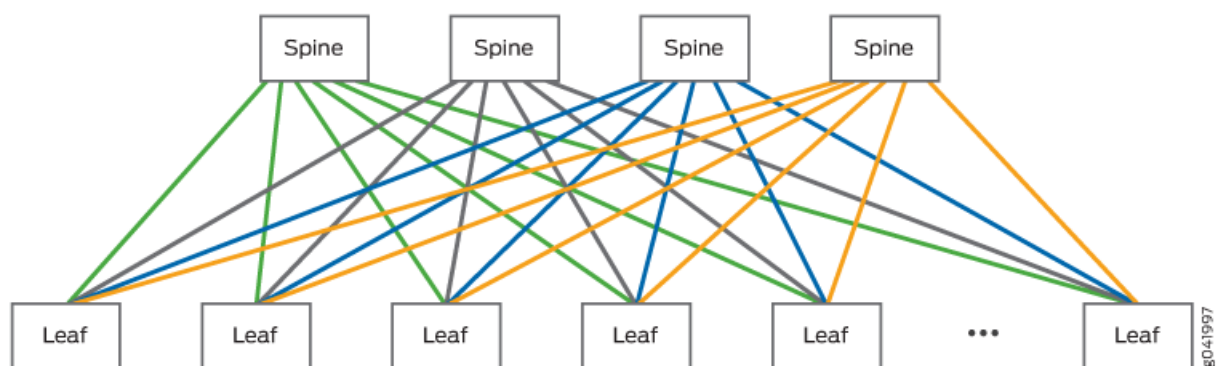
IN THIS SECTION

- [Selecting a Provisioning Method | 514](#)
- [Adding Devices | 515](#)
- [Specifying Software Packages | 519](#)
- [Reviewing the Virtual Chassis Fabric Summary and Deploying Changes | 520](#)

The Juniper Networks Virtual Chassis Fabric (VCF) provides a low-latency, high-performance fabric architecture that can be managed as a single device.

Network Director supports VCFs that are constructed using the spine-and-leaf architecture. In the spine-and-leaf architecture, each spine device is connected to one or more leaf devices. You can configure up to 20 QFX5100 switches in a VCF configuration, with two to four QFX5100s in the spine and up to 18 QFX5100s as leaf nodes. See [Figure 22 on page 514](#) for an illustration of the VCF spine-and-leaf architecture:

Figure 22: VCF Spine-and-Leaf Architecture



This topic describes how to set up a VCF using Network Director.

Before you setup a VCF, make sure at least one spine device is managed by Network Director. You can use this spine device to designate as a primary device in the following configuration.

To start setting up a VCF:

1. Select **Build** or **Deploy** in the Network Director banner.
2. Select **Network Builder > Set Up Virtual Chassis Fabric** in the Tasks pane.

The Select Provisioning Method page of the Setup Virtual Chassis Fabric wizard opens.

This topic describes:

Selecting a Provisioning Method

Use the Select Provisioning Method page to select the VCF provisioning method.

1. Select a provisioning method.

For a description of the fields on the Select Provisioning Method page, see [Table 122 on page 515](#).

Table 122: Select Provisioning Method Page

Field	Description
Select Provisioning Method Section	
Select the provisioning method.	
Auto Provision (Plug n Play)	<p>Selects automatic provisioning.</p> <p>You need to provision only the spine nodes when you use this method. All eligible nodes connected to spine nodes that are running the factory-default configuration will be added automatically to the VCF.</p>
Manual Provision	Selects manual provisioning. You must manually provision all VCF nodes when you use this method.
Enable Mixed Switch Mode	Selects a mixed mode VCF, which can contain a mix of device models. If a VCF is not in mixed mode, it can contain only QFX 5100 devices. In a spine-and-leaf mixed mode VCF, the spine nodes must be QFX 5100 devices.

2. Click **Next** to continue the wizard. The Add Devices page opens.

Adding Devices

Use the Add Devices page to assign devices to the VCF.

The following sections describe the methods of assigning devices to a VCF:

- [Assigning Devices to a Virtual Chassis Fabric by Using a Graph | 515](#)
- [Assigning Devices to a Virtual Chassis Fabric by Using a Grid | 518](#)

Assigning Devices to a Virtual Chassis Fabric by Using a Graph

On the graph, the VCF is represented by two sets of rectangles that represent devices. The top set represents spine devices, and the bottom set represents leaf devices.

Follow these steps to assign devices to a VCF by using a graph:

1. If the grid view is open, click **Switch to Graph View** on the Assign Devices page.
2. Click **Add Spines** to add spine devices. The Add Spine Switches to Virtual Chassis Fabric window opens.
3. To add devices that are managed by Network Director as spine nodes, click **Add Managed Devices as Spine Switches**.

If the VCF is in non-mixed mode, you can select only QFX 5100 devices to be members. Do the following for each device that you want to add in the Add Managed Swiches to Virtual Chassis Fabric window that opens:

- a. Select the check box next to the device.
- b. Select the node's role in the **Role** column.
- c. Select the star icon in the Designated Master column if you want to make it the Master Routing Engine node.

Only nodes with the role Routing Engine can be designated master.

- d. When you finish configuring nodes, click **OK**. The window closes and the added devices appear in the node list on the Add Devices as Spine Node(s) window.

4. To add devices that are not managed by Network Director as spine or Routing Engine nodes, click **New Devices**. Do the following in the Add New Devices window that opens:

- a. Enter the device serial number in the **Serial Number** field.
- b. Select the node's role from the **Role** list.
- c. To add another device, click **Add More** and repeat the steps to add more devices.
- d. When you are finished adding nodes, click **Add**. The window closes and the added devices appear in the node list on the Add Devices as Spine Node(s) window.

5. When you finish adding spine nodes, do the following in the Add Devices as Spine Node(s) window:

- To remove a node from the list, click the trash can icon at the end of the node's table row.
- When you finish configuring the list of spine nodes, click **Done**. The window closes and the spine nodes you configured appear in the Graph on the Setup Virtual Chassis Fabric page. The Designated Master Routing Engine node is identified with a star.

6. (Optional) If you selected the Auto Provisioning (Plug and Play) provisioning method, you can assign devices that are managed by Network Director to leaf nodes. Manually assigning a leaf node device causes the device to reset to the default configuration. When a device with the default configuration is connected to a VCF spine node, the device is added as a leaf node:

- a. Select the **Assign Leaf nodes manually** check box.
- b. Click a rectangle representing a leaf node to assign one or more devices to that role. The Add Device as Leaf Node(s) window opens.

- c. Click **Managed Devices**. The Select Managed Device(s) window opens. If the VCF is in non-mixed mode, you can select only QFX 5100 devices to be members.
- d. Select the check box for the device(s) you want to add as leaf nodes.
- e. Click **OK**. The window closes and the selected nodes are listed on the Add Device as Leaf Node(s) window.
- f. In the Add Device as Leaf Node(s) window, do the following:
 - To remove a node from the list, click the trash can icon at the end of the node's table row.
 - When you finish configuring the list of leaf nodes, click **Done**. The window closes and the leaf nodes you configured appear in the graph on the Setup Virtual Chassis Fabric page.
7. If you selected the Manual Provision provisioning method, click **Add Leaves** to add one or more devices as leaf nodes. The Add Device to Leaf Node(s) window opens.
8. To add devices that are managed by Network Director as leaf nodes:
 - a. Click **Managed Devices**. The Select Managed Devices window opens. If the VCF is in non-mixed mode, you can select only QFX 5100 devices to be members.
 - b. Select the check box next to the device(s) to add.
 - c. Click **OK**. The window closes and the added devices appear in the node list on the Add Devices as Leaf Node(s) window.
9. To add devices that are not managed by Network Director as leaf nodes:
 - a. Click **New Devices**. The Add New Devices window opens.
 - b. Enter the device serial number in the **Serial Number** field.
 - c. To add another device, click **Add More** and repeat the steps to add more devices.
 - d. When you are finished adding nodes, click **Add**. The window closes and the added devices appear in the node list on the Add Devices as Leaf Node(s) window.
10. When you finish adding leaf nodes, do the following in the Add Devices as Leaf Node(s) window:
 - To remove a node from the list, click the trash can icon at the end of the node's table row.

- When you finish configuring the list of leaf nodes, click **Done**. The window closes and the leaf nodes you configured appear in the graph on the Setup Virtual Chassis Fabric page.

11. Click **Next** to continue the wizard. The Specify Software Package page opens.

Continue with [“Specifying Software Packages” on page 519](#).

Assigning Devices to a Virtual Chassis Fabric by Using a Grid

This section describes how to assign devices to a VCF by using a grid.

To assign devices to a VCF by using a grid:

1. If the graph view is open, click **Switch to Grid View** on the Assign Devices page.
2. To add devices that are managed by Network Director as nodes, click **Managed Devices**. If the VCF is in non-mixed mode, you can select only QFX 5100 devices to be members. Do the following for each device that you want to add in the Select Managed Devices window that opens:
 - a. Select the check box next to the device.
 - b. Select the node's role in the **Role** column.
The options are Routing Engine or Line Card.
 - c. Select the node's type in the **Type** column.
The options are Spine or Leaf. Nodes with the role of Routing Engine must have the type Spine.
 - d. Select the star icon in the **Designated Master** column if you want to make the node the Master Routing Engine node.
Only nodes with the role of Routing Engine can be Master. If you select the Designated Master star for a node with the role of Line Card, the role is changed to Routing Engine, and if necessary, the type is changed to Spine.
 - e. When you are finished configuring nodes, click **OK**. The window closes and the added devices appear in the node list on the Add Devices page.
3. To add devices that are not managed by Network Director as nodes, click **New Devices**. Do the following in the Add New Devices window that opens:
 - a. Enter the device serial number in the **Serial Number** field.
 - b. Select the node's type from the **Type** list.
The type Leaf is available only if you selected the Manual provisioning method.

- c. Select the node's role from the **Role** list.

The role Routing Engine is applicable only to nodes of the type Spine. If you select Routing Engine, the type changes to Spine automatically.

- d. To add another device, click **Add More** and repeat the steps to add more devices.
 - e. When you are finished adding nodes, click **Add**. The window closes and the added devices appear in the node list on the Add Devices page.
4. When you finish adding nodes, you can remove a node from the list on the Add Devices page by clicking the trash can icon at the end of the node's table row.
 5. Click **Next** to continue the wizard. The Specify Software Package page opens.

Continue with ["Specifying Software Packages" on page 519](#).

Specifying Software Packages

You can configure a VCF to automatically upgrade new devices that are added to the VCF to a specified software package. For this functionality to work, you must specify the location of the software upgrade image files. Use the Specify Software Package page to specify the location of the image files for each device type that is included in the VCF. If the VCF is in non-mixed mode, only the QFX 5100 device type will appear on this page because it is the only device type allowed in the VCF.

To configure software packages on the VCF, do the following for each device type listed on the page:

1. Select the location of the software image by selecting a radio button:
 - Select **FTP Server** if the software image file is located on an FTP server.
 - Select **HTTP Server** if the software image file is located on an HTTP server.
 - Select **Device** if the software image file is located on the device.
 2. Enter the URL where the software image file is located in the text box:
 - If you selected **FTP Server** or **HTTP Server**, the text box is labeled Software Package URL. Enter the URL where the software image file is located on the server. Begin the URL with **ftp://** for an FTP server or **http://** for an HTTP server.
 - If you selected **Device**, the text box is labeled Software Package Pathname. Enter the pathname where the software image file is located on the device. Begin the pathname with a slash (/) to indicate the root directory.
 3. Click **Next** to continue the wizard.
- The Review page opens.

Reviewing the Virtual Chassis Fabric Summary and Deploying Changes

In the Review page, review the VCF setup. To make any changes, click **Back** to return to a previous step, or click the step in the wizard flowchart at the top of the page.

To deploy the VCF configuration, click **Deploy**.

When you deploy the VCF configuration, Network Director performs these major steps:

1. Configures the VCF on the Primary device, including the mixed/non-mixed mode setting, then reboots the Primary device.
2. Configures the VCF management interface (vme) based on the Primary device's management interface (for example, the em0 interface).
3. Provisions the VCF member devices on the Primary device.
4. Restarts all designated VCF devices, other than the Primary device, in factory-default mode.

The Primary device will remain a standalone switch in network Director until more than one VCF member device becomes operational, when it will become a VCF. Until this happens, you cannot edit the VCF using the Manage Virtual Chassis Fabric task. You can see the status of the deployment job. See [“Managing Configuration Deployment Jobs” on page 594](#) for information. If deployment to the VCF Primary device fails, the entire VCF deployment is cancelled.

Note these considerations when you create or manage a VCF using Network Director:

- If the management interface of a designated member is configured using DHCP, Network Director might lose connectivity with the member after successfully forming a VCF.
- If the em1 or em2 interface is used to manage a designated member, Network Director might lose connectivity with the member after successfully forming a VCF.

RELATED DOCUMENTATION

[Understanding Build Mode in Network Director | 77](#)

[Network Director Documentation home page](#)

Managing Virtual Chassis Fabrics

IN THIS SECTION

- [Adding Spine Device\(s\) to the Virtual Chassis Fabric | 521](#)
- [Adding Leaf Node\(s\) to the Virtual Chassis Fabric | 522](#)
- [Replacing Spine or Leaf Devices | 523](#)
- [Removing Spine or Leaf Devices | 524](#)
- [Configuring Software Package for Upgrade | 525](#)
- [Deploying VCF Configuration Changes | 525](#)

Network Director supports Virtual Chassis Fabrics (VCFs) that are constructed using the spine-and-leaf architecture. In the spine-and-leaf architecture, each spine device is connected to one or more leaf devices.

Network Director enables you to modify an existing VCF by using the Manage Virtual Chassis Fabric task. You can add, remove, or replace VCF members, and configure software packages for automatically upgrading new member devices. If the VCF is in non-mixed mode, you can add only QFX 5100 devices.

To modify a VCF:

1. In Build mode, select the VCF that you want to modify from the View pane and select **Device Management > Manage Virtual Chassis Fabric** from the Tasks pane.

The Virtual Chassis Topology page opens displaying the connectivity of the VCF in the spine-and-leaf architecture. The Designated Master node is marked with a gold star. The Backup Master node is marked with a silver star. Nodes that are not present, inactive, unprovisioned, or pre-provisioned appear in gray color. You cannot replace or remove such nodes.

2. Click **Edit VCF**. Network Director displays the Add Spines, Add Leaves, and Specify Software Package buttons.

You can perform the following tasks from the Virtual Chassis Topology page:

Adding Spine Device(s) to the Virtual Chassis Fabric

A VCF can contain up to 4 spine nodes and a total of 20 nodes.

To add a spine device:

1. Click **Add Spine**. The Add Device as Spine Node(s) window opens. If the VCF is in non-mixed mode, you can add only QFX 5100 devices, so only those devices are listed.
2. To add devices that are managed by Network Director as spine nodes, click **Managed Devices**. Do the following for each device that you want to add in the Select Managed Devices window that opens:
 - a. Select the check box next to the device.
 - b. Select the node's role in the **Role** column.
 - c. When you have selected all the nodes you want, click **OK**. The window closes and the added device appears in the node list on the Add Devices as Spine Node(s) window.
3. To add devices that are not managed by Network Director as spine nodes, click **New Devices**. Do the following in the Add New Devices window that opens:
 - a. Enter the device serial number in the **Serial Number** field.
 - b. Select the node's role from the **Role** list.
 - c. To add more devices, click **Add More** and repeat the steps to add more devices.
 - d. When you finish adding devices, click **Add**. The window closes and the added devices appear in the node list on the Add Devices as Spine Node(s) window.
4. When you finish adding spine nodes, do the following in the Add Devices as Spine Node(s) window:
 - To remove a node from the list, click the trash can icon at the end of the node's table row.
 - When you finish configuring the list of spine nodes, click **Done**. The window closes and the spine nodes you configured appear as spine devices in the Virtual Chassis Topology page. The Master Routing Engine node is identified with a gold star. The backup Routing Engine node is identified with a silver star.
5. Click **Deploy** to deploy the changes that you made to the VCF.

Adding Leaf Node(s) to the Virtual Chassis Fabric

A VCF can contain up to 4 spine nodes and a total of 20 nodes.


To add leaf nodes to the VCF:

1. Click **Add Leaves**. The Add Device as Leaf Node(s) window opens. If the VCF is in non-mixed mode, you can only add QFX 5100 devices, so only those devices are listed.
2. To add managed devices as leaf nodes, click **Managed Devices**:
 - a. Select the check box next to the devices you want to add.
 - b. Click **OK**. The window closes and the added devices appear in the node list on the Add Devices as Leaf Node(s) window.
3. To add a device that is not managed by Network Director as a leaf node, click **New Devices**. This option is not available if the VCF is configured for automatic provisioning. Do the following in the window that opens:
 - a. Enter the device serial number in the **Serial Number** field.
 - b. To add another device, click **Add More** and repeat the steps to add more devices.
 - c. When you finish adding devices, click **Add**. The window closes and the added devices appear in the node list on the Add Devices as Leaf Node(s) window.
4. When you finish adding leaf nodes, do the following in the Add Devices as Leaf Node(s) window:
 - To remove a node from the list, click the trash can icon at the end of the node's table row.
 - When you finish configuring the list of leaf nodes, click **Done**. The window closes and the leaf nodes you configured appear as leaf devices in the Virtual Chassis Topology page.
5. Click **Deploy** to deploy the changes that you made to the VCF.

Replacing Spine or Leaf Devices


Network Director enables you to replace a spine or leaf device with another VCF-compatible device. You cannot replace the Designated Master node.

To replace a spine or leaf device:

1. Mouse over the leaf or spine device that you want to replace and click . The Replace Spine Node or the Replace Leaf Node window opens depending on the type of device that you want to replace.
2. To replace a spine or the leaf node with a device that Network Director is managing, click **Managed Devices**. Do the following for the device that you want to replace in the Select Managed Devices window that opens:
 - a. Select the check box next to the device.
 - b. Select the node's role in the **Role** column.
 - c. Click **OK**. The window closes and the selected devices appear in the node list on the Replace Spine Node or the Replace Leaf Node window.
3. To replace a device that is not managed by Network Director as a spine or the leaf node, click **New Devices**. This option is not available if the VCF is configured for automatic provisioning. Do the following in the Add New Devices window that opens:
 - a. Enter the device serial number in the **Serial Number** field.
 - b. Click **Add**. The window closes and the added devices appear in the node list on the Replace Spine Node or the Replace Leaf Node window.
4. Click **Done**. The window closes and the nodes that you replaced appear in the Virtual Chassis Topology page, marked with a pencil icon.
5. Click **Deploy** to deploy the changes that you made to the VCF.

Removing Spine or Leaf Devices

Follow these steps to remove a spine or leaf device from the VCF. You cannot remove the Designated Primary device.

1. Mouse over the leaf or spine device that you want to replace and click . Network Director prompts you to confirm the deletion.
2. Click **Yes** to confirm or **No** to cancel the removal.
3. Click **Deploy** to deploy the changes that you made to the VCF.

Configuring Software Package for Upgrade

You can configure a VCF to automatically upgrade new devices that are added to the VCF to a specified software package. For this functionality to work, you must specify the location of the software upgrade image files. Use the Specify Software Package page to specify the location of the image files for each device type that is included in the VCF.

To configure software packages for upgrading VCF devices:

1. Click **Specify Software Package**. The Specify Software Package window opens.
2. To configure software image upgrades on the VCF, do the following for each device type listed on the page:
 - a. Select the location of the software image by selecting a radio button:
 - Select **FTP Server** if the software image file is located on an FTP server.
 - Select **HTTP Server** if the software image file is located on an HTTP server.
 - Select **Device** if the software image file is located on the device.
 - b. Enter the URL where the software image file is located in the text box:
 - If you selected **FTP Server** or **HTTP Server**, the text box is labeled Software Package URL. Enter the URL where the software image file is located on the server. Begin the URL with **ftp://** for an FTP server or **http://** for an HTTP server.
 - If you selected **Device**, the text box is labeled Software Package Pathname. Enter the pathname where the software image file is located on the device. Begin the pathname with a slash (/) to indicate the root directory.
3. Click **Done** to save the changes.
4. Click **Deploy** to deploy the changes that you made to the VCF.

Deploying VCF Configuration Changes

When you deploy the modified VCF configuration, Network Director performs these major steps:

1. Deletes the VCP ports that are connected to the members that are selected for deletion.
2. Recycles the deleted member.
3. Deploys the configuration to the VCF.
4. Restarts all newly added managed VCF devices in factory-default mode.

You can see the status of the deployment job. See [“Managing Configuration Deployment Jobs” on page 594](#) for information. If deployment to the VCF Primary device fails, the entire VCF deployment is cancelled.

RELATED DOCUMENTATION

[Setting Up Virtual Chassis Fabrics | 513](#)

[Understanding Build Mode in Network Director | 77](#)

[Network Director Documentation home page](#)

Configuring Overlay Networks and Tenants

IN THIS CHAPTER

- [VXLAN—EVPN Overlay Overview | 527](#)
- [Create a Layer 3 Fabric based Underlay Network | 528](#)
- [Creating and Managing Overlay Fabrics | 531](#)
- [Setting Up a VXLAN—EVPN-Based Data Center | 535](#)
- [Creating and Managing Tenants | 536](#)

VXLAN—EVPN Overlay Overview

Spanning Tree Protocol (STP), multichassis link aggregation group (MC-LAG), and Transparent Interconnection of Lots of Links (TRILL) were some of the commonly used technologies in traditional data centers. However, as the data centers started to grow exponentially, these technologies were not able to scale to meet the requirements of data centers. To cater to these requirements, data center administrators started using orchestration tools such as VMware vCenter, VMware vCenter with NSX, OpenStack, and OpenStack with NSX plug-in to orchestrate the networking needs of the tenants that a data center serves. This approach meant that the data center might require additional plug-ins to configure VLANs and gateways or that you might need to make changes to the physical network topology, to accommodate a new tenant. Most of the current day data centers handle multiple customer groups, organizations, or tenants that require a new data center architecture that decouples the underlay network from tenant overlay networks. A Layer 3 Fabric underlay coupled with a Virtual Extensible LAN (VXLAN)—Ethernet VPN (EVPN) overlay solution that uses bare metal servers and/or virtual servers, or both Network Director for management enables data center and cloud operators to deploy much larger networks than that are otherwise possible with traditional Layer 2 Ethernet-based architectures.

Some of the major advantages that VXLAN—EVPN overlay networks provide are:

- **Scalability**—Most enterprises accommodate their growth by increasing the use of cloud services, while others choose to deploy their own private and hybrid clouds. Service providers also must be able to grow rapidly to have sufficient capacity to meet the demands of the enterprises. Today's networks are often too rigid and difficult to change for scaling to meet the needs of large enterprises and service providers. But by using VXLAN—EVPN, data centers and cloud operators can have up to 16 million overlay networks in a cloud data center.

- **Operational efficiency**—As enterprises expand geographically, the physical distance between the data centers and users also increases, which makes timely maintenance and application mobility a challenge. The VXLAN—EVPN solution enables network administrators to easily migrate applications within the data center and between data centers for business continuity so that they can maintain the data center without downtime, for effective load balancing.
- **High Performance**—End users often experience poor response times and even outages of business-critical applications caused by bandwidth limitations and latency problems. Multi-pathing and control plane learning features, that are part of the VXLAN—EVPN solution, can optimize network traffic flows, rein in network faults, and ensure maximum utilization of bandwidth.

RELATED DOCUMENTATION

[Setting Up a VXLAN—EVPN-Based Data Center | 535](#)

[Creating and Managing Tenants | 536](#)

Create a Layer 3 Fabric based Underlay Network

Layer 3 Fabrics based on multi-stage Clos architecture increases resiliency and also supports new technologies such as VMware NSX that allow enterprises to deploy applications, servers, and virtual networks within seconds. The Layer 3 Fabrics can also function as an underlay network for a EVPN—VXLAN based overlay networks. Layer 3 Fabrics use BGP as the control plane protocol to advertise prefixes, perform traffic engineering, and tag traffic. Layer 3 Fabric uses spine-and-leaf topology. In the spine-and-leaf topology, all the leaf devices are connected to the spine devices in a mesh.

You can use one of the following device models as the spine device in a Layer 3 Fabric:

- QFX10008
- QFX10002-36Q
- QFX10002-72Q
- QFX5100-24Q-2P
- QFX5200-32C-32Q

You can use the following device models as leaf devices:

- QFX5100-48S-6Q
- QFX5100-96S-8Q
- QFX5100-48T-6Q

- QFX5200-32C-32Q
- QFX5100-24Q
- QFX5100-24Q-2P

NOTE: QFX5100-48T-6Q, QFX5100-48S-6Q and QFX5100-96S-8Q can be standalone or Virtual Chassis leaf devices. Whereas, QFX10008 and QFX5100-24Q are supported only as standalone devices. QFX5200-32C-32Q can be supported as both spine and leaf device, however, it cannot be a Virtual Chassis member.

You can quickly create and deploy Layer 3 Fabrics by using Network Director.

Before you begin, ensure that you have the necessary privileges on the FTP and the file server that Network Director uses for Zero Touch Provisioning. For more details, see [“User Privileges Required for the DHCP and File Server While Using Zero Touch Provisioning” on page 490](#).

To create and deploy a Layer 3 Fabric using Network Director:

1. With Build mode or Deploy mode selected, select one of these options under Views: **Logical View**, **Location View**, **Device View**, or **Custom Group View**.
2. In the Tasks pane, click **Network Builder > Manage Layer 3 Fabrics**.

The Manage Layer 3 Fabrics page opens.

3. Click **Create** to create a new Layer 3 Fabric. [Table 123 on page 529](#) describes the various pages that are available in the Layer 3 Fabric wizard and the actions that you must perform on each page.

Table 123: Layer 3 Fabric Wizard Pages

Name of the wizard page	Action
Fabric Requirements	Specify the fabric name, spine and leaf device models, number of spine and leaf devices, and settings for dynamic host link aggregation group (LAG) creation.
Device Details	Displays the number of spine and leaf devices that you are provisioning as part of the initial capacity, enables you to edit the hostname prefix for all the spine and leaf devices.
Configuration	Specify the configuration details such as the Loopback Network Address, Interconnect Network Address, VLAN Network Address, Start Management IP, Maximum Hosts or VMs per leaf, Spine-BGP Autonomous System Number, Leaf-BGP Autonomous System Number, Device Password, and Management Gateway IP address.

Table 123: Layer 3 Fabric Wizard Pages (*continued*)

Name of the wizard page	Action
Cabling	<p>Displays the recommended cabling plan for the device that you select in the left pane. If you specify all the spine and leaf devices, the cabling plan displays the exact port numbers that you must use to connect your spine and leaf devices. However, if you have not specified any leaf devices and have specified only the maximum leaf count, the plan displays all the leaf devices as unknown. The leaf devices in this case are plug-and-play and you can use any of the uplink ports on the leaf devices.</p> <p><i>You must follow the cabling plan for all the physical connections between the spine and leaf devices in your fabric.</i></p>
ZTP	<p>Zero Touch Provisioning (ZTP) enables you to provision devices in your network automatically, without manual intervention. You can use the ZTP wizard page that is part of the Layer 3 Fabric wizard. When a device is physically connected, it boots up with the factory-default configuration and auto installs a configuration file from the network. In Network Director, ZTP is used to provision Layer 3 Fabrics and all the configurations are pushed through OpenClos.</p> <p>NOTE: When you select QFX10008 as the spine device, only the leaf devices in the fabric are provisioned using ZTP. For the QFX10008 spine device, you must either copy the configuration file from Network Director or manually download it from the file server. To copy the configuration file from the file server, SSH or Telnet must be enabled on the device (QFX10008).</p>
Review	Review the Layer 3 Fabric configuration, deploy the fabric, or save and close the Layer 3 Fabric wizard.

For detailed steps for creating the Layer 3 Fabric, see [“Creating Layer 3 Fabrics” on page 493](#).

- After you have deployed the Layer 3 Fabric, you can run connectivity checks to troubleshoot any connectivity issues in the fabric. You can initiate a connectivity check by clicking **Run Connectivity Check** or **Re-run Connectivity Check** in the Cabling column in the Manage Layer 3 Fabrics page.

Network Director uses LLDP to check the connectivity of each device with the device’s neighbor and compares it with the recommended cabling plan and reports the results in the Cabling Check Results page. For more details, see [“Performing Layer 3 Fabric Connectivity Checks” on page 512](#).

5. Click **View Topology** in the Summary column of the Manage Layer 3 Fabric page to view the physical topology of the Layer 3 Fabric.
6. After you have successfully deployed a Layer 3 Fabric, you can monitor the performance, identify errors or bottlenecks in your underlay network by using the Monitoring mode of Network Director. For details, see [“Status Monitor for Layer 3 Fabrics” on page 748](#) and [“Port Status for IP Fabric Monitor” on page 743](#).

Creating and Managing Overlay Fabrics

IN THIS SECTION

- [Creating an Overlay Fabric | 532](#)
- [Modifying an Overlay Fabric | 533](#)
- [Viewing Overlay Fabric Configuration Details | 534](#)

You can create and manage overlay fabrics from the Manage Overlay Fabrics page.

Before you begin, make sure that you have at least one Layer 3 Fabric based data center in Network Director.

NOTE: Layer 3 Fabric configured using CLI and discovered in Network Director is not used to create an overlay fabric. You can use the Layer 3 fabric created in Network Director to create an overlay fabric.

To open the Manage Overlay Fabrics page:

1. Open the Datacenter View in Network Director.
2. Select the Layer 3 Fabric based data center from the View pane and select **Overlay Networks Builder > Manage Overlay Fabrics** from the Tasks pane.

The Manage Overlay Fabrics page opens. [Table 124 on page 532](#) describes the fields that are displayed in the Manage Overlay Fabrics page.

Table 124: Manage Overlay Fabrics Field Descriptions

Field	Description
Name	Name of the overlay fabric.
Description	Description for the overlay fabric.
Autonomous System No.	<p>The autonomous system number given for the overlay fabric.</p> <p>Autonomous system (AS) is a collection of devices that act as a single administrative entity, such as an overlay fabric. Each autonomous system is assigned a unique autonomous system number. In BGP routing, AS numbers are used to uniquely identify a network on the Internet.</p>
Route Reflector Network Address	<p>IP address of the route reflector device.</p> <p>Route reflector devices have the ability to readvertise routes learned from an internal peer to other internal peers. Therefore, rather than all internal peers of the fabric to be fully meshed with each other, route reflection requires that only the route reflector be fully meshed with all internal peers. The route reflector and all of its internal peers form a cluster.</p>
Layer 3 Fabrics	The Layer 3 Fabrics that is part of the overlay fabric.
Deployment State	<p>Deployment status of the overlay fabric. Deployment state can be one of the following:</p> <ul style="list-style-type: none"> • Pending Deployment—The fabric details have been saved, but not deployed. • Deployment in progress—The fabric is being deployed on the Layer 3 Fabric. • Partially deployed—The overlay fabric is successfully deployed on some devices, and the deployment failed on the other devices. • Deployed—The overlay fabric is successfully deployed on all the devices. • Deploy failed—Deployment failed on all devices. • Pending decommission—The overlay fabric is deleted. There might be a slight delay for Network Director to apply the changes to the devices; the system displays this status during this time.

From the Manage Overlay Fabrics page, you can:

Creating an Overlay Fabric

You can create an overlay fabric by following the steps given in this section. Alternatively, you can create an overlay fabric inline while you create a tenant.

To create an overlay fabric:

1. Click **Add** in the Manage Overlay Fabrics page.

The Create Overlay Fabric window appears.

2. Enter a name and description for the overlay fabric.
3. Specify the autonomous system (AS) number for the overlay fabric. You can specify the AS number in the range 1 to 4294967295.
4. Specify the IP address and the subnet mask of the route reflector device in the overlay fabric.
5. Select the Layer 3 Fabric that you want to include in the overlay fabric from the **Available** list and click the arrow to add it to the overlay network. Network Director lists all the Layer 3 Fabrics that are part of the data center and available for assignment in the Available list.

NOTE: You must have at least one Layer 3 Fabric available when you create an overlay fabric. To create a Layer 3 Fabric, follow the procedure given in [“Creating Layer 3 Fabrics” on page 493](#).

NOTE: Layer 3 Fabric configured using CLI and discovered in Network Director is not used to create an overlay fabric. You can use the Layer 3 fabric created in Network Director to create an overlay fabric.

6. Click **OK** to create the overlay fabric.
7. Click **Preview & Deploy** to deploy the overlay configuration.

After you have created or deployed the overlay configuration, you can create a tenant. See [“Creating and Managing Tenants” on page 536](#) for the procedure to create a tenant.

Modifying an Overlay Fabric

To modify an overlay fabric:

1. Select the overlay fabric in the Manage Overlay Fabrics page and click **Edit**. The Edit Overlay Fabric page appears.

2. Modify the overlay fabric details. You can modify the Description, Autonomous System Number, Route Reflector Network Address, and the Layer 3 Fabric. By modifying Layer 3 Fabric, you can assign a new Layer 3 Fabric to the overlay fabric or delete a Layer 3 Fabric from the overlay fabric.

NOTE: You cannot delete a Layer 3 Fabric that has tenants created in the overlay network. You must remove all the tenants from the Overlay Fabric before you can delete the Layer 3 Fabric.

Viewing Overlay Fabric Configuration Details

After you have successfully deployed an overlay fabric, you can view the configuration details of the overlay fabric from the View Overlay Fabric Configuration page. Click the **Deployed**, **Partially Deployed**, or the **Deploy Failed** link corresponding to the overlay fabric under the Deployment State column in the Manage Overlay Fabrics page to open this page.

The View Overlay Fabric Configuration page displays configuration details as described in [Table 125 on page 534](#).

Table 125: View Overlay Fabric Configuration Details

Field	Description
Overlay Fabric Name	Name of the overlay fabric for which the configuration details are being viewed.
Description	Description of the overlay fabric.
Devices	Device list mapped to the overlay fabric.
Element Name	Name of the instance mapped to the overlay fabric.
Configurations	Configuration details of the device for the element (for example, Overlay Fabric).
Status	Deployment status of the overlay fabric.
Error Message	Error message for a failed deployment. NOTE: The error message is visible when the user clicks the hyperlink under the status column.

RELATED DOCUMENTATION

Setting Up a VXLAN—EVPN-Based Data Center

Junos Space Network Director enables you to set up a VXLAN—EVPN based data center network by performing a few easy tasks. After you have created a VXLAN—EVPN based data center, you can monitor the performance of your data center, find bottlenecks in your data center, and troubleshoot these issues.

Setting Up a VXLAN—EVPN based data center using Network Director involves the following steps:

1. Create a Layer 3 Fabric (s) using Network Director. To know more about creating Layer 3 Fabrics, see [“Create a Layer 3 Fabric based Underlay Network” on page 528](#).

NOTE: Layer 3 Fabric configured using CLI and discovered in Network Director is not used to create an overlay fabric.

2. Create a Data Center. For more information, see *Creating Data Centers Using Network Director*.
3. Assign the Layer 3 Fabric or Fabrics to the data center. For detailed steps on assigning a fabric to a data center, see Assigning Network Devices to a Data Center section in *Creating Data Centers Using Network Director*.
4. Select the data center and create an overlay fabric from the Manage Overlay Fabrics page. You can add one or more Layer 3 Fabrics to an overlay fabric. For detailed steps, see [“Creating and Managing Overlay Fabrics” on page 531](#).

NOTE: You can skip this step and choose to create an overlay fabric when you create the tenant.

5. Create a tenant using the Create Tenant task under Manage Tenants and assign the tenant to an overlay fabric. You can either create an overlay fabric inline while creating a tenant or choose an existing overlay fabric. As part of the create tenant task, you also create one or more overlay networks, and select ports and multi-homed LAG to be part of this overlay network. You can also create new LAGs as part of this task. For detailed steps, see [“Creating and Managing Tenants” on page 536](#).

6. Preview and deploy the tenant to the data center.
7. After you have successfully deployed a tenant, you can monitor the performance and identify errors or bottlenecks in your overlay network by using the tenant monitor. For detailed steps on monitoring tenant details, see [“Monitoring Tenant Details” on page 691](#).

RELATED DOCUMENTATION

[VXLAN—EVPN Overlay Overview | 527](#)

[Creating and Managing Tenants | 536](#)

Creating and Managing Tenants

IN THIS SECTION

- [Creating a Tenant | 537](#)
- [Modifying a Tenant | 540](#)
- [Viewing Tenant Configuration Details | 541](#)

You can create and manage tenants from the Manage Tenants page.

Before you begin, make sure that you have at least one Layer 3 Fabric based data center in Network Director.

To open the Manage Tenants page:

1. Open the Datacenter View in Network Director.
2. Select a data center from the **View** pane for which you want to create a tenant and select **Overlay Networks Builder > Manage Tenants** from the **Tasks** pane.

The Manage Tenants page opens. [Table 126 on page 537](#) describes the fields that are displayed in the Manage Tenants page.

Table 126: Manage Tenants Field Descriptions

Field	Description
Tenant Name	Name of the tenant.
Description	Description given for the tenant.
Overlay Name	Name of the overlay fabric.
Deployment State	<p>Deployment status of the tenant. Deployment state can be one of the following:</p> <ul style="list-style-type: none"> • Pending Deployment—The tenant details have been saved, but not deployed. • Deployment in progress—The tenant is being deployed on the Layer 3 Fabric. • Partially deployed—The tenant is successfully deployed on some devices and the deployment failed on the other devices. • Deployed—The tenant is successfully deployed on all devices. • Deploy failed—Deployment failed on all the devices. • Pending decommission—The tenant is deleted. There might be a slight delay for Network Director to apply the changes to the devices; the system displays this status during this time.

You can perform the following tasks from the Manage Tenants page:

Creating a Tenant

To create a tenant:

1. Click **Add** in the Manage Tenants page.

The Create Tenant page opens.

2. Enter a name and description for the tenant.

You can enter a maximum of 255 characters. Underscore is the only special character that is allowed in the name field.

3. Click **Overlay Fabric Name** to select an overlay fabric for the tenant.

You can also create an overlay fabric inline by clicking **Create New**.

To create the overlay fabric:

- a. Enter a name and description for the overlay fabric.
- b. Specify the autonomous system (AS) number for the overlay fabric.

- c. Specify the IP address and the subnet mask of the route reflector network address in the overlay fabric.
- d. Select the Layer 3 Fabric that you want to include in the overlay network from the Available list and click the arrow to add it to the overlay fabric. Network Director lists all the Layer 3 Fabrics that are part of the data center and available for assignment in the Available list.
- e. Click **OK** to create the overlay fabric. The overlay fabric is selected to be used by the tenant.

4. Click **Add** in the Overlay Networks table to create an overlay network for the tenant.

The Create Overlay Network for Tenant window opens.

5. Enter the Virtual Extensible LAN Network Identifier (VNI) for the VXLAN overlay network. You can also select a VLAN from the list of existing VLANs.

VNI is a numeric value to identify a VXLAN. You can enter a number from 1 through 16777214.

6. Create a VRF instance that is to be assigned to the VXLAN overlay network or use the default VRF instance. You can also select a VRF instance from the list of existing VRF instances.

The default VRF instance is named as **default_tenant-name**.

Each tenant in an overlay fabric has its own tenant-specific routing table that contains the routing information for that tenant. To separate the route of each tenant from the other network traffic and routes, the spine device in the Layer 3 Fabric creates a separate routing table for each tenant called the Virtual Routing and Forwarding (VRF) instance. A can tenant can have multiple VRF instances.

To create a new VRF instance, click **Create**. Specify a name, description, and loopback network address for the VRF instance and click **OK**.

NOTE: The VRF loopback network address is a mandatory field for routing devices. You can modify the value in this field, but cannot delete it once it is deployed.

NOTE: If you have mapped a VRF instance to an overlay network, for example, overlay-network1 is mapped to VRF1, you cannot modify the VRF instance for this overlay network (after it is deployed) to another VRF instance, for example, to VRF2. To map the new instance, you must delete the overlay network and re-create the overlay network.

7. Specify the ID and name of the VLAN. Network Director maps this VLAN ID to the VXLAN overlay network.

The VLAN name is auto-populated, which you can edit, if needed.

8. Specify the IP address and subnet mask of the default layer 3 gateway network address. The default gateway enables the traffic to be routed between different VXLANs networks within the VRFs of the tenant.
9. Click **Auto Select Ports** to list the ports or LAGs that are connected to the devices. You can select these ports or LAG and assign it to the overlay network.
10. Click the **Ports Selection** tab to select the ports that you want to be part of the overlay network.

The Ports Selection tab displays the leaf devices of Layer 3 Fabrics that are part of the overlay fabric in the data center. When you select one or more leaf devices in the left pane, the ports associated with the device are displayed in a table in the right pane.

11. To preview a port selection, select a port and click **Preview**. You can edit the port selection listed in the preview by clicking **Edit**.
12. Select a leaf device or devices and click **Select by Range** to specify the port range that you want to include in the overlay network. You can select the port type as **Normal Ports** or **Channelized Ports**.

The normal ports are physical ports such as et, ge, and xe. Channelized ports are et ports that are 40-Gigabit, 100-Gigabit and can be converted to xe ports. For example, et-0/0/0 (40-Gigabit) can be converted to xe-0/0/0:0, xe-0/0/0:1, xe-0/0/0:2, xe-0/0/0:3. Similarly a 10-Gigabit and 100-Gigabit can be converted to 25-Gigabit.

NOTE: You can also select multiple devices from the left pane by the Shift key pressed and selecting devices. All the ports for selected devices are listed.

13. Click the **LAG Selection (Multi-homed)** tab to select the link aggregation groups (LAGs) that you want to be part of the overlay network. You can select one or more LAGs from the existing LAGs that are created in Network Director or create and add new LAGs.

NOTE: The lag you created will have a temporary system generated name with the prefix LAG; for example, LAG0 or LAG1. When you save or deploy a tenant, the available aggregate Ethernet interface name for example, ae0 or ae1 replaces this LAG name. The LAG name must be common for all the leaf devices across a host.

14. To select from a list of existing LAGs, click **Select Existing LAG**. The LAG Selection (multi-homed) window opens. Select the LAGs that you want to include in the overlay network and click **OK**. Network Director adds the LAGs that you selected to the LAG Selection (multi-homed) table.
15. Do the following to create a new LAG:
 - a. Click **Add** in the LAG Selection (multi-homed) table. The Create LAG window opens.
 - b. Select the Layer 3 Fabric on which you want to create the LAG.
 - c. Click **Add** to specify the device and port details for the LAG. Network Director adds a row to the Device and Port details table.
 - d. In the Device column, click and select the leaf device from which you want to select a port for the LAG.
 - e. In the Port column, click to select the ports that you want to add to the LAG.
 - f. Repeat steps c through e to add ports from another leaf device.
The Port column lists only those ports that are not part of any LAG or port selection tab.
 - g. To remove a device and port entry from the LAG, select the row from the Device and Port details table and click **Remove**.
 - h. Click **OK** to save the details and create the LAG. Network Director adds the LAG that you created, to the LAG Selection (multi-homed) table.
16. To delete a LAG from the overlay network, select the LAG in the LAG Selection (multi-homed) table and click **Delete**.
17. Click **OK** to save the details and add the overlay network for the tenant. Network Director lists the overlay network that you added in the Overlay Networks table in the Create Tenant page.
18. Repeat steps 4 through 17 to add more overlay networks.
19. Click **Preview & Deploy** to review the changes and deploy the tenant details to the overlay fabric.

Modifying a Tenant

To modify a tenant:

1. Select a tenant in the Manage Tenants page and click **Edit**.

The Edit Tenant page appears.

2. You can modify the tenant description, overlay network and its details, and the VRF details. You can also add a new overlay network by clicking **Add**.
3. To modify the overlay network details, select the overlay network you want to modify and click **Edit**. The Edit Overlay Network for Tenant page opens.
4. You can modify VNI, VLAN ID, L3 Gateway Network Address, VLAN Name, Port Selection, and LAG Selection (multi-homed) details and click **OK**. For details, see [“Creating a Tenant” on page 537](#).

NOTE: You cannot change a VLAN name after an overlay network is deployed.

5. To edit the VRF instance details and the loopback network address, you can either:

- Click the VRF instance name in the table.
- Select an overlay network and click **Edit VRF Details**.

The Edit VRF Instance for tenant page appears.

6. Modify the description and the loopback network address and click **OK**.

NOTE: You can create a VRF instance while you create the overlay network (this VRF instance has the tenant scope). However, you cannot delete it directly as it might be used by other overlay networks. Therefore, to enable deleting make sure you do not map this VRF instance to the VLAN while creating the tenant. You can however, edit the details of a VRF instance. If you want to delete a VRF instance, you must first delete the deployed VLANs, to which the VRF instance is associated. After you delete these VLANs, the VRF instance is removed after you save and deploy the configuration.

Viewing Tenant Configuration Details

After you successfully deploy a tenant, you can view the configuration details of the tenant from the View Tenant Configuration page. Click the **Deployed** link corresponding to the tenant name under the Deployment State column in the Manage Tenant Fabrics landing page to open this page.

The View Configuration page displays the details as described in [Table 127 on page 542](#).

Table 127: View Tenant Configuration Details

Field	Description
Tenant Name	Name of the overlay fabric for which the configuration details are being viewed.
Description	Description of the overlay fabric.
Overlay Fabric Name	Device list that is mapped to the selected overlay fabric.
View VNI	<p>Click View VNI to view the VNI details of the overlay network that are mapped to the tenant. You can view the following details:</p> <ul style="list-style-type: none"> • Ports selection, that are part of the overlay network. • Lag selection (multi-homed), that are part of the overlay network
View VRF	<p>Click View VRF, to view the VRF instance configuration details that are mapped to the tenant. You can view the following details:</p> <ul style="list-style-type: none"> • Devices—Devices that are mapped to the tenant. • Element—Name of the VRF instance mapped to the tenant. • Configurations—Configuration details of the VRF instance for the selected device. • Status—Deployment status of the tenant. • Message—The error message if any. <p>NOTE: To go back to the Overlay Network table click < Back to Overlay Networks link.</p>
VNI	VXLAN Network identifier of the tenant.
VLAN ID	VLAN ID of the tenant.
VLAN Name	VLAN name of the tenant.
L3 Gateway Network Address	IP address and subnet mask of the default L3 gateway.
Ports	Port that is part of the overlay network.
LAGs	Lag that is part of the overlay network
VRF Instance	VRF instance that is assigned to the VXLAN overlay network.
Status	Deployment status of the tenant.

RELATED DOCUMENTATION

| [VXLAN—EVPN Overlay Overview](#) | 527

Configuring VRRP Profiles

IN THIS CHAPTER


- [Understanding VRRP Profiles | 544](#)
- [Creating and Managing VRRP Profiles | 545](#)

Understanding VRRP Profiles

Virtual Router Redundancy Protocol (VRRP) enables hosts on a LAN to make use of redundant routing devices on that LAN without requiring more than the static configuration of a single default route on the hosts. The routing device on which VRRP is enabled share the IP address corresponding to the default route configured on the hosts. At any time, one of the routing devices is the primary (active) and the others are backups. If the primary fails, one of the backup routers becomes the new primary, providing a virtual default routing platform and enabling traffic on the LAN to be routed without relying on a single routing device. Using VRRP, a backup routing device can take over a failed primary router within a few seconds and without any interaction with the hosts.

Routing devices on which VRRP is enabled dynamically elect the primary and backup devices. You can also configure the assignment of the primary and the backup routers by specifying the priorities from 1 through 255 for primary-role election, with 255 being the highest priority. VRRP functions by the default primary sending advertisements to the backup devices at regular intervals. The default interval is 1 second, but you can set this interval. If a backup device does not receive an advertisement for the set period, the backup device with the next highest priority takes over as primary and begins forwarding packets. To minimize network traffic, VRRP is designed in such a way that only the device that is acting as the primary sends out VRRP advertisements at any given point in time. The backup devices do not send any advertisement until and unless they take over as the primary.

The following figure illustrates a basic VRRP topology. In this example, routers A, B, and C are running VRRP and together they function as a virtual router. The IP address of this virtual router is 10.10.0.1 (the same address as the physical interface of Router A).



Because the virtual router uses the IP address of the physical interface of router A, router A is the primary router, while routers B and C function as backup VRRP routers. Clients 1 through 3 are configured with the default gateway IP address of 10.10.0.1. As the primary router, router A forwards packets sent to its IP address. If the primary virtual router fails, the backup router configured with the higher priority becomes the primary virtual router and provides uninterrupted service for the LAN hosts. When router A recovers, it becomes the primary virtual router again.

RELATED DOCUMENTATION

[Creating and Managing VRRP Profiles | 545](#)

[Creating and Managing Port Profiles | 251](#)

[Creating and Managing VLAN Profiles | 329](#)

Creating and Managing VRRP Profiles

IN THIS SECTION

- [Managing VRRP Profiles | 546](#)
- [Creating VRRP Profiles | 547](#)
- [Specifying VRRP Settings for an EX Switching or Campus Switching ELS or Data Center Switching ELS | 547](#)

VRRP profiles enable grouping of VRRP parameters and applying them to one or more interfaces. You can configure the attributes for this profile by using the VRRP option under Profiles. You can also choose this profile as an in-line profile in a Port profile and a VLAN profile.

- VRRP on Port profile—Select VRRP in Port profile if you want to configure VRRP on a physical interface. The VRRP settings in Port profile are displayed only when you select the Service Type as Custom and Family Type as Routing. The VRRP attributes such as group ID and priority are applied to the device during the profile assignment.
- VRRP on VLAN profile—Select the VRRP in VLAN profile if you want to configure VRRP on an integrated routing and bridging (IRB) interface. The VRRP attributes such as group ID and priority are applied to the device during the profile assignment.

This topic describes:

Managing VRRP Profiles

From the Manage VRRP Profiles page, you can:

- Create a new profile by clicking **Add**.
- Modify an existing profile by selecting it and clicking **Edit**.
- View information about a profile by selecting the group and clicking **Details** or by clicking the profile name.
- Clone a profile by selecting a profile and clicking **Clone**.
- Delete profiles by selecting the profiles and clicking **Delete**.

TIP: You cannot delete profiles that are in use—that is, assigned to objects or used by other profiles. To see the current assignments for a profile, click the profile name

The following table describes the information provided about VRRP profiles on the Manage VRRP Profiles page. This page lists all VRRP profiles defined for your network, regardless of your current selected scope in Network view.

Table 128: Managing Profiles


Field	Description
Profile Name	Name given to the profile when the profile was created.
Description	<p>Description of the profile that was entered when the profile was created. If the profile was created by using the CLI and then discovered by Network Director, the description is: <i>Profile created as part of device discovery</i>.</p> <p>NOTE: To display the entire description, you might need to resize the Description column by clicking the column border in the heading and dragging it.</p>

Table 128: Managing Profiles (continued)

Field	Description
Family Type	The device family on which the profile was created: EX Series Switches or Campus Switching ELS .

Creating VRRP Profiles

To create VRRP profiles for EX Series switches or Campus Switching ELS:

1. Click  in the Network Director banner.
The Manage VRRP Profiles page opens.
2. Under Views, select one of the following views: **Logical View**, **Location View**, **Device View**, or **Custom Group View**.
3. Under **Tasks**, expand **Wired** and click **VRRP**.
The Manage VRRP Profiles page opens.
4. Click **Add**.
The Device Family Chooser appears.
5. Select **Switching (EX)** or **Campus Switching ELS**.
The Create VRRP Profile page appears for the selected family with the appropriate fields for configuring that family.

Specifying VRRP Settings for an EX Switching or Campus Switching ELS or Data Center Switching ELS

Use the Create VRRP Profile page to define a common set of VRRP attributes, which you can then apply to a group of interfaces. These directions address creating a VRRP profile for EX Series switches.

Table 129: VRRP Profile Settings

Field	Action
Profile Name	<p>Type the name of the profile.</p> <p>You can use up to 64 characters in the profile name of profiles created for wired devices. Profile name must not contain special characters or spaces. Note that profiles that are automatically created by Network Director as part of device discovery or out-of-band changes might contain the underscore (_) character in the profile name.</p>
Description	Type a description for the profile.
VRRP Configuration Settings	
Family	Select the IPv4 or IPv6 address family.
VRRP Group Identifier [0 - 255]	Select the VRRP group identifier, which identifies the virtual routing device where the packet is routed to. Each VRRP group is identified by a unique virtual identifier. MAC addresses ranging from 00:00:5e:00:01:00 through 00:00:5e:00:01:ff are reserved for VRRP groups, as defined in RFC 2338. The VRRP group number must be the decimal equivalent of the last hexadecimal byte of the virtual MAC address, which is 0 through 255.
Advertise Interval ([1-255] secs for IPv4/ [100-40000] msec for IPv6):	Configure the interval in milliseconds between VRRP IPv4 and IPv6 advertisement packets.
Fast Interval in msec [10-40950]	<p>Configure the interval, in milliseconds, between VRRP advertisement packets. All devices in the VRRP group must use the same advertisement interval.</p> <p>Range: 100 through 999 milliseconds</p> <p>Default: 1 second</p>
Authentication Type(for IPv4)	<ul style="list-style-type: none"> • simple—Use a simple password. The password is included in the transmitted packet. • md5—Use the MD5 algorithm to create an encoded checksum of the packet.
Authentication Key(for IPv4)	<p>Configure a VRRP IPv4 authentication key or password. You also must specify a VRRP authentication scheme by including the authentication-type statement. All devices in the VRRP group must use the same authentication scheme and password.</p> <p>For simple authentication, the password can contain 1 through 8 characters. For MD5 authentication, it can contain 1 through 16 characters. If you include spaces, enclose all characters in quotation marks (" ").</p>

Table 129: VRRP Profile Settings (*continued*)

Field	Action
Preempt	Determine whether or not a backup device can preempt a primary device: When no-preempt is configured, the backup device cannot preempt the primary device even if the backup device has a higher priority.
Hold Tim in secs [0 - 3600]	Set the hold time before a higher-priority backup device preempts the primary device. Range: 0 through 3600 seconds Default: 0 seconds
Accept Data	Determine whether or not an interface accepts packets destined for the virtual IP address This feature helps to debug connectivity issues by making devices respond to ping packets on virtual IP.
Virtual Link Local Address (IPv6)	Configure a virtual link local address for the VRRP IPv6 groups. You must explicitly define a virtual link local address for each group. The virtual link local address must be in the same subnet as the physical interface address.
Virtual IP Addresses (IPv4)	
IP Addresses	The addresses of the virtual routers in a VRRP IPv4 group. You can configure up to eight addresses. Do not include a prefix length. If the address is the same as the interface's physical address, the interface becomes the primary virtual device for the group.
Virtual IP Addresses (IPv6)	
IP Addresses	The addresses of the virtual routers in a VRRP IPv6 group. You can configure up to eight addresses. Do not include a prefix length. If the address is the same as the interface's physical address, the interface becomes the primary virtual device for the group.

RELATED DOCUMENTATION

[Understanding VRRP Profiles | 544](#)
[Creating and Managing Port Profiles | 251](#)
[Creating and Managing VLAN Profiles | 329](#)

Managing Network Devices

IN THIS CHAPTER

- Viewing the Device Inventory Page | 551
- Physical Topology | 553
- Viewing Profiles Assigned to a Device | 558
- Viewing the Physical Inventory of Devices | 560
- Viewing Licenses With Network Director | 561
- Viewing a Device's Current Configuration from Network Director | 563
- Assigning Devices to Logical Category | 564
- Accessing a Device's CLI from Network Director | 565
- Accessing a Device's Web-Based Interface from Network Director | 566
- Deleting Devices | 568
- Rebooting Devices | 569
- Viewing Virtual Machines | 569

Viewing the Device Inventory Page

The Device Inventory page lists devices managed by Network Director and provides basic information about the devices, such as IP address and current operating status. The Device Inventory page is available in Build and Deploy mode and is the default landing page for Build mode.

The scope you have selected in the View pane and the network view that you have selected from the View selector determines which devices are listed in the Device Inventory page. For example:

- If you are in the Logical View and select My Network, all devices managed by Network Director are listed.
- If you are in the Datacenter View and select Network Devices under a data center, only the network devices that are part of that data center are displayed.
- If you select a building in Location view, only those devices assigned to that building (including the floors and closets in the building) are listed.

The Device Inventory page provides three pie charts that summarize the status of the devices in your selected scope:

- Devices by Family—Indicates the proportion of devices in each device family.
- Connection State—Shows the proportion of devices that are up or down. In this chart, Virtual Chassis count as one device.
- Configuration State—Shows the proportion of devices in each configuration state. See the Config State entry in [Table 130 on page 551](#) for definitions of the configuration states.

Mouse over a pie segment to view the actual number of devices and the percentage represented by that pie segment.

[Table 130 on page 551](#) describes the fields in the Device Inventory table.

Table 130: Fields in the Device Inventory Table

Field	Description
Hostname	Configured name of the device or IP address if no hostname is configured.
IP Address	IP Address of the device.
Serial Number	Serial number of device chassis.
Platform	Model number of the device.
OS Version	Operating system version running on the device.

Table 130: Fields in the Device Inventory Table (*continued*)

Field	Description
Device Family	<p>Device family of the device:</p> <ul style="list-style-type: none"> • JUNOS-EX for EX Series switches • JUNOS for Campus Switching ELS • JUNOS-QFX for QFX Series switches
Device Type	<p>Type of the device:</p> <ul style="list-style-type: none"> • Switch—Standalone switch • VC—Virtual Chassis primary • VC Member—Virtual Chassis member switch • XRE—External Routing Engine for EX8200 Virtual Chassis
Connection State	<p>Connection status of the device in Network Director:</p> <ul style="list-style-type: none"> • UP—Device is connected to Network Director. • DOWN—Device is not connected to Network Director.
Config State	<p>Displays the configuration status of the device:</p> <ul style="list-style-type: none"> • In Sync—The configuration on the device is in sync with the Network Director configuration for the device. • Out Of Sync—The configuration on the device does not match the Network Director configuration for the device. This state is usually the result of the device configuration being altered outside of Network Director. You cannot deploy configuration on a device from Network Director when the device is Out Of Sync. To resolve this state, use the Resynchronize Device Configuration task in Deploy mode. • Sync failed—An attempt to resynchronize an Out Of Sync device failed. • Synchronizing—The device configuration is in the process of being resynchronized.
Manageability State	<p>Displays if the device is directly manageable or not.</p> <p>This is a hidden field. To display the Manageability State field, click any column, click the down arrow to expand the list, select Columns from the list, and then enable Manageability State.</p>

NOTE: Juniper Networks devices require a license to activate the feature. To understand more about Network Director Licenses, see, [“Viewing Licenses With Network Director” on page 561](#).

RELATED DOCUMENTATION

[Understanding the Network Director User Interface | 3](#)

[Understanding Resynchronization of Device Configuration | 611](#)

[Device Inventory Report | 820](#)

[Viewing Licenses With Network Director | 561](#)

[Network Director Documentation home page](#)

Physical Topology

At the device level, you can view the connectivity details of a device and the details of all the devices that are connected to the specified device by using the Device Connectivity task in Network Director. The Device Connectivity page displays various details about a selected device and its immediate neighbors. The level of detail that Network Director displays in the Device Connectivity page depends on the type of device that you select.

To view the connectivity details of devices:

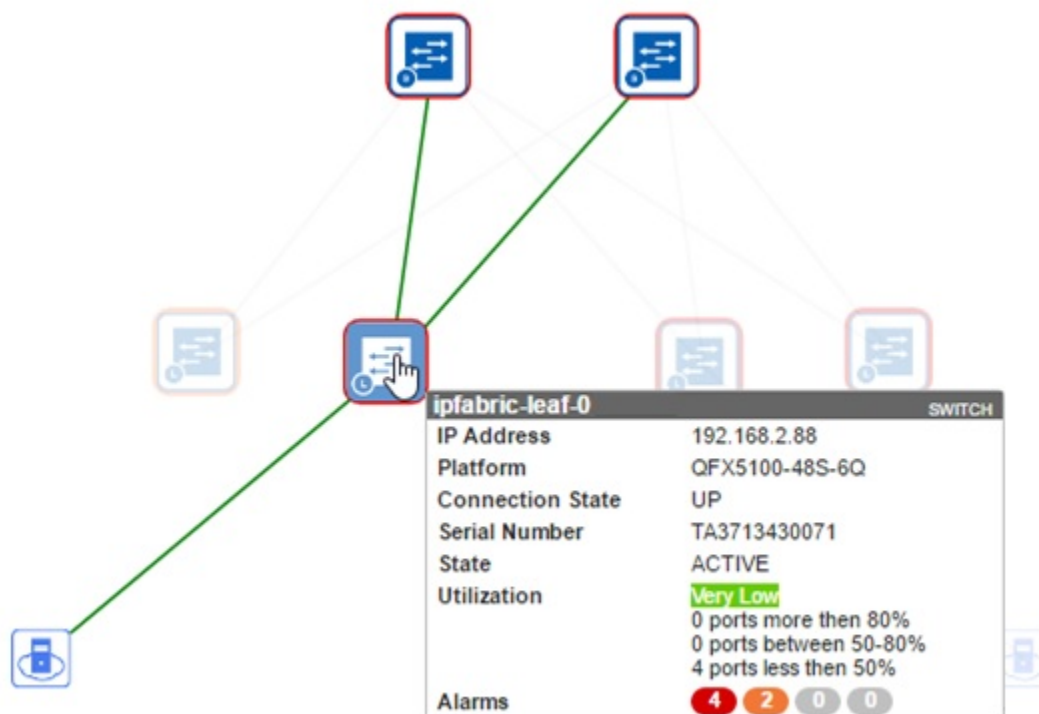
1. While in the Logical, Location, Device, Custom, or Topology View, under Build Mode, select the device for which you want to view connectivity from the View pane or the network topology (in case of Topology View) and click **Connectivity > View Device Connectivity** from the Tasks pane.

The Device Connectivity page opens. You can view the device connectivity details either in graph view or in grid view. The default view is the graph view.

In the graph view, each device and its network connectivity to all the connected devices are displayed as shown in [Figure 23 on page 554](#). Mouse over a device to select a device and view details of the device.

NOTE: If the selected device is connected to a device that is not managed by Network Director, the latter appears dimmed in the Device Connectivity page.

Figure 23: Displaying Connection Details in Graph View



If the selected device is connected to more than sixty devices, then all the connected devices are highlighted in a circular form or a grid form. If the selected device is connected to less than 60 devices, then the links between the interconnected devices are displayed.

The device details displayed include name, IP address, and the alarm state information in colored labels that provide health and reachability information. You can also view the details of the hosts or virtual machines that are connected to the devices.

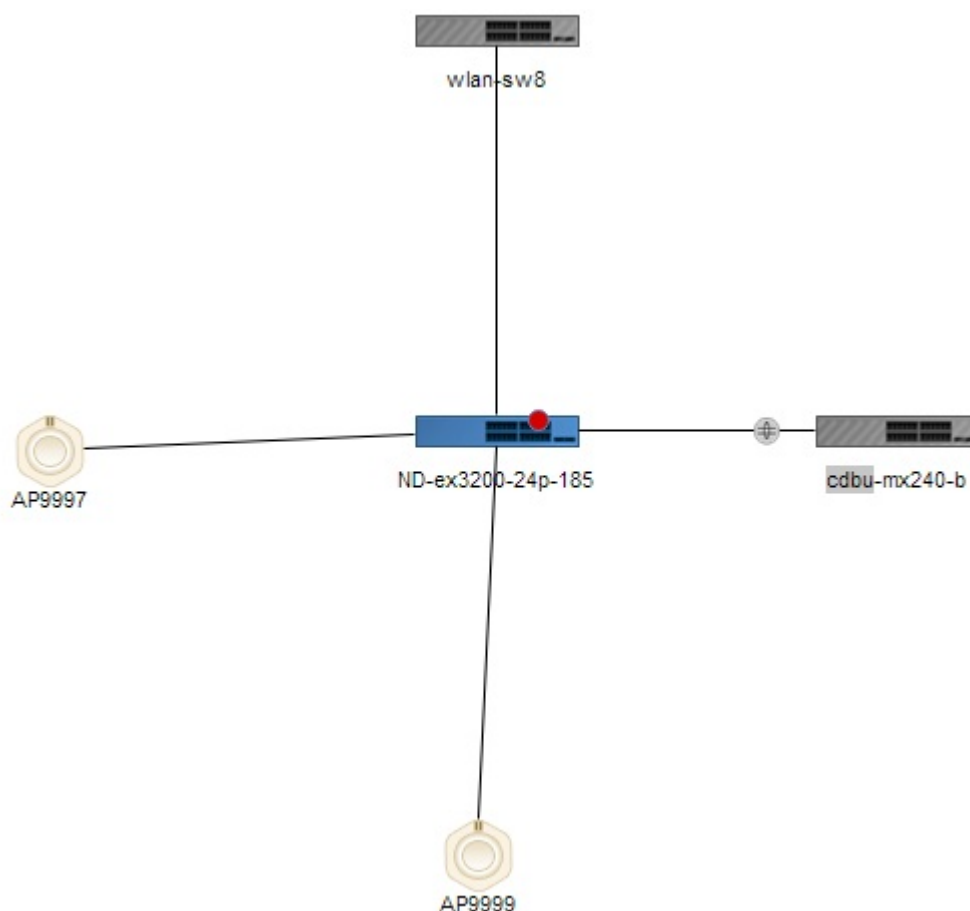
You can view the following details when you mouse over a device in the Device Connectivity—Graph view:

NOTE: The level of detail displayed depends on the type of device selected.

- **Name**—The name of the device provided while configuring the device. The device name and the device type are displayed in a label.
- **IP Address**—The IP address of the device.
- **Platform**—The device family and platform information. For example, EX4300-24P, QFX10002-36Q, and so on.
- **Connection State**—Connection state of the device. Connection state can be UP, DOWN, or N/A. Network Director updates and displays the connection state changes in real time.

- Config State—Device state. The device state can be online/provisioned or down/provisioned.
- Serial Number—Serial number of the device.
- Link status—Indicates whether the link between two devices is UP or DOWN. Network Director updates and displays the link status changes in real time. It might take up to 2 minutes for the updates to reflect.
- Utilization—Overall color-coded bandwidth utilization level and the breakup of bandwidth utilization by each port on the device.
- Alarms—Alarm details displaying the number of critical, major, minor alarms, or info for the device. Alarms details are color coded to indicate their severity level. Network Director updates and displays the alarm status changes in real time.
- Slot Number—(Applicable to Junos Fusion satellite devices) The FPC identifier of the satellite device in Junos Fusion. Slot number ranges from 65 through 255, and functions as the FPC identifier in the interface name when satellite device interfaces are being configured.
- LAG—Identifies connections that are configured as LAGs as shown in the following figure.

Figure 24: LAGs in the Device Connectivity view



You can view the following details of virtual machines (VMs) that are connected to hosts:








- Virtual Machine—Name of the virtual machine.
- Host Name—Name of the host to which the virtual machine is connected to.
- VNetwork—Name of the virtual network.
- OS—Name of the operating system on which the virtual machine is running.
- Connection State—Connection status of the virtual machine. Connection state can be UP, DOWN, or N/A.
- Power State—State of the power supply: Powered On or Powered Off.

You can view the following details of the Desktop machine:

- Host Name—Name of the host to which Desktop machine connected to.
- OS—Name of the operating system on which the Desktop machine is running.
- Connection State—Connection status of the Desktop machine. Connection state can be UP, DOWN, or N/A.

In the graph view, each networking device has a unique icon for easy identification. [Table 131 on page 556](#) describes the networking device that each of these icons indicate.

Table 131: Icons on the Device Connectivity Page

Icon	Description
	Juniper Networks switch
	Desktop computer
	Desktop IP phone
	Printer
	Satellite device cluster in a Junos Fusion system. Double-click this icon to view the connectivity of the devices that are part of the cluster.
	Hypervisor server (KVM or ESXi host).
	Bare metal server

2. You can perform the following tasks from the graph view of the Device Connectivity page:
 - Click the number adjacent to a device icon to view details about the device alarms. The Alarm Details by Severity page opens. For more details, see [“Alarm Detail Monitor” on page 771](#).

For example, in [Figure 25 on page 557](#) click the number 8 to view details of the alarms on the corresponding device.

Figure 25: Alarm Count in the Device Connectivity Page



- Double-click the satellite device cluster icon to expand and view the member satellite devices and their connectivity.
 - Click **Links** and select **Color Code Port Utilization** to view the color-coded port utilization level in the graph view. Network Director displays a port utilization legend in the upper right corner of the graph view, which you can use to identify links that are optimally used, overutilized, or underutilized and take necessary corrective actions.
 - Select **Stop Updates**, to freeze the link status changes in real time in the Device Connectivity page that might be required while the user is performing some tasks in this page.
 - Enter the device name, IP address, or the tag name of the device in the search field to quickly locate a device in the graphical view.
3. Click **Show Grid View** to view the device connectivity details in a tabular format as displayed in [Figure 26 on page 557](#). This view has two tabs: the External Links tab and the Fabric Links tab. Clicking the **External Links** tab displays the external device interface details that are connected to the fabric devices. Clicking the **Fabric Links** tab displays fabric link interface details.

Figure 26: Displaying the Connection Details in Grid View

Device Connectivity : nd-72q1-elit

External LinksFabric Links

Show Graph View

Source Device	Source Port	Source Port Bandw...	Destination Device	Destination Port	Destination Port Bandw...	Link Status
nd-72q1-elit	[LAG] ae0	NA	nd-36q1-elit	[LAG] ae0	NA	🟢
nd-72q1-elit	[LAG] ae1	NA	nd-36q1-elit	[LAG] ae1	NA	🔴
nd-72q1-elit	et-0/0/0 (ae0)	0	nd-36q1-elit	et-0/0/0 (ae0)	0	🟢
nd-72q1-elit	et-0/0/1 (ae1)	0	nd-36q1-elit	et-0/0/1 (ae1)	0	🔴
nd-72q1-elit	et-0/0/2 (ae2)	0	nd-opus-48s4	et-0/0/48	0	🟢

The following details are displayed in the grid view:

- Source Device—Name of the device specified while configuring the device.
- Source Port—Source port of the device.
- Source Port Bandwidth %—Real-time percentage of bandwidth utilized at the source port.

- Destination Device—Name of the destination device or devices the source device is connected to.
- Destination Port—The port number on the destination device to which the source device is connected to.
- Destination Port Bandwidth %—Real-time percentage of bandwidth utilized at the destination port.
- Link Status—Indicates whether the link to the device is up or down.

You can sort the details in the table in the ascending order or descending order for each column. You can also use filters to display device connectivity details for specific devices. If you type a text string and click **Go**, entries that do not contain the text string (filter criterion) are removed from the table

RELATED DOCUMENTATION

| [Setting Up the Topology View](#) | 133

Viewing Profiles Assigned to a Device

View Assigned Profile page list all the profiles associated with a selected device or with an object such as ports within that device. To view the profiles assigned to a device, you must have the profiles already assigned to the devices or ports within that device. Only those profiles that are assigned to a specified object will be displayed in the Profiles Assigned to the Device page. In addition to displaying profiles assigned to objects, the Profiles Assigned to Device page also shows link aggregation groups (LAGs) assigned to devices.

The View Assigned Profiles task is available in the Logical, Location, and Device panes for EX Series switches.

You can access the View Assigned Profiles page by selecting the object (device or port) and clicking the View Assigned profiles menu.

You can view the profiles assigned to an EX Series switch . To view the assigned profiles to a particular device:

While in Build mode, select an EX Series switch from the Switching Network cabinet under the View pane and select **View Assigned profiles** from the Tasks pane.

The Profiles Assigned to the Device page displays a list of profiles that are already assigned to the selected device. The details displayed are described in [Table 132 on page 559](#).

Table 132: Details of Assigned Profiles to a Device

Field	Description
Profile Type	The type of the profile. The profiles are grouped based on the type of the profile. The profiles that are directly deployed on the device are displayed in the list. For example, Device, VLAN, Port and so on.
Profile Name	Name of the profile that was specified at the time of creating the profile.
Object	The name of the device (EX Series switch) or port.
Object Type	Specifies whether the object is a device (EX Series switch) or a port. TIP: For an EX Series switch, the object type is the device or a port
Assignment State	The status of the profile whether it is deployed or in progress. <ul style="list-style-type: none"> • Deployed—the profile is provisioned to the device • Pending deployment—the profile is assigned to the device, but pending provisioning.

RELATED DOCUMENTATION

[Assigning Device Common Settings to Devices | 198](#)
[Assigning a VLAN Profile to Devices or Ports | 345](#)
[Network Director Documentation home page](#)

Viewing the Physical Inventory of Devices

You can view the physical inventory of all the devices in your network in the Device Physical Inventory page. The Device Physical Inventory page displays information about the slots that are available for a device and provides information about power supplies, chassis cards, fans, part numbers, and so on. Network Director displays hardware inventory by device name, based on data retrieved both from the device during discovery and resynchronizing operations, and from the data stored in the hardware catalog. For each managed device, the physical inventory page provides descriptions for field replaceable units (FRUs), part numbers, model numbers, and the pluggable locations from which empty slots are determined.

To view the Device Physical Inventory page, while in the Build mode select a standalone EX Series switch, a Virtual Chassis from the View pane and select **Device Management > Physical Inventory** from the Tasks pane.

The physical inventory page displays the model number, part number, serial number, and description for the following, depending on the device that you selected:

- For standalone EX Series switches and Virtual Chassis, the page displays details of the switch, the chassis, the Flexible PIC Concentrator (FPC), the PIC slot, the PIC installed in the PIC slot, the power supply, the fan tray, and the routing engine.

You can view the following details from the Device Physical Inventory page as described in [Table 133 on page 560](#).

Table 133: Fields in the Device Physical Inventory Table

Field	Description
Item	Name of the device and the components that are part of the device. By default, Network Director displays the device and components in an expanded tree structure. You can click a device or component to collapse or expand the sub-components.
Model Number	<ul style="list-style-type: none">• For standalone EX Series switches and Virtual Chassis, the full Junos EX Series model number of the device.
Part Number	Part number of the EX Series switch chassis component.
Serial Number	The hardware serial number of the device.
Description	The description about the component.

NOTE: Juniper Networks devices require a license to activate the feature. To understand more about Network Director Licenses, see [“Viewing Licenses With Network Director” on page 561](#). Please refer to the product Data Sheets for further details, or contact your Juniper Account Team or Juniper Partner.

RELATED DOCUMENTATION

[Understanding the Network Director User Interface | 3](#)

[Viewing Licenses With Network Director | 561](#)

[Network Director Documentation home page](#)

Viewing Licenses With Network Director

Juniper Networks devices require a license to operate some features. You can view the licenses for devices connected to Network Director.

To view the license for a Juniper Networks device on your network:

1. Select the **Build** icon in the Network Director banner.
2. In the View pane, select a wired device.
3. In the Tasks pane, select **View License Information**.

The Licenses page for that object is displayed with the fields listed in [Table 134 on page 561](#).

Table 134: Viewing Licenses with Network Director

Field	Description
Feature Name	Name of the licensed SKU or feature. It can be used to look up the license with Juniper Networks. Not all devices support this.
License Count	Number of times an item has been licensed. This value can have contributions from more than one licensed SKU or feature. Alternatively, it can be 1, no matter how many times it has been licensed.

Table 134: Viewing Licenses with Network Director (*continued*)

Field	Description
Used Count	Number of times the feature is used. For some types of licenses, the license count will be 1, no matter how many times it is used. For capacity-based licensable items, if infringement is supported, the license count can exceed the given count, which has a corresponding effect on the need count.
Need Count	Number of times the feature is used without a license. Not all devices can provide this information.
Given Count	Number of instances of the feature that are provided by default.

NOTE: If a device does not have a license, a blank page is displayed with the message, **No license is installed on this device.** If you are sure the device has a license, try resynchronizing the device before displaying the license again.

- Optionally, expand the license information by feature name to view the feature SKU information. [Table 135 on page 562](#) describes the additional fields that are displayed.

Table 135: Additional Licensing Information

Field	Description
Validity Type	Validity type can be Databased (license expires on end date), Permanent, Countdown (license expires when time remaining is zero), or Trial. If the validity type is either Databased or Countdown, more information is displayed—License Name, License Version, License State, and Time Remaining. Additional information can be added in the details grid based on the SKU type (SKU or Feature)—Start Date, End Date, or Original Time Allowed.
License Name	If the validity type is either Databased or Countdown, the identifier associated with a license key is displayed.
License Version	If the validity type is either Databased or Countdown, the version of a license is displayed. The version indicates how the license is validated, the type of signature, and the signer of the license key.
License State	If the validity type is either Databased or Countdown, the state of the license is displayed—Valid, Invalid, or Expired.

Table 135: Additional Licensing Information (*continued*)

Field	Description
Time Remaining	If the validity type is either Databased or Countdown, the remaining time left on the license is displayed. For a trial license, the number of days remaining after you installed the device is displayed. For a commercial license, the time remaining is unlimited.
Start Date	Based on the SKU type, the start date of the license can be displayed in the details grid.
End Date	Based on the SKU type, the end date of the license can be displayed in the details grid.
Original Time Allowed	Based on the SKU type, the original license timeframe can be displayed here.

RELATED DOCUMENTATION

[Resynchronizing Device Configuration | 616](#)
[Network Director Documentation home page](#)

Viewing a Device's Current Configuration from Network Director

You can view a device's current configuration from Network Director. This is a convenient way to view device configurations without leaving Network Director.

To view a device's current configuration:

1. Click **Build** or **Deploy** in the Network Director banner.
2. Select the device in the View pane.
3. Select **Device Management > Show Current Configuration** in the Tasks pane.
4. The device's current configuration displays in the main window.

RELATED DOCUMENTATION

[Understanding the Network Director User Interface | 3](#)

[Understanding the Build Mode Tasks Pane | 82](#)

[Network Director Documentation home page](#)

Assigning Devices to Logical Category

Network Director enables you to tag the available EX Series switches to different categories—Core, Aggregation, Access, or Unassigned categories. Once you tag a switch, it appears under the tagged category after refreshing the screen.

To assign a group of EX Series switches under Core, Aggregation, Access, or Unassigned category, select Logical View in Build mode (available only in Build mode) and select one of the cabinets under the Switching Network cabinet:

1. Select **Assign Device to Logical Category** under Device Management from the Tasks pane.
The Assign Device to Logical category page is displayed.
2. Click **Add** from the Selected Devices table. The Please select devices dialog box is displayed.
3. Select the device or devices that you want to assign by selecting the check box next to the hostname. Click **OK**. The selected device or devices with the details appears in the Selected Devices table.
4. Select a role from the New Role list.
The available roles are: Access, Aggregation, Core, and Unassigned.
5. Click **Done** to change the role or click **Cancel** if you do not want to change the role. The message: **Device Role successfully changed** appears if you have selected Done.

To assign an EX Series switch under Core, Aggregation, Access, or Unassigned category, select Logical View in Build mode (available only in Build mode) and select any EX Series switch from the one of the cabinets under the Switching Network cabinet:

1. Select **Assign Device to Logical Category** under Device Management from the Tasks pane.
The Assign Device to Logical category page is displayed. The page displays the name of the selected device and the current role of the device.
2. Select a role from the New Role list to change the current role.

The available roles are: Access, Aggregation, Core, and Unassigned.

3. Click **Done** to change the role or click **Cancel** if you do not want to change the role. The message: **Device Role successfully changed** appears if you have selected Done.

RELATED DOCUMENTATION

[Viewing Profiles Assigned to a Device | 558](#)

[Network Director Documentation home page](#)

Accessing a Device's CLI from Network Director

Network Director enables you to connect to the CLI for switches using SSH.

This topic describes the steps to connect to a switch by using SSH (Secure Shell). SSH is a cryptographic network protocol used for remote shell services or command execution. SSH is one of the many access services that are supported on the Juniper Networks devices. All Juniper Network devices have SSH enabled by default.

To connect to a device by using SSH:

1. Do one of the following:
 - In the View pane, select the device to which you want to connect.
 - In the Topology View, locate the device to which you want to connect.
2. Do one of the following:
 - With the device selected in the View pane, select **Build** mode and select **Tasks > Device Management > SSH to Device**.
 - While in the Topology View, select the device to which you want to launch the SSH connection and click **Device Management > SSH To Device**.

The SSH to Device dialog box appears.

3. Enter the username and password to connect to the selected device and click **Connect**.

NOTE: Ensure that you have removed Pop-Up blockers, if any, before you click Connect.

The SSH console to the switch opens in a separate browser tab or window depending on your browser settings. Refer to the [EX Series documentation](#) for more information about using the CLI for EX Series switches.

NOTE: Any configuration changes you make to a device, using the CLI qualify as out-of-band changes in Network Director. Out-of-band configuration changes can cause the configuration state of a managed device to become out of sync, which indicates that the device configuration no longer matches the Build mode configuration for the device. Use the Resynchronize Device Configuration task in Deploy mode to resynchronize the device configuration.

RELATED DOCUMENTATION

[Understanding Resynchronization of Device Configuration](#) | 611

[Accessing a Device's Web-Based Interface from Network Director](#) | 566

[Network Director Documentation home page](#)

Accessing a Device's Web-Based Interface from Network Director

Network Director enables you to connect to the switches using the device Web-based interface.

This topic describes the steps to connect to a switch by using the J-Web interface. The J-Web interface is a graphical user interface, using which you can monitor, configure, troubleshoot, and manage switches.

You can connect and configure a device by using the J-Web interface or Web View only if the device is configured to accept HTTP or HTTPS as a management service. You can configure HTTP or HTTPS as a management service using the Device Common Settings profile. For more information, see [“Creating and Managing Device Common Settings” on page 170](#).

To connect to a device using the J-Web interface or Web View:

1. Do one of the following:
 - In the View pane, select the device to which you want to connect.
 - In the Topology view, locate the device to which you want to connect.
2. Do one of the following:
 - While selecting the device in the View pane, select Build mode and select **Tasks** pane > **Device Management** > **Launch Web View**.

- While in the Topology View, select the device for which you want to launch the Web connection and click **Device Management** > **Launch Web View**.

The Web View or J-Web Login page appears.

3. Enter the username and password to connect to the selected switch and click **Login**.

If the credentials that you entered are valid, the system displays the J-Web or Web View home page for the selected device.

NOTE: Any configuration changes you make to a device using the Web interface qualify as out-of-band changes in Network Director. Out-of-band configuration changes can cause the configuration state of a managed device to become out of sync, which indicates that the device configuration no longer matches the Build mode configuration for the device. Use the Resynchronize Device Configuration task in Deploy mode to resynchronize the device configuration.

RELATED DOCUMENTATION

[Understanding Resynchronization of Device Configuration](#) | 611

[Accessing a Device's CLI from Network Director](#) | 565

[Network Director Documentation home page](#)

Deleting Devices

You can delete devices that are no longer used from Network Director. Deleting a device removes all device configuration and device inventory information from the Junos Space database. Once a device is deleted from the database, all the profiles associations, device configurations, and inventory information of the deleted device are also deleted. However, the system maintains the audit logs and monitoring data for the device even after the device is deleted.

Use the Delete Devices page to delete devices from Network Director. While in Build mode, click **Delete Devices** from the **Tasks > Device Management** menu. The Delete Devices page appears.

The Delete Devices page displays the devices contextually depending on your selection in the View pane. For example, if you select a site in Location view and click Delete Devices, Network Director displays all the devices that are assigned to the buildings or floors in the selected site in the Delete Devices page. If you select a particular switch family in Device View and click Delete Devices, only switches that belong to that switch family are displayed.

To delete devices, complete the following tasks:

1. Select the check box adjacent to the switch that you want to delete.
2. Click **Done**.

Network Director prompts you to confirm the deletion. Click **Yes** to confirm the deletion or **No** to go back and make changes to the selection.

RELATED DOCUMENTATION

[Understanding the Network Director User Interface | 3](#)

[Discovering Devices in a Physical Network | 94](#)

[Viewing the Device Inventory Page | 551](#)

[Network Director Documentation home page](#)

Rebooting Devices

Use the Reboot Devices task to immediately reboot the selected device. This task is available in all scopes when in Build mode. To reboot one or more devices immediately:

1. Select the scope in the View pane that contains the devices you want to reboot.
2. Select Reboot Devices from the Tasks pane.
3. Expand the tree on the page as needed to locate the available devices.
4. Select the check box for one or more devices.
5. Click **Done** to start the reboot or click **Cancel** to return to the Device Inventory page.

The rebooting process triggers a Cold Start Alarm that can be seen in Fault mode.

RELATED DOCUMENTATION

[Understanding the Build Mode Tasks Pane | 82](#)

[Understanding the Network Director User Interface | 3](#)

[Network Director Documentation home page](#)

Viewing Virtual Machines

EX Series switches (standalone and Virtual Chassis) and QFX Series switches in your network can be connected to one or more ESX/ESXi hosts. Each host can have one or more virtual machines running on them.

You can use the View Virtual Machine task to view details about virtual machines that are connected to a switch.

To view the virtual machines

1. While in the Logical View with Build mode selected, select the standalone switch or virtual chassis for which you want to view the connected hosts.
2. Click **Connectivity > View Virtual Machines** from the Tasks pane.

3. The View Virtual Machines table displays the details of the virtual machines that are connected to the selected switch. [Table 136 on page 570](#) describes the fields in this table:

Table 136: Manage Virtual Machines Page Field Descriptions

Field	Description
Switch Port	The switch port on the physical switch that is connected to the host.
Host	Name of the host on which the virtual machine is running.
Host NIC	The network adapter on the host that connects the physical switch to the host.
VLANs	The VLANs configured on the physical switch port.
Virtual Machines	<p>The name of the virtual machines that are running on the given host.</p> <p>Mouse over this field to view the number of virtual machines that are running on the host.</p>

RELATED DOCUMENTATION

| [Network Director Documentation home page](#)

4

PART

Working in Deploy Mode

About Deploy Mode | **572**

Deploying and Managing Device Configurations | **580**

Deploying and Managing Software Images | **632**

Managing Devices | **645**

Setting Up Zero Touch Provisioning for Devices | **649**

About Deploy Mode

IN THIS CHAPTER

- [Understanding Deploy Mode in Network Director | 572](#)
- [Understanding the Deploy Mode Tasks Pane | 577](#)

Understanding Deploy Mode in Network Director

IN THIS SECTION

- [Deploying Configuration Changes | 573](#)
- [Managing Software Images | 574](#)
- [Zero Touch Provisioning | 575](#)
- [Managing Devices | 575](#)
- [Managing Device Configuration Files | 575](#)
- [Managing Baseline Configuration | 575](#)

The Deploy mode enables you to deploy configuration changes and software upgrades to devices and perform several device management and configuration file management tasks.

NOTE: Deploy mode is enabled for devices in Dashboard View, Logical View, Location View, Device View, Custom Group View, and Topology View.

This topic describes:

Deploying Configuration Changes

When you make configuration changes in Build mode, the changes are not deployed to devices automatically. You must manually deploy the changes to devices in Deploy mode. Every time you make configuration changes in Build mode that affect a device, the device is automatically added to the list of devices with pending changes.

NOTE: The device is added to the list of devices with pending changes only when you make the device configuration changes. If you make changes to the configurations (associated with the device) that are specific only to Network Director the device is not listed with pending changes. For example, when you make changes to the profile name associated with the device, the device is not added to the list of devices with pending changes.

Configuration changes are deployed to devices at the device level. When you deploy configuration changes to a device, all pending configuration changes for that device are deployed.

You can deploy the device configurations in the following two ways:

- **Auto Approval**—In this mode, the device configuration changes are approved automatically by the system and do not require explicit (manual) approval by a configuration approver before they can be deployed. This is the default approval mode.
- **Manual Approval**—In this mode, the device configuration changes are required to be explicitly approved by a configuration approver before the changes can be deployed to the device.

For more information about enabling these modes, see [“Setting Up User and System Preferences” on page 26](#).

For manual approval, the Network Director - Configuration Approver role is available in Junos Space, which is specific to the Network Director. A user with this role reviews device configurations and proposed changes to device configurations and can either approve or reject them.

An operator performs device configurations and creates a change request for that configuration and submits it for approval to an approver. The approvers are notified by e-mail whenever a change request is created. If a configuration or a change to it is approved by an approver, then the operator is able to deploy it. If a configuration is rejected then the operator must make the necessary changes, resubmit the change request, and procure an approval before the configuration can be deployed. For more information, see [“Approving Change Requests” on page 606](#)

NOTE: You can specify any number of approvers. If you specify more than one approver while configuring the Manual Approval mode, once an approver accepts or rejects the proposed change, the change request is not listed for the other approvers and they cannot approve or reject the same change request.

You can do the following configuration deployment tasks on devices that have pending changes:

- Run configuration deployment jobs immediately or schedule them for future times.
- Approve the change requests for pending configurations, if you have selected the Manual Approval mode.
- Preview pending configuration changes before deploying the changes.
- Validate that the pending changes are compatible with the device's configuration.
- Manage configuration deployment jobs.

Configuration changes are validated for each device both in Network Director and on the device. If any part of a configuration change for a device fails validation, no configuration changes are deployed to the device. You can see the results of each validation phase separately.

Network Director will not deploy configuration to a device with a configuration that is out of sync (meaning that the device's configuration differs from Network Director's version of that device's configuration), or to a device that has uncommitted changes to its candidate configuration. Deployment to such devices will fail.

When you schedule a deployment job, that job and any profiles and devices assigned to that job are locked within Network Director. You cannot edit the job or any of its assigned profiles until the job runs or gets cancelled. This locking feature prevents you from deploying unintended configuration changes that could result from editing profiles and devices that are already scheduled to deploy. To change any properties of a scheduled job, cancel the job and create a new scheduled job with the desired properties. You cannot edit the profile assignments of a device that has scheduled pending configuration changes.

Managing Software Images

Network Director can manage software images on the nodes it manages. You can do the following software image management tasks:

- Deploy a software image stored in an image repository on the Network Director server to multiple devices with a single job.
- Track the status of software image management jobs.
- Stage and install software images as separate tasks.

- Schedule staging and installation to happen at independent future times.
- Perform several software image upgrade options, such as rebooting devices automatically after the upgrade finishes.

NOTE: Using nonstop software upgrade (NSSU) to upgrade EX Series switches is supported in Network Director.

Zero Touch Provisioning

Zero touch provisioning enables you to provision new Juniper Networks switches in your network automatically—without manual intervention. When you physically connect a switch to a network and boot it with the factory-default configuration, the switch attempts to upgrade the Junos OS software automatically and autoinstall a configuration file from the network.

Managing Devices

In Deploy mode you can perform several device management tasks, including:

- View the device inventory.
- Show a device's current configuration.
- Resynchronize the device configuration maintained in Build mode with the configuration on the device. For more information about resynchronization of device configuration, see [“Understanding Resynchronization of Device Configuration” on page 611](#)
- Enable or disable switch network ports.
- Convert QSFP+ port configuration.

Managing Device Configuration Files

You can back up device configuration files to the Network Director server. You can perform several actions on backed up configuration files, such as restoring configuration files to devices, and viewing and comparing configuration files.

Managing Baseline Configuration

You can create a baseline of the Network Director device configuration and the OS version on the Network Director server. You can perform several actions on the baseline configuration files, such as restoring configuration files to devices, and viewing and comparing configuration files.

RELATED DOCUMENTATION

Deploying Configuration to Devices	 580
Deploying Software Images	 635
Viewing the Device Inventory Page	 551
Viewing a Device's Current Configuration from Network Director	 563
Resynchronizing Device Configuration	 616
Enabling or Disabling Network Ports on Switches	 645
Managing Device Configuration Files	 622
Network Director Documentation home page	

Understanding the Deploy Mode Tasks Pane

The Tasks pane in Deploy mode lists the available tasks. All Deploy mode tasks are always available, regardless of the scope selected in the View pane.

Deploy mode tasks are divided into the following categories:

- **Configuration Deployment**—These tasks enable you to deploy configuration changes to devices and manage configuration deployment jobs. [Table 137 on page 577](#) describes the configuration deployment tasks.
- **Image Management**—These tasks enable you to manage software images on devices. [Table 138 on page 578](#) describes the image management tasks.
- **Device Management**—These tasks enable you to view the device inventory, resynchronize the configuration of out-of-sync devices, manage the administrative state of ports and convert QSFP+ port configuration. [Table 139 on page 578](#) describes the device management tasks.
- **Device Configuration File Management**—These tasks enable you manage configuration files on managed devices. [Table 140 on page 578](#) describes the device configuration file management tasks.
- **Baseline Management**—These tasks enable you to manage baseline configuration of devices. [Table 141 on page 578](#) describes the baseline management tasks.
- **Zero Touch Provisioning**—These tasks enable you to provision new Juniper Networks switches in your network automatically—without manual intervention. [Table 142 on page 579](#) describes the zero touch provisioning tasks.
- **Key Tasks**—Network Director enables you to group the tasks that you perform frequently and create a list of key tasks. You can add any task from the Tasks pane to the Key Tasks list by selecting a task and clicking the plus (+) sign that appears adjacent to the task. For some modes, you can see that Network Director has predefined some key tasks for you. You can modify this set of tasks to suit your requirements. This feature is available in Task pane irrespective of your current mode, scope, or view.

[Table 137 on page 577](#) through [Table 140 on page 578](#) describe the tasks in each task category.

Table 137: Configuration Deployment Tasks

Task	Description
Deploy Configuration Changes	Deploys pending configuration changes to devices.
Approve Change Requests	Enables a configuration approver to approve or reject a change request, which has been submitted for approval by an operator.
Set SNMP Trap Configuration	Enables SNMP traps on network devices so that Network Director can collect and manage event and error information from these devices.

Table 137: Configuration Deployment Tasks (continued)

Task	Description
View Deployment Jobs	Manages configuration deployment jobs.

Table 138: Image Management Tasks

Task	Description
Manage Image Repository	Manages the software images repository on the server.
Deploy Images to Devices	Deploys software images from the repository to devices.
View Image Deployment Jobs	Manages software image deployment jobs.

Table 139: Device Management Tasks

Task	Description
Convert Ports	Converts QSFP+ port configuration.
Manage Port Admin State	Enables or disables switch network ports.
Resynchronize Device Configuration	Resynchronizes the device configuration maintained in Build mode with the running configuration on the devices.
Show Current Configuration	Shows the selected device's current configuration.
View Inventory	Displays the device inventory of the selected node.

Table 140: Device Configuration File Management Tasks

Task	Description
Manage Device Configuration Files	Manages backup device configuration files.
View Configuration File Mgmt Jobs	Manages device configuration file management jobs.

Table 141: Baseline Management Tasks

Task	Description
Manage Baseline	Manages baseline configuration files.
View Baseline Mgmt Jobs	Manages baseline configuration file management jobs.

Table 142: Zero Touch Provisioning Tasks

Task	Description
Setup	Set up the zero touch provisioning profile to configure the DHCP server and to upload the software image and configuration file to a file server.
Monitor	View details of the devices that are provisioned using a given zero touch provisioning profile.

RELATED DOCUMENTATION

Understanding Deploy Mode in Network Director 572
Understanding the Network Director User Interface 3
Network Director Documentation home page

Deploying and Managing Device Configurations

IN THIS CHAPTER

- [Deploying Configuration to Devices | 580](#)
- [Managing Configuration Deployment Jobs | 594](#)
- [Deploy Configuration Window | 596](#)
- [Importing Configuration Data from Junos OS Configuration Groups | 598](#)
- [Enabling High-Frequency Traffic Statistics Monitoring on Devices | 602](#)
- [Configuring Network Traffic Analysis | 604](#)
- [Approving Change Requests | 606](#)
- [Enabling SNMP Categories and Setting Trap Destinations | 608](#)
- [Understanding Resynchronization of Device Configuration | 611](#)
- [Resynchronizing Device Configuration | 616](#)
- [Managing Device Configuration Files | 622](#)
- [Creating and Managing Baseline of Device Configuration Files | 626](#)

Deploying Configuration to Devices

IN THIS SECTION

- [Selecting Configuration Deployment Options | 581](#)
- [Using the Change Request Details Page | 585](#)
- [Creating a Change Request | 586](#)
- [Validating Configuration | 586](#)
- [Discarding the Pending Configurations | 587](#)
- [Viewing Pending Configuration Changes | 587](#)
- [Using the Pending Changes Window | 587](#)
- [Using the Configuration or Pending Configuration Window | 588](#)
- [Using the Deploy Configuration Errors/Warnings Window | 589](#)

- [Using the Configuration Validation Window | 589](#)
- [Deploying Configuration Changes to Devices Immediately | 589](#)
- [Scheduling Configuration Deployment | 590](#)
- [Specifying Configuration Deployment Scheduling Options | 590](#)
- [Editing Change Requests | 591](#)
- [Deleting Change Request | 592](#)
- [Resubmitting a Change Request | 592](#)
- [Performing a Rollback | 593](#)

When you make configuration changes in Build mode, the changes are not deployed to devices automatically. You must manually deploy the changes to devices in Deploy mode.

To start deploying configuration changes:

1. Click **Deploy** in the Network Director banner.
2. Select a node in the View pane that contains the devices to which you want to deploy.
3. In the Tasks pane, select **Configuration Deployment > Deploy Configuration Changes**.

Depending upon the type of approval mode you select different windows are displayed.

If you select the Auto Approval mode, the Devices with Pending Changes page opens in the main window, listing the devices within the selected node that have pending configuration changes.

If you select the Manual Approval mode, the following two sections open in the main window:

- **Devices with recent configuration changes**—This section lists the devices with pending changes (along with the details of the change) performed by the user currently logged into the system.
- **Change Requests**—This section lists the change requests created by the user currently logged into the system.

This topic describes:

Selecting Configuration Deployment Options

Based on the approval mode, you can choose to deploy the device configuration changes in the following ways:

- When you select the auto approval mode, the page Devices with Pending Changes open. From the Devices with Pending Changes page, you can:

- Deploy configuration changes immediately by selecting one or more devices and clicking Deploy Now. For more information, see [“Deploying Configuration Changes to Devices Immediately” on page 589](#).
- Schedule configuration deployment by selecting one or more devices and clicking Schedule Deploy. For more information, see [“Scheduling Configuration Deployment” on page 590](#).
- View configuration changes that are pending on a device by clicking View in the Configuration Changes column. For more information, see [“Viewing Pending Configuration Changes” on page 587](#).
- Validate that the pending changes for a device are compatible with the device’s configuration by selecting up to ten devices and clicking Validate Pending Configuration Changes. For more information, see [“Validating Configuration” on page 586](#).
- Discard the pending configuration changes. For more information, see [“Discarding the Pending Configurations” on page 587](#).

[Table 143 on page 582](#) describes the information provided in the table on the Devices with Pending Changes page. Only the subset of devices within the selected object that have pending configuration changes are listed in the table.

Table 143: Devices with Pending Changes Page

Table Column	Description
Check box	Select to perform an action on the device in that row
Name	Device name
IP Address	Device IP address
Model	Device Model
OS Version	Operating system version running on device
Connection State	<p>State of the connection to the device:</p> <ul style="list-style-type: none"> • Up—Network Director can communicate with the device. • Down—Network Director cannot communicate with the device. You cannot deploy configuration to devices that are down.

Table 143: Devices with Pending Changes Page (continued)

Table Column	Description
Configuration State	<p>Indicates whether the device's configuration is in sync with Network Director's version:</p> <ul style="list-style-type: none"> • In Sync—The configuration on the device is in sync with the Network Director configuration for the device. • Out Of Sync—The configuration on the device does not match the Network Director configuration for the device. This state is usually the result of the device configuration being altered outside of Network Director. <p>You cannot deploy configuration on a device when the device is Out Of Sync. To resolve this state, use the Resynchronize Device Configuration task in Deploy mode.</p> <ul style="list-style-type: none"> • Synchronizing—The device configuration is in the process of being resynchronized. • Sync failed—An attempt to resynchronize an Out Of Sync device failed.
Configuration Changes	<p>Click to view pending configuration changes for a device. The Pending Changes window opens.</p>

If you select the Manual Approval mode, the windows Devices with recent configuration changes and Change Requests opens.

From the Devices with recent configuration changes window, you can:

- Create a device configuration change request approval and submit it for approval. Upon submission, all device changes made by an operator are validated and all the approvers are notified of the details of the proposed change request by e-mail. For more information, see [“Creating a Change Request” on page 586](#).
- View configuration changes that are pending on a device by clicking View in the Configuration Changes column. For more information, see [“Viewing Pending Configuration Changes” on page 587](#).
- Validate that the pending changes for a device are compatible with the device's configuration . For more information, see [“Validating Configuration” on page 586](#).
- Discard the pending configuration changes. For more information, see [“Discarding the Pending Configurations” on page 587](#).

NOTE: You cannot delete a device from the Devices with Pending Changes list. To remove a device from the list, you must undo the Build mode configuration changes that placed the device on the list.

Table 144 on page 584 describes the information provided in the table on the Devices with recent configuration changes page.

Table 144: Devices with recent configuration changes

Table Column	Description
Name	Indicates the name of the device and profile node. Below each device node, a profile node is listed.
Change Type	Indicates the type of the configuration change done to the device.
Associations Added	Lists the ports that are added to that profile.
Associations Deleted	Lists the ports that are deleted from that profile.
Configuration	Click to view pending configuration changes for a device. The Pending Changes window opens.
Deployment State	Indicates the deployment state of a change request.

From the Change Requests window, you can:

- Deploy configuration changes immediately by selecting one or more devices and clicking Deploy Now. For more information, see [“Deploying Configuration Changes to Devices Immediately” on page 589](#).
- Schedule configuration deployment by selecting one or more devices and clicking Schedule Deploy. For more information, see [“Scheduling Configuration Deployment” on page 590](#).
- Resubmit for the change request for approval after making the necessary modifications. For more information, see [“Resubmitting a Change Request” on page 592](#).
- Edit or delete the change requests by selecting one or more change requests and clicking Edit or Delete respectively. For more information, see [“Editing Change Requests” on page 591](#) and [“Deleting Change Request” on page 592](#).
- Roll back the device configuration that is already deployed. For more information, see [“Performing a Rollback” on page 593](#).
- View the details of the change request created. For more information, see [“Using the Change Request Details Page” on page 585](#)

[Table 145 on page 584](#) describes the information provided in the table on the change requests submitted for the devices for which configuration changes are sought.

Table 145: Change Requests

Table Column	Description
Check Box	Select to perform an action on the device in that row.

Table 145: Change Requests (*continued*)

Table Column	Description
Change Request No	Indicates the change request number of the change request that is waiting to be deployed.
Title	Indicates the title name of the change request.
Created On	Indicates the change request creation date.
Approver	Indicates the username of the configuration approver.
Last Action On	Indicates the date on which the change request status is changed.
Approval Status	Indicates whether a change request is approved or rejected by the approver.
Deployment Status	Indicates whether a change request is deployed after the approval.
History Icon	Records the audit trail details of a change request, such as operation performed on a change request during a given period of time, username of the approver or operator, and so on.

Using the Change Request Details Page

Use the Change Request Details window to view the details of the change request before you either approve or reject a change request. This window provides you the details such as change request number, title, username of the user who created the change request, change request creation date and so on. A Devices table is also displayed showing the deployment status. [Table 146 on page 585](#) describes the fields in this table.

Table 146: Change Request Details

Column	Description
Name	Indicates the name of the device and profile node. Below each device node, a profile node is listed.
Change Type	Indicates the type of the configuration change done to the device.
Associations Added	Lists the ports that are added to that profile.
Associations Deleted	Lists the ports that are deleted from that profile.

Table 146: Change Request Details *(continued)*

Configuration	Click to view pending configuration changes for a device. The Pending Changes window opens.
Deployment Status	Indicates the deployment state of a change request.

Creating a Change Request

To create a change request for device configurations approval:

1. Click **Create Change Request** in the Devices with recent configuration changes page.

The Create Change Request page opens.

2. Enter the change request number.

You can either enter a number or retain the autogenerated number in this field.

3. Enter an appropriate title name for the change request.

4. Optionally, you can enter comments for the device configuration changes.

5. Click **Submit**.

The Create Change Request page opens, listing the change request details such as change request number, title, and comment, along with the change request submission job details. A Devices table is also displayed showing the validation status of the device and configuration generated for that device.

6. Click **Close**.

A new change request entry with the status Pending Approval is added to the Change Request section.

Validating Configuration

When you deploy configuration changes to a device, validation checks are performed to validate that the pending changes are compatible with the device. You can also perform this validation without deploying.

NOTE: You can also verify the configuration from the Build mode by clicking **Tasks > Domain Management > Validate Pending Configuration**.

To validate that the pending changes for devices are compatible with the device configuration:

1. For Auto Approval mode, select up to ten devices in the Devices with Pending Changes page.

NOTE: For Manual Approval mode, you cannot choose the devices for which validation needs to be done. All the configuration changes for all the devices are validated.

2. Click **Validate Pending Configuration Changes**.

The Configuration Validation window opens. See [“Using the Configuration Validation Window” on page 589](#) for a description of the window.

Discarding the Pending Configurations

Use the Discard Local Configuration Changes Results window to discard all the pending configurations that were made on a device. Once you discard the local configuration changes on a device, the configuration state of the device changes to In Sync or Out of Sync based on the system of record (SOR) mode set for the Junos Space Network Management Platform. If the SOR mode is set to Network as system of record (NSOR), then the configuration state changes to In Sync and if the SOR mode is set to Junos Space as system of record (SSOR), then the configuration state changes to Out of Sync.

To discard the configuration changes:

1. For Auto Approval mode, select the devices for which you want to discard the pending configuration and click **Discard Pending Configuration**.

The Discard Local Configuration Changes Results window opens displaying the status of the discard pending configuration job.

2. Click **Close** to close the Discard Local Configuration Changes Results window.

Viewing Pending Configuration Changes

To view pending configuration changes for a device, click **View** in the Pending Changes column.

The Pending Changes window opens. See [“Using the Pending Changes Window” on page 587](#) for a description of the window.

Using the Pending Changes Window

Use the Pending Changes window to view the pending Network Director changes for a device. [Table 147 on page 588](#) describes the fields in this window.

Table 147: Pending Changes Window

Field	Description
Name	Lists each selected device. Expand a device by clicking its plus sign to see its pending changes. Each pending change to a profile or other configuration object for the device is listed.
State	<p>Describes the nature of the pending change to the configuration object. These are the possible states:</p> <ul style="list-style-type: none"> • Added—The profile or configuration object was added to this device. • Removed—The profile or configuration object was removed from the device • Updated—The profile or configuration object was updated.
Configuration	<p>Click View to view the pending configuration changes for a device. The Pending Configuration window opens. See “Using the Configuration or Pending Configuration Window” on page 588 for information about the window.</p> <p>NOTE: The device configuration state must be In Sync for you to view the pending configuration changes.</p>
Close	Click to close the window.

Using the Configuration or Pending Configuration Window

Use the Pending Configuration window to view the configuration changes that will be deployed to a device when a job runs. Use the Configuration window to see changes that were deployed to a device when a completed job ran. The configuration changes are shown in these formats:

- Select the **XML View** tab to view the configuration changes in XML format. This view shows the XML-formatted configuration that will be deployed to the device’s Device Management Interface (DMI), which is used to remotely manage devices.
- Select the **CLI View** tab to view the configuration changes in CLI format. This view shows the Junos configuration statements that will be deployed to the device.

In both views, the content is color-coded for easier reading:

- Black text indicates configuration that is already active on the device, and will not be changed if you deploy.
- Green text indicates configuration that will be added if you deploy.
- Red text indicates configuration that will be removed if you deploy.

Using the Deploy Configuration Errors/Warnings Window

Use the Deploy Configuration Errors/Warnings window to view the results of deploying configuration to a device. The Errors/Warnings in validating the device configuration pane shows the results of configuration validation by Network Director. The Errors/Warnings in Updating Device configuration pane shows the results of configuration validation on the device.

Using the Configuration Validation Window

Use the Configuration Validation window to validate that the pending changes for a device are compatible with the device's configuration. [Table 148 on page 589](#) describes this window.

Table 148: Configuration Validation Window

Table Column	Description
Object name	Lists the devices you selected for validation. Click the arrow next to a device to expand it. If there are no errors or warnings, one item labeled No Validation warnings appears. If the device has errors or warnings, they appear under the device. The device contains a list of the profiles that caused errors or warnings. Expand a profile name to see the of errors and warnings it caused.
Errors/Warnings	Describes the error or warning.

Deploying Configuration Changes to Devices Immediately

To deploy configuration changes to devices immediately:

1. Select the device or devices in the Devices with Pending Changes page.
2. Click **Deploy Now**.

The Deploy Options window opens.

3. In the Deploy Options window, enter a job name in the Deployment Job Name field, then click **OK**.

The configuration deployment job runs. The Deploy Configuration window opens and shows the results of the deployment job. For a description of fields in this window, see ["Deploy Configuration Window" on page 596](#).

Scheduling Configuration Deployment

To schedule configuration deployment to devices:

- 1. Select the device or devices in the Devices with Pending Changes page.
- 2. Click **Schedule Deploy**.

The Deploy Options window opens.

- 3. Use the Deploy Options window to schedule the configuration deployment. See [“Specifying Configuration Deployment Scheduling Options” on page 590](#) for a description of the window.

Specifying Configuration Deployment Scheduling Options

Use the Deploy Options window to schedule configuration deployment jobs. [Table 149 on page 590](#) describes the actions for the fields in this window.

Table 149: Deploy Options Window

Field	Action
Deployment Job Name	Enter a job name.
Date and Time	Enter the job’s start date and time.
OK	Click to accept changes and exit the window.
Cancel	Click to cancel changes and exit the window.

Editing Change Requests

You can edit a change request to change the profile that was added to a device or delete some of the profile associations. After editing a change request, you can resubmit the change request for approval. While editing a change request, if you try to delete all the profile associations in a given change request, the system prompts a message that a change request should have at least one valid association. Deleting all the associations in a change request makes it invalid. Hence, you cannot delete all the associations in a given change request. However, you can delete a change request itself to delete all the associations for that change request.

NOTE: You are unable to delete a change request or an association of a change request if an association is in pending removal state.

You are unable to edit a change request that is in Cancelled, Deployed, Rollback Success, or Rollback Failed state.

To change a profile or delete the profile associations of a change request:

1. Select the change request in the Change Requests pending action page to edit.

2. Click **Edit**.

The Edit Change Request window opens.

3. Click the call out symbol to change the profile and choose the new profile that you want to assign the change request.

4. To delete a profile association, click **Delete**

5. Click **Save**.

The Edit Change Request window opens, listing the change request details such as change request number, title, and comment, along with the change request submission job details. A Devices table is also displayed showing the validation status of the device and configuration generated for that device.

6. Click **Close**.

Deleting Change Request

Sometimes you might need to delete a change request from the change request list. A change request is assigned with profile associations. If you delete a change request, all the associations of that change request are also deleted.

To delete a change request:

1. Select the change request or change requests in the Change Requests pending action window.
2. Click **Delete**.

The Delete Change Request window opens, displaying the message: **Are you sure you want to delete Change Request?**.

3. Click **Yes** to delete the change request; else click **No**.

If you clicked **Yes**, the message: **Change Request deleted successfully** appears.

4. Click **OK**.

Resubmitting a Change Request

You can resubmit only those change requests that are in Pending Approval, Pending Deployment, Deploy Failed, and Create Failed state. You are unable to resubmit change requests in Deployed, Cancelled, Rollback Success, or Rollback Failure state.

In certain situations, a device can go out of sync while a user is creating a change request for that device. The change request is created, but the configuration changes for that change request are not generated. You can select the change request and resubmit it after the device is in sync again, which generates the configuration for this change request. You can resubmit change requests only for devices that have pending configuration changes.

To resubmit a change request:

1. Select the change request in the Change Requests pending action window to edit.
2. Click **Resubmit**.

A warning message pops up indicating if you want to resubmit the change request.

3. Click **Yes**.

The Resubmit Change Request window opens, listing the change request details such as change request number, title, and comment, along with the change request submission job details. A Devices table is also displayed showing the validation status of the device and configuration generated for that device.

4. Click **Close**.

Performing a Rollback

In case of any misconfigurations, you can choose to roll back a configuration that has already been deployed to the device. The following conditions apply for a rollback operation:

- The maximum number of change requests that you can roll back is the rollback limit specified in Preferences.
- Change requests are rolled back in reverse chronological order; the later change requests are rolled back first. If there are any conflicting change requests, roll back is not supported. For example, assume that a user assigns port profile P1 to ge-0/0/1 and creates a change request CR1 and deploys the profile. After this, if the user edits P1, creates another change request CR2 and deploys and removes P1 from the port by assigning some other port profile and deploys device changes or configurations as part of CR3. If the user now tries to roll back CR1, an error message about the conflicting change requests CR2 and CR3 is shown. To roll back CR1, the user must roll back CR3, then CR2, and then CR1.

To roll back a device:

1. Select the device in the left navigation pane for which you want to perform the roll back operation.
2. Select **Rollback Configuration Changes** task under Configuration Deployment.
3. All the devices with previously stored configuration of the device are listed in the working area of the right pane.

You can view the stored configuration also user can view the difference of the current device configuration and stored configuration.

4. Choose the configuration for which you want to perform the roll back.

NOTE: You can choose only one rollback configuration out of the available configurations per device however you can choose multiple devices.

5. Click **Rollback**.

A rollback job is started with all the selected devices and all the devices are resynchronized after the configuration is pushed.

RELATED DOCUMENTATION

- [Deploying Configuration Changes | 573](#)
- [Managing Configuration Deployment Jobs | 594](#)
- [Approving Change Requests | 606](#)
- [Setting Up User and System Preferences | 26](#)
- [Network Director Documentation home page](#)

Managing Configuration Deployment Jobs

IN THIS SECTION

- [Selecting Configuration Deployment Job Options | 594](#)
- [Viewing Configuration Deployment Job Details | 595](#)
- [Canceling Configuration Deployment Jobs | 596](#)

When you deploy configuration changes or schedule a configuration deployment, a configuration deployment job is created.

To start managing configuration deployment jobs:

1. Click **Deploy** in the Network Director banner.
2. In the Tasks pane, select **Device Management > View Deployment Jobs**.

The Deploy Configuration page opens in the main window. The table on that page lists configuration deployment jobs.

This topic describes:

Selecting Configuration Deployment Job Options

From the Deploy Configuration page, you can:

- View the details of a configuration deployment job by clicking Show Details. See [“Viewing Configuration Deployment Job Details” on page 595](#) for more information.

- Cancel a scheduled configuration deployment job by clicking Cancel Job. See [“Canceling Configuration Deployment Jobs” on page 596](#) for more information.

[Table 150 on page 595](#) describes the information provided on the Deploy Configuration page

Table 150: Deploy Configuration Table Description

Table Column	Description
Check box	Select to perform an action on the job in that row.
Job Id	An identifier assigned to the job.
Job Name	Job name (user-created).
Percent	Percentage of the job that is complete.
Status	<p>Job status. The possible states are:</p> <ul style="list-style-type: none"> • CANCELLED—The job was cancelled by a user. • FAILURE—The job failed. This state is applied if any of the devices in the job failed. But some of the devices might have completed successfully. View the job details for the status of each device. • INPROGRESS—The job is running. • SCHEDULED—The job is scheduled but has not run yet. • SUCCESS—The job completed successfully. This state is applied if all of the devices in the job completed successfully.
Summary	Job summary.
Scheduled Start Time	Job's scheduled start time
Actual Start Time	Time when the job started.
End Time	Time when the job ended
User	User who created the job
Recurrence	This field is not used for configuration deployment jobs.

Viewing Configuration Deployment Job Details

To view the details of a configuration deployment job:

1. Select the job in the table.

- 2. Click **Show Details**. The Deploy Configuration window opens. See [“Deploy Configuration Window” on page 596](#) for a description of the window.

Canceling Configuration Deployment Jobs

To cancel a configuration deployment job:

- 1. Select the job in the table.
- 2. Click **Cancel Job**.
- 3. Click **Yes** in the confirmation window that opens.

RELATED DOCUMENTATION

Deploying Configuration Changes 573
Deploying Configuration to Devices 580
Network Director Documentation home page

Deploy Configuration Window

The Deploy Configuration window shows the results of a completed deployment job or information about a scheduled job. See [Table 151 on page 596](#) for a description of the fields in this window.

Table 151: Deploy Configuration Window

Field	Description
Job Name	Job name.
Job Start Time	Time when job started or will start.
Job End Time	Time when job ended.
Percentage Completed	Percentage of the job that is complete.
Number of Devices	Number of devices in the deployment job.
Deployed Devices table	

Table 151: Deploy Configuration Window (*continued*)

Field	Description
Name	Device name.
IP Address	Device IP address.
Deployment Status	<p>Status of configuration deployment on device:</p> <ul style="list-style-type: none"> • Scheduled—Job is scheduled for future deployment. • In Progress—Deployment is in progress. • Success—Deployment completed successfully. • Failed—Deployment failed.
Configuration	<p>Click View to see the configuration changes that were deployed to the device. See “Using the Configuration or Pending Configuration Window” on page 588 for more information.</p> <p>For a scheduled job, this column does not contain a link. See “Deploying Configuration to Devices” on page 580 for information about viewing pending configuration changes for a device.</p>
Result Details	Click View to see the results of configuration deployment for the device. See “Using the Deploy Configuration Errors/Warnings Window” on page 589 for more information.
Close	Click to close the window.

RELATED DOCUMENTATION

[Deploying Configuration Changes | 573](#)
[Deploying Configuration to Devices | 580](#)
[Managing Configuration Deployment Jobs | 594](#)
[Network Director Documentation home page](#)

Importing Configuration Data from Junos OS Configuration Groups

IN THIS SECTION

- [Enabling Import of Configuration Group Data | 598](#)
- [Viewing Configuration Group Data | 599](#)
- [Using the Configuration or Pending Configuration Window | 600](#)
- [Deploying Configuration Group Changes to Devices Immediately | 601](#)
- [Scheduling Configuration Group Change Deployment | 601](#)
- [Specifying Configuration Deployment Scheduling Options | 601](#)
- [Using the Deploy Configuration Errors/Warnings Window | 602](#)

The configuration groups feature in Junos OS enables you to create a group containing configuration statements and to direct the inheritance of that group's statements in the rest of the configuration on a device. The same group can be applied to different sections of the configuration, and different sections of one group's configuration statements can be inherited in different places in the configuration.

Configuration groups enable you to create smaller, more logically constructed configuration files, making it easier to configure and maintain Junos OS. For example, you can group statements that are repeated in many places in the configuration, such as when configuring interfaces, and thereby limit updates to just the group.


You can use the Migrate Config Groups task in Network Director to import and deploy supported configurations from these configuration groups on devices.

This topic describes:

Enabling Import of Configuration Group Data

For Network Director to be able to import configuration group data.

To enable the import of configuration group data:

1. Click  in the Network Director banner and select **Preferences**.
2. In the Preferences window, select the **Config & Deploy** tab.

3. Select the **Enable migration from Ethernet Design** check box to enable import of configuration group data.
4. Click **Save** to save and close the preferences.

Viewing Configuration Group Data

After you enable the import of configuration group data, Network Director adds a new task—Migrate from Ethernet Design—in the Deploy mode. You can use this task to flatten the active group configurations. All the devices which are discovered and managed by Network Director is listed in the Migrate from Ethernet Design page. Network Director displays the devices in this page based on you selection in the Tree view. You can view the default configuration, active configuration, and the configurations that are to be deployed on the devices from this page.

To view the configuration group data:

1. Click **Deploy** in the Network Director banner.
2. In the Tasks pane, select **Configuration Deployment > Migrate From Ethernet Design**. The Migrate Ethernet Design Configuration Groups page opens displaying all the devices that has configurations from configuration groups that are not deployed on the device yet.

[Table 152 on page 599](#) describes the information provided in the table on the Migrate Config Groups page.

Table 152: Migrate Config Groups Page

Table Column	Description
Check box	Select to perform an action on the device in that row
Name	Device name
IP Address	Device IP address
Model	Device Model
Device Family	Device family can be junos, junos-ex, or junos-qfx
OS Version	Operating system version running on device

Table 152: Migrate Config Groups Page (*continued*)

Table Column	Description
Connection State	<p>State of the connection to the device:</p> <ul style="list-style-type: none"> • Up—Network Director can communicate with the device. • Down—Network Director cannot communicate with the device. You cannot deploy configuration to devices that are down.
Configuration State	<p>Indicates whether the device's configuration is in sync with Network Director's version:</p> <ul style="list-style-type: none"> • In Sync—The configuration on the device is in sync with the Network Director configuration for the device. • Out Of Sync—The configuration on the device does not match the Network Director configuration for the device. This state is usually the result of the device configuration being altered outside of Network Director. <p>You cannot deploy configuration on a device when the device is Out Of Sync. To resolve this state, use the Resynchronize Device Configuration task in Deploy mode.</p> <ul style="list-style-type: none"> • Synchronizing—The device configuration is in the process of being resynchronized. • Sync failed—An attempt to resynchronize an Out Of Sync device failed.
Configuration Changes	<p>Click to view the configuration changes for a device. The Pending Changes window opens. For more details on using the Pending Changes window, see “Using the Configuration or Pending Configuration Window” on page 600.</p>

Using the Configuration or Pending Configuration Window

Use the Pending Configuration window to view the configuration changes that will be deployed to a device when a job runs. Use the Configuration window to see changes that were deployed to a device when a completed job ran. The configuration changes are shown in these formats:

- Select the **XML View** tab to view the configuration changes in XML format. This view shows the XML-formatted configuration that will be deployed to the device's Device Management Interface (DMI), which is used to remotely manage devices.
- Select the **CLI View** tab to view the configuration changes in CLI format. This view shows the Junos configuration statements that will be deployed to the device.

In both views, the content is color-coded for easier reading:

- Black text indicates configuration that is already active on the device, and will not be changed if you deploy.
- Green text indicates configuration that will be added if you deploy.
- Red text indicates configuration that will be removed if you deploy.

Deploying Configuration Group Changes to Devices Immediately

To deploy configuration group changes to devices immediately:

- 1. Select the device or devices in the Devices with Pending Changes page.
- 2. Click **Deploy Now**.

The Deploy Options window opens.

- 3. In the Deploy Options window, enter a job name in the Deployment Job Name field, then click **OK**.

The configuration deployment job runs. The Deploy Configuration window opens and shows the results of the deployment job. For a description of fields in this window, see [“Deploy Configuration Window” on page 596](#).

Scheduling Configuration Group Change Deployment

To schedule configuration group change deployment to devices:

- 1. Select the device or devices in the Devices with Pending Changes page.
- 2. Click **Schedule Deploy**.

The Deploy Options window opens.

- 3. Use the Deploy Options window to schedule the configuration deployment. See [“Specifying Configuration Deployment Scheduling Options” on page 601](#) for a description of the window.

Specifying Configuration Deployment Scheduling Options

Use the Deploy Options window to schedule configuration deployment jobs. [Table 153 on page 601](#) describes the actions for the fields in this window.

Table 153: Deploy Options Window

Field	Action
Deployment Job Name	Enter a job name.
Date and Time	Enter the job's start date and time.
OK	Click to accept changes and exit the window.
Cancel	Click to cancel changes and exit the window.

Using the Deploy Configuration Errors/Warnings Window

Use the Deploy Configuration Errors/Warnings window to view the results of deploying configuration to a device. The Errors/Warnings in validating the device configuration pane shows the results of configuration validation by Network Director. The Errors/Warnings in Updating Device configuration pane shows the results of configuration validation on the device.

RELATED DOCUMENTATION

[Deploying Configuration Changes | 573](#)

[Managing Configuration Deployment Jobs | 594](#)

[Network Director Documentation home page](#)

Enabling High-Frequency Traffic Statistics Monitoring on Devices

To use Network Director monitoring analytics features such as latency heat maps and congestion monitoring, you must enable high-frequency traffic statistics monitoring on the QFX devices to be monitored. You can also configure the high-frequency traffic statistics sampling rate and event thresholds on devices.

NOTE: You must specify the IP address of DLE server under **Preferences > Monitoring > Data Learning Engine Settings** before you can enable high-frequency traffic statistics monitoring as described in this topic.

To enable high-frequency traffic statistics monitoring on devices:

1. Log in to Network Director.
2. Under Views, select **Logical View**, **Location View**, **Device View**, or **Custom Group View**.
3. Click **Deploy** in the Network Director banner to open Deploy mode.
4. Select the task **Configuration Deployment > Enable High Frequency Stats** in the Tasks pane.

The Enable High Frequency Stats page opens. It contains a table listing the QFX devices that support high-frequency statistics monitoring.

5. To enable or disable high-frequency traffic statistics monitoring on a device, use the check box in the Enable column.

- 6. In the Device / Port column, specify whether you want to enable high-frequency traffic statistics monitoring on a device as a whole or on selected ports.

If you select **Device**, high-frequency traffic statistics monitoring is enabled on all ports on the device with the settings you specify. If you select **Port**, you can selectively enable/disable high-frequency traffic statistics monitoring on individual ports on the device and specify different settings for each port.
- 7. If you have selected **Port** in the previous step, click the arrow next to the device name to expose the list of ports and to enable high-frequency traffic statistics monitoring on selected ports.
- 8. To change high-frequency traffic statistics monitoring settings, double-click a current setting in the table to edit it. [Table 154 on page 603](#) describes these settings.
- 9. To deploy the settings currently configured on the page, click **Deploy**.
- 10. To reset all settings on the page to the default values and deploy those settings, click **Restore all values to default** and **Deploy**.

Table 154: High-Frequency Traffic Statistics Monitoring Settings

Setting	Description
Traffic Stats Sampling Interval (seconds)	Sets the interval for traffic statistics sampling, in seconds.
Latency Stats Sampling Interval (milli seconds)	Sets the interval for latency statistics sampling, in milliseconds.
Latency Threshold (nano seconds)	Sets the latency threshold, in nanoseconds. Monitored latency values higher than this threshold are considered congestion events.

RELATED DOCUMENTATION

Understanding Deploy Mode in Network Director 572
Understanding Monitor Mode in Network Director 658

Configuring Network Traffic Analysis

Network Traffic Analysis (NTA) is a monitoring technology for high-speed switched or routed networks. Once enabled, Network Director randomly samples network packets and sends the samples to a data learning engine (DLE) for analysis. Network traffic analysis uses packet-based sampling. Network Director samples one packet out of a specified number of packets from an interface enabled for network traffic analysis and sends the packet to the DLE. DLE uses this sampling information to create a picture of the network traffic, which includes the applications that contribute to the traffic, traffic statistics, and the top applications.

You can enable network traffic analysis on all devices.

Before you configure network traffic analysis, ensure that:

- You have configured one or more data learning engines to analyze the network traffic.
- You have specified the IP address and port number of these DLEs in the System Preferences window. For detailed steps, see the Specifying the Data Learning Engine (DLE) Settings in the [“Setting Up User and System Preferences” on page 26](#) topic.

NOTE: On devices for which you want to enable NTA, sFlow must not be configured. Before you enable NTA, make sure that sFlow is not enabled on the devices to be monitored. To verify that the sFlow is not enabled, log in to the device CLI and execute the following command and verify the output:

```
[root@user ~]# show protocols sflow | display set
```

No output is displayed when sFlow is not configured on the device.

NOTE: Network Director support the sFlow on physical port only and not in aggregation links or logical ports.

To configure network traffic analysis:

1. Log in to Network Director.
2. Under Views, select **Logical View**, **Location View**, **Device View**, or **Custom Group View**.
3. Click **Deploy** in the Network Director banner.
4. In the Tasks pane, select **Configuration Deployment > Enable Network Traffic Analysis**.

The Enable Network Traffic Analysis page opens.

5. Select the check box adjacent to **Enable Traffic Analysis on Devices when Port Utilization exceeds certain percentage** to enable network traffic analysis. The default port utilization percentage value is 90%. You can change the default value to a value that is appropriate for your network.

6. Specify the number of packets from which a packet must be sampled in the **Sample rate** field.

Network Director samples one packet out of a specified number of packets from an interface enabled for network traffic analysis and sends the packet to DLE.

For example, if you configure a sampling rate of 10, one packet is sampled from every 10 packets.

7. Click **Add Devices** to add new devices for network traffic analysis.

The Add Devices window opens.

8. In the Add Devices window, select the devices for which you want to enable network traffic analysis.

9. Click **Add**.

Network Director adds the selected devices to the list in the Enable Network Traffic Analysis page.

To remove a device, select a device from the list and click **Remove**.

10. Click **Save** to save the network traffic analysis configuration details.

Network Director initiates traffic analysis when traffic utilization on any interface of the devices added to the Enable Network Traffic Analysis page exceeds the port utilization that you specified. You can view the traffic analysis details from the **Monitor** mode > **Traffic Analysis** or the Device & Port Utilization dashboard widget.

RELATED DOCUMENTATION

[Device & Port Utilization Widget](#) | 63

[Monitoring Port Traffic Statistics](#) | 671

Approving Change Requests

NOTE: This option is available only for the users who are assigned a Configuration Approver role.

When you select the Approve Change Request option, the page Change request(s) pending approval and the page approved/declined change request(s) open in the top and bottom panels respectively.

The [Table 155 on page 606](#) shows details of the change requests that are pending for approval by the approver.

Table 155: Change request(s) pending approval

Table Column	Description
Change Request No	Indicates the change request number that was either approved or rejected by the approver.
Title	Indicates the title of the change request.
Created By	Indicates the operator name who created the change request and submitted it for approval.
Created On	Indicates the date on which the change request was created.
Age	Indicates the age of the change request, time since the change request was created.
Deployment Status	Indicates the deployment state of the change request.
History Icon	Records the audit trail details of a change request, such as operation performed on a change request during a given period of time, username of the operator, and so on.

The [Table 156 on page 606](#) shows the change requests that were approved or rejected by the currently logged in approver. The approver can also provide comments

Table 156: approved/declined change request(s)

Table Column	Description
Change Request No	Indicates the change request number that was either approved or rejected by the approver.

Table 156: approved/declined change request(s) (*continued*)

Table Column	Description
Title	Indicates the title of the change request.
Created By	Indicates the operator name who created the change request and submitted it for approval.
Created On	Indicates the date on which the change request was created.
Age	Indicates the age of the change request, time since the change request was created.
Approval Status	Indicates the approval state of the change request.
Deployment Status	Indicates the deployment state of the change request.
History Icon	Records the audit trail details of a change request, such as operation performed on a change request during a given period of time, username of the approver, and so on.

To approve or reject the change requests submitted by an operator:

1. Select **Approve Change Requests** under Configuration Deployment.
2. Select the check box against the change request and click on a change request in the change request(s) pending approval page.
The Change Request Details page opens.
3. Review details of the profile and its associations.
4. Click on the **View** link.
The Pending Configuration device name page opens.
5. Click **Close**.
6. Click **Approve** or **Reject** to approve or reject the device configuration changes respectively.
The Change Request Details page opens.
7. Type your comments and click **Approve** to approve; else click **Reject**.
After the successful approval, you can deploy the device configurations immediately or schedule the deployment for a later period.

RELATED DOCUMENTATION

[Setting Up User and System Preferences | 26](#)

[Deploying Configuration to Devices | 580](#)

[Network Director Documentation home page](#)

Enabling SNMP Categories and Setting Trap Destinations

IN THIS SECTION

- [Viewing Eligible Devices for Trap Forwarding | 608](#)
- [Enabling Trap Forwarding | 609](#)
- [Deploying SNMP Trap Configurations | 610](#)

SNMP traps must be enabled on network devices for Network Director to collect and manage event and error information from these devices.

Network Director organizes switch by categories. These categories must be enabled and deployed in order to forward trap information to Network Director.

NOTE: Network Director uses protocol port 10162 for receiving traps from devices. This port must be open on the devices.

NOTE: Network Director only supports SNMP V1 and V2C traps.

This topic describes:

Viewing Eligible Devices for Trap Forwarding

Traps are enabled on the Devices page in Deploy mode. To locate this page:

1. Select **Deploy** in the Network Director banner.

2. Select **Set SNMP Trap Configuration** in the Tasks pane. The Devices page opens. For a description of fields in the Devices page, view [Table 157 on page 609](#).

Table 157: Device Page Fields

Field	Description
Name	Either the hostname or the IP address of the device.
IP Address	Device IP address.
Model	Device model number.
OS Version	Version and release level of the operating system running on the device.
Connection State	State of connection to the device. Valid states are: <ul style="list-style-type: none"> • Up—Network Director is in communication with the device. • Down—Network Director cannot communicate with the device. You cannot enable traps on devices that are in this state.
Configuration State	Either the device's configuration is in sync or out-of-sync with Network Director's version: <ul style="list-style-type: none"> • IN_SYNC—The configuration is in-sync with the database. • OUT_OF_SYNC—The configuration is out-of-sync with the database.

Enabling Trap Forwarding

Select **Set SNMP Trap Configuration** in Deploy mode to enable your network devices to pass SNMP traps and events to Network Director. Network Director creates a target group called *networkdirector_trap_group* using target port 10162. The Community name is *public* and the access is *read-write-notify*.

Before enabling trap forwarding, complete device discovery for all the devices and ensure they are in the up state. Down devices cannot be enabled for trap forwarding.

Selecting Set SNMP Trap Configuration displays the Devices page which contains a table of all discovered switches in the network. To enable SNMP traps on switches:

1. Either select individual check boxes for devices, or select the check box next to the Name heading to select all devices. These devices must be up and in the same device family.
2. Click **Deploy Trap Configuration**. The Deploy Options window opens.
3. Fill in a new deployment job name or accept the default name of Deploy SNMP Targets.

4. Either select check boxes for individual traps, or select the check box next to the Trap Name heading to select all traps. These traps are discussed further in [“Deploying SNMP Trap Configurations” on page 610](#).

TIP: To clear an existing configuration, do not select any of the check boxes.

5. Click **Ok**. The Deploy Configuration window opens, which shows the status of deploying the configuration change.
6. Review the outcome of the deployment.

After enabling the traps, enable the alarms and establish the alarm retention period. These tasks are located in Preferences in the Network Director banner.

Deploying SNMP Trap Configurations

The Deploy Options for trap forwarding enable you to select individual traps or all traps for the selected device family.

The device family determines which traps are displayed in the Deploy Options window. The following tables map the trap to one or more MIBs being used.

- EX Series switches traps and related MIBs are shown in [Table 158 on page 610](#).

Table 158: EX Series Switches Traps

Trap	MIB
Chassis	jnxExMibRoot.mib
Link	snmpTraps.mib
Configuration	jnxCfgMgmt.mib
Authentication	jnxJsAuth.mib
Remote operations	jnxPing.mib
Routing	jnx-ipv6.mib
Startup	snmpTraps.mib

Table 158: EX Series Switches Traps (*continued*)

Trap	MIB
Rmon-alarm	jnxRmon.mib
Vrrp-events	rfc2787a.mib
Services	jnxServices.mib
Sonet-alarms	jnx-sonetaps.mib
Otn-alarms	jnxMIbs.mib
PoE-alarms	mib-rfc3621.mib

RELATED DOCUMENTATION

[Setting Up User and System Preferences | 26](#)
[Understanding Fault Mode in Network Director | 761](#)
[Network Director Documentation home page](#)

Understanding Resynchronization of Device Configuration

IN THIS SECTION

- [The Resynchronize Device Configuration Task | 612](#)
- [How Resynchronization Works in NSOR Mode | 613](#)
- [How Resynchronization Works in SSOR Mode | 614](#)
- [How Network Director Resynchronizes the Build Mode Configuration | 616](#)

In a network managed by Network Director, three separate repositories about device configuration are maintained:

- The configuration information on the devices themselves. Each switch maintains its own configuration record.
- The configuration information maintained by the Junos Space Network Management Platform. When a device is discovered, either by Junos Space or Network Director, Junos Space stores a record of the configuration on that device.

Network Director uses the configuration record maintained by Junos Space to determine what configuration commands need to be sent to the device when you deploy configuration on the device in Deploy mode.

- The configuration information maintained by Network Director in Build mode. This information takes the form of the profiles assigned to the device, plus the additional configuration, such as LAG , that you can do under device management.

In Network Director, the configuration state of a device is shown as In Sync when the configuration information in all three repositories match. If there is a conflict between the configuration information in one or more of the repositories, Network Director shows the device configuration state as Out of Sync.

An Out of Sync state is usually the result of out-of-band configuration changes—that is, configuration changes made to a device using a management tool other than Network Director. Examples of such changes include changes made by:

- Using the device CLI.
- Using the device Web-based management interface (the J-Web interface or Web View).
- Using the Junos Space Network Management Platform configuration editor.
- Restoring or replacing device configuration files.

You cannot deploy configuration on a device when the device configuration state is Out of Sync.

This topic describes how Network Director enables you to resynchronize the device configuration state. It covers:

The Resynchronize Device Configuration Task

Network Director provides a task in Deploy mode that enables you to resynchronize the repositories of configuration information. When an out-of-band configuration change is made, you can use this task to resynchronize both the Junos Space configuration record and the Build mode configuration with the configuration on the device.

How Network Director performs resynchronization depends on the system of record (SOR) mode set for the Junos Space Network Management Platform. There are two possible modes:

- Network as system of record (NSOR). This is the default mode.

- Junos Space as system of record (SSOR).

You set the mode in Junos Space under Administration > Applications > Network Management Platform > Modify Application Settings.

How Resynchronization Works in NSOR Mode

In NSOR mode, the network device is considered the system of record for device configuration, which means the configuration maintained by the device takes precedence over the configuration maintained by Junos Space and Network Director. Thus when you perform a resynchronization, the Junos Space configuration record and the Network Director Build mode configuration are updated to match the device configuration.

When an out-of-band change is made on a managed device when Junos Space is in NSOR mode:

1. Junos Space detects that a configuration change has occurred on the device and informs Network Director about the change.
2. Both Junos Space and Network Director set the device configuration state to Out of Sync.
3. Junos Space and Network Director automatically resynchronizes its configuration record to match the device configuration and set the device configuration state to In Sync when the synchronization completes. Network Director performs auto-synchronization when it is operating in the Network as System Of Record (NSOR) mode. The auto-resynchronization parameters are defined in the Preferences page. These parameters enables auto-resynchronization after the interval specified in the Preferences page. For more information see, [“Setting Up User and System Preferences” on page 26](#).
4. When the device out-of-band changes does not conflict with Network Director, Network Director automatically resynchronizes the network changes and retains the local changes in Network Directory. The configuration state of the device and the profile associated with that device remain unaffected. For example, if you modify the MTU value of the port ge-0/0/1 in Network Director and another user modifies the MTU value of port ge-0/0/2 on the same device, Network Director automatically resynchronizes the changes on ge-0/0/2 into Network Director and retains the local changes on ge-0/0/1. The profile corresponding to ge-0/0/1 continues to remain in Pending Deployment state and the profile corresponding to port ge-0/0/2 is in Deployed state.
5. When the device out-of-band changes conflict with the changes made in Network Director, Network Director does not automatically resynchronize the device changes into Network Director. The device is marked as Conflict. You must manually resynchronize the changes by using the Resynchronize Configuration task. After this, the local changes are discarded and are replaced by the latest network configuration. For example, if you modify MTU of ge-0/0/1 from Network Director and another user modifies MTU of the same port on the device, Network Director does not automatically synchronize and marks this device as Out Of Sync.

6. When a profile associated with a device is either added or removed from that device while another user tries to change the attributes corresponding to that profile, Network Director does not automatically synchronize the device and marks the device as conflict, and you must manually resynchronize the changes by using the Resynchronize Configuration task.
7. When you make local changes to profiles, the changes are merged with the new profiles if there is no conflicting configuration. If there are conflicting changes, Network Director receives an Out Of Sync message from Junos Space and you need to manually choose the appropriate profile value.

When you do not make any local changes on a profile, the device association with the profile is deleted and a new device association is created. However, when a profile has local changes, the device association of the profile is not deleted.

NOTE: Automatic resynchronization, as described in Step 3 above, is a default setting for the Junos Space Network Management Platform. If automatic resynchronization is disabled, you must manually resynchronize the Junos Space configuration with the device configuration. You can do so in two ways:

- Use the Resynchronize with Network action in Junos Space. The Junos Space configuration is synchronized with the device configuration. However, the Build mode configuration is not synchronized, so the device state in Network Director remains Out of Sync. You must use the Resynchronize Device Configuration task in Deploy mode to resynchronize the Build mode configuration.
- Use the Resynchronize Device Configuration task in Deploy mode. In this case, Network Director resynchronizes both the Junos Space configuration and the Build mode configuration with the device configuration.

How Resynchronization Works in SSOR Mode

When Junos Space is in SSOR mode, Junos Space is considered the system of record for device configuration. In this mode, when an out-of-band configuration change occurs on a device, you can choose whether to accept the change or to overwrite the change with the configuration maintained by Junos Space.

When an out-of-band change is made on a managed device when Junos Space is in SSOR mode:

1. Junos Space detects that a configuration change has occurred on the device and informs Network Director about the change.
2. Junos Space sets the device configuration state as Device Changed, and Network Director sets the device configuration state to Out of Sync.

Network Director sets the device configuration state to Out of Sync even if the configuration change does not affect configuration you can perform in Build mode. This allows you to resolve the Device Changed configuration state for Junos Space from Network Director.

3. In Network Director, use the Resynchronize Device Configuration task to accept or reject the out-of-band changes:
 - If you accept the out-of-band changes, both the Junos Space configuration record and the Network Director Build mode configuration are resynchronized to reflect the out-of-band configuration changes.
 - If you reject the out-of-band changes, the configuration on the device is overwritten by the configuration record maintained by Junos Space. The Network Director Build mode configuration remains unchanged.
4. Both Junos Space and Network Director set the device configuration state to In Sync.

The above process differs somewhat when out-of-band configuration changes are made through the Junos Space configuration editor. In this case:

1. Junos Space sets the device configuration state as Space Changed after the configuration change is saved.

At this point, the changes have been made only in the Junos Space configuration record and the changes have not yet been deployed to the device. Network Director shows the device configuration state as In Sync.

NOTE: Because the device configuration state is In Sync in Network Director, you can deploy configuration on the device from Network Director at this point. If you do so, the Network Director changes are deployed on the device, but the Junos Space changes are not. The device state in Junos Space remains Space Changed.

2. When the changes are deployed to the device from Junos Space, Junos Space changes the device state to In Sync, while Network Director changes the device state to Out of Sync.
3. In Network Director, use the Resynchronize Device Configuration task to resolve the Out of Sync state. In this case, because the Junos Space configuration record and the device configuration are in sync, you cannot reject the changes. When you resynchronize the device in Network Director, the Build mode configuration is updated to reflect the configuration changes.
4. Network Director sets the device configuration state to In Sync.

If you use Junos Space instead of Network Director to resolve out-of-band configuration changes in SSOR mode, note the following:

- If you reject an out-of-band change, the device state becomes In Sync in both Network Director and Junos Space.
- If you accept an out-of-band change that does not affect the Build mode configuration, the device state becomes In Sync in both Network Director and Junos Space.
- If you accept an out-of-band change that affects the Build mode configuration, the device state becomes In Sync in Junos Space but remains Out Of Sync in Network Director. You must use the Resynchronize Device Configuration task to resolve the Out of Sync state.

How Network Director Resynchronizes the Build Mode Configuration

When you use the Resynchronize Device Configuration task to resynchronize the Build mode configuration to the device configuration, Network Director launches a resynchronization job. This job deletes all profile assignments configured for the device. The profiles themselves are not deleted—just the assignments of the profiles to the device are deleted. It then reimports the device configuration, as if the device were a newly discovered device. It reassigns existing profiles and creates new profiles as necessary. Profiles that were originally assigned to the device will be reassigned to the device if the profiles were unaffected by the out-of-band changes. All profiles assigned to the device are in a deployed state at the end of the process. Any profile that is not reassigned to the device and is not assigned to any other device will be in a unassigned state.

RELATED DOCUMENTATION

[Resynchronizing Device Configuration | 616](#)

[Network Director Documentation home page](#)

Resynchronizing Device Configuration

IN THIS SECTION

- [The Resynchronize Device Configuration List of Devices | 618](#)
- [Resynchronizing Devices When Junos Space Is in NSOR Mode | 619](#)
- [Resynchronizing Devices When Junos Space Is in SSOR Mode | 619](#)
- [Resynchronizing Devices in Manual Approval Mode | 620](#)

- Viewing the Network Changes | 620
- Viewing Resynchronization Job Status | 621

A network managed by Network Director has three repositories of information about the configuration of a network device—the configuration stored on the device itself, the device configuration record maintained by Junos Space, and the Build mode configuration maintained by Network Director.

When the configuration contained in all three repositories match, the device configuration state is shown as In Sync in Network Director. When the repositories do not match, the configuration state is shown as Out of Sync. A common cause for this state is out-of-band configuration changes—that is, configuration changes made to a device outside of Network Director.

When a device state is Out of Sync, you cannot deploy configuration changes on the device in Deploy mode. Use the Resynchronize Device Configuration task to resynchronize the three configuration repositories and change the device configuration state back to In Sync.

How the Resynchronize Device Configuration task performs the resynchronization depends on the system of record (SOR) mode setting for the Junos Space Network Management Platform:

- When Junos Space is in network as system of record (NSOR) mode, the device is considered the system of record for configuration. When you resynchronize a device when Junos Space is in NSOR mode, both the Junos Space configuration record and the Network Director Build mode configuration are updated to reflect the device configuration—in other words, the out-of-band configuration changes are incorporated into both the Junos Space and the Network Director configuration repositories.
- When Junos Space is in Junos Space as system of record (SSOR) mode, you can choose whether accept or reject the out-of-band changes reflected in the device configuration. If you accept the changes, both the Junos Space configuration record and the Network Director Build mode configuration are updated to reflect the device configuration. If you reject the changes, the out-of-band changes are rolled back on the device so that the device configuration matches the Junos Space configuration record and the Network Director Build mode configuration.

For more information about out-of-band configuration changes, Junos Space SOR modes, and how Network Director resynchronizes device configuration, see [“Understanding Resynchronization of Device Configuration” on page 611](#).

This topic covers:

The Resynchronize Device Configuration List of Devices

The Resynchronize Device Configuration page displays a list of all devices in the selected scope whose configuration was successfully imported during device discovery and whose configuration state is now Out Of Sync. You can select devices from this list and resynchronize them.

[Table 159 on page 618](#) describes the fields in the list of devices.

Table 159: Resynchronize Device Configuration Fields

Field	Description
Name	Device hostname or device IP address.
IP address	IP address of device.
Model	Model number of the device.
OS Version	Operating system version currently running on the device.
Connection State	<p>Connection state:</p> <ul style="list-style-type: none"> • UP—Network Director is connected to the device • DOWN—Network Director cannot connect to the device
Configuration State	<p>Shows the configuration state of the device:</p> <ul style="list-style-type: none"> • Out Of Sync—The device configuration is out of sync with either the Network Director Build mode configuration or the Junos Space configuration record or both. • Resynchronizing—The device configuration is in the process of being resynchronized. • Sync Failed—The resynchronization attempt failed. <p>If the resynchronization is successful, the device is removed from the table.</p>
Local Changes	<p>Specifies whether configuration changes have been made in Build mode and are pending deployment on the device.</p> <ul style="list-style-type: none"> • None—There are no configuration changes pending deployment. • View—There are configuration changes that are pending deployment. Click View to view the changes. These changes will be lost if you resynchronize the Build mode configuration to match the device configuration. <p>NOTE: The Pending Changes window that appears when you click View allows you to see what profiles have been added, modified, or changed. However, because the device is not in sync, you cannot view the specific changes in CLI or XML format.</p>

Table 159: Resynchronize Device Configuration Fields (*continued*)

Field	Description
Network Changes	<p>Indicates whether you can view the out-of-band changes:</p> <ul style="list-style-type: none"> • None—The out-of-band changes are not available for viewing. You cannot view out-of-band changes in NSOR mode. In SSOR mode, you cannot view the out-of-band changes if they are already resolved in Junos Space—that is, the device configuration state in Junos Space is In Sync. • View—You can view the out-of-band changes made on the device. Click View to view the changes presented in XML format.

Resynchronizing Devices When Junos Space Is in NSOR Mode

To resynchronize devices when the Junos Space Network Application Platform is in NSOR mode:

1. On the Resynchronization Device Configuration page, select the device or devices that you want to resynchronize.
2. (Optional) View any pending changes to a device's configuration in Network Director by clicking **View** in the Local Changes column. These pending changes are deleted when you resynchronize the device.
3. Click **Resynchronize Configuration**.

The Resynchronize Device Configuration Results window appears. This window will be updated with status of the resynchronization when the resynchronization completes.

Resynchronizing Devices When Junos Space Is in SSOR Mode

To resynchronize devices when the Junos Space Network Management Platform is in SSOR mode:

1. On the Resynchronization Device Configuration page, select the device or devices that you want to resynchronize.
2. (Optional) View any pending changes to a device's configuration in Network Director by clicking **View** in the Local Changes column. These pending changes are deleted if you accept the out-of-band changes when you resynchronize the device.
3. (Optional) View the out-of-band configuration changes by selecting **View** in the Network Changes column. If you accept the out-of-band changes when you resynchronize the device, these changes will be reflected in the Build mode configuration. If you reject the out-of-band changes when you resynchronize the devices, these changes will be deleted from the device. For more information about viewing the out-of-band changes, see [“Viewing the Network Changes” on page 620](#).

NOTE: Out-of-band changes that were made with the Junos Space configuration editor or that were already accepted in Junos Space are not shown. Such changes also cannot be rejected.

4. Click **Resynchronize Configuration**.

5. In the Confirm dialog box:

- Click **Accept device changes** if you want to accept the out-of-band changes.
- Click **Reject device changes** if you want to reject the out-of-band changes and have the configuration that existed on the device before the out-of-band changes were made be reinstated.

click **Submit**.

The Resynchronize Device Configuration Results window appears. This window will be updated with status of the resynchronization when the resynchronization completes.

NOTE: Device changes made by the Junos Space configuration editor or device changes that have been accepted in Junos Space cannot be rejected. Even if you select Reject device changes, these changes will not be rejected and instead will be incorporated into the Build mode configuration.

Resynchronizing Devices in Manual Approval Mode

When out-of-band changes exist, device resynchronization merges the changes done by using the CLI with the local changes provided that there are no conflicts. If there are conflicting changes, the changes made using the CLI take precedence over the local changes. Therefore, configuration changes that are part of a change request might be lost. The configuration change requests that are lost are marked as Cancelled against the corresponding device. When device resynchronization is initiated for a device, a message is displayed that lists the change requests that will be lost because of conflicting CLI and local changes. All other changes remain unaffected.

Viewing the Network Changes

The Network Changes window shows the out-of-band configuration changes made to a device when Junos Space is in SSOR mode.

Not all out-of-band configuration changes are shown in this window. Configuration changes are shown only when the device configuration differs from the Junos Space configuration record—that is, when the device configuration state in Junos Space is not In Sync. For example, if the out-of-band changes were deployed from the Junos Space configuration editor or if the out-of-band changes were already accepted in Junos Space, the configuration changes will not appear in this window.

The configuration changes are shown in XML format. If there have been multiple out-of-band changes—that is, there has been more than one configuration commit, or save, on the device—the changes are grouped by each commit.

The following information is provided for each configuration commit:

- `junos:commit-seconds`—Specifies the time when the configuration was committed as the number of seconds since midnight on 1 January 1970.
- `junos:commit-localtime`—Specifies the time when the configuration was committed as the date and time in the device's local time zone.
- `xmlns:junos`—Specifies the URL for the DTD that defines the XML namespace for the tag elements.
- `junos:commit-user`—Specifies the username of the user who requested the commit operation.

Viewing Resynchronization Job Status

The Resynchronize Device Configuration Results window appears after you start a resynchronization job. This window is automatically updated with the resynchronization status for each device when the job completes.

You can also view the status of the resynchronization jobs using the Manage Jobs task in System mode. The following jobs are associated with resynchronization:

- **Resynch Network Elements**—This job runs in NSOR mode and resynchronizes the Junos Space configuration record with the device configuration.
- **Resolve OOB Changes**—This job runs in SSOR mode and resolves the out-of-band changes for Junos Space—either accepting the changes and updating the Junos Space configuration or rejecting the changes and rolling back the changes on the device.
- **Resynchronize devices**—This job runs in both NSOR and SSOR mode and resynchronizes the Build mode configuration with the device configuration.

RELATED DOCUMENTATION

[Understanding Resynchronization of Device Configuration](#) | 611

[Managing Jobs](#) | 23

Managing Device Configuration Files

You can back up device configuration files to the Network Director server. You can perform several actions on backed up configuration files, such as restoring configuration files to devices, and viewing and comparing configuration files.

To start managing device configuration files:

1. Click **Deploy** in the Network Director banner.
2. In the Tasks pane, select **Device Configuration Files > Manage Device Configuration Files**.

The Manage Device Configuration page opens in the main window. The table lists the devices that have configuration files backed up.

This topic describes:

- [Selecting Device Configuration File Management Options | 622](#)
- [Backing Up Device Configuration Files | 623](#)
- [Restoring Device Configuration Files | 624](#)
- [Viewing Device Configuration Files | 624](#)
- [Comparing Device Configuration Files | 625](#)
- [Deleting Device Configuration Files | 625](#)
- [Managing Device Configuration File Management Jobs | 625](#)

Selecting Device Configuration File Management Options

From the Manage Device Configuration page, you can:

- Back up device configuration files by clicking Backup. See [“Backing Up Device Configuration Files” on page 623](#) for more information.
- Restore backup device configuration files to devices by selecting devices and clicking Restore. See [“Restoring Device Configuration Files” on page 624](#) for more information.
- View backed up configuration files by selecting a device and clicking View Configuration File. See [“Viewing Device Configuration Files” on page 624](#) for more information.

- Compare backed up device configuration files by selecting devices and clicking Compare Config Files. See [“Comparing Device Configuration Files” on page 625](#) for more information.
- Delete backup device configuration files by selecting devices and clicking Delete. See [“Deleting Device Configuration Files” on page 625](#) for more information.

[Table 160 on page 623](#) describes the information provided in the Manage Device Configuration table.

Table 160: Manage Device Configuration Table

Table Column	Description
Device Name	Device name.
Config File Version	Version number of the backup configuration file.
First Backup on	Date when the oldest version of the backup configuration file was created.
Most Recent Backup on	Date when the configuration file was backed up most recently.

Backing Up Device Configuration Files

To back up device configuration files:

1. Click **Backup**.

The Backup Devices Configuration page opens in the main window.

2. Select the devices to back up from the device tree.

3. To back up configuration files immediately, click **Backup Now**.

The backup job runs. When it finishes, the Manage Device Configuration table shows updated information for the devices you backed up.

4. To schedule the backup to run later, click **Schedule Backup**.

The Schedule Backup window opens.

- a. Select the **Schedule at a later time** check box.
- b. Specify when the backup will run using the **Date and Time** fields.
- c. Optionally, configure the backup job to repeat by selecting the **Repeat** check box, then specifying the backup schedule using the provided fields.

Optionally, you can specify when repeated backups will stop by selecting the **End Time** check box, then specifying the last date on which the repeated backup job will run using the **Date and Time** fields.

- d. Click **Schedule Backup**.

Restoring Device Configuration Files

You can restore a backed up configuration file to the device from which it was backed up.



CAUTION: Restoring a configuration file to a device is considered an out-of-band configuration change, which can cause some unexpected results. For more information, see [“Out-of-Band Configuration Changes” on page 81](#).

To restore backed up configuration files to devices:

1. Select the devices to restore from the Manage Device Configuration list.

2. Click **Restore**.

The Restore Device Configuration File(s) window opens.

3. To restore a configuration file that is older than the most recent version, click in the **Latest Version** cell and select the version to restore.

4. Click **Restore**.

Viewing Device Configuration Files

To view the backed up configuration files for a device:

1. Select the device from the Manage Device Configuration list.

2. Click **View Configuration File**.

The Device Configuration Summary window opens, displaying the most recently backed up configuration file.

3. To view an older stored configuration file version, select a version number from the **Config File Version** list.

Comparing Device Configuration Files

To compare backed up device configuration files:

1. Select the configuration files to compare from the Manage Device Configuration list.
2. Click **Compare Configuration Files**.

The Compare Configuration Files window opens.

3. Select a source device from the **Source Device** list and a configuration file version from the **Config File Version** list.
4. Select a target device from the **Target Device** list and a configuration file version from the **Config File Version** list.
5. The configuration file versions you selected are displayed in the window. The file name and version appears at the top of each file. The differences between the configuration files are color-coded. The color-coding legend appears at the top of the window.

Deleting Device Configuration Files

When you delete a device's backed up configuration, all of the configuration file versions for the device are deleted.

To delete device configuration files:

1. Select the configuration files to delete from the Manage Device Configuration list.
2. Click **Delete**.

The Delete Device Configuration File(s) window opens.

3. Verify that the correct devices are listed, then click **Delete**.

Managing Device Configuration File Management Jobs

Each time you back up or restore device configuration files, a device configuration file management job is created.

To manage device configuration file management jobs:

1. Click **Deploy** in the Network Director banner.

2. In the Tasks pane, select **Device Configuration Files > View Configuration File Mgmt Jobs**.

The Device Configuration Jobs page opens in the main window, listing the device configuration file management jobs.

Managing these jobs is similar to managing other types of jobs using the System mode. The advantage of accessing the jobs this way is that the jobs list show only configuration file management jobs. See [“Managing Jobs” on page 23](#) for more information.

RELATED DOCUMENTATION

[Understanding the Deploy Mode Tasks Pane | 577](#)

[Understanding Deploy Mode in Network Director | 572](#)

[Managing Jobs | 23](#)

[Network Director Documentation home page](#)

Creating and Managing Baseline of Device Configuration Files

IN THIS SECTION

- [Selecting Baseline Management Options | 627](#)
- [Baselining Device Configuration Files | 628](#)
- [Restoring Baseline Device Configuration Files | 629](#)
- [Viewing Baseline Configuration Files | 629](#)
- [Comparing Baseline Configuration with Current Configuration | 629](#)
- [Deleting Baseline | 630](#)
- [Managing Baseline Management Jobs | 630](#)

You can create a baseline device configuration and the device Junos (OS) version on the Network Director server. By creating a baseline configuration file for a device you define a reference point to save the device configuration and its OS version to a particular known state and later restore the configuration to that known state. You can select the devices at the scope level, custom grouping, or for individual devices and create baseline configuration files and images for all or for the selected devices. The baseline configuration file includes the entire configuration and image files. When you restore a device configuration, you restore both the baseline configuration file and the image of the file. However, restoring image is optional.

An alarm is triggered if there are any changes to the baseline configuration. The alarm contains the delta information for later reference. For example, when you add a new device an alarm of with minor severity is generated to inform user about the device addition. When you move a device from an unassigned category to a specific category, an alarm of major severity is generated.

To start baseline file management:

1. Click **Deploy** in the Network Director banner.
2. In the Tasks pane, select **Baseline Management > Manage Baseline**.

The Manage Device Baseline page opens in the main window.

This topic describes:

Selecting Baseline Management Options

From the Manage Device Baseline page, you can:

- Create device baseline configuration files by clicking **Baseline**. See [“Baselining Device Configuration Files” on page 628](#) for more information.
- Restore baselined device configuration files to devices by selecting devices and clicking **Restore**. See [“Restoring Baseline Device Configuration Files” on page 629](#) for more information.
- View the baselined configuration files by selecting the device and clicking **View Configuration File**. See [“Viewing Baseline Configuration Files” on page 629](#) for more information.
- Compare baseline device configuration files with current configuration files by selecting devices and clicking **Compare With Current Config**. See [“Comparing Baseline Configuration with Current Configuration” on page 629](#) for more information.
- Delete baseline configuration files for devices by selecting devices and clicking **Delete**. See [“Deleting Baseline” on page 630](#) for more information.

[Table 161 on page 627](#) describes the information provided in the Manage Device Baseline table.

Table 161: Manage Device Baseline Table

Table Column	Description
Device Name	Name of baseline device.
Baseline Label	Name of the baseline configuration.
Baseline Update Time	Date and time when the baseline configuration file for a device is last updated.

Table 161: Manage Device Baseline Table (*continued*)

Table Column	Description
Baseline State	Indicates whether the baseline configuration is same as or out of sync with the current configuration.

Baselining Device Configuration Files

To create baseline configuration for devices:

1. Click **Baseline**.

The Create Baseline for Devices page opens in the main window.

2. Type a baseline label name.

3. Select the devices for which you want to create a baseline configuration files from the device tree.

4. To back up configuration files immediately, click **Create Baseline Now**.

The device baseline job runs. When it finishes, the Manage Device Baseline table shows updated information for the devices for which you performed the baselining.

5. To schedule the backup to run later, click **Schedule Baseline**.

The Schedule Baseline window opens.

- a. Select the **Schedule at a later time** check box.

- b. Specify when the baseline will run using the **Date and Time** fields.

- c. Optionally, configure the baseline job to repeat by selecting the **Repeat** check box, and then specifying the backup schedule by using the fields provided.

Optionally, you can specify when repeated backups will stop by selecting the **End Time** check box, and then specifying the last date on which the repeated backup job will run, using the **Date and Time** fields.

- d. Click **Schedule Baseline**.

Restoring Baseline Device Configuration Files

To restore the baseline configuration file and the OS image of devices:

1. Select the devices from the Manage Device Baseline list.

2. Click **Restore**.

The Restore Baseline(s) window opens.

3. Click **Restore Image** to restore the OS image and select **Reboot device after successful installation** if you want to reboot the device.

4. To view the device configuration summary, click **Click to View Config File**.

5. Click **Restore**.

The device baseline job runs. When it finishes, the Manage Device Baseline table shows updated information for the devices for which you performed the restore.

Viewing Baseline Configuration Files

To view the baseline configuration files for a device:

1. Select the device from the Manage Device Baseline list.

2. Click **View**.

The Device Configuration Summary window opens, displaying the most recently baseline configuration file.

3. To compare the baseline configuration with the current configuration, click **Compare With Current Config**.

Comparing Baseline Configuration with Current Configuration

To compare the baseline configuration with the current configuration:

1. Select the configuration files to compare from the Manage Device Baseline list.

2. Click **Compare With Current Config**.

The Compare Baseline With Current Configuration window opens.

3. The baseline versions and the current version are displayed in the window. The device name and version appear at the top of each file. The differences between the configuration files are color-coded. The color-coding legend appears at the top of the window.
4. Click **Close**.

Deleting Baseline

When you delete a baseline configuration file, all its corresponding versions are also deleted.

To delete baseline configuration file for a device:

1. Select the device name from the Manage Device Baseline list.
2. Click **Delete**.

The Delete Baseline(s) window opens.

3. Verify that the correct devices are listed, then click **Delete** to delete the device configuration.

Managing Baseline Management Jobs

Each time you create baseline configuration files, a baseline management job is also created.

To manage device configuration file management jobs:

1. Click **Deploy** in the Network Director banner.
2. In the Tasks pane, select **Device Configuration Files > View Baseline Mgmt Jobs**.

The Device Configuration Jobs page opens in the main window, listing the device configuration file management jobs.

3. To view the details of a job, select the job name and click **Show Details**.

Managing these jobs is similar to managing other types of jobs using the System mode. The advantage of accessing the jobs this way is that the jobs list shows only configuration file management jobs. See [“Managing Jobs” on page 23](#) for more information about managing jobs.

RELATED DOCUMENTATION

[Understanding the Deploy Mode Tasks Pane](#) | 577

Understanding Deploy Mode in Network Director | 572

Managing Jobs | 23

Deploying and Managing Software Images

IN THIS CHAPTER

- [Managing Software Images | 632](#)
- [Deploying Software Images | 635](#)
- [Managing Software Image Deployment Jobs | 640](#)

Managing Software Images

This topic describes how to manage software images for managed devices.

To start managing software images:

1. Click **Deploy** in the Network Director banner.
2. In the Tasks pane, select **Image Management > Manage Image Repository**.

The Device Image Repository page opens in the main window. The table lists the software images in the repository.

Starting with the Network Director 4.1R1 release, images uploaded on the Junos Space Platform from the **Images and Scripts > Images** page will also be available on the **Manage Image Repository** page.

NOTE:

- Only images uploaded on Junos Space Platform are available in Network Director. However, images uploaded in Network Director will NOT be available in Junos Space Platform.
- If you delete an image from Network Director, the image is NOT deleted in Junos Space Platform.

3. In the Tasks pane, select **Device Configuration File Management > Manage Device Configuration**.

The Manage Device Configuration page opens in the main window. The table lists the devices that have configuration files backed up software images in the repository.

This topic describes:

- [Selecting Software Image Management Options | 633](#)
- [Adding Software Images to the Repository | 633](#)
- [Using the Device Image Upload Window | 634](#)
- [Viewing Software Image Details | 634](#)
- [Using the Device Image Summary Window | 634](#)
- [Deleting Software Images | 635](#)

Selecting Software Image Management Options

From the Device Image Repository page, you can:

- Add a software image to the repository by clicking Add.
- View details about a software image by selecting it and clicking Details.
- Delete software images from the repository by selecting them and clicking Delete.

[Table 162 on page 633](#) describes the information provided in the Device Image Repository table.

Table 162: Device Image Repository Table

Table Column	Description
Check box	Select to perform an action on the software image in that row.
Name	Software image name.
Version	Software version.
Series	Device series that uses the software image.
Uploaded By	User who uploaded the software image.
Created On	Time when the software image was uploaded to the server.
Size(MB)	Size of the software image in megabytes.

Adding Software Images to the Repository

Software images are stored in a repository on the Network Director server.

To add a software image to the repository:

1. Click **Add**.
The Device Image Upload window opens.
2. Use the Device Image Upload window to upload a device software image. See [“Using the Device Image Upload Window” on page 634](#) for a description of the window.

Using the Device Image Upload Window

To use the Device Image Upload window to add a software image to the repository:

1. Click **Browse** and browse to the software image file.
2. Click **Upload** to add the file to the repository.

Viewing Software Image Details

To view details about a software image:

1. Select the software image file in the table.
2. Click **Details**.
The Device Image Summary window opens. See [“Using the Device Image Summary Window” on page 634](#) for information about this window.

Using the Device Image Summary Window

Use the Device Image Summary window to view detailed information about a software image. [Table 163 on page 634](#) describes the fields in this window.

Table 163: Device Image Summary Window

Field	Description
Name	Software image filename.
Version	Software version (release number).
Series	Device series on which the software is supported.
Supported Platforms	Platforms on which the software is supported.

Table 163: Device Image Summary Window *(continued)*

Field	Description
Uploaded By	User who uploaded the image to the server.
Created On	Date and time when the software image was uploaded.
Size (MB)	Size of the software image file, in megabytes.
OK	Click to close the window.

Deleting Software Images

To delete software image files:

1. Select the check box in the rows of the software image files that you want to delete.
2. Click **Delete**.

RELATED DOCUMENTATION

Managing Software Images 574
Deploying Software Images 635
Managing Software Image Deployment Jobs 640
Network Director Documentation home page

Deploying Software Images

This topic describes how to deploy software images to managed devices. You must upload software images to the Network Director server before you can deploy them to devices. See [“Managing Software Images” on page 632](#) for more information.

To start deploying software images:

1. Click **Deploy** in the Network Director banner.
2. Select a node in the View pane that contains the devices on which you want to deploy software images.

3. In the Tasks pane, select **Image Management > Deploy Images to Devices**.

The Select Devices page of the Deploy Images to Devices wizard opens in the main window.

This topic describes:

- [Specifying Software Deployment Job Options | 636](#)
- [Selecting Software Images To Deploy | 637](#)
- [Selecting Options for Software Deployment | 638](#)
- [Summary of Software Deployment | 640](#)

Specifying Software Deployment Job Options

To specify software deployment job options in the Select Devices page:

1. In the Job name field, enter a job name.
2. From the Device and deployment options list, select an option:
 - Select **Staging only (Download image to the device)** to download the software image to the device but not install it.
 - Select **Upgrade only (Install previously staged image on device)** to upgrade the device to a software image that was previously staged on the device.
 - Select **Staging and Upgrade (Download and Install image on device)** to download the software image and install it on the device.

Devices are not automatically rebooted after upgrade to make the device begin running the new software version. You can select the option to reboot the device automatically after the upgrade in a later wizard page.

3. Click **Next** to continue to the next page.

The Select Images page opens. Select a software image as described in [“Selecting Software Images To Deploy” on page 637](#).

Selecting Software Images To Deploy

The Select Images page includes a table listing each device group and device that you selected for deployment. See [Table 164 on page 637](#) for a description of the table columns.

If you selected the Upgrade only (Install previously staged image on device) option, only devices that contain a previously staged software image appear in the table. You cannot select a different image to install on these devices.

To select the software images to deploy, perform the following steps on the table row for each device group or individual device that you want to upgrade:

1. In the Proposed Image Version/Profile column, click **Select Image/Profile**.

The Select Image/Profile list is displayed.


2. From the Select Image/Profile list, select a software image.



TIP: To clear this field, select **Select Image/Profile** from the list.

3. After you finish selecting software images, click **Next** to continue to the next page.

The Select Options page opens.



TIP: A message notifies you if you do not select a software image for all the listed devices. This is just for your information. No action is taken on devices for which you do not select a software image. In effect, this removes those devices from the job.

Select options for software deployment as described in [“Selecting Options for Software Deployment” on page 638](#).

Table 164: Select Images for Devices Table Description

Table Column	Description
Device Family	<p>Device family to which the device belongs. Devices are grouped by family. To display the devices within a device family, click the arrow next to the device family name.</p> <p>For example, all the EX2300 devices are listed under EX2300 device group and all the Junos Fusion fabric devices are listed under <i>Fusion Enterprise</i> group. The Junos Fusion device family, when expanded displays the aggregation devices, satellite devices, and software upgrade groups, if defined.</p>

Table 164: Select Images for Devices Table Description (*continued*)

Table Column	Description
Count	Number of devices contained within a device family.
IP Address	Device's IP address.
Device Name	Device's name.
State	Device's state: <ul style="list-style-type: none"> • UP—Network Director can communicate with the device. • DOWN—Network Director cannot communicate with the device.
Running Image Version	Software version the device is running.
Proposed Image Version/Profile	Software version that is installed on the device when the job runs successfully. You can select to upgrade the software image on one or more individual device, device family, aggregation devices, satellite devices, or on a software upgrade group.

Selecting Options for Software Deployment

The options that you can configure in the Select Options page are described in [Table 165 on page 638](#). The options that are available depend on the job flow you chose in the Select Images page.

After you finish selecting options, click **Next** to continue to the next page. The Summary page opens. Review the job summary as described in [“Summary of Software Deployment” on page 640](#).

Table 165: Image Management Job Options

Option	Action
Select Options	
All Device Types	
Delete any existing image before download	Select to delete any existing software images on devices before downloading the new software image.

Table 165: Image Management Job Options (*continued*)

Option	Action
Reboot device after successful installation	<p>Select to reboot the device after the software image is installed. A reboot is required to begin running the new software version on the device.</p> <p>NOTE: This option might get disabled based on your details that you specify in the remaining fields. This indicates that for the options that you specified, the system automatically reboots the device as per the requirement during or after the image upgrade.</p>
Wired Devices	
Check compatibility with current configuration	Select to validate the software package or bundle against the current configuration as a prerequisite to adding the software package or bundle.
ISSU/NSSU	<p>Select if you want to perform an in-service software upgrade (ISSU) or a nonstop software upgrade (NSSU).</p> <p>ISSU enables you to upgrade between two different Junos OS releases with minimal disruption on the control plane and with minimal disruption of traffic.</p> <p>NSSU enables you to upgrade the software running on an EX Series switch with redundant Routing Engines or on most EX Series Virtual Chassis by using a single command and with minimal disruption to network traffic</p>
Archive data (Snapshot)	Select to take an archive snapshot of the files currently used to run the switch and copy them to an external USB storage device connected to the switch.
Copy to alternate slice	<p>Select to copy the new Junos OS image into the alternate root partition. This ensures that the resilient dual-root partitions feature operates correctly.</p> <p>This option is available only if you select Reboot device after successful installation.</p>
Select Schedule	
Stage now	Select Stage now to start staging software images to devices as soon as the job runs.
Stage later time	Select Stage later time to schedule the staging for a later time.
Staging Schedule	If you selected Stage later time, enter the date and time for staging to start.
Upgrade now	Select Upgrade now to start upgrading software images on devices as soon as staging finishes.
Upgrade later time	Select Upgrade later time to schedule the software upgrade for a later time.

Table 165: Image Management Job Options (continued)

Option	Action
Deployment Schedule	<p>If you selected Upgrade later time, enter the date and time for upgrade to start.</p> <p>If you scheduled staging, you must schedule the upgrade for at least 10 minutes after staging, to ensure that staging completes before upgrade starts.</p>

Summary of Software Deployment

On the Summary page, review the selections you made for the job. To change selections, click **Edit** in the area that you want to change. You can also click the boxes in the process flowchart above the wizard page to navigate between pages. When you are done making selections, click **Finish** on the Summary page to save the job, and run it if you configured the job to run immediately.

RELATED DOCUMENTATION

Managing Software Images 574
Managing Software Image Deployment Jobs 640
Managing Software Images 632
Network Director Documentation home page

Managing Software Image Deployment Jobs

This topic describes how to manage software image jobs. A software image job is created each time you deploy software images to devices or schedule a software image deployment. You can check the status of jobs, see job details, and cancel scheduled jobs.

To start managing software image jobs:

1. Click **Deploy** in the Network Director banner.
2. In the Tasks pane, select **Image Management > View Image Deployment Jobs**.

The Image Deployment Jobs page opens in the main window.

This topic describes:

- [Selecting Software Image Management Options | 641](#)
- [Viewing Software Image Job Details | 642](#)
- [Using the Device Image Staging Window | 642](#)
- [Canceling Software Image Jobs | 643](#)

Selecting Software Image Management Options

From the Image Deployment Jobs page, you can:

- Show deployment job details by selecting a job and clicking Show Details. See [“Viewing Software Image Job Details” on page 642](#) for more information.
- Cancel a pending job by selecting the job and clicking Cancel Job. See [“Canceling Software Image Jobs” on page 643](#) for more information.

[Table 166 on page 641](#) describes the information provided in the of the Image Deployment Jobs table.

Table 166: Image Deployment Jobs Table

Table Column	Description
Job Id	An identifier assigned to the job.
Check box	Select to perform an action on the job in that row.
Job Name	Job name.
Percent	Percentage of the job that is complete.
Status	Job status. The possible states are: <ul style="list-style-type: none">● CANCELLED—The job was cancelled by a user.● SCHEDULED—The job is scheduled but has not run yet.● INPROGRESS—The job is running.● SUCCESS—The job completed successfully. This state is applied if all of the devices in the job completed successfully.● FAILURE—The job failed. This state is applied if any of the devices in the job failed. But some of the devices might have completed successfully. View the job details for the status of each device.
Summary	Job summary.
Scheduled Start Time	Job’s scheduled start time.

Table 166: Image Deployment Jobs Table *(continued)*

Table Column	Description
Actual Start Time	Time when the job started.
End Time	Time when the job ended.
User	User who created the job.
Recurrence	This field is not used for software image management jobs.

Viewing Software Image Job Details

To view the details of a software image job:

1. Select the job in the table.
2. Click **Show Details**.

The Device Image Staging window opens. See [“Using the Device Image Staging Window” on page 642](#) for a description of the window.

Using the Device Image Staging Window

Use the Device Image Staging window to view information about software image jobs. [Table 167 on page 642](#) describes this window.

Table 167: Device Image Staging Window Description

Field	Description
Job Name	Job name.
Start Time	Job’s scheduled start time.
End Time	Time when the job ended.
% Complete	Percentage of the job that is complete.

Table 167: Device Image Staging Window Description (*continued*)

Field	Description
Status	<p>Job status. The possible statuses are:</p> <ul style="list-style-type: none"> • CANCELLED—The job was cancelled by a user. • SCHEDULED—The job is scheduled but has not run yet. • INPROGRESS—The job is running. • SUCCESS—The job completed successfully. • FAILURE—The job failed.
Host Name	Host name of device.
Status	<p>Device status. The possible statuses are:</p> <ul style="list-style-type: none"> • INPROGRESS—The job is running. • SUCCESS—The job completed successfully. • FAILURE—The job failed.
% Complete	Percentage of the job that is complete on the device.
Start Time	Time when the job started on the device.
End Time	Time when the job ended on the device.
Description	Description of the job on the device. Can include error messages for failed devices.
Close	Click to close the window.

Canceling Software Image Jobs

To cancel a software image job:

1. Select the job in the table.
2. Click **Cancel**.

RELATED DOCUMENTATION

[Managing Software Images](#) | 574

[Deploying Software Images](#) | 635

[Managing Software Images | 632](#)

[Network Director Documentation home page](#)

Managing Devices

IN THIS CHAPTER

- [Enabling or Disabling Network Ports on Switches | 645](#)
- [Converting the QSFP+ Ports on QFX Series Devices | 646](#)

Enabling or Disabling Network Ports on Switches

Network ports connect switches to the network and carry network traffic. You can enable or disable network ports of switches that are part of your network. When you enable or disable a port, the administrative status of the port changes to UP or DOWN respectively. When you disable a port, the system marks that port as administratively down, without removing the port configurations.

You can enable or disable one or more ports at a time using the Manage Port Admin State page. The status of the port is indicated by the Admin State and the Link State fields. The administrative status of a port is indicated by the Admin State field.

To enable or disable a network interface:

1. Do one of the following:
 - In the topology view, locate the device for which you want to enable or disable ports and click **Device Management > Manage Port Admin State** from the Tasks pane.
 - While in the Deploy mode, select the device for which you want to enable or disable ports in the View pane and click **Device Management > Manage Port Admin State** from the Tasks pane.

The Manage Port Admin State page appears displaying all the physical ports available on the selected device and the current status of each port. This page also displays the port mode of each interface, if any. Port mode can be access, tagged-access, or trunk mode.

2. Do one of the following:
 - Select the check box adjacent to the ports that you want to enable and click **Change Admin State UP**.
 - Select the check box adjacent to the interfaces that you want to disable and click **Change Admin State DOWN**.

3. Click **Done**. Network Director changes the administrative status of the ports and displays a confirmation message confirming the changes.

RELATED DOCUMENTATION

[Understanding the Network Director User Interface | 3](#)
[Network Director Documentation home page](#)

Converting the QSFP+ Ports on QFX Series Devices

IN THIS SECTION

- [Selecting Devices | 646](#)
- [Converting Ports | 648](#)
- [Reviewing and Deploying Port Conversions | 648](#)

The 40-Gbps QSFP+ ports on QFX Series devices can be configured to operate as four 10-Gigabit Ethernet (xe) ports, one 40-Gigabit Ethernet (xle) port.

To start converting QSFP+ ports:

1. Click **Deploy** in the Network Director banner.
2. Select the node in the View pane that contains the ports you want to convert.
3. Select the task **Device Management > Convert Ports** in the Tasks pane.

The Ports Conversion wizard opens to the Device Selection page. Continue with [“Selecting Devices” on page 646](#).

This topic describes:

Selecting Devices

Use the Device Selection page to select the devices that contain the ports you want to convert.

To select devices that contain the ports you want to convert:

1. Select the device family to filter and display devices from that device family. You can select **Data Center ELS**.
2. Select the devices that contain the ports you want to convert by selecting their check boxes.
3. Click **Next**.

The Convert Ports page opens. Continue with section [“Converting Ports” on page 648](#).

[Table 168 on page 647](#) describes the information provided about devices on the Device Selection page. This page lists all the devices in the selected scope that contain QSFP+ ports.

Table 168: Port Conversion Device Selection Page

Column	Description
Hostname	Configured name of the device or IP address if no hostname is configured.
IP Address (Standalone devices only)	IP Address of the device.
Serial Number	Serial number on device chassis.
Platform	Model number of the device.
Connection State	<p>Connection status of the device in Network Director:</p> <ul style="list-style-type: none"> • UP—Device is connected to Network Director. • DOWN—Device is not connected to Network Director. • N/A—Connection status is unavailable to Network Director.
Config State (Standalone devices only)	<p>Displays the configuration status of the device:</p> <ul style="list-style-type: none"> • In Sync—The configuration on the device is in sync with the Network Director configuration for the device. • Out Of Sync—The configuration on the device does not match the Network Director configuration for the device. This state is usually the result of the device configuration being altered outside of Network Director. <p>You cannot deploy configuration on a device from Network Director when the device is Out Of Sync. To resolve this state, use the Resynchronize Device Configuration task in Deploy mode.</p> <ul style="list-style-type: none"> • Sync failed—An attempt to resynchronize an Out Of Sync device failed. • Synchronizing—The device configuration is in the process of being resynchronized. • N/A—The device is down.

Converting Ports

Use the Convert Ports page to convert QSFP+ ports between port types.

The page contains a table in which you configure the port conversion. The Port Name (Default) column displays the default port name.

To convert QSFP+ ports:

1. To convert a port, click its **Convert to Port** column.
2. Select an option from the list that opens:
 - **No Change**—Does not change the port type.
 - **xle**—Configures the port as one 40-Gigabit Ethernet port.
 - **xe**—Configures the port as a group of 10-Gigabit Ethernet ports.
3. The Port Name (After Conversion) column displays what the port name will be if you commit the current settings.
4. When you finish making port type settings, click **Next**.

The Review page opens. Continue with [“Reviewing and Deploying Port Conversions” on page 648](#)

Reviewing and Deploying Port Conversions

Use the Review page to review settings and deploy the port conversion:

1. To change settings from the Review page, click **Back** to return to previous wizard pages.
2. When you finish making changes, click **Deploy** to deploy the port conversion to the selected devices.

RELATED DOCUMENTATION

[Understanding Deploy Mode in Network Director | 572](#)

[Network Director Documentation home page](#)

Setting Up Zero Touch Provisioning for Devices

IN THIS CHAPTER

- [Understanding Zero Touch Provisioning in Network Director | 649](#)
- [Configuring and Monitoring Zero Touch Provisioning | 650](#)

Understanding Zero Touch Provisioning in Network Director

Zero touch provisioning allows you to provision new Juniper Networks switches in your network automatically—without manual intervention. When you physically connect a switch to a network and boot it with the factory-default configuration, the switch attempts to upgrade the Junos OS software automatically and autoinstall a configuration file from the network. Use the Zero Touch Provisioning wizard to create a profile that applies all the configurations to a Dynamic Host Configuration Protocol (DHCP) server that you configure. You can apply one or more profiles to a DHCP server.

After you enable zero touch provisioning for a DHCP server that is part of a given subnet in your network, and connect a new switch to that subnet, the following series of events occurs:

1. The switch contacts the DHCP server to obtain an IP address. The DHCP server assigns an IP address to the switch. The DHCP server also passes on the location of the software image, and the configuration file to the switch. This information is passed on to the DHCP server from Network Director when you create and save a zero touch provisioning profile.
2. The switch uses this information to locate the software image, and the configuration file. These files are stored in an FTP, TFTP, or an HTTP server.
3. The switch then upgrades the operating system version by using the software image and loads the configuration file.

For more information on zero touch provisioning for switches, see [Understanding Zero Touch Provisioning](#).

RELATED DOCUMENTATION

[Configuring and Monitoring Zero Touch Provisioning | 650](#)
[Network Director Documentation home page](#)

Configuring and Monitoring Zero Touch Provisioning

IN THIS SECTION

- [Configuring Zero Touch Provisioning | 651](#)
- [Specifying the Server Details | 652](#)
- [Specifying the Software Image and Configuration Details | 654](#)
- [Reviewing and Modifying Zero Touch Provisioning Settings | 655](#)
- [What To Do Next | 655](#)
- [Configuration Statements for Custom Configuration of DHCP Server | 655](#)
- [Monitoring Zero Touch Provisioning Profiles | 656](#)

Zero touch provisioning (ZTP) allows you to provision new switches in your network automatically—without manual intervention. When you physically connect a switch to a network and boot it with the factory-default configuration, the switch attempts to upgrade the Junos OS software automatically and autoinstall a configuration file from the network.

The switch uses information that you configure on a Dynamic Host Control Protocol (DHCP) server to determine whether to perform these actions and to locate the necessary software image and configuration files on the network. You can configure the DHCP server by using a zero touch provisioning profile. If you do not configure a DHCP server, the switch boots with the preinstalled software and the default configuration.

The type of DHCP server that you want to use determines whether Network Director configures the DHCP server for you or whether you must manually configure the DHCP server. If you select CentOS or Ubuntu DHCP servers, Network Director configures the DHCP server by using the details that you specified in the zero touch provisioning profile. If you use any other DHCP server, you must manually configure the DHCP server. For such DHCP servers, you can use Network Director only to monitor the switches once they are provisioned. For details on configuring a DHCP server manually, see the DHCP server documentation.

For more information on zero touch provisioning for switches, see [Understanding Zero Touch Provisioning](#).

Before you begin, ensure that you have the necessary privileges on the FTP and the file server that Network Director uses for zero touch provisioning..

NOTE: For detailed information about DHCP and DHCP options, see RFC2131 (<http://www.ietf.org/rfc/rfc2131.txt>) and RFC2132 (<http://www.ietf.org/rfc/rfc2132.txt>). These documents refers to Internet Systems Consortium (ISC) DHCP version 4.2. For more information about this version, see <http://www.isc.org/software/dhcp/documentation>.

Configuring Zero Touch Provisioning

Before you begin:

Ensure that the switch has access to the following network resources:

- The DHCP server that provides the location of the software image and configuration files on the network
See your DHCP server documentation for configuration instructions.
- The File Transfer Protocol (anonymous FTP), the Hypertext Transfer Protocol (HTTP) server, or the Trivial File Transfer Protocol (TFTP) server on which the software image and configuration files are stored. If you are using an FTP server, ensure that the FTP server is configured to enable anonymous access. Refer to your FTP server documentation to know more about this.

NOTE: Although TFTP is supported, we recommend that you use FTP or HTTP instead, because these transport protocols are more reliable.

- (Optional) A Network Time Protocol (NTP) server to perform time synchronization on the network
- (Optional) A system log (syslog) server to manage system log messages and alerts

Identify the type of DHCP server that you will be using for zero touch provisioning:

- CentOS DHCP Server—If your DHCP server uses the following command to restart the server, then select **CentOS** as the DHCP server type:

```
service dhcpd restart
```

- Ubuntu DHCP Server—If your DHCP server uses the following command to restart the server, then select **Ubuntu** as the DHCP server type:

```
service isc-dhcp-server restart
```

- Other—If your server is not an ISC DHCP server running on Linux operating system, then you must select **Other** and configure the DHCP server manually.

NOTE: CentOS 6.10 is the supported or qualified version of CentOS for the DHCP server in Network Director 4.1 release.

For information about the CentOS and Ubuntu versions supported by Network Director Release 4.1, see the *Supported Platforms* section of the [Network Director Release Notes](#).

To configure zero touch provisioning:

1. While in the Deploy mode, select **Zero Touch Provisioning > Manage ZTP** from the Tasks pane.
The Manage ZTP Profiles page appears.
2. Specify the server details in the Server Setup wizard page as described in [“Specifying the Server Details” on page 652](#).

Specifying the Server Details

To configure the server settings:

1. Enter the settings described in [Table 169 on page 652](#). Required settings are indicated in the user interface by a red asterisk (*) that appears next to the field label.

Table 169: Server Details

Field	Description
Profile Name	Name of the zero touch provisioning profile.
DHCP Server Info	
DHCP Server Type	<p>The type of DHCP server that provides the necessary information to the switch. You can choose to configure a CentOS DHCP server, an Ubuntu DHCP server, or any other DHCP server.</p> <p>If you select Other, Network Director also selects the Manually Configure Server check box and hides all the other details except the File Server Details. You must configure the DHCP server manually.</p>
Manually Configure Server	<p>Select to indicate that you want to manually configure the DHCP server. You can configure the CentOS and Ubuntu DHCP servers manually or from Network Director.</p> <p>If you select Manually Configure Server check box, Network Director hides all the other details except the File Server Details.</p>
DHCP Server	IP address or the hostname of the DHCP server.

Table 169: Server Details (continued)

Field	Description
DHCP User	Username for the DHCP server. NOTE: This user must have write permission for the <code>dhcpd.conf</code> file.
DHCP Password	Password for the specified username.
Confirm Password	Confirm the password.
File Transfer Server Info	
File Server	The type of file server where the software images and the configuration files are to be stored. You can choose to use an FTP, HTTP, or a TFTP file server.
File Server IP	IP address or the hostname of the file server.
File Server Root Dir	The root directory of the file server.
Optional Settings	
Syslog Server IP	IP address of the system log server, if you want to perform data logging for zero touch provisioning.
NTP Server IP	IP address of the NTP server, if you want to use time synchronization.

2. Click **Next** and proceed to specify the software image, configuration file, and the IP address range to be configured on the DHCP server. For more details, see [“Specifying the Software Image and Configuration Details” on page 654](#).

Specifying the Software Image and Configuration Details

To specify the software image, configuration file, and the IP address range to be configured on the DHCP server:

1. Enter the password that you want to set for the root user on the switch, in the ZTP Devices Root User Password field and confirm the password in the Confirm Password field.

NOTE: Once the switch is successfully provisioned, Network Director uses this password for discovering the device.

2. In the Configure Settings table, click **Add** to specify details for a switch model.

Network Director adds a row to the Configure Settings table.

3. In the Device Model field, select the switch model for which you want to specify the image and configuration file details.
4. (Only for the CentOS DHCP server) In the Image File field, select the image file that you want to upload for the selected switch model. This field lists the software images that you have uploaded to Network Director from the Device Image Repository page. For details about uploading a software image, see [“Managing Software Images” on page 632](#).
5. Do one of the following to upload the configuration file to the DHCP server:
 - Select the factory-default configuration file for the selected switch model in the Config File field. Network Director ships with a factory-default configuration for all supported switch models.
 - If you want to upload a custom configuration file for the given switch model, click **Upload Config** and select a configuration file. When you upload a custom configuration file, ensure that the configurations mentioned in [“Configuration Statements for Custom Configuration of DHCP Server” on page 655](#) are included in the configuration file.
6. In the Subnet field, specify the subnet that the DHCP server caters to.
7. In the From IP and To IP fields, specify the range of IP addresses that the DHCP server can assign to new switches.
8. (Only for the CentOS or Ubuntu DHCP server) Click **Export DHCP Config** if you want to view the configuration that Network Director sends to the DHCP server.

Network Director downloads the configuration and you can view it using any text editor. If you chose to configure the DHCP server manually in the Server Details page, you can use this configuration file to complete the manual configuration.

9. Click **Next** to review the details of the zero touch provisioning profile that you created.

Reviewing and Modifying Zero Touch Provisioning Settings

From this page, you can save or make changes to a zero touch provisioning profile:

- To make changes to the profile, click the **Edit** button associated with the configuration you want to change.

Alternatively, you can click the appropriate buttons in the zero touch provisioning workflow at the top of the page that corresponds to the configuration you want to change.

When you are finished with your modifications, click **Review** to return to this page.

- To save a zero touch provisioning profile or to save modifications to the settings of an existing profile, click **Finish**.

What To Do Next

- For manual configuration, use the DHCP configuration file to manually configure the DHCP server. If you selected the DHCP server as CentOS or Ubuntu, Network Director uploads the software image to the file server that you specified. If you selected any other DHCP server, you must manually upload the software image to the file server and specify the path when you configure the DHCP server.
- (Only for the CentOS or Ubuntu DHCP servers) For automatic configuration, Network Director configures the DHCP server with the details that you specified in the zero touch provisioning profile and uploads the software image to the file server that you specified.

Configuration Statements for Custom Configuration of DHCP Server

Insert the following configuration statements to the configuration file, if you want to upload a custom configuration file to the DHCP server:

```
system {
  root-authentication {
    encrypted-password "PASSWORD"; ## SECRET-DATA
  }
}
event-options {
```

```

policy target_add_test {
  events snmpd_trap_target_add_notice;
  then {
    raise-trap;
  }
}
}
trap-group networkdirector_trap_group {
version all;
  destination-port NDPORT;
  categories {
    link;
    services;
    authentication;
  }
  targets{
    NDIP;
  }
}

```

Monitoring Zero Touch Provisioning Profiles

You can use the Monitor ZTP Profiles page to view details about the switches that were provisioned using a given zero touch provisioning profile and added successfully to the Network Director inventory.

To monitor a zero touch provisioning profile:

1. While in the Deploy mode, select **Zero Touch Provisioning > Monitor** from the Tasks pane. The Monitor ZTP Profiles page appears.
2. In the Choose ZTP Profile box, select the zero touch provisioning profile that you want to monitor.

Network Director displays the zero touch provisioning summary and details of switches that were discovered using the selected profile.

RELATED DOCUMENTATION

[Understanding Zero Touch Provisioning in Network Director | 649](#)

[Managing Software Images | 632](#)

[Network Director Documentation home page](#)

5

PART

Monitoring Devices and Traffic

About Monitor Mode | **658**

Monitoring Traffic | **670**

Monitoring Client Sessions | **701**

Monitoring Devices | **708**

Monitoring and Analyzing Fabrics | **717**

Monitoring Virtual Networks | **723**

General Monitoring | **728**

Monitor Reference | **733**

About Monitor Mode

IN THIS CHAPTER

- [Understanding Monitor Mode in Network Director | 658](#)
- [Understanding the Monitor Mode Tasks Pane | 665](#)

Understanding Monitor Mode in Network Director

IN THIS SECTION

- [Scope and Monitor Tab Availability | 659](#)
- [Monitors and Tasks | 660](#)
- [Scope and Data Aggregation | 661](#)
- [How Network Director Collects and Displays Monitoring Data | 662](#)
- [How Network Director Displays and Stores Trend Data | 662](#)
- [More About the Monitor Tabs | 663](#)

Monitor mode in Network Director provides you visibility into your network status and performance. Network Director monitors its managed devices and maintains the information it collects from the devices in a database. Monitor mode displays this information in easy-to-understand graphs and in tables that you can sort and filter, allowing you to quickly visualize the state of your network, spot trends developing over time, and find important details.

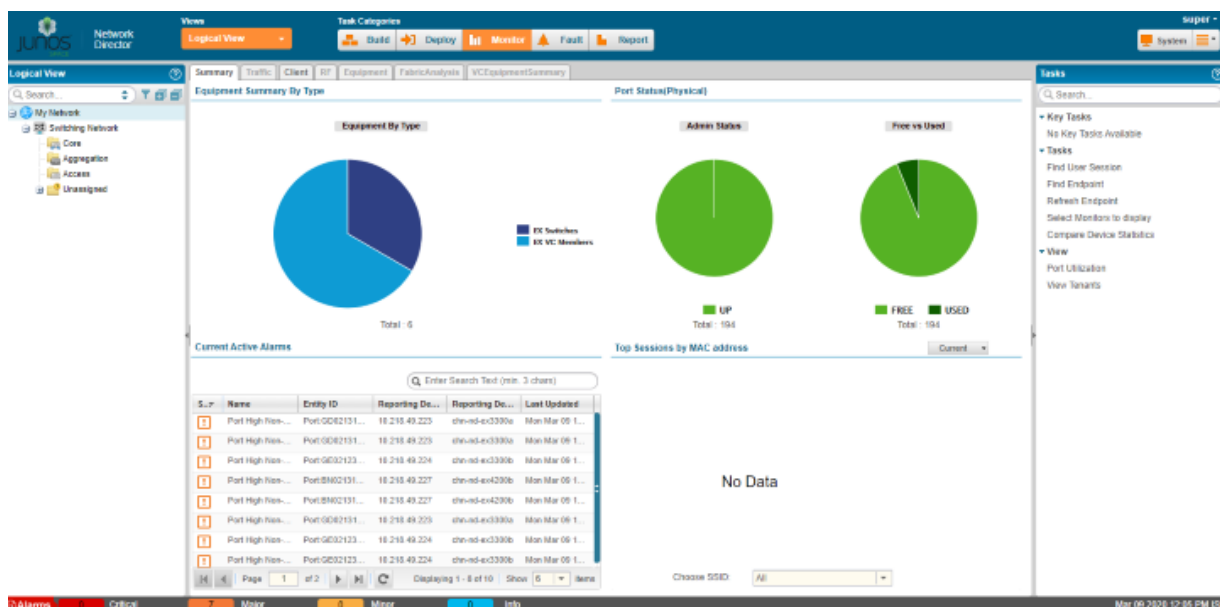
NOTE: Monitor mode for the Datacenter View displays information specific to the virtual networks, virtual switches, and hosts. To know more about Monitor mode for Datacenter View, see [“Using Monitor Mode for Virtual Devices” on page 723](#).

Monitor mode divides monitoring activity into the following categories:

- Traffic—Provides information about traffic on switches and interfaces.
- Client—Provides session information about clients connected to 802.1X authenticator switch ports.
- Equipment—Provides information about the state of switches, and interfaces.
- Fabric Analysis—Displays the results of running the Run Fabric Analyzer task on a Layer 3 fabric, Junos Fusion Fabric or Virtual Chassis Fabric. It shows information about the health, connectivity, and topology of the fabric.
- VC Equipment Summary—Provides information about the operational status of the virtual chassis devices.

You can access these categories through tabs on the Monitor mode landing page, as shown in [Figure 27 on page 659](#). An additional tab, the Summary tab, is available that provides a high-level dashboard for the scope selected in the View pane. The monitoring information displayed in the Summary tab also appears on other tabs.

Figure 27: Monitor Mode Landing Page and Tabs



This topic describes:

Scope and Monitor Tab Availability

Your current scope—that is, your view and node selection in the View pane—affects which Monitor tabs are available.

The shading of the tabs indicate whether a tab is selected, available, or not available:

- The currently selected tab has dark text on a light background.
- Tabs that are available but not selected have dark text on a dark background.

- Tabs that are not available for your current scope have light text on a light background.


When you enter Monitor mode from another mode, the Summary tab is selected for all scopes. If you have selected a tab and then change scope, the tab remains selected if it is supported in the new scope. If it is not supported in the new scope, Network Director selects a default tab for that scope.

Monitors and Tasks

When you click a Monitor tab, the landing page for that tab is displayed, which contains a set of monitors. These monitors enable you to see at a glance important information about the aspect of your network being monitored. For example, the monitors in the Client tab present high-level information about the sessions in the selected scope: the users and client sessions consuming the most bandwidth, the distribution of active sessions by type, and the trend in session count over time.

Detailed information is also available from many monitors when you click the Details icon on the monitor. If the Details icon is not visible in the title bar of a monitor, mouse over the monitor to make it visible. For example, if you click the Details icon from the Current Sessions By Type monitor, you can view detailed information about the current sessions, as shown in [Figure 28 on page 661](#).

Figure 28: Accessing Session Details from the Current Session by Type Monitor



In addition to monitors, each tab provides a set of tasks available from the Tasks pane. These tasks enable you to perform additional monitoring functions. Some tasks enable you to view more specialized monitoring data; others enable you to perform an operation, such as pinging a host. For a complete list of tasks available in Monitor mode, see [“Understanding the Monitor Mode Tasks Pane” on page 665](#).

The scope you select affects which monitors are displayed and which tasks are available. In the Equipment tab, for example, you see a different set of monitors for an EX Series switch.

Scope and Data Aggregation

Network Director enables you to monitor one or more devices. It provides a broader network view by aggregating data from devices and making that data available for viewing at higher scopes within the network.

Not all data is aggregated at higher scopes. For example, it does not make sense to provide power supply status at any higher scope than the device itself. Whenever monitors are available at a scope higher than the device scope, however, the data presented is aggregated data from all devices contained in that scope.

How Network Director Collects and Displays Monitoring Data

Network Director collects monitoring data from all its managed devices at regular intervals known as polling intervals. These polling intervals can vary according to the type of data being collected. Network Director sets default polling intervals for each type of data—you can, however, change these polling intervals in Preferences.

The polling intervals are aligned to clock time. For example, if the polling interval is set to 5 minutes, then within every hour, Network Director collects data at :00, :05, :10, :15, and so on. If the polling interval is set to 15 minutes, Network Director collects data within every hour at :00, :15, :30, and :45.

Network Director uses the Juniper Networks Device Management Interface (DMI) to the managed devices to collect the data. If you have a Junos Space fabric, Network Director balances the load of polling the managed devices across the nodes in the fabric.

When you display a monitor, the current data is from the last polling interval. Displaying or refreshing a monitor does not trigger Network Director to collect data. However, Network Director automatically refreshes monitors with new data after a polling interval completes. Each monitor displays the time that the data was last refreshed.

The detail windows for monitors are not automatically refreshed after a polling period completes. You must manually refresh them to obtain new polling data.

How Network Director Displays and Stores Trend Data

In addition to displaying current data, Network Director also displays historical data in trend graphs so that you can view trends in network performance over time.

When you display a trend graph, you can select the time period over which the data is displayed—usually 1 hour, 8 hours, 1 day, 1 week, 1 month, 3 months, 6 months, or 1 year. These predefined periods are always relative to the current time and date—that is, if you select a week, the data is from the last 7 days. You can also define a custom time period, which enables you to display data for a period between specific dates and times.

For a trend graph displaying a predefined period of 1 hour, the number of data points depends on the configured polling interval. For periods greater than an hour, the number of data points displayed depends on the time period selected and how Network Director consolidates data over time.

To allow storing of monitoring data for a long period of time, Network Director consolidates older data. Consolidation involves deriving a single value from a set of shorter term values, generally by averaging the shorter term values, and then using that value as a data point in a longer term data set. After the shorter term data is consolidated into longer term data, it is discarded to save storage space. For example, if a value is polled every 5 minutes, the set of 12 values is consolidated into a single value after an hour has passed. That value then becomes one of the 24 data points that makes up the data set for a day. Similarly, after a day has passed, data is consolidated into one data point that represents that day; after a month

has passed, data is consolidated into a one data point that represents that month. Data is not kept for more than a year. You can, however, run reports on some monitoring data in Report Mode and archive the reports to maintain a history that is longer than a year.

For all trend graphs, Network Director will not display data until it has more than two data points to display. This means that after you discover a device, trend data will not appear until three polling periods have passed.

More About the Monitor Tabs

IN THIS SECTION

- [The Summary Tab | 663](#)
- [The Traffic Tab | 663](#)
- [The Client Tab | 664](#)
- [The Equipment Tab | 664](#)
- [The Fabric Analysis Tab | 664](#)
- [VC Equipment Summary Tab | 664](#)

The following sections provide more information about each tab in Monitor mode.

The Summary Tab

The Summary tab is displayed whenever you enter Monitor mode. It serves as a high-level dashboard for the current selected scope in the View pane.

The monitors displayed in the Summary tab can belong to any of the Monitor categories. Each scope has a predefined set of monitors that are displayed.

When you select an individual device in the View pane, the Summary tab itself displays an arrow that indicates whether the device is up (green up arrow) or down (red down arrow).

For the My Network scope, you can customize what monitors appear on Summary tab, giving you the ability to view at a glance those aspects of network health and performance that are most important to you.

The Traffic Tab

The Traffic tab provides information for analyzing traffic on switches. The four monitors provide an aggregated view of all network traffic on a device, such as proportion of current proportion of multicast, unicast, broadcast traffic or the trend in packet errors. Tasks provide more detailed looks at traffic, such as traffic statistics for individual ports or the degree in which a port's bandwidth is being used.

The Client Tab

The Client tab provides information about clients and sessions on the network. A client is any device that is connected to the network through an access port on a switch that is an 802.1X authenticator port. Examples of clients include VoIP phones, laptops, printers, security cameras, and so on. When a client connects to the network, a session starts, which is uniquely identified by the MAC address of the client.

The Client tab monitors provide a view of overall client session activity in the selected scope. They show the total number of sessions, sessions consuming the most bandwidth, and trends in the number of sessions. Detailed views provide information about each client, such as MAC address, IP address, username, and client VLAN, the client is connected to. You can also search for a particular client session or sessions using a variety of search criteria and view client history.

NOTE: Because traffic information is unavailable for sessions connected to access ports on switches, monitors that show session traffic are not displayed for scopes that contain switches only.

The Equipment Tab

The Equipment tab provides information about the operational status of individual devices. Monitors display CPU and memory use, power supply and fan status, port status, and general device information for switches. Additional information provided by this tab includes the state of logical Ethernet switching interfaces on standalone switches, the topology of Virtual Chassis.

The Fabric Analysis Tab

The Fabric Analysis tab displays the results of running the Fabric Analyzer on Layer 3 fabric or Junos Fusion Fabric, or Virtual Chassis Fabric (VCF). It shows information about the health, connectivity, and topology of the fabric. For information about analyzing fabrics, see [“Analyzing Virtual Chassis Fabrics” on page 718](#)

VC Equipment Summary Tab

The VC Equipment Summary tab displays the operational status of the virtual chassis members associated with Logical View, Location View, Device View and Custom Group View.

NOTE: The VC Equipment Summary tab is enabled only for members that are at the virtual chassis container level, virtual chassis level or virtual chassis device level in the applicable view panes.

When you select any virtual chassis or virtual chassis members in the View pane, the VC Equipment Summary tab displays a pie chart that indicates the total number of virtual chassis members and the connection status of the virtual chassis members. Mouse over a pie segment to view the actual number of VC members and the percentage represented by that pie segment. The connection states are as follows:

- UP (green)—Device is connected to Network Director
- DOWN (red)—Device is not connected to Network Director

To see detailed information about the virtual chassis members, click the Details icon on the top right corner of the VC Equipment Member Status title bar. The details shown depends on the type of the virtual chassis device selected. In the VC Equipment Member Status window, the connection status is shown as green up arrow when the virtual chassis members are up and red down arrow when the virtual chassis members are down.

RELATED DOCUMENTATION

[Understanding the Monitor Mode Tasks Pane | 665](#)

[Understanding the Network Director User Interface | 3](#)

[Network Director Documentation home page](#)

Understanding the Monitor Mode Tasks Pane

The Tasks pane in Monitor mode displays a list of tasks that are available for the currently selected Monitor tab. These tasks provide monitoring functions in addition to the monitors available under each tab.

The tasks listed in the Tasks pane vary according to the selected tab—that is, Summary, Traffic, Client, or Equipment—and the scope you have selected in the View pane. For example, the VC Protocols Statistics task is available only when you select the Traffic tab and a Virtual Chassis or Virtual Chassis member in the View pane.

For each Monitor mode tab, the following tables list each task and provide a short description of the task:

- [Table 170 on page 666](#): Summary Tab Tasks
- [Table 171 on page 666](#): Traffic Tab Tasks
- [Table 172 on page 667](#): Client Tab Tasks
- [Table 173 on page 668](#): Equipment Tab Tasks
- [Table 174 on page 669](#): VC Equipment Summary Tab Tasks
- Key Tasks—Network Director enables you to group the tasks that you perform frequently and create a list of key tasks. You can add any task from the Tasks pane to the Key Tasks list by selecting a task and clicking the plus (+) sign that appears adjacent to the task. For some modes, you can see that Network Director has predefined some key tasks for you. You can modify this set of tasks to suit your requirements. This feature is available in Task pane irrespective of your current mode, scope, or view.

Table 170: Summary Tab Tasks

Task	Description
Compare Device Statistics	Compares statistics from multiple devices in real time.
Find End Point	Finds end points that match supplied attributes and returns information about the location of the end points on the network.
Find User Session	Finds client sessions that match supplied attributes and returns information about the sessions.
View Tenants	Shows details of tenants and overlay networks.
View Congestion Events	Shows latency congestion information based on high-frequency statistics data.
Ping To a Host	From the selected device, pings the host you specify and returns the results.
Port Utilization	Displays port utilization trend information for the devices in the selected scope. You can view overall port utilization for the selected device or can view individual port utilization.
Refresh End Point	Refreshes end point location information for the Find End Point task.
Select Monitors to display	Selects the monitors that are displayed in the Summary tab
Show ARP Table	Shows Address Resolution Protocol (ARP) table information for a device.
Show PoE Interfaces	Shows details such as interface name, admin status, maximum power limit, power consumption and priority, of PoE interfaces in satellite devices of Junos Fusion Enterprise and EX devices.
Show Routing Instances	Shows the routing instances configured on an MX Series router.
Traceroute To a Host	From the selected device, traces the route to the host you specify and returns the results.

Table 171: Traffic Tab Tasks

Task	Description
Compare Device Statistics	Compares statistics from multiple devices in real time.
Find End Point	Finds end points that match supplied attributes and returns information about the location of the end points on the network.

Table 171: Traffic Tab Tasks (*continued*)

Task	Description
Find User Session	Finds client sessions that match supplied attributes and returns information about the sessions.
L3 VLAN Statistics	Displays packet in and out statistics for Layer 3 VLANs on the selected device.
View Congestion Events	Shows latency congestion information based on high-frequency statistics data.
Ping To a Host	From the selected device, pings the host you specify and returns the results.
Port Statistics	Displays packet and error statistics for all ports on the selected device.
Port Utilization	Displays port utilization trend information for the devices in the selected scope. You can view overall port utilization for the selected device or can view individual port utilization.
Show ARP Table	Shows Address Resolution Protocol (ARP) table information for a device.
Show PoE Interfaces	Shows details such as interface name, admin status, maximum power limit, power consumption and priority, of PoE interfaces in satellite devices of Junos Fusion Enterprise and EX devices.
Show Routing Instances	Shows the routing instances configured on an MX Series router.
Traceroute To a Host	From the selected device, traces the route to the host you specify and returns the results.
VC Protocol Statistics	Displays Virtual Chassis Control Protocol (VCCP) statistics for the selected Virtual Chassis or Virtual Chassis member, such as the kind and number of protocol data units (PDUs) sent and received.

Table 172: Client Tab Tasks

Task	Description
Compare Device Statistics	Compares statistics from multiple devices in real time.
Find End Point	Finds end points that match supplied attributes and returns information about the location of the end points on the network.
Find User Session	Finds client sessions that match supplied attributes and returns information about the sessions.

Table 172: Client Tab Tasks (*continued*)

Task	Description
View Congestion Events	Shows latency congestion information based on high-frequency statistics data.
Ping To a Host	From the selected device, pings the host you specify and returns the results.
Refresh End Point	Refreshes end point location information for the Find End Point task.
Show ARP Table	Shows Address Resolution Protocol (ARP) table information for a device.
Show PoE Interfaces	Shows details such as interface name, admin status, maximum power limit, power consumption and priority, of PoE interfaces in satellite devices of Junos Fusion Enterprise and EX devices.
Show Routing Instances	Show the routing instances configured on an MX Series router.
Traceroute To a Host	From the selected device, traces the route to the host you specify and returns the results.

Table 173: Equipment Tab Tasks

Task	Description
Compare Device Statistics	Compares statistics from multiple devices in real time.
Find End Point	Finds end points that match supplied attributes and returns information about the location of the end points on the network.
Find User Session	Finds client sessions that match supplied attributes and returns information about the sessions.
View Congestion Events	Shows latency congestion information based on high-frequency statistics data.
Logical Interfaces	Displays the status of the Ethernet switching interfaces on the device, including aggregated Ethernet interfaces. Information includes VLAN membership, STP state, and port mode.
Ping to a Host	From the selected device, pings the host you specify and returns the results.
Show ARP Table	Shows Address Resolution Protocol (ARP) table information for a device.

Table 173: Equipment Tab Tasks (*continued*)

Task	Description
Show PoE Interfaces	Shows details such as interface name, admin status, maximum power limit, power consumption and priority, of PoE interfaces in satellite devices of Junos Fusion Enterprise and EX devices.
Show Routing Instances	Show the routing instances configured on an MX Series router.
Traceroute To a Host	From the selected device, traces the route to the host you specify and returns the results.

Table 174: VC Equipment Summary Tab Tasks

Task	Description
Find User Session	Finds client sessions that match supplied attributes and returns information about the sessions.
Compare Device Statistics	Compares statistics from multiple devices in real time.

RELATED DOCUMENTATION

[Understanding Monitor Mode in Network Director | 658](#)
[Understanding the Network Director User Interface | 3](#)
[Network Director Documentation home page](#)

Monitoring Traffic

IN THIS CHAPTER

- [Monitoring Traffic on Devices | 670](#)
- [Monitoring Port Traffic Statistics | 671](#)
- [Monitoring Traffic on Layer 3 VLANs | 674](#)
- [Monitoring Routing Instances | 676](#)
- [Monitoring Port Utilization | 687](#)
- [Monitoring Tenant Details | 691](#)
- [Monitoring Virtual Chassis Protocol Statistics | 697](#)
- [Viewing Congestion Events | 699](#)

Monitoring Traffic on Devices

The monitors on the Traffic tab provide information about the traffic traversing switches, routers, Virtual Chassis, Virtual Chassis Fabrics (VCFs) and Layer 3 Fabrics.

To monitor traffic on a device:

1. Click **Monitor** in the Network Director banner.
2. Select the device in the View pane that contains the traffic you want to monitor.
3. Select the **Traffic** tab to open the traffic monitors.
4. To get help for a monitor, click the Help button in its title bar.

The available monitors include:

- [“Unicast vs Broadcast/Multicast Monitor” on page 753](#): shows the distribution of unicast, broadcast, and multicast traffic entering and leaving the device.
- [“Unicast vs Broadcast/Multicast Trend Monitor” on page 754](#): shows trend data about the distribution of unicast, broadcast, and multicast traffic entering and leaving the device.

- [“Traffic Trend Monitor” on page 753](#): shows trend data about the amount of traffic entering and leaving the device.
- [“Error Trend Monitor” on page 736](#): shows trend data about the amount of errors on the device.
- [“Top Talker - Wired Devices Monitor” on page 752](#) shows the hosts that are using the most bandwidth.

RELATED DOCUMENTATION

[Understanding Monitor Mode in Network Director | 658](#)

[Monitoring Port Traffic Statistics | 671](#)

[Monitoring Virtual Chassis Protocol Statistics | 697](#)

[Monitoring Traffic on Layer 3 VLANs | 674](#)

[Network Director Documentation home page](#)

Monitoring Port Traffic Statistics

IN THIS SECTION

- [Procedure for Monitoring Port Traffic Statistics | 671](#)
- [Port on Device Window | 672](#)
- [Port Traffic Stats Window | 673](#)

This topic describes how to monitor port traffic statistics on a device. You can monitor port traffic statistics for a switch, router, Virtual Chassis, Virtual Chassis Fabric, or Layer 3 Fabric.

This topic describes:

Procedure for Monitoring Port Traffic Statistics

1. Click **Monitor** in the Network Director banner.
2. Do one of the following:
To view the port statistics:
 - a. Select a node in the View pane that contains the port traffic you want to monitor.

- b. Select the **Traffic** tab.

- c. In the Tasks pane, select **Port Statistics**.

The Port Traffic Stats window opens. For information about this window, click the Help button in the title bar of the window or see [“Port Traffic Stats Window” on page 673](#).

To view the port status and analyze network traffic for devices that are configured for network traffic analysis:

- a. Select a device from the view pane.

- b. In the Tasks pane, select **Traffic Analysis**.

The Port on Device window opens. For information about this window and the tasks that you can perform from this window, see [“Port on Device Window” on page 672](#).

Port on Device Window

Port on Device window displays the details of all the ports on a device. [Table 175 on page 672](#) describes the fields that are displayed in the Port on Device window.

Table 175: Port on Device table field descriptions

Field Name	Description
Port Name	Identification of the port.
Admin State	The administrative state of the port: enabled (UP) or disabled (DOWN).
Operational State	The operational status—link up (UP) or link down (DOWN).
Max Bandwidth	The actual bandwidth available on the port, in megabits (Mb).
Negotiated Bandwidth	The negotiated bandwidth based on the speed that is configured or auto-negotiated for the interface.

To view more details about the traffic on any port, select the port and click View Traffic. The Traffic on Port window opens. For more details and field descriptions, see [“Device & Port Utilization Widget” on page 63](#).

Port Traffic Stats Window

The Port Traffic Stats window displays information about the port traffic on the node you selected in the View pane. It contains the following elements:

- **Port Traffic Trend graph**—This line graph shows trends in the data and error rates on the port selected in the ports table below it. The horizontal axis shows the times when samples were taken. The vertical axis shows the data rate on the left side (in packets per second) and the error rate on the right side (in errors per second).

To display traffic for a different port, select the port from the table below the graph. To change the time period over which to display the traffic trends, select a time period from the list in the upper right corner.

NOTE: After a device is discovered, trend data does not appear immediately—it appears only after three polling periods have elapsed.

To highlight a line on the graph, mouse over the line legend. To remove or restore a line, click the line legend. To see numeric values, mouse over where a data line intersects with a dotted vertical grid line.

- **Ports table** (on the lower left side of the window)—This table provides information about the ports as described in [Table 176 on page 673](#). Selecting a port from this table updates the Port Traffic Trend graph to display traffic information about the selected port.
- **Counter selection table** (on the lower right side of the window)—This table enables you to select which counters to display on the Port Traffic Trend graph. It includes separate tabs for packet counters and error counters. Select the check box in the Show column of each counter that you want to display on the graph. The Per/Sec column shows the rate per second of that row's counter.

Table 176: Port Traffic Window

Table Column	Description
Serial Num	Serial number of the device to which the port belongs.
Port Name	Port number. Channelized ports are indicated by the port number followed by :<channelized port number>. For example, xe/0/0/1:2 indicates that this channelized port is a part of the xe/0/0/1 port with a channelized port number of 2.
Port Usage Type	Port mode—either ACCESS or UPLINK.

Table 176: Port Traffic Window (*continued*)

Table Column	Description
MAC Addresses	Port MAC address.
Link Type	Full duplex, half duplex, or unspecified.
In Packets/Sec.(Current)	Current rate of inbound packets.
Out Packets/Sec.(Current)	Current rate of outbound packets.

RELATED DOCUMENTATION

[Understanding the Monitor Mode Tasks Pane | 665](#)

[Network Director Documentation home page](#)

Monitoring Traffic on Layer 3 VLANs

IN THIS SECTION

- [Procedure for Monitoring Layer 3 VLAN Traffic Statistics | 674](#)
- [L3 VLAN Traffic Stats Window | 675](#)

This topic describes how to monitor Layer 3 VLAN traffic statistics on a device. You can monitor Layer 3 VLAN statistics for a switch, router, Virtual Chassis, Virtual Chassis Fabric, Layer 3 Fabric, and the aggregation devices in a Junos Fusion fabric.

This topic describes:

Procedure for Monitoring Layer 3 VLAN Traffic Statistics

1. Click **Monitor** in the Network Director banner.
2. Select a node in the View pane that contains the Layer 3 VLAN traffic you want to monitor.

3. Select the **Traffic** tab.
4. In the Tasks pane, select **View > L3 VLAN Statistics**.

The L3 VLAN Traffic Stats window opens. For information about this window, click the Help button in the title bar of the window or see [“L3 VLAN Traffic Stats Window” on page 675](#).

L3 VLAN Traffic Stats Window

The L3 VLAN Traffic Stats window displays information about the Layer 3 VLAN traffic on the node you selected in the View pane. You can monitor Layer 3 VLAN statistics for a switch, router, Virtual Chassis, Virtual Chassis Fabric, Layer 3 Fabric, and the aggregation devices in a Junos Fusion fabric.

NOTE: For a Junos Fusion fabric, you must select an aggregation device to view the Layer 3 VLAN statistics. The L3 VLAN Statistics option is not available if you select a Junos Fusion satellite device.

The L3 VLAN Traffic Stats window contains two panes:

- **VLAN Traffic line graph**—This graph shows the data transmission rate on the Layer 3 VLAN selected in the table beneath the graph. The horizontal axis shows the times when samples were taken. The vertical axis shows the data rate, in bytes per second.

To show a Layer 3 VLAN on the VLAN Traffic line graph, select the Layer 3 VLAN from the table beneath the graph. To highlight a line on the graph, mouse over the line legend. To remove or restore a line, click the line legend. To see numeric values, mouse over a data point.

- **Layer 3 VLAN traffic statistics table**—This table provides information about the Layer 3 VLANs as described in [Table 177 on page 675](#). Selecting a Layer 3 VLAN from this table updates the VLAN Traffic graph to display the traffic information for the selected Layer 3 VLAN.

Table 177: Layer 3 VLAN Traffic Statistics Table

Table Column	Description
L3 Interface	Layer 3 interface assigned to the VLAN.
SerialNo	The serial number of the device containing the Layer 3 VLAN.
VLAN Name	VLAN name.
VLAN ID	VLAN ID.
Description	VLAN description.

Table 177: Layer 3 VLAN Traffic Statistics Table (*continued*)

Table Column	Description
In Packet	Number of packets entering the VLAN.
Out Packet	Number of packets leaving the VLAN.

RELATED DOCUMENTATION

[Understanding the Monitor Mode Tasks Pane | 665](#)
[Network Director Documentation home page](#)

Monitoring Routing Instances

IN THIS SECTION

- [Procedure for Monitoring Routing Instances | 677](#)
- [Show Routing Instances Window | 677](#)
- [Show Interfaces Window | 678](#)
- [Show Bridge Domains Window | 679](#)
- [Show Connections | 680](#)
- [Show Routing Tables | 683](#)
- [Show MAC Table | 686](#)

This topic describes how to monitor VPN routing instances on MX Series routers by using Network Director. Using Network Director, you can determine which interfaces and bridge domains belong to the routing instances and view traffic statistics for those interfaces and bridge domains. You can also display connection information for Layer 2 VPN and virtual private LAN service (VPLS) routing instances.

Network Director can be used to monitor the following types of Layer 2 routing instances:

- Default routing instance
- Ethernet VPN (EVPN)
- Layer 2 VPN

- VPLS
- Virtual switch

Network Director can be used to monitor the following types of Layer 3 routing instances:

- Layer 3 VPN

This topic describes:

Procedure for Monitoring Routing Instances

Use the options in the Show Routing Instances window to monitor routing instances.

1. Click **Monitor** in the Network Director banner.
2. Select an MX Series router in the View pane that contains the port traffic you want to monitor.
3. In the Tasks pane, select **Tasks > Show Routing Instances**.

The Show Routing Instances window opens. For information about this window, click the Help button in the title bar of the window or see [“Show Routing Instances Window” on page 677](#).

Show Routing Instances Window

The Show Routing Instances window lists the routing instances configured on a selected device. Use this window to display the interfaces or bridge domains belonging to a routing instance and obtain traffic statistics for the interfaces. You can also display information about the VPLS and Layer 2 VPN connections. [Table 178 on page 677](#) describes the fields in this window.

Table 178: Fields in the Show Routing Instances Window

Field	Description
Routing Instance Name	Name of the routing instance.
	The default routing instance is named <i>default-switch</i> .



Table 178: Fields in the Show Routing Instances Window (*continued*)

Field	Description
Type	<p>Identifies the routing instance type:</p> <ul style="list-style-type: none"> • EVPN • L2VPN • L3VPN • Virtual Switch <p>The default routing instance is of this type.</p> <ul style="list-style-type: none"> • VPLS • VRF (L3VPN)
Details	<p>Provides the following information (if configured for the routing instance):</p> <ul style="list-style-type: none"> • Route Distinguisher—Used to identify all routes that are part of the VPN. The route distinguisher makes IP addresses globally unique, so that the same IP address prefixes can be used for different VPNs. • Target—Extended BGP community used to match routes for import and export.
Interfaces	<p>Displays the number of interfaces belonging to the routing instance. Click the number to open the Show Interfaces window, described in “Show Interfaces Window” on page 678.</p>
Bridge Domains	<p>Displays the number of bridge domains belonging to the routing instance. Click the number to open the Show Bridged Domains window, described in “Show Bridge Domains Window” on page 679.</p>
Actions	<ul style="list-style-type: none"> • Click Show Connections to display information about Layer 2 VPN and VPLS connections. The information described in “Show Connections” on page 680 is displayed. This link is available only for Layer 2 VPN and VPLS routing instances. • Click Show MAC Table to display the MAC table for the selected routing instance. For details, see “Show MAC Table” on page 686. • Click Show Routing Table to view the routing table information for the selected routing instance. For details, see “Show Routing Tables” on page 683.

Show Interfaces Window

The Show Interfaces window lists the logical interfaces configured on the routing instance and provides the information about the interfaces as described in [Table 179 on page 679](#).

Table 179: Show Interfaces Information

Field	Description
Interface Name	The interface name.
Port Mode	Indicates one of two modes—Access or Trunk: <ul style="list-style-type: none"> • Access—The interface can be in a single VLAN only. • Trunk—The interface can be in multiple VLANs and accept tagged packets from multiple devices.
Interface State	Indicates whether the interface is  UP or  DOWN.
STP State	Indicates whether the interface is in a discarding (blocked) or in forwarding (unblocked) state. (Not shown for interfaces belonging to Layer 2 VPN and Layer 3 VPN routing instances.)
Local IP Address	Local IP address. (Shown only for interfaces belonging to Layer 2 VPN and Layer 3 VPN routing instances.)
Remote IP Address	Remote IP address. (Shown only for interfaces belonging to Layer 2 VPN and Layer 3 VPN routing instances.)
Actions	<ul style="list-style-type: none"> • Click View Statistics to display traffic statistics for the interface. The Show Interface Statistics window opens, which charts the number of input and output packets and the number of input and output bytes. • Click Show MAC Table to display the MAC table for the interface. For more details, see “Show MAC Table” on page 686.



Show Bridge Domains Window

The Show Bridge Domains window lists the bridge domains configured on the routing instance. To display information about the VLAN IDs and interfaces configured on a bridge domain, select the bridge domain. [Table 180 on page 679](#) describes the information provided in the Show Bridge Domains window.

Table 180: Show Bridge Domains Information

Field	Description
Bridge Domains	The bridge domain name.
Actions	Click Show MAC Table to display the MAC table for the selected bridge domain. For details, see “Show MAC Table” on page 686 .

Table 180: Show Bridge Domains Information (continued)

Field	Description
VLAN ID	The VLAN ID or IDs assigned to the bridge domain.
Interface Name	The name of a logical interface assigned to the VLAN ID.
Port Mode	Indicates one of two modes—access or trunk: <ul style="list-style-type: none"> • Access—The interface can be in a single VLAN only. • Trunk—The interface can be in multiple VLANs and accept tagged packets from multiple devices.
Interface State	Indicates whether the interface is  UP or  DOWN.
STP State	Indicates whether the interface is in a discarding (blocked) or in forwarding (unblocked) state.
Actions	<ul style="list-style-type: none"> • Click View Statistics to display traffic statistics for the interface. The Show Interface Statistics window opens, which charts the number of input and output packets and the number of input and output bytes. • Click Show MAC Table to display the MAC table for the interface. For details, see “Show MAC Table” on page 686.

Show Connections

The Show Connections window provides information about the VPN connections for Layer 2 VPN and VPLS routing instances as described in [Table 181 on page 680](#).

Table 181: Show Connections Information



Field	Description
Local Site Name	Name of the local site.
Local Site ID	Identifier for the local site.
Local Interface Name	Name of the local interface.
Interface Status	Indicates whether the local interface is  UP or  DOWN.
Remote Site ID	Identifier for the remote site.
Remote IP	IP address of the remote provider edge device (PE device).

Table 181: Show Connections Information (*continued*)

Field	Description
Connection Status	

Table 181: Show Connections Information (*continued*)

Field	Description
	<p>Status of the connection:</p> <ul style="list-style-type: none"> • EI—The local VPN interface is configured with an encapsulation that is not supported. • EM—The encapsulation type received on this connection from the neighbor does not match the local connection interface encapsulation type. • VC-Dn—The virtual circuit is currently down. • CM—The two routers do not agree on a control word, which causes a control word mismatch. • CN—The virtual circuit is not provisioned properly. • OR—The label associated with the virtual circuit is out of range. • OL—No advertisement has been received for this virtual circuit from the neighbor. There is no outgoing label available for use by this virtual circuit. • LD—All of the CE-facing interfaces to the local site are down. Therefore, the connection to the local site is signaled as down to the other PE routers. No pseudowires can be established. • RD—All the interfaces to the remote neighbor are down. Therefore, the remote site has been signaled as down to the other PE routers. No pseudowires can be established. • LN—The local site has lost path selection to the remote site and therefore no pseudowires can be established from this local site. • RN—The remote site has lost path selection to a local site or to a remote site and therefore no pseudowires are established to this remote site. <p>In a multihoming configuration, one multihomed PE site displays the state LN, and the other multihomed PE site displays the state RN in the following circumstances:</p> <ul style="list-style-type: none"> • The multihomed links are both configured to be the backup site. • The two multihomed PE routers have the same site ID, but have a peering relationship with a route reflector (RR) that has a different site ID. <ul style="list-style-type: none"> • XX—The connection is down for an unknown reason. This is a programming error. • MM—The MTUs for the local site and the remote site do not match. • BK—The router is using a backup connection. • PF—Profile parse failure. • RS—The remote site is in a standby state. • NC—The interface encapsulation is not configured as an appropriate CCC (circuit cross-connect), TCC (translational cross-connect), Layer 2 VPN, or VPLS encapsulation. • WE—The encapsulation configured for the interface does not match with the encapsulation configured for the associated connection within the routing instance. • NP—The router detects that interface hardware is not present. The hardware might be offline, a PIC might not be of the compatible type, or the interface might be configured in a different routing instance. • ->—Only the outbound connection is up.

Table 181: Show Connections Information (*continued*)

Field	Description
	<ul style="list-style-type: none"> • <--—Only the inbound connection is up. • Up—The connection is operational. • Dn—The connection is down. • CF—The router cannot find enough bandwidth to the remote router to meet the connection bandwidth requirement. • SC—The local site identifier is the same as the remote site identifier. No pseudowire can be established between these two sites. You must configure different values for the local and remote site identifiers. • LM—The local site identifier is not the minimum designated, which means it is not of the lowest value. There is another local site with a lower value for site identifier. Pseudowires are not being established to this local site and the associated local site identifier is not being used to distribute Layer 2 VPN or VPLS label blocks. However, this is not an error state. Traffic continues to be forwarded to the PE router interfaces connected to the local sites when the local sites are in this state. • RM—The remote site identifier is not the minimum designated, which means it is not the lowest. There is another remote site connected to the same PE router which has lower site identifier. The PE router cannot establish a pseudowire to this remote site and the associated remote site identifier cannot be used to distribute VPLS label blocks. However, this is not an error state. Traffic continues to be forwarded to the PE router interface connected to this remote site when the remote site is in this state. • IL—The incoming packets for the connection have no MPLS label. • MI—The configured mesh group identifier is in use by another system in the network. • ST—The router has switched to a standby connection. • PB—Profile is busy. • SN—The neighbor is static.
Time Last Up	The time when the connection was last in the Up condition.

Show Routing Tables

The Routing Tables window enables you view the routing table information for the selected virtual routing instance. For L3VPN and EVP services, you can determine which LSPs or tunnels are being used by looking at the routing tables.

- **Routing Tables**—The Routing Tables table shows the routing tables associated with the virtual instance and the number of active routes in each table. Click on a routing table to display the actual contents of the routing table.

- **Details**—The Details table shows the contents of the selected routing table. [Table 182 on page 684](#) displays the fields that are displayed in the Details table.

Table 182: Show Routing Table Field Descriptions

Name	Description
Routing Instance	Name of the routing instance.
Number of Destinations	Number of destinations for which there are routes in the routing table.
Active Routes	Number of routes that are active.
Hidden Routes	Number of routes that are not used because of routing policy.
Hold-down Routes	Number of routes that are in the hold-down state before being declared inactive.
Total Routes	Total number of routes.
Destination Prefix	<p>Route destination (for example:10.0.0.1/24). Sometimes the route information is presented in another format, such as:</p> <ul style="list-style-type: none"> • MPLS-label (for example, 80001). • interface-name (for example, ge-1/0/2). • neighbor-address:control-word-status:encapsulation type:vc-id:source (Layer 2 circuit only. For example, 10.1.1.195:NoCtrlWord:1:1:Local/96): <ul style="list-style-type: none"> • neighbor-address—Address of the neighbor. • control-word-status—Whether the use of the control word has been negotiated for this virtual circuit: NoCtrlWord or CtrlWord. • encapsulation type—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport. • vc-id—Virtual circuit identifier. • source—Source of the advertisement: Local or Remote.
State	State of the route.
Protocol	Name of the protocol from which the route was learned. For example, OSPF , RSVP , and Static .
Protocol Preference	Preferred protocol for this routing instance. Junos OS uses this preference to choose which routes become active in the routing table.
Age	Displays how long since the route was learned.

Table 182: Show Routing Table Field Descriptions (*continued*)

Name	Description
Metric	Cost value of the indicated route. For routes within an AS, the cost is determined by the IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.
BGP Local Preference	A metric used by BGP sessions to indicate the degree of preference for an external route. The route with the highest local preference value is preferred.
Route Learned From	Interface from which the route was received.
AS Path	<p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> • I—IGP • E—EGP • ?—Incomplete; typically, the AS path was aggregated.
Validation State	<p>(BGP-learned routes) Validation status of the route:</p> <ul style="list-style-type: none"> • Invalid—Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message exceeds the maximum length permitted in the database. • Unknown—Indicates that the prefix is not among the prefixes or prefix ranges in the database. • Valid—Indicates that the prefix and autonomous system pair are found in the database.
Next Hop Type	<p>Next hop to the destination. An angle bracket (>) indicates that the route is the selected route.</p> <p>If the destination is Discard, traffic is dropped.</p>
Local Interface	The local interface used to reach the next hop.
Address	IP address of the interface.
Via Interface	Interface used to reach the next hop. If there is more than one interface available to the next hop, the interface that is actually used is followed by the word Selected.
MPLS Label	MPLS label and operation occurring at the next hop. The operation can be pop (where a label is removed from the top of the stack), push (where another label is added to the label stack), or swap (where a label is replaced by another label).

Show MAC Table

The Show MAC table window displays the MAC table for the selected routing instance.

[Table 183 on page 686](#) describes the fields that are displayed in the Show MAC Table window.

Table 183: Show MAC Table Field Descriptions

Field Name	Description
Routing Instance	Name of the routing instance.
Type	Identifies the routing instance type: <ul style="list-style-type: none"> • EVPN • L2VPN • L3VPN • Virtual Switch <p>The default routing instance is of this type.</p> <ul style="list-style-type: none"> • VPLS • VRF (L3VPN)
Bridge Domain	Name of the bridging domain.
VLAN ID	VLAN ID of the routing instance or bridge domain in which the MAC address was learned.
MAC Address	MAC address or addresses learned on a logical interface.
MAC Flags	Status of MAC address learning properties for each interface: <ul style="list-style-type: none"> • S—Static MAC address is configured. • D—Dynamic MAC address is configured. • L—Locally learned MAC address is configured. • C—Control MAC address is configured. • SE—MAC accounting is enabled. • NM—Non-configured MAC. • R—Remote PE MAC address is configured.
Logical Interface	Name of the logical interface.

RELATED DOCUMENTATION

[Understanding the Monitor Mode Tasks Pane | 665](#)

Monitoring Port Utilization

IN THIS SECTION

- [How to Access the Port Utilization Task | 687](#)
- [Port Utilization Details Window | 688](#)
- [Utilization for Device Window | 688](#)
- [Utilization for Fabric Devices Window | 690](#)

Network Director provides information about port utilization in either one of two places, depending on the node you select in the View pane:

- Port Utilization monitor—This monitor, available in the Summary tab, provides a bar chart that shows the aggregate utilization of the ports on a device or devices over a period of time that you select. For more information about using the Port Utilization monitor, see [“Port Utilization Monitor” on page 743](#).
- Port Utilization task—This task, available from **View > Port Utilization** in the Tasks pane of the Summary or Traffic tabs, provides a bar chart similar to the Port Utilization monitor bar chart. Unlike the Port Utilization monitor, it also enables you to obtain information on individual port utilization over time when you have selected an individual device or Layer 3 Fabric in the View pane.

This topic describes the Port Utilization task. It describes:

How to Access the Port Utilization Task

1. Click **Monitor** in the Network Director banner.
2. Select a node in the View pane that contains the ports whose utilization you want to monitor.
3. Select the **Summary** or **Traffic** tab.
4. In the Tasks pane, select **View > Port Utilization**.

If you have selected a node that contains more than one device, the Port Utilization Details window opens. For information about this window, see [“Port Utilization Details Window” on page 688](#).

If you have selected an individual device, the Utilization for Device window opens. For information about this window, see [“Utilization for Device Window” on page 688](#).

If you have selected a fabric device such as Junos Fusion, Layer 3 Fabric, or Virtual Chassis Fabric, the Utilization for Fabric Device window opens. For information about this window, see [“Utilization for Fabric Devices Window” on page 690](#).

Port Utilization Details Window

This window provides a bar chart showing the aggregate port utilization trend for the devices within the selected scope.

Each bar in the bar chart represents the overall port utilization for all the devices at a polling interval. The vertical axis shows the number of ports polled. The horizontal axis shows the time when each poll was taken.

Each bar is divided into the following colored sections to indicate the distribution of port utilization at the polling interval:

- Green indicates ports that operated at less than 50% of negotiated speed.
- Yellow indicates ports that operated at between 50% and 80% of negotiated speed.
- Red indicates ports that operated at more than 80% of negotiated speed.

You can perform the following actions on the bar chart:

- Change the time period over which to display the trend data by selecting a time period from the list in the upper right corner.
- Remove or restore a utilization category (bar color) by clicking its legend.
- Display a numeric value by mousing over a bar.

Utilization for Device Window

The Utilization for Device window shows the port utilization trend for individual devices and ports. It is available when you select a individual device in the View pane.

The Utilization for Device window provides two views of port utilization:

- **Device**—This view provides a trend chart of overall port use on the device over time.
- **Port**—This view provides a heat map of all the ports on the device, enabling you to view the utilization of individual ports. You can choose a port from the heat map to view a utilization trend chart for that particular port.

The Device view is the default view. Click **Port** to change to the Port view.

Device View

The Device view provides a bar chart that shows the trend of overall port use on the device. Each bar represents the overall port utilization at a polling interval. The vertical axis shows the number of ports polled. The horizontal axis shows the time when each poll was taken.

Each bar is divided into the following colored sections to indicate the distribution of port utilization at the polling interval:

- Green indicates ports that operated at less than 50% of negotiated speed.
- Yellow indicates ports that operated at between 50% and 80% of negotiated speed.
- Red indicates ports that operated at more than 80% of negotiated speed.

You can perform the following actions in Device view:

- Change the time period over which to display the trend data by selecting a time period from the list in the upper right corner.
- Remove or restore a utilization category (bar color) by clicking its legend.
- Display a numeric value by mousing over a bar.

Port View

The Port view provides utilization heat maps of the ports on the device—one heat map for access ports and another for uplink ports. In the heat maps, each port on the device is represented by a box that is color-coded to indicate the level of port utilization. Cooler colors (for example, green) indicate lower port utilization, while hotter colors (for example, red) indicate higher port utilization.

Click a port box to display a utilization trend chart for that individual port.

You can perform the following actions in the Port view:

- On a heat map:
 - Mouse over a port box to see more information about the port such as the port utilization percent, port type, MAC address of the port, duplex mode, device serial number, admin status and operational status of the port, negotiated speed, and the last flap time.
 - Change the time period over which the port utilization percentage is derived.
 - Click a port box to display the utilization trend chart for that port.
 - Use the percentage slider under the port heat map to display only those ports for which utilization falls within a certain percentage range.
- On the port utilization trend chart:
 - Change the time period over which to display the trend data.
 - Display the percentage utilization and polling time by mousing over a data point.

Utilization for Fabric Devices Window

The Utilization for Fabric Devices window provides information about port utilization for the devices and ports within a fabric device such as a Junos Fusion, Layer 3 Fabric, or Virtual Chassis Fabric. It is available when you select a fabric device from the Fabrics container in the View pane.

The top part of the Utilization for Fabric Devices window displays a heat map of the devices in the fabric. Each device in the fabric is shown as either a spine or leaf device (for Layer 3 Fabric or Virtual Chassis Fabric) or aggregation device and satellite device (for Junos Fusion devices) and is color-coded to show the overall port utilization on the device.

You can interact with this fabric-level heat map as follows:

- Mouse over a box representing a device. Information about that device is displayed, such as IP address, model, overall port utilization, and a list of the five ports with the highest utilization.
- Click a box representing a device. The information in the remainder of the window is changed to reflect the port utilization of the selected device. For example, for a Junos Fusion device, clicking an aggregation device displays the heat map for the cascade ports and the uplink ports whereas clicking the satellite device displays the heat map for the extended ports.

You can select two different views of the port utilization on the device:

- **Device**—This view provides a trend chart of overall port use on the device over time.
- **Port**—This view provides a heat map of all the ports on the device, enabling you to view the utilization of individual ports. You can choose a port from the heat map to view a utilization trend chart for that particular port.

The Device view is the default view. Click **Port** to change to the Port view.

Device View

The Device view provides a bar chart that shows the trend of overall port use on the selected device. Each bar represents the overall port utilization at a polling interval. The vertical axis shows the number of ports polled. The horizontal axis shows the time when each poll was taken.

Each bar is divided into the following colored sections to indicate the distribution of port utilization at the polling interval:

- Green indicates ports that operated at less than 50% of negotiated speed.
- Yellow indicates ports that operated at between 50% and 80% of negotiated speed.
- Red indicates ports that operated at more than 80% of negotiated speed.

You can perform the following actions on the bar chart:

- Change the time period over which to display the trend data by selecting a time period from the list in the upper right corner.

- Remove or restore a utilization category (bar color) by clicking its legend.
- Display a numeric value by mousing over a bar.

Port View

The Port view provides utilization heat maps of the ports on the device—one heat map for access ports and another for uplink ports. In the heat maps, each port on the device is represented by a box that is color-coded to indicate the level of port utilization. Cooler colors (for example, green) indicate lower port utilization, while hotter colors (for example, red) indicate higher port utilization.

You can perform the following actions on the device heat map:

- Mouse over a port box to see more information about the port, such as the port utilization percent, port type, MAC address of the port, duplex mode, device serial number, admin status and operational status of the port, negotiated speed, and the last flap time.
- Change the time period over which the port utilization percentage is derived.
- Use the percentage slider under the port heat map to display only those ports whose percent utilization falls within a certain range.
- Click a port box to display the utilization trend chart for that port.

The port utilization trend chart shows the utilization trend for the selected port. You can:

- Change the time period over which to display the trend data.
- Display the percentage utilization and polling time by mousing over a data point.

RELATED DOCUMENTATION

[Port Utilization Monitor | 743](#)

[Network Director Documentation home page](#)

Monitoring Tenant Details

IN THIS SECTION

- [Viewing the List of Tenants | 693](#)
- [View Port Details of Tenants | 694](#)
- [View Endpoints | 694](#)
- [View the Port Utilization Trend for a VXLAN Port | 696](#)

This topic describes how to monitor details about the tenants that are part of your overlay network. You can create overlay networks and tenants in Network Director by using Layer 3 Fabrics that are created and managed from Network Director. You can monitor the tenant details for the entire network, for a specific Layer 3 Fabric that acts as the underlay for the tenant overlay network, or for a data center that uses a Layer 3 Fabric.

To view the tenant details:

1. Click **Monitor** in the Network Director banner.
2. Select the **Summary** tab.
3. For the Monitor life cycle mode, select a combination of view and a device or container as shown in [Table 184 on page 692](#).

Table 184: Scopes that You Can Use to View Tenant Details

View Selector	Selection from the View Pane
Logical or Device	My Network
Logical or Device	My Network > Fabric > Layer 3 Fabric My Network > Fabrics > Layer 3 Fabric > Spine Device My Network > Fabrics > Layer 3 Fabric > Leaf Device
Logical or Device	My Datacenters > Data center My Datacenters > Data center > Layer 3 Fabric My Datacenters > Data center > Layer 3 Fabric > Spine Device My Datacenters > Data center > Layer 3 Fabric > Leaf Device

From the Tasks pane, select **Tasks > View > View Tenants**.


The View Tenants window opens, displaying a list of tenants.

NOTE: The View Tenant task is enabled for devices that are at the virtual chassis level.

You can use the filters available on this window to sort ports by **Port Utilization** percentage or **Port Status**.

Use the breadcrumbs at the top of the window to navigate to the various views within the View Tenant window. For example, in [Figure 29 on page 693](#), from the Endpoint view, you can click **All Ports** to view details of all the ports that Tenant 1 uses, click **Tenant 1** to view the port details summary for Tenant 1, or click **All Tenants** to view the details of all the tenants in the View Tenants window.

Figure 29: Breadcrumbs on the View Tenants window



You can perform the following tasks from the View Tenants window:

Viewing the List of Tenants

The View Tenants window enables you to view details about the tenants, the number of ports used by a tenant, and the status of the ports. The level of information that Network Director displays in this window depends on the scope that you select. If you select a Layer 3 Fabric, this window displays the tenants that are part of that fabric, whereas if you select from My Network, this window displays all the tenants that are part of the network—tenants from multiple data centers and fabrics.

The View Tenant details table displays the details of all the tenants and their port status and utilization. See [Table 185 on page 693](#) for a description of the fields in this table.

Table 185: View Tenant Details Table Field Descriptions

Field	Description
Tenants	Name of the tenant.
VXLAN ID	VXLAN ID of the overlay networks for the tenant.
Total Ports	Number of ports that the tenant uses.
Number of Ports with Status	Displays the number of ports that are up and ports that are down in separate columns.
Ports with Utilization (%)	Displays the number of ports that utilize high, medium, and low bandwidth in separate columns.

View Port Details of Tenants

To view more details about the ports that are used by a tenant:

1. In the View Tenants window, click the number of ports field corresponding to a tenant for which you want to view more details.

The Port Details view opens.

[Table 186 on page 694](#) describes the fields of the Port Details view.

NOTE: You can click the number of ports in any of fields—Total Ports, Number of Ports with Status, or Ports with Utilization (%)—to open the Ports Details view. Network Director filters the ports based on the column that you clicked. For example, if you click the port number in the Total Ports column, the Port Details window displays all the ports that the tenant uses; If you click the port number in the Ports with Utilization (%) > High, the Port Details view filters and displays only the ports that have high utilization.

Table 186: Port Details View Field Descriptions

Field	Description
Device	The device on which the given port exists.
Port	The port number.
Port Status	Indicates whether the port is up or down.
Port Utilization %	Utilization (percentage) of the selected port.
Actions	Click to view details about the VXLAN endpoints. The View All Endpoints window opens.

View Endpoints


Endpoints are the hosts on which a tenant network terminates. You can view the endpoint details for each VXLAN overlay network.

To view endpoint details:

From the Port Details view, click **Show Endpoints** in the Actions column to view the end point details for a tenant on a specific device and port. The Show Endpoints window opens displaying the port number and the corresponding MAC address on a host.

View the Port Utilization Trend for a VXLAN Port

To view the port utilization trend for a port:


1. From the Port Details view, click  in the Port Utilization % column to view the utilization trend for the port.

The Port Utilization Trend window is divided into three sections—packets details graph, port details, packet counter and error count table.

The packet details graph displays the number of packets that passed through the port per second plotted during a certain time period as shown in [Figure 30 on page 696](#). The vertical axis on the left shows the number of packets per second. The horizontal axis shows the time when the packet count was taken. The dotted lines indicate the number of errors. You can read the error count for any time period by using the vertical axis on the right.

NOTE: The default polling interval for the packet details graph is 1 hour. You can modify this by selecting an appropriate value from the polling interval drop down list that is displayed above the packet details graph.

Figure 30: Packet Details Graph



The port details section displays the port name, port usage, MAC address of the port, and the number of packets that passed through the port.

The packet counter and error count table consists of two sub-tabs. The Packet Counter sub-tab displays the distribution of unicast, broadcast, and multicast traffic types on the port you selected. The port traffic is classified as *Unicast In*, *Broadcast In*, *Multicast In*, *Unicast Out*, *Broadcast Out*, and *Multicast Out*.

The Error Count sub-tab displays the count of packet errors under each major error category.

You can select or clear each counter to view or hide details about a specific counter in the packet details graph.

RELATED DOCUMENTATION

[Understanding the Monitor Mode Tasks Pane | 665](#)

[VXLAN—EVPN Overlay Overview | 527](#)

[Network Director Documentation home page](#)

Monitoring Virtual Chassis Protocol Statistics

IN THIS SECTION

- [Procedure for Monitoring Virtual Chassis Protocol Statistics | 697](#)
- [Virtual Chassis Protocol Statistics Window | 697](#)

This topic describes how to monitor Virtual Chassis protocol statistics on a device. You can monitor Virtual Chassis protocol statistics for a Virtual Chassis node in any view.

This topic describes:

Procedure for Monitoring Virtual Chassis Protocol Statistics

1. Click **Monitor** in the Network Director banner.
2. Select a node in the View pane that contains the Virtual Chassis protocol traffic you want to monitor.
3. Select the **Traffic** tab.
4. In the Tasks pane, select **View > VC Protocol Statistics**.

The Virtual Chassis Protocol Statistics window opens. For information about this window, click the Help button in the title bar of the window or see [“Virtual Chassis Protocol Statistics Window” on page 697](#).

Virtual Chassis Protocol Statistics Window

The Virtual Chassis Protocol Statistics window displays information about the Virtual Chassis protocol statistics on the Virtual Chassis node you selected in the View pane. It contains these panes:

- The top pane of the window lists the Virtual Chassis members and provides the information about each member that is described in [Table 187 on page 698](#).

Select a member's table row to see information about that member in the other panes.

- The middle and bottom panes provide the information described in [Table 188 on page 698](#).

Table 187: Virtual Chassis Protocol Statistics Window Top Pane

Table Column	Description
Member	Virtual Chassis member's ID.
Role	Member's Virtual Chassis role. Roles include Primary, Backup, and LineCard.
FPC Slot	Member's FPC slot in the Virtual Chassis.
Member Serial Number	Member's serial number.

Table 188: Virtual Chassis Protocol Statistics Window Middle and Bottom Panes

Field or Table Column	Description
System Name	Member system name.
Purges initiated	Number of purges that the system initiated. A purge is initiated if the software determines that a link-state PDU must be removed from the network.
Shortest-path-first runs	Number of shortest-path-first (SPF) calculations that have been performed.
Link-state PDUs queue length	Number of link-state PDUs waiting in the queue for processing. This value is almost always 0.
Link-state PDU fragments computed	Number of link-state PDU fragments that the local system has computed.
Link-state PDUs regenerated	Number of link-state PDUs that have been regenerated. A link-state PDU is regenerated when it is nearing the end of its lifetime and it has not changed.
Protocol data unit type (PDU)	Protocol data unit type.
PDUs Received	Number of PDUs received since VCCP started or since the statistics were set to zero.
PDUs Processed	Number of PDUs received minus the number of PDUs dropped.
PDUs Dropped	Number of PDUs dropped.

Table 188: Virtual Chassis Protocol Statistics Window Middle and Bottom Panes (*continued*)

Field or Table Column	Description
PDU's Transmitted	Number of PDU's transmitted after VCCP started or after the statistics were set to zero.
PDU's Retransmitted	Number of PDU's retransmitted after VCCP started or after the statistics were set to zero.

RELATED DOCUMENTATION

[Understanding the Monitor Mode Tasks Pane | 665](#)

[Network Director Documentation home page](#)

Viewing Congestion Events

This topic describes how to view congestion events on a device. A congestion event occurs when congestion on a device port exceeds the configured threshold.

You can view congestion events only for devices that support Cloud Analytics Engine and that have the high-frequency traffic statistics feature enabled in Network Director.

To view congestion events on a device, you must first do the following:

- Configure the Data Learning Engine (DLE) settings under **Preferences > Monitoring > Data Learning Engine Settings**. The DLE is a component of Cloud Analytics Engine. For information on configuring the DLE settings, see Setting Up User and System Preferences in [“Specifying the Data Learning Engine \(DLE\) Settings” on page 34](#).
- Enable high-frequency traffic statistics on the device and optionally configure thresholds. For information, see [“Enabling High-Frequency Traffic Statistics Monitoring on Devices” on page 602](#).

To view congestion events on a device:

1. In the View pane, select a device on which the high-frequency traffic statistics feature is enabled.
2. Click **Monitor** in the Network Director banner to open Monitor mode.
3. In the Tasks pane, select **Tasks > View Congestion Events**. The View Congestion Events window opens.

The View Congestion Events window lists congestion events that occurred on the device during the time span of 1 minute. The table column headings are the seconds within the selected minute. Each row represents a device interface. Each cell represents the activity on that interface during that second. When congestion events occurred during that second, a bubble appears in the cell. The size of the bubble indicates how many congestion events occurred during that second. The color of the bubble indicates the severity of the congestion during that second: cooler colors indicate lower severity, and hotter colors indicate higher severity.

You can perform these actions in the View Congestion Events window :

- Use the Select Hour and Select Minute lists to select the minute in which to display congestion events and then click **Submit**.
- Mouse over a port name to change the bubbles in its row into the number of congestion events that occurred during each second.
- Click a bubble to open a bar chart that shows detailed information about the congestion events that occurred during that second.

RELATED DOCUMENTATION

[Enabling High-Frequency Traffic Statistics Monitoring on Devices | 602](#)

[Understanding Monitor Mode in Network Director | 658](#)

[Network Director Documentation home page](#)

Monitoring Client Sessions

IN THIS CHAPTER

- Finding User Sessions | 701
- Finding End Points | 704
- Monitoring Client Sessions | 706

Finding User Sessions

IN THIS SECTION

- Procedure for Finding User Sessions | 701
- Search User Session Window | 702

This topic describes how to find user sessions on the network. You can search for sessions based on several session attributes. When you find a session, you can view its current and historical bandwidth usage.

This topic describes:

Procedure for Finding User Sessions

1. Click **Monitor** in the Network Director banner.

You can search for user sessions in any tab in Monitor mode.

2. In the Tasks pane, select **Tasks > Find User Session**.

The Search User Session window opens. For information about this window, click the Help button in the title bar of the window or see [“Search User Session Window” on page 702](#).

Search User Session Window

The Search User Session window enables you to search for and view information about user sessions. The search scope is the entire managed network, regardless of which node is selected in the View pane. You can view current and historical session information.

To find user sessions:

1. Enter search text in the text box. The search looks for the search text in these session attributes:
 - MAC address
 - IP address (IPv4 or IPv6)
 - User name

2. Click **Search**.

The found user sessions appear in a table. See [Table 189 on page 702](#) for a description of this table.

3. To view more information about a session, click its table row.

Detailed information about the session appears. The MAC address appears at the top of the page. The page contains these sections:

- Current Session Information—Displays information about the current session. [Table 190 on page 703](#) describes the information shown for sessions connected to the network by a wired connection.
 - Past Session Information—Displays information about the MAC addresses' past sessions. This information is not shown for sessions connected to the network by a wired connection. You can select the time period to view from the list above the table. [Table 191 on page 703](#) describes the information shown. You can expand the record of a past session to see more information about it by using the plus and minus buttons in the left column.
4. When you are done viewing a session's details, to return to the search results, click **Back** in the top left corner of the window.

Table 189: User Session Details Table for Found User Sessions

Table Column	Description
MAC Address	MAC address of the connected device.
Client IPv4	IPv4 address of the connected device.
User Name	User name of the connected user.
Session Type	Shows whether the session is connected by wired connection.
Client IPv6	IPv6 address of the connected device.

Table 189: User Session Details Table for Found User Sessions (*continued*)

Table Column	Description
Link-local	Link-local address.

Table 190: Current Session Information for Found Wired User Sessions

Field	Description
Username	Username of the connected user.
Device IP	IP address of the device.
Authentication Type	Type of authentication used to authenticate the session.
VLAN	Name of the VLAN the session is using.
Device Serial	Device's serial number.
Port	Port to which the device is connected.

Table 191: Past Session Information for Found User Sessions

Table Column	Description
Session Start Time	Time when the current session started.
Elapsed Time	Length of time the session has been active.
Client IPv4	IPv4 address of the connected device.
Client IPv6	IPv6 address of the connected device.
Link-local	Link-local address.
VLAN	VLAN to which the client is connected.
RSSI	Received signal strength indication (RSSI), in decibels referred to 1 milliwatt (dBm). A higher value indicates a stronger signal.
SNR	Signal-to-noise ratio (SNR),. A measure of the level of a desired signal against the level of background noise, measured in decibels (dB).
RxUniKBytes Value	Unicast bytes received by the session.

Table 191: Past Session Information for Found User Sessions (continued)

Table Column	Description
RxMultiKBytes Value	Unicast bytes transmitted by the session.
TxUniKBytes Value	Multicast bytes received by the session.
Data Usage (KBytes)	Data transmitted and received by the client, in kilobytes.

RELATED DOCUMENTATION

[Understanding the Monitor Mode Tasks Pane | 665](#)

[Network Director Documentation home page](#)

Finding End Points

IN THIS SECTION

- [Procedure for Finding End Points | 704](#)
- [Find End Point Window | 705](#)
- [Refreshing End Point Information | 705](#)

This topic describes how to find end points on the network. End points are computing devices that are connected to the network. You can search for end points based on several attributes. When you find an end point, you can see its last known location in the network.

This topic describes:

Procedure for Finding End Points

1. Click **Monitor** in the Network Director banner.
You can search for end points in any tab in Monitor mode.
2. In the Tasks pane, select **Tasks > Find Endpoint**.

The Find End Point window opens. For information about this window, click the Help button in the title bar of the window or see [“Find End Point Window” on page 705](#).

Find End Point Window

The Find End Point window enables you to search for end points and see their last known location in the network. The search scope is the entire managed network, regardless of which node is selected in the View pane.

To find end points:

1. Enter search text in the text box. The search looks for the search text in these end point attributes:
 - MAC address
 - IP address
2. Click **Search**.

The found end points appear in the Search Results table. See [Table 192 on page 705](#) for a description of this table.

Table 192: Table of Found End Points

Table Column	Description
MAC Address	MAC address of the connected end point.
IP Address	IP address of the connected end point.
Device Name	Name of the networking device that last saw the end point on the network.
Interface Name	Name of the device interface that last saw the end point on the network.
VLAN	VLAN on which the end point was last seen.
Last Seen	When the end point was last seen on the network.
Actions	Click Verify Current Location to verify the information shown for the end point. If any information changed since the last poll, it is updated in the table.

Refreshing End Point Information

Information about all end points connected to the managed network is polled automatically once every 24 hours. You can refresh this information manually.

NOTE: Refreshing end point information can consume significant system resources and take several minutes to complete, depending on the size of the network and the number of connected end points.

To refresh information about all end points connected to the managed network:

1. Click **Monitor** in the Network Director banner.
2. In the Tasks pane, select **Tasks > Refresh Endpoint**.

A confirmation window opens, listing the job ID of the refresh job.

The endpoint refresh runs as a job. You can monitor the job status in System mode by selecting **Tasks > Manage Jobs** in the Task pane.

RELATED DOCUMENTATION

[Understanding the Monitor Mode Tasks Pane | 665](#)

[Network Director Documentation home page](#)

Monitoring Client Sessions

The Client tab in Monitoring mode provides information about clients and sessions on the network. It is available when the node you select in the View pane contains client and session data. The types of available monitoring data vary depending on the node or node type selected.

To monitor client sessions:

1. Click **Monitor** in the Network Director banner.
2. Select a node in the View pane that contains the client sessions you want to monitor.
3. Select the **Client** tab.
4. To get information about a monitor, click the Help button in its title bar.

The Client monitors include:

- [“Current Sessions by Type Monitor” on page 735](#) shows the current active sessions by their type.

RELATED DOCUMENTATION

[Understanding Monitor Mode in Network Director | 658](#)

[Finding User Sessions | 701](#)

[Network Director Documentation home page](#)

Monitoring Devices

IN THIS CHAPTER

- [Comparing Device Statistics | 708](#)
- [Showing ARP Table Information | 710](#)
- [Viewing PoE Information | 711](#)
- [Monitoring the Status of Logical Interfaces | 713](#)
- [Monitoring the Status of a Virtual Chassis | 715](#)
- [Monitoring the Status of Virtual Chassis Members | 716](#)

Comparing Device Statistics

IN THIS SECTION

- [Procedure for Comparing Device Statistics | 708](#)
- [Compare Interfaces Window | 709](#)

This topic describes how to compare statistics from multiple network devices and interfaces in real time. You select which devices, interfaces, and counters to compare, and how often to poll for new statistics.

This topic describes:

Procedure for Comparing Device Statistics

1. Click **Monitor** in the Network Director banner.

You can compare device statistics in any tab in Monitor mode.

2. In the Tasks pane, select **Tasks > Compare Device Statistics**.

The Compare Interfaces window opens. For information about this window, click the Help button in the title bar of the window or see [“Compare Interfaces Window” on page 709](#).

Compare Interfaces Window

The Compare Interfaces window enables you to compare statistics from multiple device interfaces in real time. The search scope is the entire managed network, regardless of which node is selected in the View pane.

To compare device statistics:

1. Select the devices to compare from the device tree in the Select Devices section.
2. Select a device in the Selected Devices section to select which of its interfaces to compare.
The Select Interfaces section lists the device's interfaces. You can select up to two interfaces per device.
3. Select an Interface in the Select Interfaces section to select which of its counters to compare.
The Select Counters section lists the interface's counters.
4. Select the counters to compare in the Select Counters section.
5. Repeat the process of selecting devices, interfaces, and counters to compare until you are finished selecting what to compare.
6. Select how often the data will be refreshed from the **Data Collection Frequency** list.
7. Click the **Compare** button to start comparing information.
A page opens containing a line graph for each counter you selected. Each graph displays all the interfaces for which its counter is selected.
8. To pause data collection, click the **Pause** button. To resume data collection, click the **Resume** button.
9. To change data collection settings, click the **Back** button.

RELATED DOCUMENTATION

[Understanding the Monitor Mode Tasks Pane | 665](#)

[Network Director Documentation home page](#)

Showing ARP Table Information

IN THIS SECTION

- Procedure for Showing ARP Table Information | 710
- Show ARP Table Information Window | 710

This topic describes how to show Address Resolution Protocol (ARP) table information for a device. ARP table information is collected from the selected device when this task runs. You can search for ARP table records.

Procedure for Showing ARP Table Information

To show ARP table information for a device:

1. Click **Monitor** in the Network Director banner.
2. Select the device in the View pane that you want to monitor.
3. Select **Tasks > Show ARP Table** in the Task pane.

The Show ARP Table Information window opens. For information about this window, click the Help button in the title bar of the window or see [Table 193 on page 710](#). You can click the Refresh button below the table to refresh the data from the device.

Show ARP Table Information Window

The Show ARP Table Information Window shows information from the selected device's ARP table.

Table 193: Show ARP Table Information Window

Control or Column	Description
Search controls	Search for ARP table records. Enter search text in the text box. The table of ARP records displays only matching records. Click the X button to clear the search and display all records.
MAC Address	MAC address.
IP Address	IP address.

Table 193: Show ARP Table Information Window (*continued*)

Control or Column	Description
Interface Name	Interface name.
Expiring in (sec)	Number of seconds until the record expires from the ARP table.

RELATED DOCUMENTATION

[Understanding the Monitor Mode Tasks Pane | 665](#)
[Network Director Documentation home page](#)

Viewing PoE Information

IN THIS SECTION

- [Procedure for Viewing PoE Information | 711](#)
- [Show PoE Information Window | 712](#)

This topic describes how to view Power over Ethernet (PoE) information for EX devices.

Procedure for Viewing PoE Information

To view PoE information for a device:

1. Click **Monitor** on the Network Director banner.
2. Select the device in the View pane that you want to monitor.
3. Select **Tasks > Show PoE Interfaces** in the Tasks pane.

The Show PoE Information window opens. For information about this window, click the Help button in the title bar of the window or see [Table 194 on page 712](#). You can click the **Refresh** button below the table to refresh the data from the device.

Show PoE Information Window

The Show PoE Information window shows information about the PoE ports of the selected device.

Table 194: Show PoE Information Window

Control or Column	Description
Search controls	Search for PoE records. Enter search text in the text box. The PoE table displays only the matching records. Click the X button to clear the search and display all records.
Interface Name	Name of the PoE interface.
Admin Status	Administrative state of the PoE interface—Enabled or Disabled. If the PoE interface is disabled, it can provide network connectivity, but it cannot provide power to connected devices.
Operational Status	Operational status of the PoE interface. The operational status can have one of the following values: <ul style="list-style-type: none"> • ON—The interface is currently supplying power to a powered device. • OFF—PoE is enabled on the interface, but the interface is not currently supplying power to a powered device. • FAULT—PoE interface is in the OFF state due to a fault condition. • Disabled—PoE is disabled on the interface.
Max Power Limit	Maximum power that can be provided by the interface.
Priority	Interface power priority—High or Low.
Power Consumption	Amount of power being used by the interface.
Class	Class of the powered device—IEEE 802.3af (PoE) or IEEE 802.3at (PoE+). Class 0 is the default class and is used when the class of the powered device is unknown. If no powered device is connected, this column displays Not Applicable.
LLDP Negotiation Priority	Interface power priority negotiated by LLDP.
LLDP Negotiation Power	Amount of power negotiated by LLDP, to be used by the interface.

RELATED DOCUMENTATION

Monitoring the Status of Logical Interfaces

IN THIS SECTION

- [Locating Information about Logical Interfaces | 713](#)
- [Show Logical Interface Information Table | 713](#)

Network Director provides real-time statistics on logical Ethernet switching interfaces for switches, routers, Virtual Chassis, and Layer 3 Fabrics.

This topic describes:

Locating Information about Logical Interfaces

Real-time logical interface statistics, including VLAN information, are available from the Show Logical Interfaces window in Monitoring mode. To find this information:

1. Select **Monitor** in the Network Director banner.
2. Select a node in the View pane that contains the logical interface you want to monitor.
3. Select the Equipment tab.
4. Click **Logical Interfaces** in the Tasks pane to open the Show Logical Interface Information table in main window.

Show Logical Interface Information Table

The Show Logical Interface Information table provides interface, VLAN, and spanning-tree status for an interface. The information is presented in a tabular format. The fields in the Show Logical Interface Information table are described in [Table 195 on page 714](#).

Table 195: Show Logical Interface Information Fields

Field	Description
Logical Interface Name	The logical interface name.
Serial Number	The serial number of the device to which the logical interface belongs.
VLAN Membership ID	The VLAN to which the interface belongs.
Bridge Domain Membership	The bridge domain to which the interface belongs (for only devices that do not use the Enhanced Layer 2 Software (ELS)).
802.1Q Tag	The IEEE 802.1Q identifier for the VLAN.
Tagging	Indicates whether the packets entering the port are tagged or untagged.
Logical Interface State	Indicates whether the logical interface is up or down.
STP State	Indicates whether the interface is discarding (blocked) or forwarding (unblocked).
Port Mode	<p>Indicates one of three modes: access, tagged-access, or trunk.</p> <ul style="list-style-type: none"> • Access—The interface can be in a single VLAN only. • Tagged-access—The interface can accept tagged packets from one access device. • Trunk—The interface can be in multiple VLANs and accept tagged packets from multiple devices.

RELATED DOCUMENTATION

[Monitoring Traffic on Devices | 670](#)
[Monitoring Traffic on Layer 3 VLANs | 674](#)
[Network Director Documentation home page](#)

Monitoring the Status of a Virtual Chassis

When you select a Virtual Chassis from the network tree in any view, four monitors are displayed that give at-a-glance information about the status and performance of the Virtual Chassis. Use this information to monitor the chassis as a whole, without reviewing each switch independently.

To locate the Virtual Chassis monitors:

1. Click **Monitor** in the Network Director banner to ensure that you are in Monitor mode.
2. Expand the network tree to expose the Virtual Chassis node.
3. Select the Virtual Chassis in the network tree.
4. Click the **Equipment** tab to display the four monitors.
5. Click the Help icon on the monitor learn more about the purpose or fields on a monitor.

The four monitors are:

- [“Resource Utilization Monitor for Switches, Routers, Virtual Chassis, and Virtual Chassis Fabrics” on page 745](#), that provides information about the composition of the chassis, its members, and the location of neighboring switches.
- [“Status Monitor for Virtual Chassis” on page 749](#), that provides information about the uptime, IP address, and hostname of the Virtual Chassis.
- [“Port Status Monitor” on page 740](#), that provides port level status information.

RELATED DOCUMENTATION

[Monitoring the Status of Virtual Chassis Members | 716](#)

[Resource Utilization Monitor for Switches, Routers, Virtual Chassis, and Virtual Chassis Fabrics | 745](#)

[Status Monitor for Virtual Chassis | 749](#)

[Port Status Monitor | 740](#)

[Network Director Documentation home page](#)

Monitoring the Status of Virtual Chassis Members

When you select a member of a Virtual Chassis in the any view, Network Director displays four monitors. At the member node level, the information is highly specific to the equipment.

To locate these monitors:

1. Click **Monitor** in the Network Director banner to ensure that you are in Monitor mode.
2. Expand the network tree to expose the member of the Virtual Chassis node.
3. Select the Virtual Chassis in the network tree.
4. Click the **Equipment** tab to display the monitors.
5. Click the Help icon on the monitor to learn more about the purpose or fields on a monitor.

The monitors at this level are:

- [“Port Status Monitor” on page 740](#), that provides summary and detailed information about the status of the physical network interfaces for the selected node in the View pane.
- [“Power Supply and Fan Status Monitor” on page 744](#), that provides a graphical representation of the operating condition for this member. These graphs also show the ratio of filled slots to available power and fan slots.
- [“Status Monitor for Virtual Chassis Members” on page 751](#), that provides status information for this member of the Virtual Chassis.

RELATED DOCUMENTATION

[Monitoring the Status of a Virtual Chassis | 715](#)

[Port Status Monitor | 740](#)

[Power Supply and Fan Status Monitor | 744](#)

[Status Monitor for Virtual Chassis Members | 751](#)

[Network Director Documentation home page](#)

Monitoring and Analyzing Fabrics

IN THIS CHAPTER

- [Monitoring Junos Fusion Fabric Systems and Components | 717](#)
- [Analyzing Virtual Chassis Fabrics | 718](#)

Monitoring Junos Fusion Fabric Systems and Components

This topic describes how to monitor Junos Fabric systems and their components, and which Junos Fusion-specific monitors are available.

To monitor a Junos Fusion system or its components:

1. Click **Monitor** in the Network Director banner.
2. Select the fabric container or a specific Junos Fusion fabric that you want to monitor in the View pane.

The tabs and monitors that are available depends on your selection.

NOTE: You cannot select the Aggregation Device or Satellite Device nodes within a fusion fabric in the View pane. To monitor a fabric, select the fusion fabric parent node.

3. To get information about a monitor, click the Help button in its title bar, or refer to [Table 196 on page 718](#) for information about the Junos Fusion-specific monitors that are available for each node type.

Table 196: Monitors Available for Junos Fusion Systems and Components

Selection in the View Pane	Monitoring Tab	Available Junos Fusion-Specific Monitors
Junos Fusion system	Summary	<ul style="list-style-type: none"> • Equipment Summary By Type Monitor on page 739 • Port Status Monitor on page 740 • Current Active Alarms Monitor on page 776
	Traffic	<ul style="list-style-type: none"> • Unicast vs Broadcast/Multicast Monitor on page 753 • Unicast vs Broadcast/Multicast Trend Monitor on page 754 • Traffic Trend Monitor on page 753 • Error Trend Monitor on page 736
	Equipment	<ul style="list-style-type: none"> • Status Monitor for Junos Fusion Systems on page 747 • Port Status Monitor on page 740
Fabrics node	Summary	<ul style="list-style-type: none"> • Equipment Summary By Type Monitor on page 739 • Port Status Monitor on page 740 • Current Active Alarms Monitor on page 776

RELATED DOCUMENTATION

[Understanding Monitor Mode in Network Director | 658](#)

[Network Director Documentation home page](#)

Analyzing Virtual Chassis Fabrics

IN THIS SECTION

- [Procedure for Analyzing a Virtual Chassis Fabric | 719](#)
- [Using the Fabric Health Check Tab | 719](#)
- [Using the Topology Tab | 719](#)

This topic describes how to analyze Virtual Chassis Fabrics (VCFs). The Run Fabric Analyzer task analyzes a VCF and provides information about its health, connectivity, and topology. The Fabric Analyzer works on VCFs configured in spine-and-leaf mode.

This topic describes:

Procedure for Analyzing a Virtual Chassis Fabric

Fabric analysis is automatically performed on managed VCFs. You do not need to manually run the analyzer.

To see the results of the analysis of a VCF:

1. Click **Monitor** in the Network Director banner.
2. Select the VCF to analyze in the View pane.
3. Select the **Fabric Analysis** tab.

For information about using the tabs within the Fabric Analysis tab, see the following sections:

- [Using the Fabric Health Check Tab on page 719](#)
- [Using the Topology Tab on page 719](#)

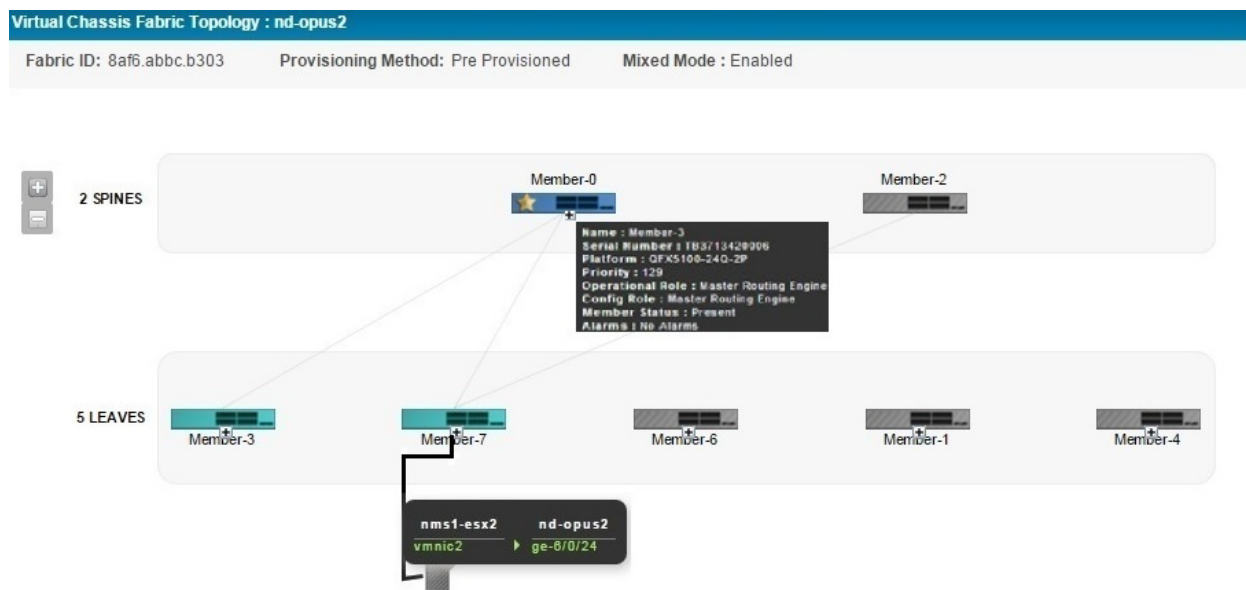
Using the Fabric Health Check Tab

To check the health of a VCF, select the **Fabric Health Check** tab on the Fabric Analysis tab in Monitor mode. The Fabric Health Check tab shows connectivity information about the VCF. Each leaf device is listed in the Leaves column, and each spine device is a table column. Each cell indicates the connectivity status between the leaf device and the spine device.

Using the Topology Tab

The Topology tab displays a topology diagram of the VCF as shown in [Figure 31 on page 720](#). Mouse over a member to view details of that member.

Figure 31: Virtual Chassis Fabric Topology



The details that are displayed in the top panel of the Virtual Chassis Topology view are described in [Table 197 on page 720](#).

Table 197: Common Details for the Virtual Chassis and Virtual Chassis Fabric

Details	Description
Fabric ID VC ID	All members of a Virtual Chassis configuration share one Virtual Chassis identifier (VCID). This identifier is derived from internal parameters.
Provisioning Method	<p>The provisioning mode of the member. Provision mode can be <i>autoprovisioned</i>, <i>preprovisioned</i> or <i>not preprovisioned</i>.</p> <p>Autoprovisioning a Virtual Chassis Fabric (VCF) enables you to “plug and play” devices into your VCF after minimal initial configuration.</p> <p>In a VCF, you can have two to four members configured in the Routing Engine role. Of this, one member acts as the primary Routing Engine and another member acts as the backup Routing Engine. In a preprovisioned configuration, the selection of which member functions as the primary Routing Engine and which as the backup Routing Engine is determined by the software based on the primary-role election algorithm.</p> <p>In a configuration that is not preprovisioned, the selection of the primary and backup is determined by the primary-role priority value and secondary factors in the primary-role election algorithm.</p>
VC Mode	Indicates whether the Virtual Chassis is mixed or not.

The following details are displayed in the Virtual Chassis Topology view for each member depending on the role of the member as shown in [Table 198 on page 721](#).

Table 198: Details of VC/VCF Members

Details	Description	Role
Name	Name of the member switch provided while configuring the device. The device name is displayed as a label.	Primary Backup Line Card
Serial Number	Serial number of the member switch.	Primary Backup Line Card
Platform	Platform of the device. Platform can be QFX3500, QFX3600, QFX5100 or QFX5110.	Primary Backup Line Card
Priority	The primary-role priority value. This is the most important factor in determining the role of the member switch within the Virtual Chassis configuration.	Primary Backup Line Card
Operational Role	Operational role of the device. A device might be configured for a particular role, but can operate in the same or a different role. For example, a spine device configured with a Routing Engine role might operate as a line card. Therefore, the operational role of this device is Line Card. Operational role can be Routing Engine or Line Card.	Primary Backup Line Card
Config Role	The configured role of the device. This can be Routing Engine or Line card.	Primary Backup Line Card
Member Status	Displays the status of each member device: <ul style="list-style-type: none"> ● Present—The device is connected and working fine. ● Not Present—The device is not connected to the VC or VCF. ● Inactive—The device is connected, but not running. ● Non Provisioned—A configuration in which the roles of the members are assigned automatically; not configured statically (preprovisioned). ● Pre Provisioned—A configuration that allows you to deterministically control the member ID and role assigned to a member by associating the member with its serial number. 	Primary Backup Line Card

RELATED DOCUMENTATION

[Understanding the Monitor Mode Tasks Pane | 665](#)

[Network Director Documentation home page](#)

Monitoring Virtual Networks

IN THIS CHAPTER

- [Using Monitor Mode for Virtual Devices | 723](#)
- [Viewing vMotion History in Network Director | 726](#)

Using Monitor Mode for Virtual Devices

IN THIS SECTION

- [Current Active Alarms Monitor | 724](#)

The Monitor mode for virtual devices in your network enables you to view details about your virtual network using the following tabs:

- **Summary**—Displays the status of the virtual network, virtual machine, or virtual switch, active alarms, and the number of hosts and the version of VMware ESXi that is running on each host.
- **vMotion History**—vSphere vMotion is a feature that enables live migration of running virtual machines from one host to another with zero downtime and continuous network availability. You can view the status of the history of all the vMotions for your virtual network in the vMotion History tab. For more details, see [“Viewing vMotion History in Network Director” on page 726](#).

Your current scope—that is, your view and node selection in the View pane—affects which Monitor widgets are available. For example, if you select a virtual switch, Network Director displays the status and the active alarms for the selected virtual switch.

This topic describes:

Current Active Alarms Monitor

The Current Active Alarms monitor shows any active alarm that has not yet been cleared. It is one of the four standard monitors available in Alarm mode. Current Active Alarms is a table that has four fields and appear by default. However, nine fields are available for selection. See [Table 199 on page 724](#) for a description of the table.

Table 199: Current Active Alarms Monitor

Table Column	Description	Shown in Summary by Default	Shown in Detailed View by Default
Name	The alarm name.	Yes	Yes
ID	A system and sequentially generated identification number.	No	No
Assigned To	If assigned to an individual, it shows the name of the person assigned; otherwise, it shows System to mark that the alarm is still unassigned.	No	Yes
Severity	The severity of the alarm. Severity levels are: <ul style="list-style-type: none"> • Critical—A critical condition exists; immediate action is necessary. • Major—A major error has occurred; escalate or notify as necessary. • Minor—A minor error has occurred; notify or monitor the condition. • Info—An informational message; no action is necessary. 	Yes	Yes
Entity ID	The identification of the entity responsible for causing this alarm. The Entity ID is the key for correlation of events into an alarm. The Entity ID can be the IP address of the device.	Yes	Yes
Reporting Device IP	The IP address of the device or network element that generated the alarm. The SNMP agent is located at the source IP. In most cases, the source IP is the IP address of the switch.	Yes	Yes
Reporting Device	The hostname or IP address of the reporting device.	Yes	Yes
Creation Date	The date and time the alarm was first reported.	No	No

Table 199: Current Active Alarms Monitor *(continued)*

Table Column	Description	Shown in Summary by Default	Shown in Detailed View by Default
Last Updated	The date and time that the information for the alarm was last modified.	Yes	Yes
Updated By	Either the system or the last user who modified the alarm.	No	No

Clicking the Details icon opens Alarm Details where you can sort and disposition alarms by state (Acknowledged, Clear, or Active).

RELATED DOCUMENTATION

[Viewing vMotion History in Network Director](#) | 726

[Network Director Documentation home page](#)

Viewing vMotion History in Network Director

vSphere vMotion is a feature that enables live migration of running virtual machines from one host to another with zero downtime and continuous network availability. vMotion is a key feature that enables the creation of a dynamic, automated and self-optimizing Network Director.

If a vMotion happens in any of the virtual machines that are under the management of Network Director, then Network Director initiates a job to track the vMotion and the corresponding changes to orchestration. You can view the status of the history of all the vMotions for your virtual network in the vMotion History page. You can also view the status of the orchestration job that is initiated because of this vMotion by clicking the Orchestration Job ID field.

After the orchestration job is completed successfully, you must manually resynchronize the physical switch's configuration by using Network Director. If the system of record (SOR) mode set for the Junos Space Network Management Platform is:

- Network as system of record (NSOR), then performing a resynchronization ensures that Junos Space automatically resynchronizes its configuration record to match the device configuration and sets the device configuration state to In Sync when the synchronization completes. For more details, see [“Resynchronizing Devices When Junos Space Is in NSOR Mode” on page 619](#).
- Junos Space as system of record (SSOR), then you must perform a resynchronization and accept the out-of-band changes. Both the Junos Space configuration record and the Network Director Build mode configuration are resynchronized to reflect the out-of-band configuration changes. For more details, see [“Resynchronizing Devices When Junos Space Is in SSOR Mode” on page 619](#).

To view the vMotion history:

1. While in the Monitor mode, select a Virtual Network.
2. Select the **vMotion History** tab.

The vMotion History page appears. You can view the details shown in [Table 200 on page 726](#) in the vMotion History page.

Table 200: View vMotion History fields

Field	Description
VM Name	Name of the virtual machine that had undergone a vMotion.
vNetwork	Name of the virtual network.
Source Host	The host from which the virtual machine moved.
Destination Host	The host to which the virtual machine moved.

Table 200: View vMotion History fields (*continued*)

Field	Description
Started On	Time when the vMotion started.
Completed On	Time when the vMotion completed.
Status	Indicates the status of the vMotion.
Source Switches	Host name of the physical switch to which the host was connected before the vMotion.
Source Switch Port	Port on the source physical switch to which the host was connected.
Destination Switches	Host name of the physical switch to which the host is connected after the vMotion.
Destination Switch Port	Port on the destination physical switch to which the host is connected.
Orchestration Job IDs	<p>The ID of orchestration job that was initiated as a result of the given vMotion.</p> <p>Click a job ID to view details about the orchestration job that got initiated as a result of the vMotion.</p>
MAC Address	MAC address of the virtual machine.

3. You can check the status of the orchestration job corresponding to a given vMotion by clicking the orchestration job ID. Network Director opens the vMotion Orchestration window displaying job details such as the name, percentage complete, status, start and end time, and the summary of the job.

RELATED DOCUMENTATION

[Network Director Documentation home page](#)

General Monitoring

IN THIS CHAPTER

- [Selecting Monitors To Display on the Summary Tab | 728](#)
- [Changing Monitor Polling Interval and Data Collection | 729](#)
- [Pinging Host Devices | 729](#)
- [Troubleshooting Network Connections Using Traceroute | 731](#)

Selecting Monitors To Display on the Summary Tab

When you select the My Network node in the View pane, the Summary tab in Monitor mode enables you to select which monitors to display. If you select more than four monitors, a scroll bar appears to allow you to scroll to the additional monitors.

To select monitors to display on the Summary tab:

1. Click **Monitor** in the Network Director banner.
2. Select the **My Network** node in the View pane (the top node in the tree).
3. To select which monitors to display on the Summary tab:

- a. Click **Select Monitors to Display** in the Tasks pane.

The Select Monitors window opens. The monitors that are already selected to display are listed in the Selected list. The other available monitors are listed in the Available list.

- b. To move a monitor from one list to the other list, click the monitor name, and then click the right or left arrow button, as appropriate.

- c. To change the order in which the selected monitors appear in the tab, select a monitor name and move it in the list using the up and down arrow buttons. The arrow buttons at the top and bottom of the stack of buttons move the selected monitor to the top or bottom of the list, respectively.
 - d. Click **Save** to save your changes, or click **Cancel** to cancel your changes.
4. To get information about a monitor, click the Help button in its title bar.

RELATED DOCUMENTATION

[Understanding Monitor Mode in Network Director | 658](#)

[Network Director Documentation home page](#)

Changing Monitor Polling Interval and Data Collection

Network Administrators can change the default polling interval for monitors. The default polling period varies by monitor category. You can change these values in Preferences, found in the Network Director banner. You can also enable or disable the data collection processes used by monitors in Preferences.

RELATED DOCUMENTATION

[Setting Up User and System Preferences | 26](#)

[Network Director Documentation home page](#)

Pinging Host Devices

Use the Ping Host task in Monitor mode to determine whether an EX Series host can be reached over the network from the device selected in the network tree. Entering a hostname or an address creates a periodic ping task that sends a series of Internet Control Message Protocol (ICMP) echo (ping) requests to the specified host. The output of the task displays in the Response Console.

The Ping from Device to a Host task is available only for EX Series switches and QFX Series switches in your network.

1. Select either IP or HostName in the Remote Host Details box.

2. Type the IP address or hostname for the device that you want to reach.
3. Click **Ping** to use the default settings and start the requests or select the plus (+) symbol to use the Advanced Search Criteria. The fields in Advanced Search Criteria are described in [Table 201 on page 730](#).

Table 201: Ping Host Advanced Search Criteria Field Descriptions

Field	Description	Default
Count	Indicates the number of ping requests to send. Valid values are 1 through 24.	5
Type of Service	Sets the type-of-service (ToS) field in the IP header of the ping packets. The range of values is 0 through 255. If the routing platform does not support ToS, the field is ignored.	0
Time To Live	Indicates the time-to-live hop count for the ping request packet. Valid values are 0 through 255.	0
Wait Interval	Indicates the amount of time in seconds between ping requests. Valid values are 0 through 24; a 0 value sends the request immediately.	1
Packet Size	Indicates the size of the ping request packet in bytes. The routing platform adds 8 bytes of ICMP header to this size before sending the request packet.	56
Interface	Sends the ping requests on the interface you specify. If you do not specify this option, ping requests are sent on all interfaces.	All
Source	Uses the source address that you specify in the ping request packet.	None

RELATED DOCUMENTATION

[Understanding Monitor Mode in Network Director | 658](#)

[Network Director Documentation home page](#)

Troubleshooting Network Connections Using Traceroute

Traceroute is a diagnostic tool that enables you to display the route that a packet takes to reach the destination and measure transit delays of packets across an Internet Protocol (IP) network. You can use traceroute to troubleshoot and identify points of failure in your switching network. In traceroute, the source device sends three Internet Control Message Protocol (ICMP) echo request packets to the destination device. This is done sequentially till the source receives an ICMP echo reply message from the destination device. The time-to-live (TTL) value is used in determining the number of intermediate devices that the packets traverse before reaching the destination device.

You can use traceroute for EX Series switches and QFX Series switches.

To start a traceroute from the selected device to another device in your network:

1. Select either IP or HostName in the Remote Host Details box.
2. Type the IP address or hostname for the device to which you want to start a traceroute.
3. Click **Trace** to use the default settings and start the traceroute or select the plus (+) symbol to use the Advanced Options. The fields in Advanced Options are described in [Table 202 on page 731](#).

Table 202: Traceroute Advanced Options Field Descriptions

Field	Description	Default
Interface	Sends the Internet Control Message Protocol (ICMP) echo request packets on the interface you specify. If you do not specify this option, ICMP packets are sent on all interfaces.	Select a value from the list.
Time To Live	Indicates the time-to-live hop count for the ICMP echo request packets. Default value is 30. Valid values are 1 through 255.	30
Wait Interval	Indicates the amount of time in seconds between echo requests. Default value is 5. Valid values are 1 through 24.	5
Type of Service	Sets the type-of-service (ToS) field in the IP header of the echo packets. The range of values is 0 through 255. If the routing platform does not support ToS, the field is ignored.	0

RELATED DOCUMENTATION

[Understanding Monitor Mode in Network Director | 658](#)

[Network Director Documentation home page](#)

Monitor Reference

IN THIS CHAPTER

- 802.11 Packet Errors Monitor | 734
- Access vs. Uplink Port Utilization Trend Monitor | 734
- Current Sessions Monitor | 735
- Current Sessions by Type Monitor | 735
- Error Trend Monitor | 736
- Equipment Summary By Type Monitor | 739
- Node Device Summary Monitor | 740
- Port Status Monitor | 740
- Port Status for IP Fabric Monitor | 743
- Port Utilization Monitor | 743
- Power Supply and Fan Status Monitor | 744
- Resource Utilization Monitor for Switches, Routers, Virtual Chassis, and Virtual Chassis Fabrics | 745
- Status Monitor for Junos Fusion Systems | 747
- Status Monitor for Layer 3 Fabrics | 748
- Status Monitor for Switches and Routers | 748
- Status Monitor for Virtual Chassis | 749
- Status Monitor for Virtual Chassis Members | 751
- Top Talker - Wired Devices Monitor | 752
- Traffic Trend Monitor | 753
- Unicast vs Broadcast/Multicast Monitor | 753
- Unicast vs Broadcast/Multicast Trend Monitor | 754
- User Session Details Window | 755
- Virtual Chassis Topology Monitor | 756
- VC Equipment Summary By Type Monitor | 758

802.11 Packet Errors Monitor

The 802.11 Packet Errors monitor displays the number of packet errors experienced by the object selected in the View pane. The object is polled and plotted at the standard polling rate.

You can perform the following actions on this graph:

- Change the time period over which to display the percentage of retransmitted packets by selecting a time period from the list in the upper right corner.
- Display a numeric value by mousing the cursor where a vertical grid line bisects a data line.

Packet error data is available when you select a floor, a building, or a site in any view. You must configure floors, buildings, and sites. See [“Creating a Site” on page 111](#), [“Configuring Buildings” on page 112](#), and [“Configuring Floors” on page 114](#) for details.

RELATED DOCUMENTATION

[Creating a Site | 111](#)

[Configuring Buildings | 112](#)

[Configuring Floors | 114](#)

[Network Director Documentation home page](#)

Access vs. Uplink Port Utilization Trend Monitor

The Access vs. Uplink Port Utilization Trend monitor shows trends in the bandwidth utilization of access and uplink ports within the selected node device. It is available on the Summary tab in Monitor mode.

NOTE: After a device is discovered, trend data does not appear immediately—it appears only after three polling periods have occurred.

The information is shown in a line graph. The vertical axis shows bandwidth utilization percentage. The horizontal axis shows the times when data was polled. At each poll, the bandwidth utilization percentage of each port type (access and uplink) is indicated by a dot. The dots are connected by lines to show the trend over time.

You can perform the following actions on the line graph:

- Change the time period over which to display the traffic trends by selecting a time period from the list in the upper right corner.
- Highlight a line in the graph by mousing over the line's legend.
- Remove or restore a line by clicking its legend.
- Display a numeric value by mousing over where a vertical grid line crosses a data line.

RELATED DOCUMENTATION

[Understanding Monitor Mode in Network Director | 658](#)
[Network Director Documentation home page](#)

Current Sessions Monitor

Current Sessions monitor displays the number of active sessions at any given point of time. This monitor is available at the My Network level from the Summary tab.

This monitor displays the number of current active sessions on a pie chart.

RELATED DOCUMENTATION

[Current Sessions by Type Monitor | 735](#)
[Network Director Documentation home page](#)

Current Sessions by Type Monitor

IN THIS SECTION

- [Current Sessions by Type | 736](#)
- [Current Session Details | 736](#)

The Current Sessions by Type monitor provides summary and detailed information about the active sessions within the node selected in the View pane. This monitor is available in the Client tab.

NOTE: If the selected scope is a single switch, this monitor is named Current Sessions by VLAN, and shows the distribution of current sessions by VLAN.

Current Sessions by Type

The summary view of the Current Sessions by Type monitor shows a pie chart of the active sessions within the node selected in the View pane. The chart shows the distribution of sessions by the session type. To change the session type shown in the monitor, select from the **Choose Sessions By Type** list.

Current Session Details

To see detailed information about the sessions in the Current Sessions monitor, click the **Details** button in the monitor title bar. The User Session Details window opens. For information about this window, see [“User Session Details Window” on page 755](#).

RELATED DOCUMENTATION

[Monitoring Client Sessions | 706](#)

[Network Director Documentation home page](#)

Error Trend Monitor

IN THIS SECTION

- [Error Trend | 737](#)
- [Error Trend Details | 737](#)

The Error Trend monitor displays inbound and outbound error trends on the node you selected in the View pane. This monitor is available in the Traffic tab.

This topic describes:

Error Trend

A line graph shows the rate inbound and outbound errors over time. The horizontal axis shows the times when samples were taken. The vertical axis shows errors per second.

NOTE: After a device is discovered, trend data does not appear immediately—it appears only after three polling periods have elapsed.

You can perform the following actions on the line graph:

- Change the time period over which to display the traffic trends by selecting a time period from the list in the upper right corner.
- Highlight a line in the graph by mousing over the line's legend.
- Remove or restore a line by clicking its legend.
- Display a numeric value by mousing over a data point.

Error Trend Details

The Error Trend details window displays detailed information about errors on the node you selected in the View pane. It contains the following elements:

- A line graph shows the rate of errors over time. The horizontal axis shows the times when samples were taken. The vertical axis shows errors.

NOTE: After a device is discovered, trend data does not appear immediately—it appears only after three polling periods have elapsed.

You can perform the following actions on the line graph:

- Change the time period over which to display the traffic trends by selecting a time period from the list in the upper right corner.
- Highlight a line in the graph by mousing over the line's legend.
- Remove or restore a line by clicking its legend.
- Display a numeric value by mousing over a data point.

- Error Trend Details table—Shows detailed information about the data gathered at each sample. For information about this table, see [Table 203 on page 738](#)
- Error Trend Additional Details table—Shows additional error trend details and enables you to display them on the graph. For information about this table, see [Table 204 on page 738](#).

Table 203: Error Trend Details Table

Column	Description
Time	Time when a data sample was taken from devices.
Errors In	Number of inbound errors reported in the sample.
Errors Out	Number of outbound errors reported in the sample.
CRC Errors In	Number of inbound cyclic redundancy check (CRC) errors reported in the sample.
CRC Errors Out	Number of outbound CRC errors reported in the sample.

Table 204: Error Trend Additional Details Table

Column	Description
Series Name	Name of the data series.
Series Value	Value of the data series.
Show	Select the check box to display the series on the graph. Clear the check box to remove the series from the graph.

RELATED DOCUMENTATION

[Monitoring Traffic on Devices | 670](#)

[Network Director Documentation home page](#)

Equipment Summary By Type Monitor

IN THIS SECTION

- [Equipment Summary By Type | 739](#)
- [Equipment Summary By Type Details | 739](#)

The Equipment Summary By Type monitor provides summary and detailed information about the type and number of devices in the scope selected in the View pane. This monitor is available on the Summary tab in Monitor mode.

Equipment Summary By Type

The summary view of the Equipment Summary By Type monitor shows the distribution of device types in the selected scope. Switches in a Virtual Chassis are counted separately from standalone switches. Similarly, the count of satellite devices and aggregation devices in a Junos Fusion system are displayed separately in a pie chart.

Mouse over a segment of the pie chart to see the actual number of devices of that type. Click the details icon to open the Equipment Summary By Type Detail View window.

Equipment Summary By Type Details

The Equipment Summary By Type Detail View window provides details about the distribution of device types in the selected scope. Each table row represents a device type. Device types are defined by the combination of a device family, platform, and operating system version (for some device types). See [Table 205 on page 739](#) for a description of the table columns.

Table 205: Equipment Summary By Type Detail View

Table Column	Description
Device Family	Device family.
Platform	Device platform.
OS Version	Operating system version running on the device.
Device Type	Device type.

Table 205: Equipment Summary By Type Detail View *(continued)*

Table Column	Description
Count	Number of devices of this platform in the selected scope.

RELATED DOCUMENTATION

[Selecting Monitors To Display on the Summary Tab | 728](#)

[Network Director Documentation home page](#)

Node Device Summary Monitor

The Node Device Summary monitor displays information about the port utilization of the nodes within the selected container within a fabric. It is on the Summary tab in Monitor mode. The information is presented in a bar chart. The vertical axis shows node names. The horizontal axis shows the number of ports. Ports are categorized based on the percentage of allocated bandwidth they use: over 80%, between 50-80%, and below 50%. The bar color codes for the categories are shown in the legend below the chart. The five nodes that are using the highest percentage of their bandwidth are shown on the monitor.

RELATED DOCUMENTATION

[Understanding Monitor Mode in Network Director | 658](#)

[Network Director Documentation home page](#)

Port Status Monitor

IN THIS SECTION

- [Port Status Summary | 741](#)
- [Port Status Details | 741](#)

The Port Status monitor provides summary and detailed information about the status of the physical network interfaces for the selected node in the View pane.

If the selected node represents an individual device, the monitor displays data specific to the ports on the device. If the selected node contains multiple devices, the monitor displays data aggregated from all the ports on all the devices.

This topic describes:

Port Status Summary

The summary view of the Port Status monitor displays two pie charts:

- **Admin Status**—Of the interfaces on the selected node, shows the proportion of interfaces that are administratively enabled and that are administratively disabled.
- **Free vs Used**—Of the network interfaces that are administratively enabled, shows the proportion of interfaces that are in use (operationally up) and that are not in use (operationally down).

Mouse over a pie segment to view the actual number of ports. Click the details icon to open the Port Status Details window.

Port Status Details





The Port Status Details table provides details about the physical network interfaces for the selected node, as shown in [Table 206 on page 741](#).

NOTE: You must have a transceiver installed in an SFP, SFP+, or XFP port for information about the port to appear.

Table 206: Port Status Details Table

Field	Description
Port Name	The name of the physical interface.
MAC Address	<p>For standalone EX Series switches, the first five groups of hexadecimal digits are determined when the switch is manufactured. The switch then assigns a unique MAC address to each interface by assigning a unique identifier as the last group of hexadecimal digits.</p> <p>For Virtual Chassis members, the first four groups of hexadecimal digits are determined when the switch is manufactured. The fifth group of hexadecimal digits reflects the role of the member in the chassis, such as primary or linecard.</p>

Table 206: Port Status Details Table (*continued*)

Field	Description
Serial Number	The hardware serial number of the device.
Host Name	The hostname of the device.
Description	A text description of the physical interface.
Current Negotiated Speed (Mbps)	The actual operating speed of the port, in megabits per second (Mbps). Depending on the results of autonegotiation, this speed might be less than the maximum speed supported by the port as indicated by port type.
Configured Speed	The speed configured for the port. If the speed is configured to be determined by autonegotiation, the configured speed is shown as Auto.
Duplex Mode	The duplex mode: full (full-duplex), half (half-duplex), or auto (autonegotiation).
Port Type	The port type (for example, 1 Gigabit Ethernet or 10 Gigabit Ethernet interface).
Admin Status	Indicates the administrative state of the port as  UP or  DOWN.
Operational Status	Indicates the operational status of the port as  UP or  DOWN.
PoE	Indicates whether the PoE traps are enabled for the port. Possible values are: <ul style="list-style-type: none"> • Enabled—PoE traps are generated for the port • Disabled—PoE traps are not generated for the port • N/A—PoE traps are not applicable for the port
Last Flap Time	Date and time at which the advertised link became unavailable, and then, available again.

RELATED DOCUMENTATION

[Monitoring the Status of Virtual Chassis Members | 716](#)
[Network Director Documentation home page](#)

Port Status for IP Fabric Monitor

The Port Status for IP Fabric monitor provides summary information about the status of the physical network interfaces in the Layer 3 Fabric selected in the View pane. The status information the monitor provides is organized by the device role in the fabric (spine or leaf) and by port role (access or uplink).

To use the monitor, first select either **Spines** or **Leaves** and then **Access Ports** or **Uplink Ports**. The information displayed is based on your selections and consists of the following:

- A table that lists the number of each type of port—for example, the number of 10 Gigabit Ethernet ports.
- Two pie charts that display:
 - The administrative status of the physical interfaces—that is the proportion of interfaces that are administratively enabled versus administratively disabled.
 - The operational status of the physical interfaces—that is, of the interfaces that are administratively enabled, the proportion of interfaces that are in use (operationally up) versus not in use (operationally down).

RELATED DOCUMENTATION

[Status Monitor for Layer 3 Fabrics | 748](#)

[Network Director Documentation home page](#)

Port Utilization Monitor

The Port Utilization Monitor displays a bar chart with information about the port traffic utilization on the node selected in the View pane. Each bar in the chart represents the port traffic utilization data gathered at a polling interval. The vertical axis shows the number of ports polled. The horizontal axis shows the time when each poll was taken. The data shown in the graph is aggregated from all the ports contained in the node selected in the View pane.

Each bar is divided into the following colored sections to indicate the distribution of port traffic utilization at the polling interval:

- Green indicates ports that operated at less than 50% of negotiated speed.
- Yellow indicates ports that operated at between 50% and 80% of negotiated speed.
- Red indicates ports that operated at more than 80% of negotiated speed.

You can perform the following actions on the bar chart:

- Change the time period over which to display the data by selecting a time period from the list in the upper right corner.
- Remove or restore a utilization category (bar color) by clicking its legend.
- Display a numeric value by mousing over a bar.

RELATED DOCUMENTATION

[Monitoring Port Utilization | 687](#)

[Network Director Documentation home page](#)

Power Supply and Fan Status Monitor

IN THIS SECTION

- [Power Supply and Fan Status | 744](#)
- [Power Supply and Fan Status Details | 745](#)

The Power Supply and Fan Status monitor provides information about the availability and status of power supplies and cooling fans for the node you select in the View pane.

This monitor is available when you select a switch, router, or a Virtual Chassis member in any view. It appears on the Equipment tab when in Monitor mode.

This topic describes:

Power Supply and Fan Status

The summary view of the Power Supply and Fan Status monitor displays two pie charts:

- **Power Supply Units**—On the node selected, shows the proportion of power supplies that are detected as absent against those that are present in the bay. Those units that are present indicate their operating status as OK, Check, or Failed. The totals shown below the title indicate the total number of power supplies that the device is capable of holding.

- **Fans**—On the node selected, shows the proportion of fans that are detected against those that are present in the bay. Those units that are present indicate their operating status as OK, Check, or Failed. The totals shown below the title indicate the total number of fans that the device is capable of holding.

Mouse over the pie segments to view the number of power supplies or fans in each segment. The total number of units is shown at the bottom of the graph. Click the details icon to open the Power Supply and Fan Status Details window.

Power Supply and Fan Status Details

The Power Supply and Fan Status Details window provides a tabular status view of each power supply and fan in the device.

The top table lists all of the power supplies available in the device. The chart shows the individual status of each power supply as OK, Absent, Check, or Failed.

The lower table lists the fans in the device. The chart shows the status of each fan unit as OK, Absent, Check, or Failed.

RELATED DOCUMENTATION

[Monitoring the Status of a Virtual Chassis | 715](#)

[Monitoring the Status of Virtual Chassis Members | 716](#)

[Network Director Documentation home page](#)

Resource Utilization Monitor for Switches, Routers, Virtual Chassis, and Virtual Chassis Fabrics

IN THIS SECTION

● [Resource Utilization Summary | 746](#)

● [Resource Utilization Details | 746](#)

The Resource Utilization monitor shows a line chart for CPU and memory use for a switch, router, and Virtual Chassis and Virtual Chassis Fabric. The vertical axis shows the percentage of the resource being

consumed. The horizontal axis shows the times when samples were taken. The time period that the chart represents can be selected from a list.

This monitor appears on the Equipment tab in Monitor mode.

This topic describes:

Resource Utilization Summary

The summary view of the Resource Utilization monitor shows a line chart representing memory and CPU use. There are six possible categories shown on the chart, depending on which device type is selected:

CPU User—(Orange) the percentage of time that the CPU uses to run user processes, such as the database.

CPU System—(Green) the percentage of time that the CPU uses on all processes for the system.

CPU Background—(Light Blue) the percentage of time that the CPU uses on background processes.

CPU Interrupt —(Red) the percentage of time that the CPU uses for interrupt handling.

CPU Idle—(Purple) the percentage of time that the CPU is available for work.

5 min Mem avg—(Dark Blue) the amount of memory being used over a 5-minute average.

You can interact with the chart to manipulate the data being displayed by:

- Mouse over a line's legend to highlight the line.
- Removing or restoring a line by clicking the legend item.
- Displaying specific chart values by mousing over a datapoint.
- Changing the time period that the chart covers.

If you select Custom, an additional dialog box opens, enabling you to select a starting and ending date and time.

Resource Utilization Details

In Resource Utilization Details, you can view utilization rates for memory and CPU over different periods of time. You can select the time period from a list or specify a custom value. If you select Custom, an additional dialog box, enabling you to select a starting and ending date and time.

The data is presented in two line charts or graphs, one for memory utilization and the other for CPU utilization. Select a timeframe for analysis from the list to update both graphs. Depending on the timeframe selected, the charts refresh to display time increments proportional to the timeframe. For example, if you select 1 Hour, the time increments are 5 minutes apart; if you select 1 Day, the time increments are 1 hour apart. Mouse over a data point to view the precise value at that point.

RELATED DOCUMENTATION

[Network Director Documentation home page](#)

Status Monitor for Junos Fusion Systems

The Status monitor for Junos Fusion systems provides key information about the status of equipment in a Junos Fusion system and is available on the Equipment tab in Monitor mode.

[Table 207 on page 747](#) describes the fields in this monitor.

Table 207: Status Monitor for Junos Fusion System Fields

Field	Function
Power Supply Status	Indicates the aggregated power supply status for devices in the Junos Fusion system. Power supplies are categorized as absent , OK , check , or failed and the number of power supplies in each category are given.
FAN Status	Indicates the aggregated fan status for devices in the Junos Fusion system. Fans are categorized as absent , OK , check , or failed and the number of fans in each are given.
Used MAC Addresses	Indicates the number of MAC addresses in use in the Junos Fusion system.
Used VLANs	Indicates the number of VLANs in use in the Junos Fusion system.
Status	Indicates the number of devices in the Junos Fusion system that are up and that are down.
Alarm Severity Status	Displays the highest severity of any alarms active on any device in the Junos Fusion system and the number of alarms at that severity level.

RELATED DOCUMENTATION

[Network Director Documentation home page](#)

Status Monitor for Layer 3 Fabrics

The Status monitor for Layer 3 Fabrics provides key information about the status of equipment in a Layer 3 Fabric and is available on the Equipment tab in Monitor mode.

[Table 208 on page 748](#) describes the fields in this monitor.

Table 208: Status Monitor for Layer 3 Fabric Fields

Field	Function
Power Supply Status	Indicates the aggregated power supply status for devices in the Layer 3 Fabric. Power supplies are categorized as absent, OK, check, or failed and the number of power supplies in each category are given.
FAN Status	Indicates the aggregated fan status for devices in the Layer 3 Fabric. Fans are categorized as absent, OK, check, or failed and the number of fans in each are given.
Used MAC Addresses	Indicates the number of MAC addresses in use in the Layer 3 Fabric.
Used Vlans	Indicates the number of VLANs in use in the Layer 3 Fabric.
Status	Indicates the number of devices in the Layer 3 Fabric that are up and that are down.
Alarm Severity Status	Displays the highest severity of any alarms active on any device in the Layer 3 Fabric and the number of alarms at that severity level.

RELATED DOCUMENTATION

[Port Status for IP Fabric Monitor | 743](#)

[Network Director Documentation home page](#)

Status Monitor for Switches and Routers

This monitor provides key information about the status for a standalone switch or a router when the device is selected in any of the views. This monitor is on the Equipment tab in Monitor mode.

[Table 209 on page 749](#) describes the fields in this monitor.

Table 209: Status Monitor Fields

Field	Function
Serial Number	Indicates the hardware serial number of the device.
IP Address	Indicates the IP address of the device.
Uptime	Indicates the amount of time since the last boot of the unit in days, hours, minutes, and seconds.
Status	Indicates whether the device is up or down.
Used MAC Addresses	Indicates the number of MAC addresses in use on the device.
Used VLANs	Indicates the number of VLAN memberships for this device.
Last Configured Time	Indicates the date and time when the device was last configured.
Temperature (°C)	Indicates the ambient temperature (in degrees Celsius).
Junos Version	Indicates the version and release level of Junos OS running on the device.

RELATED DOCUMENTATION

[Network Director Documentation home page](#)

Status Monitor for Virtual Chassis

This monitor provides status information, including power supply and fan information, for a Virtual Chassis. It is on the Equipment tab in Monitor mode.

The Summary view shows key status fields in a table format. Power supply and fan data is represented as small bar chart entries in the table. The Details view also shows the same status information, but expands the power supply and fan information. [Table 210 on page 750](#) displays these fields and their location in the monitor.

Table 210: Virtual Chassis Status Monitor Fields

Field	Function	Location
Serial Number	Indicates the hardware serial number of the primary member.	Summary Detailed
IP Address	Indicates the IP address of the primary member.	Summary Detailed
Uptime	Indicates the amount of time since the last boot of the system in days, hours, minutes, and seconds.	Summary Detailed
Status	Indicates whether the Virtual Chassis is up or down.	Summary Detailed
Used MAC Addresses	Indicates the number of MAC addresses in use on the Virtual Chassis.	Summary Detailed
Used VLANs	Indicates the VLAN memberships for the Virtual Chassis.	Summary Detailed
Last Configured Time	Indicates the date and time since the Virtual Chassis was last configured.	Summary Detailed
Temperature Range (°C)	Indicates the temperature of the coolest and hottest devices in the Virtual Chassis (in degrees Celsius).	Summary Detailed
Junos Version	Indicates the version and release level of Junos OS running on the device.	Summary Detailed
Power Supply Status	Indicates the number of power supplies that are detected as absent or present in the bay. The graphic bar and total count for missing and present power supplies is shown as OK, Check, or Failed.	Summary
Power Supply Status	In the Power Supply and Fan Status table, there are separate table entries for each power supply state with the totals for that state.	Detailed
Fan Status	Indicates the number of cooling fans that are detected as absent or present in the bay. The graphic bar and total count for missing and present fans is shown as OK, Check, or Failed.	Summary
Fan Status	In the Power Supply and Fan Status table, there are separate table entries for each power supply state with the totals for that state.	Detailed

RELATED DOCUMENTATION

Monitoring the Status of a Virtual Chassis 715
Monitoring the Status of Virtual Chassis Members 716
Network Director Documentation home page

Status Monitor for Virtual Chassis Members

Use the Member Status monitor to view key information about the status of Virtual Chassis members. It is displayed on the Equipment tab in Monitor mode when you select a Virtual Chassis member.

[Table 211 on page 751](#), describes the fields in this monitor.

Table 211: Status Monitor for Members Fields

Field	Description
Serial Number	Indicates the hardware serial number of the member.
Member ID	Identifies by number a member switch in a Virtual Chassis.
Member Serial Number	Indicates the hardware serial number of the member.
FPC Slot	Identifies the Flexible PIC Concentrator (FPC) slot number for the member: same as Member Slot.
Member Model	The model number of the member.
Member Mixed Mode	Indicates whether the switch is configured to run in mixed member mode. Valid fields are true or false.
Member Role	Indicates the function and responsibility of the switch in the Virtual Chassis. Possible values are primary, backup, and linecard.

RELATED DOCUMENTATION

Monitoring the Status of Virtual Chassis Members 716
Network Director Documentation home page

Top Talker - Wired Devices Monitor

IN THIS SECTION

- [Top Talker - Wired Devices Summary | 752](#)
- [Top Talker - Wired Devices Details | 752](#)

The Top Talker - Wired Devices monitor provides summary and detailed information about the wired devices that are using the most bandwidth.

Top Talker - Wired Devices Summary

The summary view of the Top Talker - Wired Devices monitor has a bar chart that shows summary information about the wired devices that are using the most bandwidth. Device names or addresses are listed on the vertical axis. Data usage in kilobytes is shown on the horizontal axis. You can mouse over a bar to see more information about that device.

Top Talker - Wired Devices Details

To see detailed information about the top talkers, click the **Details** button in the monitor title bar. The Top Talker - Wired Devices monitor details window has a table containing detailed information about the devices that are using the most bandwidth. [Table 212 on page 752](#) describes the columns in the table. To close the details page, click the **Minimize** button in the title bar.

Table 212: Top Hosts Monitor Details

Column	Description
Host Name	Host's host name.
MAC Address	Host's MAC address
Data Usage (KBytes)	Data used by the host, in kilobytes.
Device Serial Number	Device's serial number.

RELATED DOCUMENTATION

Traffic Trend Monitor

The Traffic Trend monitor displays inbound and outbound traffic trends on the node you selected in the View pane. This monitor is available in the Traffic tab. A line graph shows the rate of each type of traffic over time. The horizontal axis shows the times when samples were taken. The vertical axis shows the data rate, in packets per second.

NOTE: After a device is discovered, trend data does not appear immediately—it appears only after three polling periods have elapsed.

You can perform the following actions on the line graph:

- Change the time period over which to display the traffic trends by selecting a time period from the list in the upper right corner.
- Highlight a line in the graph by mousing over the line's legend.
- Remove or restore a line by clicking its legend.
- Display a numeric value by mousing over a data point.

RELATED DOCUMENTATION

Unicast vs Broadcast/Multicast Monitor

The Unicast vs Broadcast/Multicast monitor displays a pie chart of the current distribution of unicast, broadcast, and multicast traffic types on the node you selected in the View pane. This monitor is available in the Traffic tab.

The traffic is divided into these categories:

- Unicast inbound
- Unicast outbound
- Broadcast inbound
- Broadcast outbound
- Multicast inbound
- Multicast outbound

Mouse over a pie segment to view the actual number of packets.

RELATED DOCUMENTATION

[Monitoring Traffic on Devices | 670](#)

[Network Director Documentation home page](#)

Unicast vs Broadcast/Multicast Trend Monitor

The Unicast vs Broadcast/Multicast Trend monitor displays trends in the data rates of unicast, broadcast, and multicast traffic on the node you selected in the View pane. This monitor is available on the Traffic tab. A line graph shows the rate of each type of traffic over time. The horizontal axis shows the times when samples were taken. The vertical axis shows the data rate, in packets per second.

NOTE: After a device is discovered, trend data does not appear immediately—it appears only after three polling periods have elapsed.

The traffic is divided into these categories:

- Unicast inbound
- Unicast outbound
- Broadcast inbound
- Broadcast outbound
- Multicast inbound
- Multicast outbound

You can perform the following actions on the line graph:

- Change the time period over which to display the traffic trends by selecting a time period from the list in the upper right corner.
- Highlight a line in the graph by mousing over the line's legend.
- Remove or restore a line by clicking its legend.
- Display a numeric value by mousing over a data point.

RELATED DOCUMENTATION

[Monitoring Traffic on Devices | 670](#)

[Network Director Documentation home page](#)

User Session Details Window

The User Session Details window provides information about the active sessions within the node selected in the View pane. To open this window, click the **Details** button in the Current Sessions or Current Sessions by Type monitors.

The following table describes the columns that appear in the user session details table:

Table 213: User Session Details Table

Table Column	Description
User Name	Client's user name
MAC Address	Client's MAC address.
Device Type	Client's device type.
Device Group	Client's device group.
Device Profile	Client's device profile.
VLAN	VLAN to which the client is connected.
Client IP	Client's IP address.
Auth Type	Authorization type used for the client.

Table 213: User Session Details Table (*continued*)

Table Column	Description
B/w[Kbps]	Bandwidth used by the client.
Data Usage (KBytes)	Data transmitted and received by the client, in kilobytes.
Elapsed Time	Length of time the session has been active.
Sample Time	Time when the most recent sample was taken.
RSSI	Received signal strength indication (RSSI). Specified in decibels referred to 1 milliwatt (dBm). A higher value indicates a stronger signal.

TIP: Some table columns are hidden by default. To select which columns to display, mouse over any column heading, click the arrow that appears, mouse over **Columns** in the drop-down menu, and then select the columns to display from the list.

RELATED DOCUMENTATION

[Monitoring Client Sessions | 706](#)

[Network Director Documentation home page](#)

Virtual Chassis Topology Monitor

The Virtual Chassis Topology monitor provides a fast and simple way to view the members and their relationships. It is available on the Equipment tab in Monitor mode. View [Table 214 on page 757](#) for a description of the fields in the monitor.

The summary shows up to five available fields that you can configure to be displayed or hidden. The details page shows an expanded version with up to eleven fields that can also be tailored.

Table 214: Virtual Chassis Topology Fields

Field	Function	Default in Topology Monitor
Member	Identifies by member ID a member switch in a Virtual Chassis.	Summary (hidden) Details (hidden)
Member Role	Indicates the function and responsibility of the switch in the Virtual Chassis. Possible values are primary, backup, and linecard.	Summary Details
Member ID	Same as Member.	Summary Details
FPC Slot	Identifies the Flexible PIC Concentrator (FPC) slot number for the member.	Summary Details
Member Status	Identifies whether the member is present in the Virtual Chassis or Not Present.	Details
Member Serial Number	Identifies the hardware serial number of the switch.	Details
Member Model	Specifies the Juniper model number of the switch.	Details
Member Location	Identifies the wiring closet for the switch.	Details (hidden)
Member Mixed Mode	Indicates whether the switch is configured to run in mixed member mode. Valid fields are true or false.	Details
Neighbor ID	Identifies the neighbors by the member ID.	Summary Details
Neighbor Interface	The Virtual Chassis Port of the neighbor.	Details

RELATED DOCUMENTATION

[Virtual Chassis Topology Monitor | 756](#)
[Network Director Documentation home page](#)

VC Equipment Summary By Type Monitor

IN THIS SECTION

- [VC Equipment Summary By Type | 758](#)
- [VC Equipment Member Status By Type Details | 758](#)

The VC Equipment Summary By Type displays the operational status of the virtual chassis members associated with Logical View, Location View, Device View and Custom Group View.

VC Equipment Summary By Type

The summary view of the VC Equipment Summary By Type monitor shows the total number of virtual chassis members and the connection status of the virtual chassis members in the selected scope.

NOTE: The VC Equipment Summary tab is enabled only for members that are at the virtual chassis container level, virtual chassis level or virtual chassis device level in the applicable view panes.

Mouse over a segment of the pie chart to see the actual number of VC devices of that type. The connection states are as follows:

- UP (green)—Device is connected to Network Director
- DOWN (red)—Device is not connected to Network Director

Click the details icon to open the VC Equipment Member Status By Type Detail View window.

VC Equipment Member Status By Type Details

The VC Equipment Member Status By Type Detail View window provides details about the VC members in the selected scope. See [Table 215 on page 758](#) for the description of the table columns.

Table 215: VC Equipment Member Status By Type Detail View

Table Column	Description
Hostname	Configured name of the device or IP address if no hostname is configured.

Table 215: VC Equipment Member Status By Type Detail View (*continued*)

Table Column	Description
IP Address	IP Address of the device.
Serial Number	Serial number of device chassis.
Platform	Model number of the device.
OS Version	Operating system version running on the device.
Device Family	Device family of the device: <ul style="list-style-type: none"> • JUNOS-EX for EX Series switches • JUNOS for Campus Switching ELS
Device Type	Type of the device: <ul style="list-style-type: none"> • VC—Virtual Chassis primary • VC Member—Virtual Chassis member switch
Connection State	Connection status of the device in Network Director. <ul style="list-style-type: none"> • UP (green up arrow)—Device is connected to Network Director • DOWN (red down arrow)—Device is not connected to Network Director
Configuration State	Displays the configuration status of the device: <ul style="list-style-type: none"> • In Sync—The configuration on the device is in sync with the Network Director configuration for the device. • Out Of Sync—The configuration on the device does not match the Network Director configuration for the device. This state is usually the result of the device configuration being altered outside of Network Director. You cannot deploy configuration on a device from Network Director when the device is Out Of Sync. To resolve this state, use the Resynchronize Device Configuration task in Deploy mode • Sync failed—An attempt to resynchronize an Out Of Sync device failed. • Synchronizing—The device configuration is in the process of being resynchronized • N/A—The device is down.

RELATED DOCUMENTATION

[Selecting Monitors To Display on the Summary Tab | 728](#)
[Network Director Documentation home page](#)



Using Fault Mode

About Fault Mode | **761**

Using Fault Mode | **766**

Fault Reference | **771**

About Fault Mode

IN THIS CHAPTER

- [Understanding Fault Mode in Network Director | 761](#)
- [Understanding the Fault Mode Tasks Pane | 765](#)

Understanding Fault Mode in Network Director

IN THIS SECTION

- [What Are Events and Alarms? | 761](#)
- [Alarm Severity | 762](#)
- [Alarm Classification | 762](#)
- [Alarm State | 764](#)
- [Alarm Notifications | 764](#)
- [Threshold Alarms | 764](#)

The Fault mode shows you information about the health of your network and changing conditions of your equipment. Use Fault mode to find problems with equipment, pinpoint security attacks, or to analyze trends and categories of errors.

This topic describes:

What Are Events and Alarms?

Activity on a network device consists of a series of *events*. A software component on the network device, called an *entity*, is responsible for running the Simple Network Management Protocol (SNMP) to log and monitor these events. When certain types of events are persistent, or when the condition causing the event crosses a threshold, SNMP sends a notification, also called a *trap* to Network Director. Network

Director correlates traps, describing a condition, into an *alarm*. For example, multiple power supply traps coming from a device are correlated into a single power supply alarm for the device.

There are many types of alarms. An alarm can be as routine as when the device changes state or as serious as when a power supply has failed. When an alarm is sent, or *raised*, it stays raised until the triggering condition is resolved or *cleared*. The system can clear the alarm when the state changes again or an administrator can clear it manually, which indicates that the condition is now resolved.

SNMP also plays another role in Network Director. Enabling devices for SNMP with the appropriate read-only V1/V2/V3 credentials, can speed up device discovery.

Alarm Severity

Alarms are ranked by their impact to the network. The following list shows the ranking of alarms in Network Director from alarms that have the most impact to alarms that have the least impact on the network. It also shows the color scheme associated with each level of severity that is reflected in related graphs.

Critical (Red)—A critical condition exists; immediate action is necessary.

Major (Orange)—A major error has occurred; escalate or notify as necessary.

Minor (Yellow)—A minor error has occurred; notify or monitor the condition.

Info (Blue)—An informational message; no action is necessary. Informational alarms do not necessarily indicate an error. It could indicate that a device or entity has changed state.

Administrators can override the default severity of an alarm and set the severity to match their inhouse guidelines. Changing the severity level for an alarm is done on the Fault tab of System Preferences.

Alarm Classification

Network Director organizes alarms into categories so you can view trends in the types of errors occurring on a network. These categories, shown in [Table 216 on page 762](#) are derived from the SNMP Management Information Base (MIB) that is the information database or module containing the trap information for the event.

Table 216: Network Director Alarm Classifications

Category	Description
BFD	Indicates alarms for Bidirectional Forwarding Detection sessions. These alarms are generated from EX Series switches.
BGP	Indicates alarms for BGP4.
Chassis	Indicates alarms for switch hardware, in this case, EX Series switches.

Table 216: Network Director Alarm Classifications (*continued*)

Category	Description
Configuration	Indicates alarms for configuration management.
CoS	Indicates class of service alarms.
DHCP	Indicates local server DHCP alarms.
DOM	Indicates Digital Optical Monitoring alarms that are generated from optical interfaces.
FlowCollection	Indicates alarms generated when collecting and exporting traffic flows.
General	Indicates alarms that are common to all network devices, such as link up/down or authentication.
GenericEvent	Indicates an alarm that is generated from an Op script or event policies.
L2ALD	Indicates MAC address alarms generated from the Layer 2 Address Learning Daemon (L2ALD).
L2CP	Indicates alarms generated by Layer 2 Control Protocol features.
MACFDB	Indicates an alarm for when MAC addresses are learned or removed from the forwarding database of the monitored device.
Misc	Indicates alarms that do not fit into the other categories.
PassiveMonitoring	Indicates alarms that occur on a passive monitoring interface.
Ping	Indicates alarms that are generated during a Ping request.
RMon	Indicates RMON alarms
SONET	Indicates a SONET or SDH alarm on an interface.
SONET APS	Indicates alarms generated on a SONET interface that participates in Automatic Protection Switching (APS).
VirtualChassis	Indicates alarms generated from Virtual Chassis members regarding member or port status.
VNetwork	Indicates virtual networking alarms.

Alarm State

Once an alarm is active, it has one of these states:

- **Active**—Alarms that are current and not yet acknowledged or cleared.
- **Cleared**—Alarms that are resolved and the device or entity has returned to normal operation.

Some alarm states go directly from active to cleared state and require little to no administrative effort. However, other alarms with a high severity should be acknowledged and investigated.

In addition to acknowledging and clearing an alarm, you can assign an alarm to someone and you can append a note or annotation to an alarm. Annotations are helpful for documenting the resolution of an alarm or time estimates for a fix. Changes to an alarm's state are made through the Alarm State monitor in Fault mode.

Alarm Notifications

Alarms can be enabled for email notification. When an alarm with notification enabled is generated, an email is sent to a set of specified addresses. There is a list of global email addresses that receive notifications from all alarms with notification enabled. Each alarm type can also have a list of addresses that receive notification when that alarm type is generated. Administrators can enable notification for alarm types and specify addresses to receive email notifications. These tasks are done on the Fault tab of System Preferences.

Threshold Alarms

Threshold alarms are alarms that are generated when a monitored value crosses the configured threshold. They provide enhanced visibility into potential issues on the network. Administrators configure and manage threshold alarms the same way as other alarms, and can set the threshold level of individual threshold alarms on the Fault tab of System Preferences.

RELATED DOCUMENTATION

[Setting Up User and System Preferences | 26](#)

[Alarms by Severity Monitor | 778](#)

[Alarms by Category Monitor | 778](#)

[Current Active Alarms Monitor | 776](#)

[Alarms by State Monitor | 779](#)

[Network Director Documentation home page](#)

Understanding the Fault Mode Tasks Pane

The Tasks pane in Fault mode provides you with a set of tools for effectively managing alarms on your system.

From the Tasks pane, you can:

- Filter known alarms to locate a specific alarm or error condition by clicking Search Alarms. Use this task to isolate alarms that occurred during a known time-frame or that have annotations associated with them. Although each of the Fault mode monitors can sort the alarms, Search Alarms enable you to submit multiple search and sort arguments as part of your search query.

In addition, Network Director enables you to group the tasks that you perform frequently and create a list of key tasks. You can add any task from the Tasks pane to the Key Tasks list by selecting a task and clicking the plus (+) sign that appears adjacent to the task. For some modes, you can see that Network Director has predefined some key tasks for you. You can modify this set of tasks to suit your requirements. This feature is available in Task pane irrespective of your current mode, scope, or view.

RELATED DOCUMENTATION

[Searching Alarms | 766](#)

[Changing Alarm State | 769](#)

[Understanding Fault Mode in Network Director | 761](#)

[Network Director Documentation home page](#)

Using Fault Mode

IN THIS CHAPTER

- [Customizing Alarms | 766](#)
- [Searching Alarms | 766](#)
- [Changing Alarm State | 769](#)

Customizing Alarms

Ensure that all devices are enabled for SNMP trap forwarding. This task, Set SNMP Trap Configuration, is found in Deploy mode.

Network Director enables you to tailor alarms by:

- Enabling or disabling individual alarms.
- Setting the amount of time alarms are retained in the system.

You can customize alarms using Preferences in the Network Director banner.

RELATED DOCUMENTATION

[Setting Up User and System Preferences | 26](#)

[Network Director Documentation home page](#)

Searching Alarms

Use Search Alarms, available from the Tasks pane, to filter and isolate information about a specific alarm. Use this page to specify complex sorting and filtering criteria for all alarms.

Each field in the Search Alarm window helps narrow the current list of alarms. The more search items you specify, the more specific your results. All fields are optional.

1. Select or type the known descriptors for the alarm. These fields are described in [Table 217 on page 767](#).
2. Click **Search** to run the query. The Alarms Details page opens with the results of your search.
3. Review the alarm. From this page you can change the state of the alarm, annotate, or assign the alarm to personnel. For more information about changing the state of an alarm, view [“Changing Alarm State” on page 769](#).

Table 217: Alarm Search Fields

Search Criteria	Description
State	<p>Use the list to select which alarm states to search for:</p> <ul style="list-style-type: none"> • All—Alarms of all states. • Active—Alarms that are current and not yet acknowledged or cleared. • Clear—Alarms that are resolved and the device or entity has returned to normal operation.

Table 217: Alarm Search Fields (*continued*)

Search Criteria	Description
Category	<p>Fill in one of the available alarm categories:</p> <ul style="list-style-type: none"> • BFD • BGP • Chassis • ClientAndUserSession • Configuration • CoS • DHCP • DOM • FlowCollection • GENERAL • GenericEvent • L2ALD • L2CP • MACFDB • Misc • PassiveMonitoring • Ping • RMon • SONET • SONETAPS • VirtualChassis • VNetwork
Severity	<p>Pull down the list to select the severity level. Not all possible alarm severities are listed. Only the severity levels of your current active alarms are shown. Possible selections are:</p> <ul style="list-style-type: none"> • Critical • Major • Minor • Info
Advanced Search Criteria	
(from) Date	Pull down the calendar and select the starting date of the search.

Table 217: Alarm Search Fields (*continued*)

Search Criteria	Description
(from) Time	Pull down the list to select the starting time of the search. Search times are in military (24-hour) clock format in 30 minute intervals.
(to) Date	Pull down the calendar and select the ending date of the search.
(to) Time	Pull down the list to select the ending time of the search. Search times are in military (24-hour) clock format in 30 minute intervals.
Notes	Enter any keywords or phases that were listed in an existing annotation.

RELATED DOCUMENTATION

[Understanding Fault Mode in Network Director | 761](#)

[Network Director Documentation home page](#)

Changing Alarm State

When an alarm becomes active, it remains active until either the system determines that the condition is resolved or system personnel change the status. Critical alarms always need immediate attention and seldom resolve on their own, but informational messages are often expected actions and results. When a condition is severe or persistent and needs attention, follow these steps:

1. Locate the alarm.
 - a. Click **Fault** in the Network Director banner to enter Fault mode.
 - b. Click the Alarm Details icon on any of the monitors to open the Alarm Details page. Scroll or sort the alarms to find the alarm in question. As an alternate method, click **Search Alarms** in the Tasks pane and filter the active alarm list.
 - c. Select the alarm.
2. Review the Event Details that triggered the trap for the alarm. These events provide insight into the cause or location of the problem.

3. Click **Acknowledge** to indicate that the problem is now known. You should receive a message saying the alarm is acknowledged.
4. Depending whether you can resolve the alarm with the information at hand or not, either assign the alarm to a member of your staff or clear the alarm. Click **Clear** to clear the alarm or click **Assign** and fill in the assignee's name.

At any time in the life cycle of an alarm, you can attach information about the alarm to the alarm record by clicking **Annotate**. Fill in your name in the **Notes By** field and add the note description in the **Notes** field. Click **Add** to record the annotation.

RELATED DOCUMENTATION

[Alarm Detail Monitor | 771](#)

[Network Director Documentation home page](#)

Fault Reference

IN THIS CHAPTER

- [Alarm Detail Monitor | 771](#)
- [Current Active Alarms Monitor | 776](#)
- [Alarms by Category Monitor | 778](#)
- [Alarms by Severity Monitor | 778](#)
- [Alarms by State Monitor | 779](#)
- [Alarm Trend Monitor | 780](#)

Alarm Detail Monitor

IN THIS SECTION

- [Finding Specific Alarms | 772](#)
- [Sorting Alarms | 773](#)
- [Reading Events | 774](#)
- [Investigating Event Attributes | 775](#)
- [Changing the Alarm State | 775](#)

Use the Alarm Detail monitor to sort alarms, view an alarm in depth, and to assign a disposition to an alarm.

By clicking the **Details** icon, you can access the Alarm Detail monitor from any of the four alarm monitors available on the main page in Fault mode (Severity, Category, Current, or State). It is also available from the Current Active Monitors available from the Summary tab in Monitor mode.

This topic describes:

Finding Specific Alarms

Use the Alarm Detail monitor to locate a specific alarm, research the events causing the alarm, and to assign a disposition to the alarm. When an alarm is highlighted in the sorting sequence, the events contributing to the alarm are listed in Event Details and the variable settings are shown in Event Attribute Detail.

To locate an alarm and to assign a disposition to the alarm:

1. Sort the list using the Display list. Sorting choices vary depending on how you arrived here. View [“Sorting Alarms” on page 773](#) for details on sorting options.
2. Review the sorted list. Each entry shows a minimum of one to a maximum of nine fields. These fields are described in [Table 218 on page 772](#).
3. Examine the events and event attributes that contributed to sending the alarm. Events and event attributes are discussed in [“Reading Events” on page 774](#) and [“Investigating Event Attributes” on page 775](#).

Table 218: Alarm Detail Fields

Field	Value	Shown in Detailed View by Default
Name	The alarm name.	Yes
ID	A system and sequentially-generated identification number.	No
Assigned To	If assigned to an individual, it shows the name of the person assigned; otherwise, it shows System to mark that the alarm is still unassigned.	Yes
Severity	The severity of the alarm. Severity levels are: <ul style="list-style-type: none"> • Critical—A critical condition exists; immediate action is necessary. • Major—A major error has occurred; escalate or notify as necessary. • Minor—A minor error has occurred; notify or monitor the condition. • Info—An informational message; no action is necessary. 	Yes
Acknowledged	Indicates if the alarm has been acknowledged.	Yes
Entity ID	The identification of the entity responsible for causing this alarm. The Entity ID is the key for correlation of events into an alarm. The Entity ID could be an IP address of the device.	Yes

Table 218: Alarm Detail Fields (*continued*)

Field	Value	Shown in Detailed View by Default
Reporting Device IP	The IP address of the reporting device that generated the alarm.	Yes
Reporting Device Name	The hostname of the reporting device.	Yes
Creation Date	The date and time the alarm was first reported.	No
Last Updated	The date and time that the information for the alarm was last modified.	Yes
Updated By	Either the system or the last user who modified the alarm.	No

NOTE: You can enter the alarm name, entity ID, or reporting device IP in the search field and press enter to perform searches and display only those searched alarms in the Alarm Detail table.

You can also use Searching Alarms in the Tasks pane to perform searches using multiple arguments. With multiple arguments, you can isolate a single alarm from a long alarm list. For more information, see [“Searching Alarms” on page 766](#).

Sorting Alarms

Depending on the monitor you chose to access Alarm Detail, your sorting options change to reflect the summary monitor. The different sort options are listed in [Table 219 on page 773](#).

Table 219: Sort Options for Alarms

Alarms by Severity Sort	Alarms by Category Sort	Alarms by State and Current Active Alarms Sort
Info	BFD	Clear
Minor	BGP	
Major	Chassis	
Critical	ClientandUserSession	
	Config	

Table 219: Sort Options for Alarms (*continued*)

Alarms by Severity Sort	Alarms by Category Sort	Alarms by State and Current Active Alarms Sort
	CoS	
	DHCP	
	DOM	
	FlowCollection	
	GENERAL	
	GenericEvent	
	L2ALD	
	L2CP	
	MACFDB	
	Misc.	
	PassiveMonitoring	
	Ping	
	RMon	
	SONET	
	SONETAPS	
	VirtualChassis	
	VNetworkS	

Reading Events

When you select an alarm in Alarm Detail, the Event Detail table updates with information about the events that are associated with the alarm. [Table 220 on page 775](#) lists the fields in Event Detail.

Table 220: Event Detail Fields

Field	Value
Name	The event name; also known as the SNMP trap name.
ID	A system-generated, hexadecimal code that uniquely identifies the event.
Description	If the event is an SNMP event, it is shown as a system-generated event.
Type	The type of event, either fault or system alert.
Category	<p>The category of the event message. The category corresponds to the alarm categories shown in the Alarms by Category monitor and the Alarm Settings window. These categories are:</p> <ul style="list-style-type: none"> • General • Chassis • BFD • Core • Misc.
Source	The identification of the entity that is the cause of this event ; it is not necessarily the ID of the event that generated the event.
Originator	The identification of the entity that generated this event, for example, the switch IP address.
Time Updated	The date and time of the last update to the event.

Investigating Event Attributes

The Event Attribute Detail window reflects the variables set during the event. In SNMP terminology, these attributes are known as variable bindings or varbinds. These attributes can provide key information about triggers. For example, if a fan fails, the attribute field could indicate the location of the fan in the chassis.

Changing the Alarm State

When an alarm is first reported, it is considered an active alarm. To change the alarm state, to assign the alarm to a person, or simply to record notes about the alarm, use the buttons on Alarm Details. These buttons are:

- **Acknowledge**—Use this button to acknowledge or record that the alarm is known and is being addressed.
- **Clear**—Use this button to clear or remove the alarm. The clear state says that the issue sending the alarm has been resolved and no longer requires attention.
- **Annotate**—Use this button to record actions taken to resolve the alarm.
- **Assign**—Use this button to assign active or acknowledged alarms to staff.

RELATED DOCUMENTATION

[Alarms by Category Monitor | 778](#)

[Alarms by Severity Monitor | 778](#)

[Alarms by State Monitor | 779](#)

[Current Active Alarms Monitor | 776](#)

[Network Director Documentation home page](#)

Current Active Alarms Monitor

The Current Active Alarms monitor shows any active alarm that has not yet been cleared. It is one of the four standard monitors available in Alarm mode. Current Active Alarms is a table that has four fields and appears by default. However, nine fields are available for selection. View [Table 221 on page 776](#) for a description of the table.

Table 221: Current Active Alarms Monitor

Table Column	Description	Shown in Summary by Default	Shown in Detailed View by Default
Name	The alarm name.	Yes	Yes
ID	A system and sequentially-generated identification number.	No	No
Assigned To	If assigned to an individual, it shows the name of the person assigned; otherwise, it shows System to mark that the alarm is still unassigned.	No	Yes

Table 221: Current Active Alarms Monitor (*continued*)

Table Column	Description	Shown in Summary by Default	Shown in Detailed View by Default
Severity	<p>The severity of the alarm. Severity levels are:</p> <ul style="list-style-type: none"> • Critical—A critical condition exists; immediate action is necessary. • Major—A major error has occurred; escalate or notify as necessary. • Minor—A minor error has occurred; notify or monitor the condition. • Info—An informational message; no action is necessary. 	Yes	Yes
Entity ID	The identification of the entity responsible for causing this alarm. The Entity ID is the key for correlation of events into an alarm. The Entity ID could be an IP address of the device.	Yes	Yes
Reporting Device IP	The IP address of the device or network element that generated the alarm. The SNMP agent is located at the source IP. In most cases, the source IP is the IP address of the switch.	Yes	Yes
Reporting Device	The hostname or IP address of the reporting device.	Yes	Yes
Creation Date	The date and time the alarm was first reported.	No	No
Last Updated	The date and time that the information for the alarm was last modified.	Yes	Yes
Updated By	Either the system or the last user who modified the alarm.	No	No

Clicking the Details icon opens Alarm Details where you can sort and disposition alarms by state (Acknowledged, Clear, Active).

RELATED DOCUMENTATION

[Alarm Detail Monitor | 771](#)

[Understanding Fault Mode in Network Director | 761](#)

Alarms by Category Monitor

Alarms by Category is a table of all active alarms sorted by category. Use this monitor to view where errors are trending. These categories are the same categories shown in the Alarm Settings page.

This monitor is available in all views in the main window when in Fault mode.

The table shows the active categories and the number of alarms per category. Clicking the Details icon on Alarms by Category opens Alarm Details where you can sort these categories and change the state of the alarms.

To create a similar report for a specific period of time, use the Alarm Summary report in Report mode.

RELATED DOCUMENTATION

[Alarm Detail Monitor | 771](#)

[Understanding the Fault Mode Tasks Pane | 765](#)

[Setting Up User and System Preferences | 26](#)

[Alarm Summary Report | 815](#)

[Network Director Documentation home page](#)

Alarms by Severity Monitor

Alarms by Severity is a pie-chart that shows the breakdown of all alarms since the last system restart. It is available on the main page when in Fault mode.

If you mouse over each segment, the total number of alerts for those alarms is shown. Alarm severity levels are:

- Critical (Red)—A critical condition exists; immediate action is necessary.
- Major (Orange)—A major error has occurred; escalate or notify as necessary.
- Minor (Yellow)—A minor error has occurred; notify or monitor the condition.
- Info (Wedgewood Blue)—An informational message; no action is necessary. Informational alarms do not necessarily indicate an error. It could indicate that a device or entity has changed state.

Clicking the Details icon on Alarms by Severity opens Alarm Details where you can sort and disposition individual.

RELATED DOCUMENTATION

[Alarm Detail Monitor | 771](#)

[Understanding the Fault Mode Tasks Pane | 765](#)

[Setting Up User and System Preferences | 26](#)

[Network Director Documentation home page](#)

Alarms by State Monitor

The Alarms by State monitor is a pie-chart representation of the states of an alarm: active and cleared. Use this graph to get an overall perspective of the amount of alarms that are active compare to those that are cleared. The Alarms by State monitor is on the main pane when in Fault mode.

Mouse over each segment of the pie-chart shows the number of alarms in these states:

- Active—Alarms that are current and not yet cleared.
- Cleared—Alarms that are resolved and the device or entity has returned to normal operation.

You can create an Alarms by State report for a specified node or a period of time using the Alarms Summary Report in Repot mode.

Changing the state of an alarm using Network Director is performed on the Alarm Detail page. Clicking the Details icon on Alarms by State opens Alarm Details where you can sort and set the disposition of the alarms.

RELATED DOCUMENTATION

[Alarm Detail Monitor | 771](#)

[Understanding the Fault Mode Tasks Pane | 765](#)

[Setting Up User and System Preferences | 26](#)

[Alarm Summary Report | 815](#)

[Network Director Documentation home page](#)

Alarm Trend Monitor

The Alarm Trend monitor provides trend information about alarms. The trend information is shown on a line chart, where each alarm severity is shown as a colored line. The legend for the line colors is displayed below the chart. The alarm count is shown on the vertical axis. The time of the data samples is shown on the horizontal axis. This monitor includes tabs that show alarm trend information for active alarms and for new alarms. You can select the time period to display from the list in the title bar.

RELATED DOCUMENTATION

[Understanding Fault Mode in Network Director | 761](#)

[Network Director Documentation home page](#)

7

PART

Working in Report Mode

[About Report Mode](#) | **782**

[Creating and Managing Reports](#) | **787**

[Report Reference](#) | **810**

About Report Mode

IN THIS CHAPTER

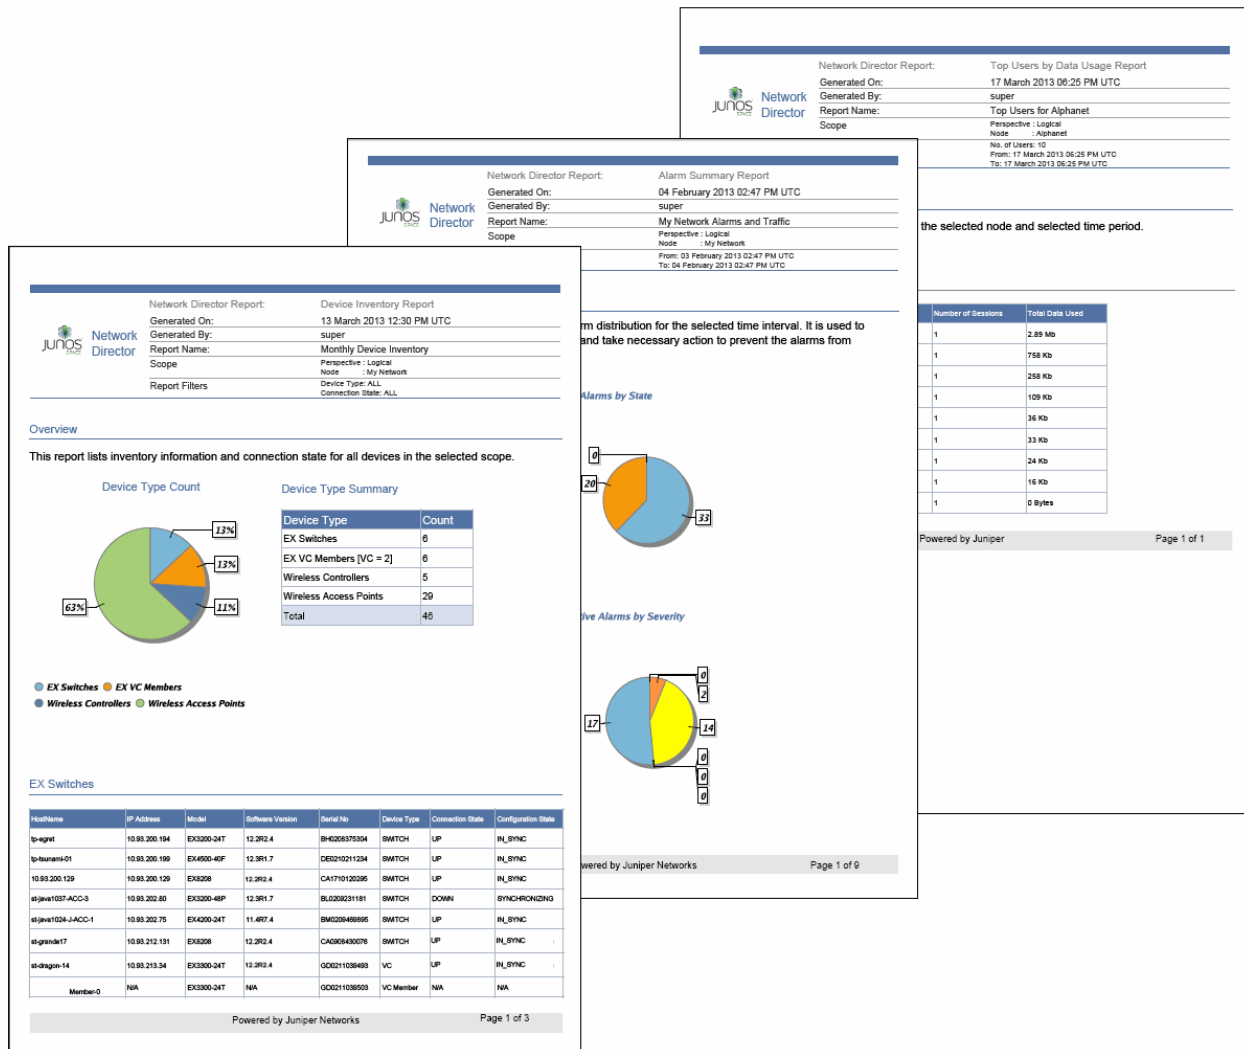
- [Understanding Report Mode in Network Director | 782](#)
- [Understanding the Report Mode Tasks Pane | 784](#)
- [Understanding the Types of Reports You Can Create | 786](#)

Understanding Report Mode in Network Director

In Report mode in Junos Space Network Director, you can create standardized reports from the monitoring and fault data collected by Network Director. An essential part of the network management lifecycle, reporting provides administrators and management insight into the network for maintenance, troubleshooting, trend and capacity analysis, and provides records that can be archived for compliance requirements.

Network Director provides reports in PDF and HTML formats that use graphs and tables to clearly convey data. Reports are also available in CSV format for importing into spreadsheets. [Figure 32 on page 783](#) shows some examples of PDF reports.

Figure 32: Examples of Network Director Reports



In addition to choosing the formats for your reports, you can:

- Run reports on-demand or schedule them to run at a specific time or on a recurring schedule.
- Select the portion of network you want the report to cover by selecting a scope in the View pane when you create a report definition. For example, you can run a Device Inventory report on your entire network, on all devices in a wiring closet, or on all EX2300 switches.
- Select the report options—for example, the historical time frame you want an Audit Trail report to cover or the type of devices you want to include in a Device Inventory report.
- Have reports sent to an e-mail address or automatically archived on a file server.

The process for generating reports is simple. Select a scope in the View pane and then create a report definition by using the Create Report Definition wizard. When you complete the report definition, the reports are immediately scheduled to run according to the scheduling choices you have made.

RELATED DOCUMENTATION

[Understanding the Types of Reports You Can Create | 786](#)

[Understanding the Report Mode Tasks Pane | 784](#)

[Network Director Documentation home page](#)

Understanding the Report Mode Tasks Pane

Network Director has built-in reporting features to create standardized reports from your network data. You can schedule these reports to run either in real time or in batch to gain insight into the network for ensuring compliance, performing maintenance, or troubleshooting.

The Report mode analyzes data from different perspectives and filters the data based on the node selected in the network tree.

From the Reports Tasks pane, you can:

- Set up a new report or change how an existing report is run by clicking Report Definition. From this page, you can launch a wizard that guide you through the process of defining a report or changing a report definition file. The report definition file is based on the report content on the view and the node you select in the network tree. The Filter option in the View pane does not affect the report content.
- View the summary details of the last run of a report, export a report, or to delete a report output by clicking Manage Generated Reports. This page is also the default Reports page. After a report definition is created and a report is generated from that definition, it is shown in the Generated Reports page.

Reports are stored on the application server on which Network Director is running. However, because reports can be large, the report is delivered in a compressed or *zipped* format. and can be stored offline or on a Secure Copy Protocol (SCP) server.

- Set or change the path to an SCP server for report storage. You can also test the connectivity to the server by clicking Test Connection on the Manage SCP Servers page.
- Set or change the path to an SMTP server for e-mail notifications of alerts or for mailing reports to administrators. You can test the connectivity to the server by clicking Test Connection on the Manage SMTP Servers page.
- Create or change a schedule that is used by one or more reports by clicking Manage Schedules. Unless you want to run the report immediately, you need to create a schedule and associate it with the report definition file. Create the schedule before you create the report definition file.

For example, you might want to run several reports that run on the weekend and are available first thing on Monday morning. You could create a single schedule that runs at midnight on Saturday and is delivered to you through e-mail.



- Add frequently performed tasks to Key tasks list. You can add any task from the Tasks pane to the Key Tasks list by selecting a task and clicking the plus (+) sign that appears adjacent to the task. For some modes, you can see that Network Director has predefined some key tasks for you. You can modify this set of tasks to suit your requirements. This feature is available in Task pane irrespective of your current mode, scope, or view.

RELATED DOCUMENTATION

Managing Reports in Network Director 787
Managing Generated Reports 801
Managing Reports on SCP Servers 804
Mailing Reports 806
Scheduling Reports 796
Understanding the Types of Reports You Can Create 786
Network Director Documentation home page

Understanding the Types of Reports You Can Create

The Report mode enables you to create standard reports from your network information. Reports are based on a report definition that can either be global or granular. You control this global or granular scope of the report definition by your selections in the View pane (the selected view and network tree node).

For example, if you want to run your reports against all switches, you could select the Logical view and the Switching Network node in the network tree. Or if you wanted to run reports on all the devices on a floor of a building, you would select the Location view and navigate to the floor node of a building in the network tree. To pinpoint the performance on a particular switch, you would select the Device view and the individual switch node in the network tree.

TIP: When naming your report definition, include the scope in the name. You cannot tell the scope from the report definition after you have created the definition; you can, however, determine the scope from the generated report.

The reports generated from the report definition file are either formatted and sent to you through e-mail or sent using Secure Copy Protocol (SCP) to a designated repository.

RELATED DOCUMENTATION

[Creating Reports | 790](#)

[Understanding Report Mode in Network Director | 782](#)

[Understanding the Report Mode Tasks Pane | 784](#)

[Network Director Documentation home page](#)

Creating and Managing Reports

IN THIS CHAPTER

- [Managing Reports in Network Director | 787](#)
- [Creating Reports | 790](#)
- [Scheduling Reports | 796](#)
- [Managing Generated Reports | 801](#)
- [Retaining Reports | 803](#)
- [Managing Reports on SCP Servers | 804](#)
- [Mailing Reports | 806](#)

Managing Reports in Network Director

IN THIS SECTION

- [How to Locate and Manage Reports | 788](#)
- [Managing Report Definitions | 788](#)

Reports are generated from a report definition. These definitions establish the type of report, when it is run, and how the report output is presented and preserved. You create, modify, or delete these report definitions from the Manage Report Definition page.

This topic describes:

How to Locate and Manage Reports

The Manage Report Definition page is available from the Report Tasks pane while the Report mode is selected. To locate this page:

1. Select **Report** in the Network Director banner. The Report Tasks pane opens displaying the tasks available in the Report mode.
2. Select **Manage Report Definitions** in the Tasks pane. The Manage Report Definition page opens displaying all existing report definition files.
3. Use the Manage Report Definition page to review existing report definitions, create new definitions, or change a definition.
 - Create a new report definition by clicking **Add**. See [“Creating Reports” on page 790](#) for help using the Report wizard.
 - Modify an existing report definition by selecting a report definition in the table and clicking **Edit**.
 - Delete an existing report definition by selecting a report type in the table and clicking **Delete**.
 - View details of the report composition, the scope, and perspective of the report definition by clicking **Details**.

Managing Report Definitions

Use the Manage Report Definition page to review existing report definitions, or follow the Report wizard to create new report definitions, delete definitions, or see report details.

Existing report definitions are listed on the page in the format discussed in [Table 222 on page 788](#). The reports created from these definitions are found under the Manage Generated Reports task.

Table 222: Manage Report Definition Fields

Field	Description
Report Definition	The name of the report definition. Specify a name that indicates the purpose of the report.

Table 222: Manage Report Definition Fields (*continued*)

Field	Description
Format	<p>The format or file extension of the report output; the final rendering of the output. Valid values are:</p> <ul style="list-style-type: none"> • PDF—(Portable Definition Format) is used for output that is either viewed in a reader or printed. • CSV—(Comma Separated Format) is used for output that is exported into a spreadsheet. • HTML—(Hypertext Markup Language) is used for output that is viewed in a Web browser.
Reporting Mode	<p>(Optional) Where the generated report is sent. Valid values are:</p> <ul style="list-style-type: none"> • Email—Sends a zipped file of the report to an e-mail address. • SCP—Sends a zipped file to a secure server.
Schedule	(Optional) When the report is scheduled to run.
Last Updated By	The userid of the last person to modify the report definition.
Last Updated Time	Time when the report definition was last updated.
Execute Report	Click Run Now to run the report.

RELATED DOCUMENTATION

[Creating Reports | 790](#)
[Managing Generated Reports | 801](#)
[Understanding the Types of Reports You Can Create | 786](#)
[Understanding the Report Mode Tasks Pane | 784](#)
[Retaining Reports | 803](#)
[Network Director Documentation home page](#)

Creating Reports

IN THIS SECTION

- [How to Create a Report Definition | 790](#)
- [Creating a Report Definition | 792](#)
- [Setting Report Options | 794](#)
- [Reviewing the Report Definition | 795](#)
- [Changing a Report Definition | 796](#)

Network Director has built-in reporting features to create standardized reports from your network data. You can schedule these reports either to run in real time or in batch to provide insight into the network for compliance, maintenance, or troubleshooting. To define a new report, you select from a number of preconfigured report types and set the scheduling and output options.

This topic describes:

How to Create a Report Definition

You create new reports from the Report Definition page while in the Report mode. To locate this page:

1. Select **Report** in the Network Director banner. The Report Tasks pane opens, displaying the tasks available in the Report mode.
2. (Optional) Select the node on which to run the report in the View pane. Some reports are designed to run at a specific scope in the network tree. For example, if you select an EX Series switch node and attempt to run a Network Neighborhood report that reports on RF strength, the report runs, but is empty.
3. Select **Manage Report Definitions** in the Tasks pane.

[Table 223 on page 791](#) describes the information provided about report definitions on the Manage Report Definition page.

Table 223: Manage Report Definition Fields

Field	Description
Report Definition	The name of the report definition. Specify a name that indicates the purpose of the report.
Format	<p>The format or file extension of the report output; the final rendering of the output. Valid values are:</p> <ul style="list-style-type: none"> • PDF—(Portable Definition Format) is used for output that is either viewed in a reader or printed. • CSV—(Comma Separated Format) is used for output that is exported into a spreadsheet. • HTML—(Hypertext Markup Language) is used for output that is viewed in a Web browser.
Reporting Mode	<p>(Optional) Where the generated report is sent. Valid values are:</p> <ul style="list-style-type: none"> • Email—Sends a zipped file of the report to an e-mail address. • SCP—Sends a zipped file to a secure server.
Schedule	(Optional) When the report is scheduled to run.
Last Updated By	The user ID of the last person to modify the report definition.
Last Updated Time	Time when the report definition was last updated.
Execute Report	Click Run Now to run the report.

Creating a Report Definition

A report definition defines the properties that are used to generate one or more reports. It includes these properties:

- Name of the report definition
- Report type(s)
- Reporting filters
- Scheduling options
- Output format

To create a report definition:

1. Click **Add** on the Manage Report Definition main page to open the Create Report Definition wizard Basic Settings page.
2. Type a name for the report in the Report Definition Name field. After the report runs, you can find a report by this name in the Generated Reports list. Names can contain letters, numbers, spaces, dashes (-), and underscores (_).
3. Select the report types for the report definition in the Select Report Type area:
 - To add one or more report types:
 - a. Click **Add**. The Assign Report Types window opens.
 - b. Select one or more report types from the list in the Assign Report Types window.
 - c. Click **OK**.

The report types you added appear in the Select Report Type list. [Table 224 on page 793](#) describes the information about report types that is available in the Select Report Type table.

TIP: When adding multiple reports types, be sure all of the reports you select are supported for the node type selected in the view pane.

- To delete one or more report types, select their check boxes in the Select Report Type list, then click **Delete**.

4. (Optional) Edit the report type options for the added report types by clicking **Edit Report Options** in the Customize Report Options column. Configure report type options in the Filter Options window, then click **OK**. [Table 225 on page 793](#) describes the available report type options.
5. Click **Next** or **Report Options** to set up the report options. You can also click **Cancel** to exit the wizard. For details on report options, see [“Setting Report Options” on page 794](#).

Table 224: Select Report Type Table Columns

Column Heading	Description
Type	The Report name.
Category	The general classification of the report.
Scope	Shows the scopes that are applicable for the report type. (Appears only in the Assign Report Types window that opens when you click the Add button.)
Description	A description of the use or purpose of the report.
Report Option	Lists the applied report type options.
Customize Report Options	Click the link to change the report type options for that report type.

Table 225: Report Type Options for Data Filtration

Filter Option	Description
Classification Reason	For the Rogue Summary report, filters the rogue devices included in the report to only those that are classified as rogue for the selected reason.
Connection State	Limits the report to devices in this state.
Device Types	Limits the report to this type of device.
Number of Users	Customizes the report to the specified number of users.
Percentage Utilization Exceeding	Specifies the utilization percentage threshold for the report. Only results that exceed the threshold will appear in the report.
Search Parameter	Specifies search parameters. Only results that match the parameters will appear in the report. The search parameters are compared to these properties of results to filter the results that appear in the report: IP address, MAC address, username. Separate multiple search parameters with commas (,).

Table 225: Report Type Options for Data Filtration (*continued*)

Filter Option	Description
Time Interval	Limits the report to the indicated time period. If you select Custom, the From and To fields become available, enabling you to set a specific reporting period.
Top N Count	Sets the number of items to include in reports that show a fixed number of items. For example, the Traffic and Congestion Summary report includes the top <i>N</i> number of devices that have the highest port utilization and latency. If the scope is a single device, the top <i>N</i> number of ports on the device are included in the report.

Setting Report Options

This page establishes the report schedule and the output format of the report.

1. Choose from the following scheduling options:

- Run the report now
- Select or create a schedule for the report
- Select to both run the report now and to run the report by a schedule

Options for report scheduling are shown in [Table 226 on page 794](#).

Table 226: Schedule Options for Reports

Field	Action
Run Now	Select this option to immediately run the report one time.
Select Schedule	Select this option to either create a schedule so that it is run at regular intervals, or to select an already established schedule. <ul style="list-style-type: none"> • The Add Schedule link enables you to create a new schedule. • The Select button opens Choose Schedule window that displays the currently configured schedules. Select the check box to choose a schedule to use for the report. To associate the schedule to your report, click OK.

2. Establish the report output format and destination.

Field	Options
Select Format	<p>A report is available in these formats:</p> <ul style="list-style-type: none"> • PDF—Choose this format If you want to print the report. Portable Definition Format (PDF) enables the report to be printed from any operating system with the same formatting results. • CSV—Choose this format if you want to export the report data to a spreadsheet or other business application. The Comma-Separated Values (CSV) format takes the raw data from the report and delineates the fields with commas so that it imports into popular spreadsheet programs. • HTML—Choose this format if you want to view the report in a browser. <p>NOTE: Because reports can be quite large, they are initially delivered as a zipped (compressed) file.</p>
Mode	<p>Reports can be sent to your e-mail address, to a secure server, or to both.</p> <ul style="list-style-type: none"> • Select EMAIL and type the e-mail address to have the report sent through e-mail. Network Director uses SMTP server settings for e-mail routing. You can configure an SMTP server from the Tasks pane. • Select SCP to send the report to the secure server that is marked as active, using Secure Copy Protocol. The settings for secure servers are available in Tasks > Manage SCP Servers.

3. Click **Next** or **Summary** to review the report definition.

Reviewing the Report Definition

The Report wizard guides you to the Summary Page where you can review your report configuration and make any changes before you run the report.

1. Review your Report Name and Report Type in basic settings. If you want to change either of these settings, click **Edit** to return to the Basic Settings page.
2. Review your Report Options. If you want to change these settings, click **Edit** to return to the Report Options page.
3. Click **Finish** when you are done with the report configuration and to exit the wizard.

Changing a Report Definition

You can change an existing report definition file from the Manage Report Definition page.

To change a report definition:

1. Select the check box for the report definition.
2. Click **Edit** to reopen the report definition in the Report wizard. The system returns you to the Summary page, where you can make changes to the report definition.
3. Click **Details** to review the details of the report definition or click **Delete** to remove the report definition. To remove all of the report definitions, select the check box in the header next to Report Definition to select all of the report definitions and click **Delete**.

RELATED DOCUMENTATION

[Managing Generated Reports | 801](#)

[Understanding the Types of Reports You Can Create | 786](#)

[Managing Reports on SCP Servers | 804](#)

[Mailing Reports | 806](#)

[Scheduling Reports | 796](#)

[Understanding Report Mode in Network Director | 782](#)

[Network Director Documentation home page](#)

Scheduling Reports

IN THIS SECTION

- [How to Create or Manage Schedules | 797](#)
- [Managing Schedules | 797](#)
- [Creating New Schedules | 798](#)
- [Editing Schedules | 800](#)
- [Deleting Schedules | 800](#)

You can run Network Director reports as needed or you can automate reports to run in batch by creating a schedule. You can associate a single schedule with one or more reports when you create the report definition. Although you can create a schedule during the report definition process, it is helpful to have the schedules configured before defining the report. To create a new schedule you name the schedule and set the time and frequency of the run.

This topic describes:

How to Create or Manage Schedules

You create a schedule from the Manage Schedules page. You can display this page from the Report mode Tasks pane.

- Select **Report** in the Network Director banner. The Report mode Tasks pane displays the tasks available in the Report mode. Reports run in any network view (Logical, Location, or Device).
- Select **Manage Schedules** in the Tasks pane. The Manage Schedules page opens, displaying all existing report schedules.

From the Manage Schedules page, you can:

- Create a new schedule
- Edit an existing schedule
- See the details of a schedule
- Delete a schedule

Managing Schedules

Use the Manage Schedules page to administer report schedules. From this page, you can create new schedules and view, modify, and delete existing schedules. On the Manage Schedules page, a table of existing report schedules appears. You can sort and customize this table to exclude fields that might not be relevant to your needs. The fields in the table are defined in [Table 227 on page 797](#).

Table 227: Manage Report Schedules

Field	Description
Schedule Name	The name of the schedule. Indicate the purpose of the schedule in the name.
Schedule Type	<p>Schedules are either One-Time or Recurring.</p> <ul style="list-style-type: none"> • One-Time—These schedules are helpful for running a non-repeating report in batch-mode, such as running it at 3:30 a.m. • Recurring—These schedules are helpful for routine reports and trend analysis.

Table 227: Manage Report Schedules (*continued*)

Field	Description
Recurrence Pattern	How often the pattern repeats. The pattern applies only to recurring schedules. Valid values are: <ul style="list-style-type: none"> • Hourly • Daily • Weekly • Monthly
Description	Details of the reoccurrence pattern.
Status	Either Active or blank. Active indicates the schedule is running.

From this page you can:

- Display a summary of all the parameter settings of a schedule by selecting the schedule and clicking Details. The Report Schedule Summary opens.
- Create a new schedule by clicking Add. The Create Schedule window opens.
- Select a schedule and change the settings by clicking Edit. The Edit Schedule window opens.
- Select a schedule and click **Delete** to remove a schedule.



CAUTION: Take care when deleting a schedule. Network Director enables you to delete a schedule even if it is active.

Creating New Schedules

Use the Create Schedule window to create a one-time or a recurring schedule.

To create a new schedule:

1. Click **Add** on the Manage Schedules page. The Create Schedule window opens.
2. Enter the name for the schedule in **Schedule Name**. Indicate the purpose of the schedule in the name.
3. Choose a one-time run-option or a recurring schedule from the list.
 - One-time schedule options are described in [Table 228 on page 799](#).
 - Recurring schedule options are described in [Table 229 on page 799](#).

Depending on the schedule range selected, these settings change dynamically.

Table 228: One-Time Schedule Options

Field	Action
Execute Start Date	Select the date when the schedule is run. You can either fill in a date directly or click the calendar icon to pick a date from a traditional calendar.
Execute Start Time	Select the time the report is run. Time is shown in 24-hour clock format in increments of 15 minutes.

Table 229: Recurring Schedule Options

Field	Action
Hourly	Run the report associated with the schedule, hourly between these hours at intervals of x minutes. Start and end times are shown in 24-hour clock format, in increments of 15 minutes. For example, if you want to schedule a report to run from 1 am to 3 pm at 30 minute intervals, your settings would be: between 13:00 and 15:00 at 30 min(s).
Daily	Run the report associated with the schedule either every weekday or on the specified number of sequential days. Use the up and down arrows to set the number of sequential days or click Every weekday .
Weekly	Run the report associated with the schedule on one or more days of the week. Set the schedule to repeat the run in the specified number of weeks. Use the up and down arrows to set the weekly frequency of how often the schedule is repeated. Click one or more of the days when the report is run.
Monthly	Run the report associated with the schedule either on a certain day of the month or on a specified day and week for the specified number of months. For example, if your organization has a congestion spike at the end of the fiscal quarter, you might want to run reports on the last Friday every 4 months.

- Specify when to start and end implementation of this schedule as described in [Table 230 on page 799](#).

Table 230: Range of Recurrence Fields

Field	Action
Start Time:	Select the time of day. The clock is in 24-hour format, in increments of 15 minutes, when the schedule begins to run.

Table 230: Range of Recurrence Fields (*continued*)

Field	Action
Start Date:	Select the date when the schedule is first implemented. Format is yyyy-mm-dd.
No end date	Select to indicate to continue to use this schedule until it is modified or deleted.
End After: x occurrence	Select to run the schedule for a specified number of times. Use the up and down arrow keys to specify the number of times to run the schedule.
End by:	Select a time and date to stop running the schedule. Date format is yyyy-mm-dd. Select the time from the list; clock times are in increments of 15 minutes.

5. Click **Add** to finish and to validate the schedule.

Editing Schedules

To change an existing schedule:

1. Select a schedule from the list in the Manage Schedules page.
2. Click **Edit** to reopen the schedule settings.
3. Change the settings based on the values described in [Table 228 on page 799](#).
4. Click **Edit** to save the revised settings.

Deleting Schedules



CAUTION: Network Director enables you to delete an active schedule. Be aware that deleting a schedule that is active will cause reports not to run.

You can also permanently remove a schedule by selecting the schedule in the Manage Schedule page and clicking Delete.

RELATED DOCUMENTATION

[Creating Reports](#) | 790

Managing Generated Reports

IN THIS SECTION

- [Reviewing Generated Reports | 801](#)
- [Viewing Report Details | 802](#)
- [Exporting Reports | 803](#)
- [Deleting Generated Reports | 803](#)

After a report definition is created and the initial report is run, Network Director populates the Generated Reports page with summary information about the run of the report.

From the Generated Reports page you can view the report details, export the report to view or store in a new location, or delete the report.

This topic describes:

Reviewing Generated Reports

After a reports runs, information about the report is recorded on the Generated Reports page, as shown in [Table 231 on page 801](#).

Table 231: Fields in the Generated Reports Page

Field	Description
Report Definition	The name assigned at report creation time.
Executed By	The userid of the report owner who ran the report.
Start Time	When the report began to run.
End Time	When the report ended.
Format	The chosen report format, possible values are PDF, CSV, or HTML.

Table 231: Fields in the Generated Reports Page (*continued*)

Field	Description
Generated Report	Links to view or download the report.

From the Generated Reports page, you can either select the check box in the heading to select all of the reports or select the check box for the individual reports and:

- View details about the running of the report by clicking Report Details.
- Export the report by clicking Export.
- Delete the report by clicking Delete.

Viewing Report Details

Use the Generated Report Details window to see information about the report composition. Viewing the Report Details is helpful when a report comprises many smaller reports. See [Table 232 on page 802](#) for these field descriptions.

Table 232: Generated Report Details

Field	Description
Report Definition Name	The name assigned at report creation time.
Executed By	The userid of the report owner who ran the report.
Start Time	When the report began to run.
End Time	When the report ended.
Format	The chosen report format, possible values are PDF, CSV, or HTML.
Generated Report	This link opens the Report Details window where you can view or download the report.

Reports are kept on the application server until either you delete them or they are deleted by the system. The amount of time a report is saved on the system depends on the report retention settings for Network Director. Network Administrators can globally set the report retention period for reports in the system Preferences, located in the Network Director banner.

Exporting Reports

Use the Export Report window to store multiple reports to a file location. Because reports can be large, they are delivered as compressed *zipped* files for both viewing or storing. After you choose a report to export, you are prompted to select a program to view the report or to download the file.

You can either unzip the report for viewing or save the report to a file location.

TIP: If you choose to save the file, you might want to give a unique name to the unzipped file. After unzipping the report, the name of the report reverts to the type of report you selected. It does not retain the name of the report.

Deleting Generated Reports

Use Delete to discard unneeded or outdated reports. When you click **Delete**, you are prompted to confirm the file deletion. Because deleted reports cannot be recovered, save the report offline before deleting them from the system.

RELATED DOCUMENTATION

[Creating Reports | 790](#)

[Managing Reports in Network Director | 787](#)

[Retaining Reports | 803](#)

[Network Director Documentation home page](#)

Retaining Reports

Reports are stored on the Junos Space server where Network Director is running; however, you set the report retention time in Network Director. The default for the report retention period is 30 days. Because accumulating old reports could eventually impact system performance, you might want to consider changing this setting. The report retention period is set in Preferences on the Report tab.

Another option is to store the report to another location such as a Secure Copy Protocol (SCP) server and then delete the report from the Network Director application server.

RELATED DOCUMENTATION

[Setting Up User and System Preferences | 26](#)

[Managing Generated Reports | 801](#)

[Managing Reports on SCP Servers | 804](#)

[Network Director Documentation home page](#)

Managing Reports on SCP Servers

IN THIS SECTION

- [How to Configure SCP Servers | 804](#)
- [Managing SCP Servers | 805](#)

If your organization requires reports be stored on a secure server using Secure Copy Protocol (SCP), you can set one or more servers as a reports repository. A reports repository enables you to keep reports long-term for compliance requirements or for your organizational needs.

This topic describes:

How to Configure SCP Servers

SCP servers are used as report repositories for Network Director reports. You can set up or manage a secure server for Network Director reports from the Manage SCP Servers page in the Report mode. To access this page:

1. Select **Report** in the Network Director banner. The Report Tasks pane opens, displaying the tasks available in the Report mode.
2. Select **Manage SCP Servers** in the Report Tasks pane. The Manage SCP Servers page opens, displaying all existing report schedules.

Use the Manage SCP Servers page to:

- View existing SCP server settings
- Set up new SCP servers for reports
- Edit SCP server settings

- Test the connection to an SCP server
- Make an SCP server active
- Delete an SCP server setting

Managing SCP Servers

Use the Manage SCP Servers page to view and manage SCP server settings.

- The Manage SCP Servers page lists any existing server settings. The fields in the Manage SCP Server page are described in [Table 233 on page 805](#).

Table 233: Managing SCP Server Fields

Field	Description
Server Name	The name you are using to identify the SCP server.
IP Address	The IP address or hostname of the SCP server.
Port Number	The forwarding port number. Default port number for SCP is 22.
Active	Either yes or no to indicate whether it is the active server. Only one server can be active at a time.
Base Path	The path on the server where the reports are to be stored.

- Create new or edit existing server settings:
 1. Establish a new server definition by clicking **Add** or edit an existing definition by clicking **Edit**. Either an Add SCP Settings or Edit SCP Settings page opens.
 2. Fill in the settings described in [Table 234 on page 805](#).

Table 234: Defining an SCP Server

Field	Action
Server Name	Type a name for this SCP server.
IP Address/Host Name	Type the IP address of SCP server.
Port Number	Type the forwarding port number. Default port number for SCP is 22.
User Name	Type the account name accessing the server.

Table 234: Defining an SCP Server (*continued*)

Field	Action
Password	Type the password twice for the account on the secure server.
Default Path	Type the file path to the server where the reports are to be stored.
Set Active	Select if you want this server to be the active server. While many servers can be set up as SCP servers for reports, only one server is marked as active.

3. Click **Done** to complete the process.
4. Click **Test Connection** to ensure your server is set up correctly. Network Director attempts to connect to the SCP server and tells you whether the connection could be established.
5. Select a server and click **Set Active** to make that server available for secure services.

You can also delete any SCP server definition from use by Network Director reports by clicking **Delete**.

RELATED DOCUMENTATION

[Understanding the Report Mode Tasks Pane | 784](#)

[Managing Generated Reports | 801](#)

[Network Director Documentation home page](#)

Mailing Reports

IN THIS SECTION

- [How to Configure SMTP Servers | 807](#)
- [Managing SMTP Servers | 807](#)
- [Adding or Editing SMTP Server Settings | 809](#)

You can set up one or more electronic mail servers to send reports to e-mail addresses. These servers use the Simple Mail Transfer Protocol (SMTP) to forward the reports. While you can configure many servers as SMTP servers, you can only designate one as the primary mail server.

This topic describes:

How to Configure SMTP Servers

An SMTP server is responsible for sending e-mails. Network Director uses the SMTP server to send reports to users. Under most circumstances, you need only one SMTP server. However, you might want to configure more than one SMTP server if you need a server with a distinct SMTP server configuration. In this case, you would configure multiple SMTP servers and mark the server you want to use as Active.

You can set up or manage SMTP servers from the Manage SMTP Servers page in the Report mode. To access this page:

1. Select **Report** in the Network Director banner. The Report Tasks pane opens, displaying the tasks available in the Report mode. The Generated Reports page loads in the main window.
2. Select **Manage SMTP Servers** in the Report Tasks pane. The Manage SMTP Servers page opens in the main window, displaying all existing SMTP servers configured for Network Director.

Use the Manage SMTP Servers page to:

- View existing SMTP server settings
- Set up new SMTP servers
- Edit existing SMTP server settings
- Test the connection to a SMTP server
- Set an SMTP server as the active server
- Delete an SMTP configuration
- See details of the SMTP configuration

Managing SMTP Servers

Use the Manage SMTP Servers page to view and manage SMTP server settings. The Manage SMTP Servers page lists any existing SMTP server settings. The fields in the Manage SMTP Server page are described in [Table 235 on page 808](#).

Table 235: Managing SMTP Server Fields

Field	Description	Hidden or Displayed by Default
Name	The name you are using to identify the SMTP server.	Displayed
Host Address	The IP address or hostname of the SMTP server.	Displayed
Port	The forwarding port number. Default port number for SMTP is 587.	Displayed
Active	Either yes or no to indicate whether it is the active server. Only one server can be active at a time.	Displayed
User Auth	Indicates whether SMTP authentication is required for the server. This field is either yes or no.	Displayed
Use TLS	Indicates whether Transport Layer Security (TLS) protocol is used to provide shared-secret encryption.	Displayed
User Name	Indicates the username when user credentials are required for SMTP authentication.	Hidden
From E-mail Address	The e-mail account that sends the report.	Hidden

1. Establish a new server definition by clicking **Add** or edit an existing definition by selecting the server and clicking **Edit**. Either an Add SMTP Settings or Edit SMTP Settings page opens. See [“Adding or Editing SMTP Server Settings” on page 809](#) for details on setting up or changing server settings.
2. Click **Done** to complete the process.
3. Click **Test Connection** to ensure your server is set up correctly. Network Director tells you whether the attempted connection to the SMTP server could be established.
4. Select a server and click **Set Active** to make that server responsible for sending e-mail.

You can also delete any SMTP server definition from use by Network Director reports by clicking **Delete**.

Adding or Editing SMTP Server Settings

The process of establishing a new SMTP server or to changing the values on an existing server is straightforward. Simply enter or change the values in the fields in the Add SMTP Server or Edit SMTP Server page. These fields are described in [Table 236 on page 809](#).

Table 236: Defining an SMTP Server

Field	Action
Server Name	Type a name for this SMTP server.
Host Address	Type the IP address of SMTP server.
Port Number	Type the forwarding port number. Default port number for SMTP is 587.
From Email Address	Type the e-mail address used to send the notification.
Set as Active Server	Checking this box sets the server as the Active server. If there is only one server, you cannot clear this box.
Use SMTP Authentication	Checking this box requires the server to use SMTP authentication. You must provide user credentials to use SMTP Authentication.
User Name	Type the account name accessing the server for SMTP authentication.
Password	Type the password twice that is used for authentication.
Use TLS	Select if you want this server to use TLS protocol on the SMTP server.

RELATED DOCUMENTATION

[Understanding the Report Mode Tasks Pane | 784](#)

[Network Director Documentation home page](#)

Report Reference

IN THIS CHAPTER

- Active User Sessions Report | 810
- Alarm History Report | 812
- Alarm Summary Report | 815
- Audit Trail Report | 817
- Client Devices Report | 819
- Device Inventory Report | 820
- Fabric Analyzer Report | 822
- Network Device Traffic Report | 824
- Port Bandwidth Utilization Report | 826
- Top Users by Data Usage Report | 827
- Traffic and Congestion Summary Report | 828

Active User Sessions Report

The Active User Sessions report is a standardized report generated in Network Director to show the activity level of current users on a specified node. There are two portions of the report: a report header and the report body. The contents of the report header are found in [Table 237 on page 810](#).

Table 237: Active User Session Report Header

Field	Description
NETWORK DIRECTOR REPORT:	The type of report; In this case, the Active User Sessions report.
Generated On:	The date and time the report ran. The date format is: [Day of the Month] [Month] [Year]. The time format follows the Coordinated Universal Time (UTC). For example: 07 December 2013 11:35 PM PST.
Generated By:	The username of the user that generated the report.

Table 237: Active User Session Report Header (*continued*)

Field	Description
Report Name	The name of the report assigned when the report was created.
Scope	<p>The report is generated from this view and node in the network.</p> <ul style="list-style-type: none"> • <i>Perspective</i>—Can be a Logical view, Location view, or a Device view of the network. • <i>Node</i>—Represents the selected object on which the report is based.

The base report is a table with the fields described in [Table 238 on page 811](#).

Table 238: Active User Session Report Fields

Field	Description
User Name	The name that identifies the user to the network.
Client MAC Address	The MAC address of the user.
Total Bytes	The total number of bytes for the session. Bytes are shown in system international (SI) notation. For example, terabytes (TB), gigabytes (GB), megabytes (MB), and Kilobytes (KB).
Client IP	The IP address of the client.
Auth Type	The authentication type.
Bandwidth (KBps)	The transfer speed in Kilobytes per second.
VLAN	The VLAN name.
Session Elapsed Time	The amount of time after the user started the session.

RELATED DOCUMENTATION

[Understanding the Types of Reports You Can Create | 786](#)

[Creating Reports | 790](#)

[Managing Generated Reports | 801](#)

[Managing Reports on SCP Servers | 804](#)

[Network Director Documentation home page](#)

Alarm History Report

IN THIS SECTION

- [Alarm History Header | 812](#)
- [Alarm History Tables | 813](#)

The Alarm History report is a standardized report generated in Network Director. It shows all active, acknowledged, and cleared alarms that occurred within a specified period of time for the indicated node. The report has two portions: a report header and the report body.

This topic describes:

Alarm History Header

The Alarm History report header provides file creation information about the report. The contents of the report header are described in [Table 239 on page 812](#).

Table 239: Alarm History Report Header

Field	Description
NETWORK DIRECTOR REPORT	The type of report—in this case, the Alarm History report.
Generated On	The date and time the report ran. The date format is: [Day of the Month] [Month] [Year]. The time format follows the Coordinated Universal Time (UTC). For example: 07 December 2012 11:35 PM PST.
Generated By	The username of the user that generated the report.
Report Name	The name of the report assigned when the report was created.
Scope	The report is generated from this view and node in the network. Perspective can be Logical view, Location view, or Device view of the network. Node represents the selected object upon which the report is based.
Report Filters	The period of time specified for data collection.

Alarm History Tables

The body of the Alarm History report is a set of tables: one table for each alarm state (active, acknowledged, and cleared). In each table the alarms are listed by severity level. The fields of the tables have a common format, which is described in [Table 240 on page 813](#).

Table 240: Active Alarm History Fields

Field	Description
Alarm Name	The SNMP alarm name.
Severity	One of six levels that indicate the gravity of the alarm based on the impact on the system: <ul style="list-style-type: none">• Critical• Major• Minor• Warning (customer defined)• Alert (customer defined)• Notice• Info

Table 240: Active Alarm History Fields (*continued*)

Field	Description
Category	<p>One of twenty-four functional areas:</p> <ul style="list-style-type: none"> • BFD • BGP • Chassis • ClientAndUserSession • Configuration • CoS • DHCP • DOM • FlowCollection • General • GenericEvent • L2ALD • L2CP • MACFDB • Misc. • PassiveMonitoring • Ping • RFDetect • RMon • SONET • SonetAPS • VirtualChassis
Description	An indication of what caused the alarm.
Source	The Entity ID of the network device sending the trap.
Acknowledged	Whether or not the alarm has been acknowledged.
Updated On	The date when the alarm was last updated (assigned, annotated, or acknowledged) otherwise, the date the alarm was created.

RELATED DOCUMENTATION

[Understanding the Types of Reports You Can Create | 786](#)

[Creating Reports | 790](#)

[Managing Generated Reports | 801](#)

[Managing Reports on SCP Servers | 804](#)

[Network Director Documentation home page](#)

Alarm Summary Report

IN THIS SECTION

- [Alarm Summary Header | 815](#)
- [Alarm Summary Charts | 816](#)

The Alarm Summary Report is a standardized report generated in Network Director. It shows a graphical summary of the alarms that occurred within a specified period of time, node, and network view. The report has two portions: a report header and a report body.

This topic describes:

Alarm Summary Header

The report header provides file creation information about the report. The contents of the report header are described in [Table 241 on page 815](#).

Table 241: Alarm Summary Report Header

Field	Description
NETWORK DIRECTOR REPORT	The type of report—in this case, the Alarm Summary report.
Generated On	The date and time the report ran. The date format is: [Day of the Month] [Month] [Year]. The time format follows the Coordinated Universal Time (UTC). For example: 07 December 2012 11:35 PM PST.
Generated By	The username of the user that generated the report.
Report Name	The name of the report assigned when the report was created.

Table 241: Alarm Summary Report Header (*continued*)

Field	Description
Scope	The report is generated from this view and node in the network. Perspective can be Logical view, Location view, or Device view of the network. Node represents the selected object upon which the report is based.
Report Filters	The period of time specified for data collection.

Alarm Summary Charts

The body of the Alarm Summary report contains a series of colored charts that provide insight into the trends or distribution of alarms in the network. The first chart summarizes the proportion of active and clear alarms. The rest of the charts are divided into two identical sets: one set for active alarms and one set for clear alarms. If there are no alarms in an active or in a clear state, the set of charts for that state are omitted.

The charts are:

- **Alarms by State**—Shows the proportion of alarms in active and clear states.
- **Active Alarms by Severity (Clear Alarms by Severity)**—Shows the proportion of alarms in each severity classification. The default severity levels are:

Critical (Red)— A critical condition exists; immediate action is necessary.

Major (Orange)— A major error has occurred; escalate or notify as necessary.

Minor (Yellow)— A minor error has occurred; notify or monitor the condition.

Info (Light Blue) —An informational message; no action is necessary.

In Preferences, administrators can override the default severity of an alarm and set the severity to match their inhouse guidelines.

- **Active Alarm Summary Chart (Clear Alarm Summary Chart)**—Shows alarms by category and severity.
- **Active Alarms by Category (Clear Alarms by Category)**—Shows the number of alarms in each alarm category and, within category, the number of alarms that are acknowledged or unacknowledged.
- **Active Alarms by Type (Clear Alarms by Type)**—Shows the number of alarms of each specific type. The types are color-coded by severity.
- **Top 10 Sources of Active Alarms (Top 10 Sources of Clear Alarms)**—Identifies the top 10 devices that are generating the most alarms. Shows the number of alarms each device is generating by severity.
- **Active Alarms by Timestamp (Clear Alarms by Timestamp)**—This section of the report contains two graphs that plot alarms by their created timestamp:

- Alarms by timestamp per severity—Plots each alarm of given severity against the time the alarm was created. For example, if during the time period covered by the report there are 10 alarms of major severity, the graph shows 10 orange data points that are plotted against the time the alarms were created.
- Active alarms by timestamp for top 10 sources (Clear alarms by timestamp for top 10 sources)—For each source, plots the alarms generated by the source against the time they were created.

RELATED DOCUMENTATION

[Understanding the Types of Reports You Can Create | 786](#)

[Creating Reports | 790](#)

[Managing Generated Reports | 801](#)

[Managing Reports on SCP Servers | 804](#)

[Network Director Documentation home page](#)

Audit Trail Report

The Audit Trail report is a standardized report generated in Network Director. It shows a history of users accessing the system, modifications to applications, and network management activities for a specific period of time. The report is defined and generated from the Report mode in Network Director.

There are two portions of the report: a report header and the report body. The contents of the report header are shown in [Table 242 on page 817](#).

Table 242: Audit Trail Report Header Information

Field	Description
NETWORK DIRECTOR REPORT:	The type of report; In this case, the Audit Trail report.
Generated On:	The date and time the report ran. The date format is: [Day of the Month] [Month] [Year]. The time format follows the Coordinated Universal Time (UTC). For example: 07 December 2012 11:35 PM PST.
Generated By:	The username of the user that generated the report.
Report Name	The name of the report assigned when the report was created.

Table 242: Audit Trail Report Header Information (continued)

Field	Description
Scope	The report is generated from this view and node in the network. Perspective can be Logical view, Location view, or Device view of the network. Node represents the selected object upon which the report is based.
Report Filters	The period of time specified for data collection.

The body of the report comprises the Audit Logs by Task Type chart, the Audit Logs by User chart, and a Audit Detail table.

The Audit Logs by Task Type is a bar chart that lists all of the user and system activities over the specified time period for all users.

The Audit Logs by Users is a pie chart that graphically represents all the active users in a specified time period.

The fields in the report body table are shown in [Table 243 on page 818](#).

Table 243: Audit Trail Report Fields

Field	Description
User Name	The userid of the individual associated with the activity.
User IP	The IP address of the client.
Task	A short summary of the activity, such as Backup or Login.
Description	The description of the system activity being logged. Examples of common logging activities include logging in and out of the system, modifying application settings, creating SMTP or SCP servers, or database backups.
Result	The result of the system activity: whether it is successful or not. Regularly scheduled events, such as backups, show as recurring.
Job ID	The system generated identification for applicable tasks.
Timestamp	The date and time of the activity. The date is shown in the format: [Day of the Month] [Month] [Year], while the time is shown in standard 12-hour clock format.

RELATED DOCUMENTATION

[Understanding the Types of Reports You Can Create | 786](#)

[Creating Reports | 790](#)

[Managing Generated Reports | 801](#)

[Managing Reports on SCP Servers | 804](#)

[Viewing Audit Logs From Network Director | 22](#)

[Network Director Documentation home page](#)

Client Devices Report

The Client Devices report is a standardized report generated in Network Director to show the distribution of user sessions by device types, device groups, and device profiles. There are two portions of the report: a report header and the report body. The contents of the report header are described in [Table 244 on page 819](#).

Table 244: Client Devices Report Header

Field	Description
NETWORK DIRECTOR REPORT:	The type of report; in this case, the Client Devices report.
Generated On:	The date and time the report ran. The date format is: [Day of the Month] [Month] [Year]. The time format follows the Coordinated Universal Time (UTC). For example: 07 December 2013 11:35 PM PST.
Generated By:	The username of the user that generated the report.
Report Name	The name of the report assigned when the report was created.
Scope	<p>The report is generated from this view and node in the network.</p> <ul style="list-style-type: none"> • <i>Perspective</i>—Can be a Logical View, Location View, or a Device View of the network. • <i>Node</i>—Represents the selected object on which the report is based.

The report body contains sections for device types, device groups, and device profiles. Each section shows the following information:

- A pie chart of the distribution of sessions by that section's category.

The definitions of the pie chart sections are listed below the chart. The session count appears next to each chart section.

- A table of the data shown in the pie chart.

RELATED DOCUMENTATION

Understanding the Types of Reports You Can Create 786
Creating Reports 790
Managing Generated Reports 801
Managing Reports on SCP Servers 804
Network Director Documentation home page

Device Inventory Report

The Device Inventory report is a standardized report generated in Network Director to show all devices that are visible to Network Director. There are two portions of the report: a report header and the report body. The contents of the report header are found in [Table 245 on page 820](#).

Table 245: Device Inventory Report Header

NETWORK DIRECTOR REPORT:	The type of report; In this case, the Device Inventory report.
Generated On:	The date and time the report ran. The date format is: [Day of the Month] [Month] [Year]. The time format follows the Coordinated Universal Time (UTC). For example: 07 December 2012 11:35 PM PST
Generated By:	The username of the user that generated the report.
Report Name:	The name of the report assigned by the user when the report was created.
Scope	<p>The report is generated from this view and node in the network.</p> <ul style="list-style-type: none"> • <i>Perspective</i>—Can be Logical view, Location view, or Device view of the network. • <i>Node</i>—Represents the selected object that the report is based.

Table 245: Device Inventory Report Header (*continued*)

Report Filters	<p>The report specified these device and connection state filters.</p> <p>Device Type—Supported device types are:</p> <ul style="list-style-type: none"> • EX Series switches • QFX • All (supported device types) <p>Connection State—Device connection states for filtering are:</p> <ul style="list-style-type: none"> • Up • Down • N/A • All (connection states)
----------------	--

The body of the report comprises:

- Device Type Count, which is a pie chart that graphically represents the network composition. Each segment represents:
 - EX Series switches
- Device Type Summary, which gives the total count of each type of device covered in the node.
- Device details by logical groups.

Following the pie chart, details for each device segment are listed by device type. For descriptions of the device fields see [Table 246 on page 821](#).

Table 246: Inventory Report Fields

Field	Description
HostName	The device label.
IP Address	Either the IPv4 or the IPv6 address.
Model	The full model number of the EX Series switch.
Software Version	The Junos software version and release number.
Serial No	The hardware serial number of the device.
Device Type	The type of hardware, such as switches. Switches can also be designated as NORMAL for standalone or VC for Virtual Chassis.

Table 246: Inventory Report Fields (*continued*)

Field	Description
Connection State	The connection state of the switch.
Configuration State	The administrative or operational state of the device.

RELATED DOCUMENTATION

[Understanding the Types of Reports You Can Create | 786](#)
[Creating Reports | 790](#)
[Managing Generated Reports | 801](#)
[Managing Reports on SCP Servers | 804](#)
[Viewing the Device Inventory Page | 551](#)
[Network Director Documentation home page](#)

Fabric Analyzer Report

The Fabric Analyzer report is a standardized report generated in Network Director to show information about a Virtual Chassis Fabric (VCF) or Virtual Chassis. There are two portions of the report: a report header and the report body. The contents of the report header are described in [Table 247 on page 822](#).

Table 247: Fabric Analyzer Report Header

Field	Description
NETWORK DIRECTOR REPORT:	The type of report; in this case, the Fabric Analyzer report.
Generated On:	The date and time the report ran. The date format is: [Day of the Month] [Month] [Year]. The time format follows the Coordinated Universal Time (UTC). For example: 07 December 2013 11:35 PM PST.
Generated By:	The username of the user that generated the report.
Report Name	The name of the report assigned when the report was created.

Table 247: Fabric Analyzer Report Header (*continued*)

Field	Description
Scope	<p>The report is generated from this view and node in the network.</p> <ul style="list-style-type: none"> • <i>Perspective</i>—Can be a Logical View, Location View, or a Device View of the network. • <i>Node</i>—Represents the selected object on which the report is based.

The contents of the report body depend on the fabric type:

For a VCF, the report body contains these sections:

- Virtual Chassis Connectivity Status—Shows summary and status information about the VCF and its members.
- Port Bandwidth Utilization—Shows information about the bandwidth used by each link between member devices.
- VCF Health Check—Shows the connection status between each leaf device and the spine devices.

For a Virtual Chassis, the report body contains these sections:

- Virtual Chassis Connectivity Status—Shows summary and status information about the Virtual Chassis and its members.
- Port Bandwidth Utilization—Shows information about the bandwidth used by each link between member devices.

RELATED DOCUMENTATION

[Understanding the Types of Reports You Can Create | 786](#)

[Creating Reports | 790](#)

[Managing Generated Reports | 801](#)

[Managing Reports on SCP Servers | 804](#)

[Network Director Documentation home page](#)

Network Device Traffic Report

IN THIS SECTION

- [Network Device Traffic Report Header | 824](#)
- [Network Device Traffic Charts | 825](#)

The Network Device Traffic report is a standardized report generated in Network Director to show the device traffic for a device. There are two portions of the report: a report header and the report body.

This topic describes:

Network Device Traffic Report Header

The contents of the report header are found in [Table 248 on page 824](#).

Table 248: Network Device Traffic Report Header

Field	Description
NETWORK DIRECTOR REPORT:	The type of report; In this case, the Network Device Traffic report.
Generated On:	The date and time the report ran. The date format is: [Day of the Month] [Month] [Year]. The time format follows the Coordinated Universal Time (UTC). For example: 07 December 2012 11:35 PM PST.
Generated By:	The username of the user that generated the report.
Report Name	The name of the report assigned by the user when the report was created.
Scope	<p>The report is generated from this view and node in the network.</p> <ul style="list-style-type: none"> ● <i>Perspective</i>—Can be Logical view, Location view, or a device view of the network. ● <i>Node</i>—Represents the selected object upon which the report is based. <p>NOTE: If you select a device that is down, the report might not contain any data.</p>
Report Filters	The period of time specified for data collection.

Network Device Traffic Charts

The body of the Network Device Traffic report is a series of four colored charts that show a comparison of data or trend information about the packets. The charts are:

- Unicast Vs Non-unicast

This pie-chart shows the percentage totals for packets over the specified period of time at the node:

- Inbound unicast packets
- Inbound non-unicast (such as broadcast and multicast) packets
- Outbound unicast packets
- Outbound non-unicast packets

The percentage of non-unicast packets is normally less than that of unicast packets. If the percentage of non-unicast packets is as high or higher than that of the unicast percentage, it means that too many non-unicast packets are being sent in the network.

- Unicast Vs Non-Unicast Trend

This line chart shows the trend in unicast and non-unicast packets over the specified period of the report. The x axis shows the polling period; the y axis shows the number of packets. Use this chart to see if the plots are symmetric or asymmetric. It can also be useful for identifying unusual patterns.

- Traffic Trend

This line chart shows the overall trend of all packets over the specified period of time. The x axis shows the polling period; the y axis shows the number of packets. Use this chart to find abnormalities in the traffic trend.

- Error Trend

This line chart shows the errors over the specified period of time. An error is caused by a missing packet. Missing packets can be a result of: packet loss in the network, uncorrectable packet out of sequence, packet length error, jitter buffer overflow, or jitter buffer underflow. Use this chart to see the overall trend in errors. The x axis shows the polling period; the y axis shows the number of errors.

RELATED DOCUMENTATION

[Understanding the Types of Reports You Can Create | 786](#)

[Creating Reports | 790](#)

[Managing Generated Reports | 801](#)

[Managing Reports on SCP Servers | 804](#)

[Network Director Documentation home page](#)

Port Bandwidth Utilization Report

The Port Bandwidth Utilization report is a standardized report generated in Network Director to show data for the last polled interval. There are two portions of the report: a report header and the report body. The contents of the report header are found in [Table 249 on page 826](#).

Table 249: Port Bandwidth Utilization Report Header

Field	Description
NETWORK DIRECTOR REPORT:	The type of report; In this case, the Port Bandwidth Utilization report.
Generated On:	The date and time the report ran. The date format is: [Day of the Month] [Month] [Year]. The time format follows the Coordinated Universal Time (UTC). For example: 07 December 2013 11:35 PM PST.
Generated By:	The username of the user that generated the report.
Report Name	The name of the report assigned when the report was created.
Scope	<p>The report is generated from this view and object in the network.</p> <ul style="list-style-type: none"> • <i>View</i>—Can be a Logical view, Location view, or a Device view of the network. • <i>Object</i>—Represents the selected object on which the report is based.
Report Filters	Shows the Percentage Utilization Exceeding value specified for the report. Only ports that exceed this percentage of their allocated bandwidth appear in the report.

The report contains a table for each device that contains ports that exceeded the specified percentage of their allocated bandwidth. The fields in these tables are described in [Table 250 on page 826](#).

Table 250: Port Bandwidth Utilization Report Fields

Field	Description
Host Name	Host name of the device.
IP Address	IP address of the device.
Device Type	Device type.
Port Name	Port name.
Percentage Utilization	Percentage of allocated bandwidth the port used.

RELATED DOCUMENTATION

Understanding the Types of Reports You Can Create 786
Creating Reports 790
Managing Generated Reports 801
Managing Reports on SCP Servers 804
Network Director Documentation home page

Top Users by Data Usage Report

IN THIS SECTION

- [Top Users by Data Usage Header | 827](#)
- [Top Users of Data Table | 828](#)

The Top Users by Data Usage report is a standardized report generated in Network Director. Use this report to identify the users with the highest data usage at the specified node during the specified time frame.

This topic describes:

Top Users by Data Usage Header

The contents of the report header are found in [Table 251 on page 827](#).

Table 251: Top 10 Users by Data Usage Report Header

Field	Description
NETWORK DIRECTOR REPORT:	The type of report; In this case, the Top Users by Data Usage report.
Generated On:	The date and time the report ran. The date format is: [Day of the Month] [Month] [Year]. The time format follows the Coordinated Universal Time (UTC). For example: 07 December 2012 11:35 PM PST.
Generated By:	The username of the user that generated the report.

Table 251: Top 10 Users by Data Usage Report Header (*continued*)

Field	Description
Report Name	The name of the report assigned by the user when the report was created.
Scope	The report is generated from this view and node in the network. Perspective can be Logical view, Location view, or Device view of the network. Node represents the selected object that the report is based.
Report Filters	The count specified for the report generation. The default count is 10 users.

Top Users of Data Table

The body of the Top Users by Data Usage report is a snapshot of the users with the highest data usage. The number of users analyzed and the time interval is determined by the Reporting Options set during report definition. The report sorts users in the table from the user with the highest bandwidth usage to the least highest. The key fields of the table are described in [Table 252 on page 828](#).

Table 252: Top 10 Users by Bandwidth Report Fields

Field	Description
User Name	The User ID with the highest bandwidth usage during the specified period.
Client MAC Address	The MAC address of the client.
Number of Sessions	The total sessions during this time period.
Total Data Used	The bandwidth used in megabytes.

RELATED DOCUMENTATION

[Network Director Documentation home page](#)

Traffic and Congestion Summary Report

The Traffic and Congestion Summary report is a standardized report generated in Network Director to show information about latency and port utilization on network devices. It provides detailed and trended

information about latency and port utilization at the network, device and port level. This report supports any scope that includes devices that support high-frequency statistics.

There are two portions of the report: a report header and the report body. The contents of the report header are described in [Table 253 on page 829](#).

Table 253: Traffic and Congestion Summary Report Header

Field	Description
NETWORK DIRECTOR REPORT:	The type of report; in this case, the Traffic and Congestion Summary report.
Generated On:	The date and time the report ran. The date format is: [Day of the Month] [Month] [Year]. The time format follows the Coordinated Universal Time (UTC). For example: 07 December 2013 11:35 PM PST.
Generated By:	The username of the user that generated the report.
Report Name	The name of the report assigned when the report was created.
Scope	<p>The report is generated from this view and node in the network.</p> <ul style="list-style-type: none"> • <i>Perspective</i>—Can be a Logical View, Location View, or a Device View of the network. • <i>Node</i>—Represents the selected object on which the report is based.
Report Filters	Shows any filters that were applied to the report definition.

The report body shows port utilization and latency information. If the report scope is My Network, the report contains sections that summarize port utilization and latency for all devices, then a section that shows more information about the top devices with the highest port utilization and latency. If the report scope is one device, the report shows more detailed information about that device, including port-level utilization and latency.

RELATED DOCUMENTATION

[Understanding the Types of Reports You Can Create | 786](#)

[Creating Reports | 790](#)

[Managing Generated Reports | 801](#)

[Managing Reports on SCP Servers | 804](#)

[Network Director Documentation home page](#)

8

PART

Working with Network Director Mobile

About Network Director Mobile | **831**

Getting Started with Network Director Mobile | **832**

Working in the Network Director Mobile Dashboard Mode | **835**

Working in the Network Director Mobile Devices Mode | **839**

About Network Director Mobile

IN THIS CHAPTER

- [Overview of Network Director Mobile | 831](#)

Overview of Network Director Mobile

Network Director Mobile is a Network Director user interface that is optimized to run in a mobile browser. It enables you to use Network Director monitoring features on a mobile device.

Network Director Mobile provides a Dashboard View that summarizes information about your entire network. It also enables you to drill down into individual devices for detailed information about those devices and the devices and sessions they manage.

RELATED DOCUMENTATION

[Network Director Mobile System Requirements | 832](#)

[Logging Into Network Director Mobile | 833](#)

[Understanding the Network Director Mobile User Interface | 833](#)

Getting Started with Network Director Mobile

IN THIS CHAPTER

- [Network Director Mobile System Requirements | 832](#)
- [Logging Into Network Director Mobile | 833](#)
- [Understanding the Network Director Mobile User Interface | 833](#)
- [Configuring Network Director Mobile Settings | 834](#)

Network Director Mobile System Requirements

Network Director Mobile runs in a mobile web browser on tablet devices. It has the following system requirements:

- On Apple iPad2, iPad3, and iPad mini devices:
 - Operating system versions 6.1.3 and 7.0.
 - Apple Safari browser versions included with the supported operating system versions.
- On Android tablet devices:
 - Android version 4.1.
 - Google Chrome browser version 29.0.1547.59 and higher.

RELATED DOCUMENTATION

[Overview of Network Director Mobile | 831](#)

[Logging Into Network Director Mobile | 833](#)

Logging Into Network Director Mobile

Network Director Mobile runs in a mobile browser. To log in to the Network Director server, navigate to this URL:

https://<server>/networkdirector/mobile, where <server> is the IP address or hostname of the Network Director server. Log in using your Network Director username and password.

RELATED DOCUMENTATION

[Network Director Mobile System Requirements | 832](#)

Understanding the Network Director Mobile User Interface

IN THIS SECTION

- [Dashboard Mode | 833](#)
- [Devices Mode | 833](#)

Network Director Mobile is a Network Director user interface that is optimized to run in a mobile browser. It enables you to use Network Director monitoring and fault management features on a mobile device.

The user interface has two modes: Dashboard and Devices. Buttons to access the modes are always available at the bottom of the page. When you log in to the server, Dashboard mode is open by default.

On any page that has a Back button, select the **Back** button to return to the previous page.

These sections describe the modes:

Dashboard Mode

Dashboard mode contains monitors that show information about your entire network.

Devices Mode

Devices mode enables you to drill down into individual devices for detailed information about these devices and the devices and sessions they manage.

RELATED DOCUMENTATION

[Monitoring Network-Wide Activity Using Network Director Mobile | 835](#)

[Locating a Device and Viewing Device Properties Using Network Director Mobile | 839](#)

[Configuring Network Director Mobile Settings | 834](#)

[Overview of Network Director Mobile | 831](#)

Configuring Network Director Mobile Settings

To configure Network Director Mobile settings:

1. Select the settings button in the main banner.

A dialog box opens.

2. Select the **General** tab to configure general settings:

- Refresh Interval—Select how often the application refreshes its data from the Network Director server.

3. Select the **Sessions** tab to configure sessions settings:

- Session Timeout—Select how long the application will wait before it logs off the session automatically if there is no user activity.

RELATED DOCUMENTATION

[Overview of Network Director Mobile | 831](#)

[Understanding the Network Director Mobile User Interface | 833](#)

Working in the Network Director Mobile Dashboard Mode

IN THIS CHAPTER

- [Monitoring Network-Wide Activity Using Network Director Mobile | 835](#)
- [Network Director Mobile Dashboard Reference | 835](#)

Monitoring Network-Wide Activity Using Network Director Mobile

Use Dashboard mode to monitor network-wide activity. To open Dashboard mode, select the **Dashboard** button that is always available at the bottom of the page.

RELATED DOCUMENTATION

| [Network Director Mobile Dashboard Reference | 835](#)

Network Director Mobile Dashboard Reference

IN THIS SECTION

- [Network Summary Monitor | 836](#)
- [Alarms Monitor | 836](#)
- [Top Sessions Monitor | 837](#)
- [Ports Monitor | 838](#)
- [Session Count Monitor | 838](#)
- [Session Trend Monitor | 838](#)

The Dashboard contains monitors that show information about your entire managed network:

Network Summary Monitor

The Network Summary monitor contains these pie charts:

- **Devices By Family**—Shows the distribution of devices based on device family.
- **Connection State**—Shows the distribution of devices based on the status of the device’s connection to the Network Director server. The possible connection states are:
 - **UP**—Device is connected to Network Director.
 - **DOWN**—Device is not connected to Network Director.
 - **N/A**—Device connection state is not available.
- **Configuration State**—Shows the distribution of devices based on whether the Network Director configuration is in sync with the device configuration. The possible configuration states for a device depend on its connection state:
 - When connection state is **UP**, the configuration state can be **Out of Sync**, **Synchronizing**, **In Sync**, or **Sync Failed**.
 - When connection state is **DOWN**, the configuration state is **N/A**.

Alarms Monitor

The Alarms monitor provides a quick summary of the critical, major, minor, and info alarms currently active in the network.

Select the **Expand** button in the right corner of the title bar to see detailed information about the alarms on the Alarms page.

On the Alarms page, you have the following options:

- Select the **Graph** or **List** buttons to view the information in those formats.
For a description of the information presented in the list view, see [Table 254 on page 836](#).
- In the graph view, you can select **By Severity**, **By Category**, or **By State** to see the distribution of active alarms by those properties.

Select **Back** to return to the Dashboard.

Table 254: Network Director Mobile Alarm Details Fields

Field	Value
Name	The alarm name.

Table 254: Network Director Mobile Alarm Details Fields (*continued*)

Field	Value
Alarm Severity	<p>The severity of the alarm. Severity levels are:</p> <ul style="list-style-type: none"> • Critical—A critical condition exists; immediate action is necessary. • Major—A major error has occurred; escalate or notify as necessary. • Minor—A minor error has occurred; notify or monitor the condition. • Info—An informational message; no action is necessary.
Entity ID	The identification of the entity responsible for causing this alarm. The Entity ID is the key for correlating events into an alarm. The Entity ID could be an IP address of the device.
Assigned to	If the alarm is assigned to an individual, it shows the name of that person; otherwise, it shows System to mark that the alarm is still unassigned.
Last Updated On	The date and time that the information for the alarm was last modified.

Top Sessions Monitor

The Top Sessions monitor contains a bar chart showing the eight user sessions that are currently using the most bandwidth.

Select the **Expand** button in the top right corner of the title bar to see detailed information on the Top Sessions page. On the Top Sessions page, you have the following options:

- Select **Top Sessions By User** to see sessions that are identified by their user.
- Select **Top Sessions By MAC** to see sessions that are identified by their MAC address.
- Select the **Graph** or **List** buttons to view the information in those formats.

For a description of the information presented in the list view, see [Table 255 on page 837](#).

- Select the time period to display from the list in the title bar.

Select **Back** to return to the Dashboard.

Table 255: Network Director Mobile Top Session Details

Table Column	Description
Username	Client's user name.
Total Data Usage (KBytes)	The session's total data usage. Appears when any time period other than Current is selected.

Table 255: Network Director Mobile Top Session Details (continued)

Table Column	Description
Number of Sessions	Number of sessions.

Ports Monitor

The Ports monitor shows information about the network’s device ports:

- Admin Status—Shows the number of ports that are up and down.
- Free Vs Used—Shows the number of ports that are free and the number that are used.

Select the **Expand** button in the top right corner of the title bar to see detailed information on the Ports page. On the Ports page, you have the option to view ports by admin status or by free versus used status.

Select **Back** to return to the Dashboard.

Session Count Monitor

The Session Count monitor shows the number of active user sessions on the network.

Session Trend Monitor

The Session Trend monitor contains a line graph that shows the number of active user sessions on the network over time.

RELATED DOCUMENTATION

Working in the Network Director Mobile Devices Mode

IN THIS CHAPTER

- [Locating a Device and Viewing Device Properties Using Network Director Mobile | 839](#)
- [Monitoring Sessions on a Device Using Network Director Mobile | 840](#)

Locating a Device and Viewing Device Properties Using Network Director Mobile

You can locate a device and view its properties by searching or by browsing.

To locate a device and view its properties:

1. Select the **Devices** button at the bottom of the page to open Devices mode.
2. To locate the device by searching:
 - a. Enter search text in the search box (it contains the text Hostname or IP until you enter text).
 - b. Select the search button.
 - c. Locate the device in the list of search results.
 - d. For information about the device properties shown, see [Table 256 on page 840](#).
3. To locate the device by browsing:
 - a. Select the device type (**Switches**).
 - b. If you selected Switches, select the device family from the list, then locate the device in the list of devices that opens.

- c. For information about the device properties shown, see [Table 256 on page 840](#).

Table 256: Device Properties Shown in Network Director Mobile Inventory

Field	Description
Hostname	Configured name of the device.
Device Family	Device family of the device. For example, MSS, EX, or WLC. Shown only on inventory pages created by searching.
Platform	Model number of the device. Shown only on inventory pages created by searching.
Model	Type of the device. Shown only on inventory pages that you browse to, not on search results pages.
Mgmt IP	IP Address of the device.
Mgmt Status	Displays whether the device is directly manageable or not.
Connection Status	Connection status of the device in Network Director: <ul style="list-style-type: none"> • UP—Device is connected to Network Director. • DOWN—Device is not connected to Network Director. • N/A—Device connection status is not available.
Serial Number	Serial number on device chassis.

RELATED DOCUMENTATION

[Monitoring Sessions on a Device Using Network Director Mobile](#) | 840

Monitoring Sessions on a Device Using Network Director Mobile

To monitor session activity on a device:

1. Locate the device as described in [“Locating a Device and Viewing Device Properties Using Network Director Mobile” on page 839](#).
2. Select the device from the list.

3. Select **Session Details**.

The Session Details page for the device opens. You can select the **Graph** or **List** buttons to view the information in those formats.

4. To see historical session data, select **Session Trend**.

You can select the time period to view by selecting a time period from the list in the page's title bar.

For a description of the information presented in the list view, see [Table 257 on page 841](#).

5. To see current session data, select **Current Sessions**.

For a description of the information presented in the list view, see [Table 258 on page 841](#).

Table 257: Session Trend Details

Table Column	Description
Time	Time when a poll occurred.
Min Session Count	Minimum session count for the time period.
Avg Session Count	Average session count.
Max Session Count	Maximum session count for the time period.

Table 258: Current Session Details

Table Column	Description
Username	Client's user name
MAC Address	Client's MAC address.
Incremental Data Usage (KBytes)	The session's current incremental data usage.

RELATED DOCUMENTATION

[Locating a Device and Viewing Device Properties Using Network Director Mobile](#) | 839