



---

Junos<sup>®</sup> Space

Edge Services Director User Guide

Release

1.1



---

Modified: 2019-05-28

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<sup>®</sup> Space Edge Services Director User Guide*

1.1

Copyright © 2019 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 . . . . .	xxvii
	Documentation and Release Notes . . . . .	xxvii
	Documentation Conventions . . . . .	xxvii
	Documentation Feedback . . . . .	xxix
	Requesting Technical Support . . . . .	xxx
	Self-Help Online Tools and Resources . . . . .	xxx
	Creating a Service Request with JTAC . . . . .	xxxi
<b>Part 1</b>	<b>Overview</b>	
<b>Chapter 1</b>	<b>Edge Services Director Overview . . . . .</b>	<b>3</b>
	Understanding the Need for Edge Services Director . . . . .	3
	Understanding Edge Services Director User Administration . . . . .	5
	Understanding the Edge Services Director User Interface . . . . .	6
	View Pane . . . . .	8
	Displaying Devices in Various Network Views . . . . .	8
	Expanding or Collapsing Nodes in the Network Tree . . . . .	9
	Searching the Network Tree . . . . .	9
	Tasks Pane . . . . .	10
	Alarms . . . . .	10
	Main Window or Workspace . . . . .	10
	Tables in Edge Services Director . . . . .	11
	Moving and Resizing Columns . . . . .	11
	Navigating Pages . . . . .	11
	Displaying the Column Drop-Down Menu . . . . .	11
	Sorting on a Column . . . . .	12
	Hiding and Exposing Columns . . . . .	12
	Searching Table Contents . . . . .	13
	Filtering Table Contents . . . . .	14
	Understanding Edge Services Director and the Management Lifecycle Modes . . . . .	15
	Service Delivery Gateway Overview . . . . .	17
	Carrier-Grade NAT . . . . .	17
	Firewalls and Intrusion Prevention System . . . . .	18
	Traffic Direct . . . . .	18
	Load Balancing and Adaptive Services . . . . .	18
	Edge Services Director Overview . . . . .	19

<b>Chapter 2</b>	<b>Getting Started . . . . .</b>	<b>21</b>
	Understanding How to Use the Edge Services Director Interface to View System Information . . . . .	22
	Getting Started Assistant in Junos Space Platform Overview . . . . .	23
	Changing Your Password for Edge Services Director . . . . .	24
	Logging In to Edge Services Director . . . . .	25
	Logging Out of Edge Services Director . . . . .	27
	Quickly Accessing Important Monitoring and Troubleshooting Details . . . . .	28
<b>Chapter 3</b>	<b>Tasks Pane . . . . .</b>	<b>31</b>
	Understanding the Build Mode Tasks Pane . . . . .	31
	Understanding the Deploy Mode Tasks Pane . . . . .	38
	Understanding the Fault Mode Tasks Pane . . . . .	40
	Understanding the Monitor Mode Tasks Pane . . . . .	41
	Understanding the Report Mode Tasks Pane . . . . .	42
<b>Chapter 4</b>	<b>Dashboard . . . . .</b>	<b>45</b>
	Understanding the Dashboard . . . . .	45
	Working with the Dashboard . . . . .	45
	SDG Views . . . . .	46
	Service Delivery Gateway Alarms . . . . .	47
	Filters . . . . .	47
	Specifying KPI Template and Alarm Filters . . . . .	48
	Service Delivery Gateways Count by Severity . . . . .	49
	Service Gateway Ticker Updates . . . . .	50
	Service Delivery Gateway Health Status Trend . . . . .	50
	Using Dashboard Widgets . . . . .	50
	Alarm Severities and States Overview . . . . .	51
	Alarm Severity . . . . .	51
	Alarm State . . . . .	52
	Viewing the Detailed Status of KPI Templates Applied to Devices . . . . .	52
<b>Part 2</b>	<b>System Administration</b>	
<b>Chapter 5</b>	<b>Handling Administrative Tasks . . . . .</b>	<b>57</b>
	Understanding Edge Services Director User Administration . . . . .	57
	Viewing Audit Logs From Edge Services Director . . . . .	58
	Managing Jobs . . . . .	59
	Collecting Logs for Troubleshooting . . . . .	60
<b>Part 3</b>	<b>Gateway View of Build Mode</b>	
<b>Chapter 6</b>	<b>About Gateway View of Build Mode . . . . .</b>	<b>65</b>
	Understanding Build Mode in Gateway View of Edge Services Director . . . . .	65
	Discovering Devices . . . . .	65
	Configuring Devices . . . . .	66
	Deploying Device Configurations . . . . .	66
	Importing Device Configurations . . . . .	67
	Out-of-Band Configuration Changes . . . . .	67
	Viewing the Devices Inventory . . . . .	67

	Service Delivery Gateway Groups . . . . .	68
	KPI Templates . . . . .	68
	Understanding Resynchronization of Device Configuration . . . . .	69
	The Resynchronize Device Configuration Task . . . . .	70
	How Resynchronization Works in NSOR Mode . . . . .	70
	How Resynchronization Works in SSOR Mode . . . . .	71
	How Edge Services Director Resynchronizes the Build Mode Configuration . . . . .	73
	Importing Devices . . . . .	74
	Device Discovery Overview . . . . .	77
	Unmanaged Devices Overview . . . . .	78
	Working With Managed Devices . . . . .	80
	Working With Unmanaged Devices . . . . .	80
	Working With Discovered Devices . . . . .	81
	Managing Jobs as a System Task . . . . .	81
<b>Chapter 7</b>	<b>Managing Service Delivery Gateways and Groups . . . . .</b>	<b>85</b>
	Discovering Devices . . . . .	85
	Preparing MX Series Devices for Discovery . . . . .	86
	Specifying a Discovery Profile and the Target Devices . . . . .	87
	Specifying SNMP Probes . . . . .	89
	Specifying Credentials . . . . .	91
	Comparing Configuration Settings of Devices . . . . .	92
	Exporting Managed Device Details to a CSV File . . . . .	94
	Changing an Unmanaged Device to a Managed Device . . . . .	95
	Modifying the SDG Group and KPI Templates for a Device . . . . .	97
	Scheduling the Discovery of Devices . . . . .	98
	Creating Service Gateway Groups . . . . .	99
	Managing Service Gateway Groups . . . . .	101
	Viewing the Service Gateway Details . . . . .	102
	Searching Unmanaged Devices . . . . .	104
	Viewing the List of Discovered, Managed, and Unmanaged Devices . . . . .	106
	Changing a Managed Device to an Unmanaged Device . . . . .	112
	Modifying Discovery Profiles . . . . .	113
	Deleting Discovery Profiles . . . . .	114
	Systems of Record in Junos Space Overview . . . . .	115
	Systems of Record . . . . .	115
	Implications on device management . . . . .	116
	Resynchronizing Managed SDGs with the Network . . . . .	117
<b>Chapter 8</b>	<b>Managing KPI Templates . . . . .</b>	<b>119</b>
	Understanding Measurement Points, Key Performance Indicators, and Baseline Values . . . . .	119
	Measurement Points . . . . .	119
	Basic Key Performance Indicators . . . . .	120
	Setting Baselines . . . . .	120
	Cloning a KPI Template . . . . .	121
	Deleting KPI Templates . . . . .	128
	Managing KPI Templates . . . . .	129
	Viewing KPI Templates . . . . .	130

	Modifying a KPI Template Associated with a Service Gateway . . . . .	131
<b>Chapter 9</b>	<b>Viewing the Device Inventory . . . . .</b>	<b>133</b>
	Viewing the Device Inventory Page . . . . .	133
	Viewing Device Statistics . . . . .	141
	Viewing the Number of Devices by Platform . . . . .	142
	Viewing Connection Status for Devices . . . . .	142
	Viewing Devices by Junos OS Release . . . . .	143
	Viewing Configuration Details of Services on Devices . . . . .	144
	Viewing Discovery Logs . . . . .	146
	Viewing Discovery Profiles . . . . .	147
<b>Part 4</b>	<b>Location and Device Views of Build Mode</b>	
<b>Chapter 10</b>	<b>Location View Configuration . . . . .</b>	<b>151</b>
	Understanding Build Mode in Location and Device Views of Edge Services	
	Director . . . . .	151
	Discovering Devices . . . . .	151
	Building the Location and Custom Views . . . . .	152
	Configuring Devices . . . . .	153
	Deploying Device Configurations . . . . .	153
	Importing Device Configurations . . . . .	153
	Out-of-Band Configuration Changes . . . . .	153
	Managing Devices . . . . .	154
	Understanding the Location View . . . . .	154
	Assigning and Unassigning Devices to a Location . . . . .	155
	How to Assign or Unassign Devices . . . . .	156
	Assigning Devices . . . . .	157
	Changing the Location of a Device . . . . .	157
	How to Move a Device to a New Location . . . . .	158
	Changing the Location of a Device . . . . .	158
	Configuring Buildings . . . . .	159
	How to Add or Edit a Building . . . . .	159
	Adding or Editing a Building for a Location . . . . .	159
	Configuring Floors . . . . .	160
	How to Add or Edit a Floor . . . . .	160
	Adding or Editing a Building Floor for a Location . . . . .	161
	Configuring Outdoor Areas . . . . .	162
	How to Configure an Outdoor Area . . . . .	162
	Configuring an Outdoor Area . . . . .	162
	Creating a Site . . . . .	163
	How to Add or Edit a Location Site . . . . .	163
	Creating or Editing a Site . . . . .	163
	Deleting Sites, Buildings, Floors, Wiring Closets, and Devices . . . . .	164
	How to Delete a Location Object . . . . .	164
	Deleting Sites . . . . .	165
	Deleting Buildings . . . . .	165
	Deleting Floors . . . . .	165
	Deleting Closets . . . . .	165
	Deleting Devices . . . . .	165

	Setting Up Closets . . . . .	166
	How to Add or Edit a Closet . . . . .	166
	Adding or Editing a Wiring Closet . . . . .	167
	Setting Up the Location View . . . . .	168
<b>Chapter 11</b>	<b>Device Management . . . . .</b>	<b>173</b>
	Accessing a Device's CLI from Edge Services Director . . . . .	173
	Deleting Devices from Edge Services Director . . . . .	174
	Rebooting Devices from Edge Services Director . . . . .	175
	Viewing the Device Inventory Page in Device View of Edge Services Director . . . . .	176
	Viewing the Physical Inventory of Devices . . . . .	178
	Viewing a Device's Current Configuration from Edge Services Director . . . . .	179
<b>Part 5</b>	<b>Service View of Build Mode</b>	
<b>Chapter 12</b>	<b>About Build Mode in Service View . . . . .</b>	<b>183</b>
	Understanding Build Mode in Service View of Edge Services Director . . . . .	183
	Service Designer . . . . .	183
	Services Inventory . . . . .	184
	Object Builder . . . . .	184
<b>Chapter 13</b>	<b>Using the Service Designer . . . . .</b>	<b>185</b>
	Object Builder Overview . . . . .	185
	Planning and Deployment of Service Templates Overview . . . . .	187
	Planning Workflow for Service Templates . . . . .	187
	Deployment Workflow for Service Templates . . . . .	187
	Service Templates Overview . . . . .	189
	Filtering Service Templates . . . . .	189
	Restoring Service Template Configurations . . . . .	190
	Viewing Service Templates . . . . .	192
	Viewing the Services Inventory Page . . . . .	193
	Using the Actions Menu on the Service Template and Service Edit Pages . . . . .	195
	Publishing a Service Template . . . . .	196
	Unpublishing a Service Template . . . . .	197
	Exporting a Service to a CSV File . . . . .	198
	Cloning a Service Template . . . . .	199
	Creating a Deploy Plan and Provisioning Services Immediately . . . . .	200
	Viewing a Graphical Statistic of Service Templates . . . . .	202
	Creating and Managing ADC Service Templates . . . . .	203
	Creating an ADC Service Template . . . . .	204
	Importing an ADC Service Template . . . . .	207
	Creating a Deployment Plan . . . . .	209
	Creating a Real Server . . . . .	210
	Creating a Group for Real Servers . . . . .	212
	Load-Balancing Methods for Real-Server Groups . . . . .	214
	Creating a Client-Facing Interface and Routing Instance . . . . .	216
	Creating a Server-Facing Interface and Routing Instance . . . . .	218
	Creating a Services PIC for an ADC Service Template . . . . .	219
	Creating a Health Check for an ADC Service Template . . . . .	221
	Creating a Custom Health Check for an ADC Instance . . . . .	222

Creating a Virtual Service for an ADC Service Template . . . . .	225
Creating a Virtual Server for an ADC Service Template . . . . .	228
Creating a Firewall Rule for an ADC Service Template . . . . .	229
Modifying ADC Service Templates . . . . .	232
Creating and Managing CGNAT Service Templates . . . . .	234
Creating a CGNAT Service Template . . . . .	235
Modifying CGNAT Service Templates . . . . .	238
Creating a Deployment Plan . . . . .	240
Importing a CGNAT Service Template . . . . .	241
Creating a Service Set . . . . .	243
Creating a Syslog . . . . .	247
Creating a Rule . . . . .	249
Creating a Rule Set . . . . .	250
Creating a Pool . . . . .	251
Creating and Managing SFW Service Templates . . . . .	252
Creating an SFW Service Template . . . . .	253
Modifying SFW Service Templates . . . . .	256
Creating a Deployment Plan . . . . .	258
Importing an SFW Service Template . . . . .	259
Creating a Service Set . . . . .	261
Creating an Application . . . . .	265
Creating an Application Set . . . . .	268
Creating a Syslog . . . . .	269
Creating a Rule . . . . .	271
Creating a Rule Set . . . . .	272
Creating a Services PIC for an SFW Service Template . . . . .	274
Creating and Managing TLB Service Templates . . . . .	275
Creating a TLB Service Template . . . . .	276
Creating a Deployment Plan . . . . .	279
Modifying TLB Service Templates . . . . .	280
Importing a TLB Service Template . . . . .	282
Creating a Real Server . . . . .	284
Creating a Group for Real Servers . . . . .	285
Creating a Services PIC for a TLB Service Template . . . . .	288
Creating a Network Monitor Profile for a TLB Service Template . . . . .	289
Creating a Command for Script-Based Health Checks . . . . .	291
Creating a Server Bypass Filter . . . . .	292
Creating a Virtual Service for a TLB Service Template . . . . .	293
Creating a Client-Facing Interface and Routing Instance . . . . .	296
Creating a Server-Facing Interface and Routing Instance . . . . .	298
Modifying Individual Service Instances and Deploying to Devices . . . . .	300
Modifying Service Instances . . . . .	300
Creating a Deployment Plan . . . . .	302
<b>Chapter 14</b>	
<b>Using the Object Builder . . . . .</b>	<b>305</b>
Understanding the Object Builder . . . . .	305
Importing All Types of Objects . . . . .	306
Importing SFW Rule Sets . . . . .	308
Importing SFW Rules . . . . .	310

	Importing Real Server Settings . . . . .	312
	Importing CGNAT Rule Sets . . . . .	313
	Importing CGNAT Rules . . . . .	315
	Importing CGNAT Pools . . . . .	316
	Importing Applications . . . . .	318
	Importing Application Sets . . . . .	319
<b>Chapter 15</b>	<b>Managing Packet Analyzers . . . . .</b>	<b>321</b>
	Packet Analyzer Overview . . . . .	321
	Pre-Service Filtering of Traffic for Service Processing . . . . .	322
	Postservice Filtering of Returning Service Traffic . . . . .	323
	Creating and Viewing Service Analyzers . . . . .	323
	Configuring the Traffic Analyzer Filter . . . . .	323
	Managing Service Analyzer Filter Instances . . . . .	326
	Viewing Service Analyzer Instance Details . . . . .	328
	Viewing the Service Analyzer Statistics in Grid Format and Graph . . . . .	330
<b>Part 6</b>	<b>Deploy Mode</b>	
<b>Chapter 16</b>	<b>About Deploy Mode . . . . .</b>	<b>335</b>
	Understanding Deploy Mode in Gateway and Service Views of Edge Services	
	Director . . . . .	335
	Deploying Configuration Changes . . . . .	335
	Transactions . . . . .	336
	Modify the Association of SDG Details and Rule Terms for a Policy	
	Filters . . . . .	336
	View Service Object Statistics . . . . .	337
	Service Edit . . . . .	337
	Policy and Filter Management . . . . .	337
	Understanding Deploy Mode in Location and Device Views of Edge Services	
	Director . . . . .	338
	Managing Software Images . . . . .	338
	Managing Devices . . . . .	338
	Managing Device Configuration Files . . . . .	338
<b>Chapter 17</b>	<b>Device Management . . . . .</b>	<b>339</b>
	Viewing the Device Inventory Page in Device View of Edge Services Director . . .	340
	Resynchronizing Device Configuration . . . . .	342
	The Resynchronize Device Configuration List of Devices . . . . .	343
	Resynchronizing Devices When Junos Space Is in NSOR Mode . . . . .	344
	Resynchronizing Devices When Junos Space Is in SSOR Mode . . . . .	344
	Resynchronizing Devices in Manual Approval Mode . . . . .	345
	Viewing the Network Changes . . . . .	345
	Viewing Resynchronization Job Status . . . . .	346
<b>Chapter 18</b>	<b>Configuration File Management . . . . .</b>	<b>347</b>
	Managing Device Configuration Files . . . . .	347
	Selecting Device Configuration File Management Options . . . . .	347
	Backing Up Device Configuration Files . . . . .	348
	Restoring Device Configuration Files . . . . .	349

	Viewing Device Configuration Files . . . . .	349
	Comparing Device Configuration Files . . . . .	350
	Deleting Device Configuration Files . . . . .	350
	Managing Device Configuration File Management Jobs . . . . .	350
	Managing Jobs . . . . .	351
<b>Chapter 19</b>	<b>Software Image Management . . . . .</b>	<b>353</b>
	Managing Software Images . . . . .	353
	Selecting Software Image Management Options . . . . .	353
	Adding Software Images to the Repository . . . . .	354
	Using the Device Image Upload Window . . . . .	354
	Viewing Software Image Details . . . . .	355
	Using the Device Image Summary Window . . . . .	355
	Deleting Software Images . . . . .	355
	Deploying Software Images . . . . .	356
	Specifying Software Deployment Job Options . . . . .	356
	Selecting Software Images To Deploy . . . . .	357
	Selecting Options for Software Deployment . . . . .	358
	Summary of Software Deployment . . . . .	359
	Managing Software Image Deployment Jobs . . . . .	359
	Selecting Software Image Management Options . . . . .	360
	Viewing Software Image Job Details . . . . .	361
	Using the Device Image Staging Window . . . . .	361
	Canceling Software Image Jobs . . . . .	362
<b>Chapter 20</b>	<b>Viewing and Editing Service Instances and Packet Filters Across All Gateways . . . . .</b>	<b>363</b>
	Viewing Service Object Statistics . . . . .	363
	Modifying Service Instances . . . . .	365
	Modifying Packet Filter Policies . . . . .	367
<b>Chapter 21</b>	<b>Enhanced Editing of Services and Packet Filters . . . . .</b>	<b>369</b>
	Enhanced Editing of Service Policies and Policy Filters Overview . . . . .	369
	Modifying the Association of SDG Details and Service Components for a Packet Filter Policy . . . . .	370
	Modifying the Association of SDG Details and Service Components for a Service Policy Filter . . . . .	372
<b>Chapter 22</b>	<b>Managing Service Instance and Policy Rule Definitions . . . . .</b>	<b>375</b>
	Policy and Filter Management Overview . . . . .	375
	States and Transitions of Policies or Filters . . . . .	376
	User Roles . . . . .	377
	Packet and Service Filters Overview . . . . .	378
	Filtering Traffic Before Accepting Packets for Service Processing . . . . .	379
	Postservice Filtering of Returning Service Traffic . . . . .	380
	Searching for CGNAT Policies . . . . .	381
	Searching for Packet Filters . . . . .	384
	Searching for SFW Policies . . . . .	386
	Managing Service and Policy Locks . . . . .	387
	Unlocking Locked Services and Policies . . . . .	389

Viewing Policy and Filter Instances . . . . .	390
Creating and Managing CGNAT Policy and Filter Instances . . . . .	395
Creating a NAT Policy . . . . .	396
Creating a Service Set . . . . .	399
Creating a Syslog . . . . .	403
Creating a Rule . . . . .	405
Creating a Rule Set . . . . .	406
Creating Addresses . . . . .	408
Creating Address Groups . . . . .	409
Address and Address Groups Overview . . . . .	409
Creating a NAT Rule Term . . . . .	410
Associating an Application and Application Set with a NAT Rule . . . . .	414
Creating a NAT Pool . . . . .	414
Associating Service Sets and Rule Sets With a NAT Rule . . . . .	415
Modifying NAT Policies . . . . .	416
Creating a Deployment Plan . . . . .	417
Creating and Managing Packet Filter Policy Instances . . . . .	419
Creating a Packet Filter Policy . . . . .	420
Creating Addresses . . . . .	422
Creating Address Groups . . . . .	423
Address and Address Groups Overview . . . . .	424
Creating a Packet Filter Rule Term . . . . .	424
Creating an Application and Application Set . . . . .	428
Associating Service Sets and Rule Sets With a Packet Filter Rule . . . . .	428
Associating Interfaces With a Packet Filter Rule . . . . .	429
Modifying Packet Filter Policies . . . . .	429
Creating a Deployment Plan . . . . .	432
Creating and Managing SFW Policy and Filter Instances . . . . .	434
Creating an SFW Policy . . . . .	435
Creating a Service Set . . . . .	438
Creating a Syslog . . . . .	442
Creating a Rule . . . . .	445
Creating a Rule Set . . . . .	446
Creating Addresses . . . . .	447
Creating Address Groups . . . . .	449
Address and Address Groups Overview . . . . .	449
Creating an SFW Rule Term . . . . .	449
Creating an Application and Application Set . . . . .	452
Associating Service Sets and Rule Sets With an SFW Rule . . . . .	452
Modifying SFW Policies . . . . .	453
Creating a Deployment Plan . . . . .	454
Viewing CGNAT Service Templates . . . . .	456
Viewing SFW Service Templates . . . . .	457
Viewing and Modifying ADC Service Instances . . . . .	459
Viewing ADC Service Instances . . . . .	459
Modifying ADC Service Instances . . . . .	461
Creating a Deploy Plan and Provisioning Services Immediately . . . . .	463
Filtering ADC Service Instances . . . . .	465
Managing ADC Service Instance Locks . . . . .	467

	Unlocking Locked ADC Service Instances . . . . .	469
	Viewing and Modifying TLB Service Instances . . . . .	471
	Viewing TLB Service Instances . . . . .	472
	Modifying TLB Service Instances . . . . .	473
	Creating a Deploy Plan and Provisioning Services Immediately . . . . .	476
	Filtering TLB Service Instances . . . . .	478
	Managing TLB Service Instance Locks . . . . .	480
	Unlocking Locked TLB Service Instances . . . . .	481
	Using the Actions Menu on the Service Policy and Packet Filter Pages . . . . .	483
	Creating a Deployment Plan . . . . .	483
	Discarding Changes Made to a Service Policy or Packet Filter Policy . . . . .	484
	Tagging Junos Space Network Management Platform Objects . . . . .	485
	Creating a Tag . . . . .	486
	Tagging an Object . . . . .	489
	Untagging an Object . . . . .	490
<b>Chapter 23</b>	<b>Managing Packet Analyzers . . . . .</b>	<b>493</b>
	Packet Analyzer Overview . . . . .	493
	Pre-Service Filtering of Traffic for Service Processing . . . . .	494
	Postservice Filtering of Returning Service Traffic . . . . .	495
	Creating and Viewing Service Analyzers . . . . .	495
	Configuring the Traffic Analyzer Filter . . . . .	495
	Managing Service Analyzer Filter Instances . . . . .	498
	Viewing Service Analyzer Instance Details . . . . .	500
	Viewing the Service Analyzer Statistics in Grid Format and Graph . . . . .	502
<b>Part 7</b>	<b>Deploy Mode</b>	
<b>Chapter 24</b>	<b>About Deploy Mode . . . . .</b>	<b>507</b>
	Understanding Deploy Mode in Gateway and Service Views of Edge Services	
	Director . . . . .	507
	Deploying Configuration Changes . . . . .	507
	Transactions . . . . .	508
	Modify the Association of SDG Details and Rule Terms for a Policy	
	Filters . . . . .	508
	View Service Object Statistics . . . . .	509
	Service Edit . . . . .	509
	Policy and Filter Management . . . . .	509
<b>Chapter 25</b>	<b>Configuration File Management . . . . .</b>	<b>511</b>
	Managing Device Configuration Files . . . . .	511
	Selecting Device Configuration File Management Options . . . . .	511
	Backing Up Device Configuration Files . . . . .	512
	Restoring Device Configuration Files . . . . .	513
	Viewing Device Configuration Files . . . . .	513
	Comparing Device Configuration Files . . . . .	514
	Deleting Device Configuration Files . . . . .	514
	Managing Device Configuration File Management Jobs . . . . .	514
	Managing Jobs . . . . .	515

<b>Chapter 26</b>	<b>Software Image Management . . . . .</b>	<b>517</b>
	Managing Software Images . . . . .	517
	Selecting Software Image Management Options . . . . .	517
	Adding Software Images to the Repository . . . . .	518
	Using the Device Image Upload Window . . . . .	518
	Viewing Software Image Details . . . . .	519
	Using the Device Image Summary Window . . . . .	519
	Deleting Software Images . . . . .	519
	Deploying Software Images . . . . .	520
	Specifying Software Deployment Job Options . . . . .	520
	Selecting Software Images To Deploy . . . . .	521
	Selecting Options for Software Deployment . . . . .	522
	Summary of Software Deployment . . . . .	523
	Managing Software Image Deployment Jobs . . . . .	523
	Selecting Software Image Management Options . . . . .	524
	Viewing Software Image Job Details . . . . .	525
	Using the Device Image Staging Window . . . . .	525
	Canceling Software Image Jobs . . . . .	526
<b>Chapter 27</b>	<b>Deploying Configurations to Devices . . . . .</b>	<b>527</b>
	Planning and Deployment of Service Templates Overview . . . . .	527
	Planning Workflow for Service Templates . . . . .	527
	Deployment Workflow for Service Templates . . . . .	528
	Viewing Deployment Plans . . . . .	529
	Creating and Assigning a Deployment Plan to Devices . . . . .	533
	Creating a Deployment Plan . . . . .	534
	Publishing a Deploy Plan . . . . .	537
	Viewing Deploy Plans and Policies . . . . .	538
	Approving a Deploy Plan and Policies . . . . .	539
	Unpublishing a Deploy Plan and Policies . . . . .	540
	Deploying a Deploy Plan and Policies Immediately . . . . .	540
	Scheduling Deployment of Services and Policies . . . . .	541
	Rejecting a Deploy Plan and Policies . . . . .	542
	Changing a Deploy Plan Action or Decommissioning a Deploy Plan . . . . .	543
	Discarding a Deploy Plan and Policies . . . . .	544
	Modifying the Association of SDG Details and Service Components for a Packet Filter Policy . . . . .	544
	Modifying the Association of SDG Details and Service Components for a Service Policy Filter . . . . .	547
<b>Chapter 28</b>	<b>Viewing Transactions Associated with Deployment Jobs . . . . .</b>	<b>551</b>
	Transactions Overview . . . . .	551
	Viewing Transactions . . . . .	552
<b>Part 8</b>	<b>Monitor Mode</b>	
<b>Chapter 29</b>	<b>About Monitor Mode . . . . .</b>	<b>559</b>
	Understanding Monitor Mode in Edge Services Director . . . . .	559
	General Monitoring . . . . .	559
	Packet Analyzer . . . . .	560

	Fault Management . . . . .	560
	Performance Management . . . . .	560
<b>Chapter 30</b>	<b>Using Fault Management Monitors . . . . .</b>	<b>561</b>
	Understanding Fault Management . . . . .	561
	Viewing the Fault Management Details . . . . .	562
	Viewing Charts of Alarms and Syslogs . . . . .	563
	Viewing Alarms, Events, and Syslogs . . . . .	563
	Changing the Alarm State . . . . .	566
	Searching Alarms . . . . .	566
	Searching Events . . . . .	567
	Searching System Log Messages . . . . .	568
	Acknowledging Alarms . . . . .	568
	Clearing Alarms . . . . .	569
	Escalating Alarms . . . . .	569
	Unacknowledging Alarms . . . . .	570
<b>Chapter 31</b>	<b>Using Performance Management Utilities . . . . .</b>	<b>571</b>
	Performance Management . . . . .	571
	The Need and Benefits of Performance Manager . . . . .	571
	Performance Manager View After a Context-Switch from the Monitoring Page . . . . .	576
<b>Chapter 32</b>	<b>General Monitoring . . . . .</b>	<b>577</b>
	Monitoring Capabilities Overview . . . . .	577
	Viewing the Monitoring Page in Gateway View . . . . .	578
	Viewing the ADC Service Details . . . . .	583
	Viewing the TLB Service Details . . . . .	585
	Viewing the CGNAT Service Details . . . . .	588
	Viewing the SFW Service Details . . . . .	591
<b>Part 9</b>	<b>Fault Mode</b>	
<b>Chapter 33</b>	<b>About Fault Mode . . . . .</b>	<b>597</b>
	Understanding Fault Mode in Edge Services Director . . . . .	597
	What Are Events and Alarms? . . . . .	597
	Alarm Severity . . . . .	598
	Alarm State . . . . .	598
	Threshold Alarms . . . . .	598
	Understanding the Fault Mode Tasks Pane . . . . .	598
<b>Chapter 34</b>	<b>Viewing and Managing Alarms . . . . .</b>	<b>601</b>
	Changing Alarm State . . . . .	601
	Searching Alarms . . . . .	601
<b>Chapter 35</b>	<b>Alarm Monitor Reference . . . . .</b>	<b>605</b>
	Alarms by State Monitor . . . . .	605
	Alarms by Severity Monitor . . . . .	605
	Current Active Alarms Monitor . . . . .	606

	Alarms by Service Type Monitor . . . . .	607
	Alarm Detail Monitor . . . . .	607
	Finding Specific Alarms . . . . .	608
	Sorting Alarms . . . . .	609
	Reading Events . . . . .	609
	Investigating Event Attributes . . . . .	610
	Changing the Alarm State . . . . .	610
<b>Part 10</b>	<b>System Mode</b>	
<b>Chapter 36</b>	<b>About System Mode . . . . .</b>	<b>613</b>
	Understanding the System Tasks Pane . . . . .	613
	Audit Logs Overview . . . . .	613
<b>Part 11</b>	<b>Appendix</b>	
<b>Chapter 37</b>	<b>Services Overview . . . . .</b>	<b>617</b>
	Adaptive Services Overview . . . . .	617
	Junos Address Aware Network Addressing Overview . . . . .	619
	Packet Flow Through the Adaptive Services or Multiservices PIC . . . . .	620
	ADC Overview . . . . .	622
	Service Instances . . . . .	623
	Installing and Configuring the ADC Software . . . . .	624
	Application-Based Health Checks . . . . .	624
	SSL Server Health Checks . . . . .	624
	DNS Health Checks . . . . .	624
	Ping Health Checks . . . . .	624
	HTTP Health Checks . . . . .	624
	Script-Based Health Checks . . . . .	625
	Script Formats . . . . .	625
	Sample IPv6 Transition Scenarios . . . . .	626
	Example 1: IPv4 Depletion with a Non-IPv6 Access Network . . . . .	627
	Example 2: IPv4 Depletion with an IPv6 Access Network . . . . .	627
	Example 3: IPv4 Depletion for Mobile Networks . . . . .	628
	Understanding Services PICs . . . . .	628
	Adaptive services and Multiservices PICs . . . . .	629
	Encryption Services (ES) PIC . . . . .	629
	Multilink Services and Link Services PICs . . . . .	630
	Monitoring Services PICs . . . . .	630
	Tunnel Services PIC . . . . .	630
	Multiservices MIC and Multiservices MPC . . . . .	630
	TLB Overview . . . . .	631
	TLB Application Description . . . . .	631
	TLB Topology . . . . .	632
	TLB Key Characteristics . . . . .	632
	TLB Application Components . . . . .	633
	Servers and Server Groups . . . . .	633
	Server Health Monitoring — Single Health Check and Dual Health Check . . . . .	633

Virtual Services .....	634
TLB Configuration Limits .....	634
Installing and Configuring TLB Using the CLI Interface .....	634
Configuring a TLB Instance .....	635
Configuring Interface and Routing Information .....	635
Configuring Servers .....	637
Configuring Network Monitoring Profiles .....	638
Configuring Server Groups .....	639
Configuring Virtual Services .....	640
Stateful Firewall Overview for Junos OS Extension-Provider Packages .....	642
Stateful Firewall Support for Application Protocols .....	643
Stateful Firewall Anomaly Checking .....	643
Network Address Translation Configuration Overview .....	645
Configuring Source and Destination Addresses Network Address Translation Overview .....	645
Configuring Pools of Addresses and Ports for Network Address Translation Overview .....	646
Configuring NAT Pools .....	646
Preserve Range and Preserve Parity .....	647
Specifying Destination and Source Prefixes without Configuring a Pool .....	647
Configuring Address Pools for Network Address Port Translation (NAPT) Overview .....	648
Round-Robin Allocation for NAPT .....	648
Sequential Allocation for NAPT .....	649
Preserve Parity and Preserve Range for NAPT .....	649
Address Pooling and Endpoint Independent Mapping for NAPT .....	650
Port Block Allocation for NAPT .....	651
Deterministic Port Block Allocation for NAPT .....	651
Comparision of NAPT Implementation Methods .....	656
Network Address Translation Rules Overview .....	656
Configuring Match Direction for NAT Rules .....	657
Configuring Match Conditions in NAT Rules .....	658
Configuring Actions in NAT Rules .....	658
Configuring Translation Types .....	659
Configuring Service Sets for Network Address Translation .....	661
Junos OS CGNAT Implementation Overview .....	663
Network Address Translation Overview for MS-DPC, MS-MPC, and MS-MIC Line Cards .....	664
Types of NAT .....	664
Inline Network Address Translation Overview for MPC Types 1, 2, and 3 ...	668
CGNAT Implementations Feature Comparison for Junos Address Aware by Type of Interface Card .....	669
ALGs Available by Default for Junos OS Address Aware NAT .....	671
Service Redundancy Daemon Overview .....	674
Introduction to the Service Redundancy Daemon .....	674
Service Redundancy Daemon Components .....	674
Service Redundancy Daemon Constraints .....	675

Service Redundancy Daemon Operation . . . . .	676
Configuring the Service Redundancy Daemon . . . . .	676
Configuring Redundancy Events . . . . .	677
Configuring Redundancy Policies . . . . .	678
Configuring Redundancy Set and Group . . . . .	680
Configuring Routing Policies Supporting Redundancy . . . . .	681
Configuring Service Sets . . . . .	682
Application Layer Gateways Overview . . . . .	683
Supported ALGs . . . . .	683
ALG Support Details . . . . .	684
Basic TCP ALG . . . . .	685
Basic UDP ALG . . . . .	685
BOOTP . . . . .	686
DCE RPC Services . . . . .	686
DNS . . . . .	686
FTP . . . . .	686
H323 . . . . .	687
ICMP . . . . .	687
IIOP . . . . .	688
IP . . . . .	688
NetBIOS . . . . .	688
NetShow . . . . .	688
ONC RPC Services . . . . .	688
PPTP . . . . .	689
RealAudio . . . . .	689
Sun RPC and RPC Portmap Services . . . . .	690
RTSP . . . . .	691
SIP . . . . .	691
SNMP . . . . .	692
SQLNet . . . . .	692
TFTP . . . . .	692
Traceroute . . . . .	692
UNIX Remote-Shell Services . . . . .	693
Winframe . . . . .	693
Juniper Networks Defaults . . . . .	693
Examples: Referencing the Preset Statement from the Junos Default Group . . . . .	704



# List of Figures

<b>Part 1</b>	<b>Overview</b>	
<b>Chapter 1</b>	<b>Edge Services Director Overview</b>	<b>3</b>
	Figure 1: The Edge Services Director User Interface Components	6
	Figure 2: Edge Services Director Banner	7
	Figure 3: Performing Search on the View Pane	10
	Figure 4: Column Drop-Down Menu	12
<b>Chapter 3</b>	<b>Tasks Pane</b>	<b>31</b>
	Figure 5: Alarms Page in Fault Mode	41
<b>Chapter 4</b>	<b>Dashboard</b>	<b>45</b>
	Figure 6: Dashboard Page	46
<b>Part 3</b>	<b>Gateway View of Build Mode</b>	
<b>Chapter 7</b>	<b>Managing Service Delivery Gateways and Groups</b>	<b>85</b>
	Figure 7: Compare Configuration View Page	93
	Figure 8: Service Gateway Groups Page	100
	Figure 9: Service Gateway Details Page	103
<b>Chapter 8</b>	<b>Managing KPI Templates</b>	<b>119</b>
	Figure 10: Network Entry Points	120
	Figure 11: Clone KPI Template Window	122
	Figure 12: KPI Templates Page	130
<b>Chapter 9</b>	<b>Viewing the Device Inventory</b>	<b>133</b>
	Figure 13: Device Inventory Page	134
	Figure 14: Device Count by Platform Report	142
	Figure 15: Device Status Report	143
	Figure 16: Device Count by OS Report	144
<b>Part 4</b>	<b>Location and Device Views of Build Mode</b>	
<b>Chapter 11</b>	<b>Device Management</b>	<b>173</b>
	Figure 17: Device Inventory Page	176
<b>Part 5</b>	<b>Service View of Build Mode</b>	
<b>Chapter 13</b>	<b>Using the Service Designer</b>	<b>185</b>
	Figure 18: Services Inventory Page	194
	Figure 19: Service Template Statistics Page	202
	Figure 20: Create ADC Service Template Window	205

	Figure 21: Select Reference Config Dialog Box . . . . .	233
	Figure 22: Create CGNAT Service Template Window . . . . .	236
	Figure 23: Select Reference Config Dialog Box . . . . .	239
	Figure 24: Create SFW Service Template Window . . . . .	254
	Figure 25: Select Reference Config Dialog Box . . . . .	257
	Figure 26: Create TLB Service Template Window . . . . .	277
	Figure 27: Select Reference Config Dialog Box . . . . .	281
	Figure 28: Select Reference Config Dialog Box . . . . .	301
<b>Chapter 14</b>	<b>Using the Object Builder . . . . .</b>	<b>305</b>
	Figure 29: Object Builder Page . . . . .	307
	Figure 30: Add to Object Builder Dialog Box . . . . .	309
<b>Chapter 15</b>	<b>Managing Packet Analyzers . . . . .</b>	<b>321</b>
	Figure 31: Service Analyzer Instances Page . . . . .	327
<b>Part 6</b>	<b>Deploy Mode</b>	
<b>Chapter 17</b>	<b>Device Management . . . . .</b>	<b>339</b>
	Figure 32: Device Inventory Page . . . . .	340
<b>Chapter 20</b>	<b>Viewing and Editing Service Instances and Packet Filters Across All Gateways . . . . .</b>	<b>363</b>
	Figure 33: Service Edit Page with Pie Charts of Configured Service Types . . . . .	364
<b>Chapter 21</b>	<b>Enhanced Editing of Services and Packet Filters . . . . .</b>	<b>369</b>
	Figure 34: Enhanced Edit Page for Packet Filters . . . . .	371
	Figure 35: Enhanced Edit Page for Service Policy Rules . . . . .	373
<b>Chapter 22</b>	<b>Managing Service Instance and Policy Rule Definitions . . . . .</b>	<b>375</b>
	Figure 36: CGNAT Services Listing Page . . . . .	393
	Figure 37: Stateful Firewall Services Listing Page . . . . .	395
	Figure 38: Create a CGNAT Rule Window . . . . .	397
	Figure 39: Create a Packet Filter Rule Term Window . . . . .	424
	Figure 40: Lock Failure Error Message for the Second User . . . . .	429
	Figure 41: Inactivity Timeout Error . . . . .	430
	Figure 42: Policy Lock Expired Message . . . . .	430
	Figure 43: Packet Filter Policy: Unsaved Changes Message . . . . .	430
	Figure 44: Packet Filter Policy: Policy Unlock by Admin Message . . . . .	430
	Figure 45: Packet Filter Policy Lock Release Message . . . . .	431
	Figure 46: Create SFW Policy Window . . . . .	436
<b>Chapter 23</b>	<b>Managing Packet Analyzers . . . . .</b>	<b>493</b>
	Figure 47: Service Analyzer Instances Page . . . . .	499
<b>Part 7</b>	<b>Deploy Mode</b>	
<b>Chapter 27</b>	<b>Deploying Configurations to Devices . . . . .</b>	<b>527</b>
	Figure 48: Deployment Plans Page . . . . .	531
	Figure 49: Create Deployment Plan Page . . . . .	535
	Figure 50: Enhanced Edit Page for Packet Filters . . . . .	545
	Figure 51: Enhanced Edit Page for Service Policy Rules . . . . .	548

<b>Chapter 28</b>	<b>Viewing Transactions Associated with Deployment Jobs . . . . .</b>	<b>551</b>
	Figure 52: Transactions Page . . . . .	553
<b>Part 8</b>	<b>Monitor Mode</b>	
<b>Chapter 32</b>	<b>General Monitoring . . . . .</b>	<b>577</b>
	Figure 53: Monitoring Page . . . . .	579
	Figure 54: Monitoring Page for ADC Service . . . . .	584
	Figure 55: Monitoring Page for TLB Service . . . . .	586
	Figure 56: Monitoring Page for CGNAT Service . . . . .	589
	Figure 57: Monitoring Page for SFW Service . . . . .	592
<b>Part 9</b>	<b>Fault Mode</b>	
<b>Chapter 33</b>	<b>About Fault Mode . . . . .</b>	<b>597</b>
	Figure 58: Alarms Page in Fault Mode . . . . .	599
<b>Part 11</b>	<b>Appendix</b>	
<b>Chapter 37</b>	<b>Services Overview . . . . .</b>	<b>617</b>
	Figure 59: Packet Flow Through the Adaptive Services or MultiServices PIC . . . .	621
	Figure 60: IPv4 Depletion Solution - IPv4 Access Network . . . . .	627
	Figure 61: IPv4 Depletion Solution - IPv6 Access Network . . . . .	628
	Figure 62: TLB Topology . . . . .	632
	Figure 63: Dynamic NAT Flow . . . . .	667
	Figure 64: Stateful NAT64 Flow . . . . .	668
	Figure 65: Supported Inline NAT Types . . . . .	669



# List of Tables

	<b>About the Documentation</b> . . . . .	<b>xxvii</b>
	Table 1: Notice Icons . . . . .	xxviii
	Table 2: Text and Syntax Conventions . . . . .	xxviii
<b>Part 1</b>	<b>Overview</b>	
<b>Chapter 1</b>	<b>Edge Services Director Overview</b> . . . . .	<b>3</b>
	Table 3: Numerical Sorts and Lexical Sorts . . . . .	12
<b>Chapter 3</b>	<b>Tasks Pane</b> . . . . .	<b>31</b>
	Table 4: Device Discovery Tasks . . . . .	33
	Table 5: Inventory Tasks . . . . .	33
	Table 6: Service Gateway Tasks . . . . .	33
	Table 7: Service Analyzer Tasks . . . . .	33
	Table 8: Device Management Tasks . . . . .	34
	Table 9: Location Management Tasks . . . . .	34
	Table 10: Service Template Tasks . . . . .	36
	Table 11: Object Builder Tasks . . . . .	36
	Table 12: Configuration Deployment Tasks . . . . .	38
	Table 13: Image Management Tasks . . . . .	39
	Table 14: Device Management Tasks . . . . .	39
	Table 15: Device Configuration File Management Tasks . . . . .	39
	Table 16: Service Deployment Tasks . . . . .	39
	Table 17: Service Edit Tasks . . . . .	40
	Table 18: Policy and Filter Tasks . . . . .	40
<b>Part 2</b>	<b>System Administration</b>	
<b>Chapter 5</b>	<b>Handling Administrative Tasks</b> . . . . .	<b>57</b>
	Table 19: Audit Logs Page Fields . . . . .	58
	Table 20: Job Management Page Fields . . . . .	60
	Table 21: Log Files in the troubleshooting.zip File . . . . .	61
<b>Part 3</b>	<b>Gateway View of Build Mode</b>	
<b>Chapter 6</b>	<b>About Gateway View of Build Mode</b> . . . . .	<b>65</b>
	Table 22: Job Management Page Fields . . . . .	82
<b>Chapter 7</b>	<b>Managing Service Delivery Gateways and Groups</b> . . . . .	<b>85</b>
	Table 23: Fields on the Service Gateway Details Page . . . . .	103
	Table 24: Fields in the Last Execution Status Dialog Box . . . . .	107
<b>Chapter 8</b>	<b>Managing KPI Templates</b> . . . . .	<b>119</b>

	Table 25: ADC Tab . . . . .	123
	Table 26: TLB Tab . . . . .	124
	Table 27: CGNAT Tab . . . . .	124
	Table 28: SFW Tab . . . . .	125
	Table 29: Chassis Tab . . . . .	125
	Table 30: HA Tab . . . . .	128
	Table 31: KPI Templates View . . . . .	131
<b>Chapter 9</b>	<b>Viewing the Device Inventory . . . . .</b>	<b>133</b>
	Table 32: Managed Status Pie Chart . . . . .	134
	Table 33: Fields Under the Gateway Tab . . . . .	135
	Table 34: Fields Under the Hardware Tab . . . . .	136
	Table 35: Fields Under the Interface Tab . . . . .	140
<b>Part 4</b>	<b>Location and Device Views of Build Mode</b>	
<b>Chapter 10</b>	<b>Location View Configuration . . . . .</b>	<b>151</b>
	Table 36: Contents of Selected Device Details . . . . .	158
	Table 37: Add or Edit Building Fields . . . . .	159
	Table 38: Floor Field Descriptions . . . . .	161
	Table 39: Outdoor Area Fields . . . . .	162
	Table 40: Site Creation Fields . . . . .	163
	Table 41: Devices that can be Assigned to each Location Component . . . . .	168
<b>Chapter 11</b>	<b>Device Management . . . . .</b>	<b>173</b>
	Table 42: Fields in the Device Inventory Table . . . . .	177
	Table 43: Fields in the Device Physical Inventory Table . . . . .	178
<b>Part 5</b>	<b>Service View of Build Mode</b>	
<b>Chapter 13</b>	<b>Using the Service Designer . . . . .</b>	<b>185</b>
	Table 44: Service Designer View . . . . .	193
	Table 45: Fields on the Services Page . . . . .	194
	Table 46: Hash Keys Supported for AMS for Service Applications . . . . .	245
	Table 47: Hash Keys Supported for AMS for Service Applications . . . . .	263
<b>Part 6</b>	<b>Deploy Mode</b>	
<b>Chapter 17</b>	<b>Device Management . . . . .</b>	<b>339</b>
	Table 48: Fields in the Device Inventory Table . . . . .	341
	Table 49: Resynchronize Device Configuration Fields . . . . .	343
<b>Chapter 18</b>	<b>Configuration File Management . . . . .</b>	<b>347</b>
	Table 50: Manage Device Configuration Table . . . . .	348
	Table 51: Job Management Page Fields . . . . .	351
<b>Chapter 19</b>	<b>Software Image Management . . . . .</b>	<b>353</b>
	Table 52: Device Image Repository Table . . . . .	354
	Table 53: Device Image Summary Window . . . . .	355
	Table 54: Select images for devices Table . . . . .	357
	Table 55: Image Management Job Options . . . . .	358

	Table 56: Image Deployment Jobs Table . . . . .	360
	Table 57: Device Image Staging Window Description . . . . .	361
<b>Chapter 21</b>	<b>Enhanced Editing of Services and Packet Filters . . . . .</b>	<b>369</b>
	Table 58: Service Edit > Packet Filter Page . . . . .	371
	Table 59: Services – CGNAT and SFW Page . . . . .	373
<b>Chapter 22</b>	<b>Managing Service Instance and Policy Rule Definitions . . . . .</b>	<b>375</b>
	Table 60: Service Edit > ADC Page . . . . .	391
	Table 61: TLB Service Edit Page . . . . .	391
	Table 62: CGNAT Policy and Filter Page . . . . .	392
	Table 63: Packet Filter Page . . . . .	393
	Table 64: SFW Policy and Filter Page . . . . .	394
	Table 65: Hash Keys Supported for AMS for Service Applications . . . . .	401
	Table 66: Hash Keys Supported for AMS for Service Applications . . . . .	440
	Table 67: CGNAT Service Edit Page . . . . .	457
	Table 68: SFW Service Edit Page . . . . .	458
	Table 69: ADC Service Edit Page . . . . .	460
	Table 70: Fields in the Manage Instance Locks Dialog Box . . . . .	470
	Table 71: TLB Service Edit Page . . . . .	473
	Table 72: Fields in the Manage Instance Locks Dialog Box . . . . .	482
<b>Part 7</b>	<b>Deploy Mode</b>	
<b>Chapter 25</b>	<b>Configuration File Management . . . . .</b>	<b>511</b>
	Table 73: Manage Device Configuration Table . . . . .	512
	Table 74: Job Management Page Fields . . . . .	515
<b>Chapter 26</b>	<b>Software Image Management . . . . .</b>	<b>517</b>
	Table 75: Device Image Repository Table . . . . .	518
	Table 76: Device Image Summary Window . . . . .	519
	Table 77: Select images for devices Table . . . . .	521
	Table 78: Image Management Job Options . . . . .	522
	Table 79: Image Deployment Jobs Table . . . . .	524
	Table 80: Device Image Staging Window Description . . . . .	525
<b>Chapter 27</b>	<b>Deploying Configurations to Devices . . . . .</b>	<b>527</b>
	Table 81: Service Edit > Packet Filter Page . . . . .	546
	Table 82: Services – CGNAT and SFW Page . . . . .	548
<b>Part 8</b>	<b>Monitor Mode</b>	
<b>Chapter 30</b>	<b>Using Fault Management Monitors . . . . .</b>	<b>561</b>
	Table 83: Alarms Tab Fields . . . . .	564
	Table 84: Events Tab Fields . . . . .	564
	Table 85: Syslogs Tab Fields . . . . .	565
<b>Part 9</b>	<b>Fault Mode</b>	
<b>Chapter 34</b>	<b>Viewing and Managing Alarms . . . . .</b>	<b>601</b>
	Table 86: Alarm Search Fields . . . . .	602

<b>Chapter 35</b>	<b>Alarm Monitor Reference . . . . .</b>	<b>605</b>
	Table 87: Current Active Alarms Monitor . . . . .	606
	Table 88: Alarm Detail Fields . . . . .	608
	Table 89: Sort Options for Alarms . . . . .	609
	Table 90: Event Detail Fields . . . . .	609
<b>Part 10</b>	<b>System Mode</b>	
<b>Chapter 36</b>	<b>About System Mode . . . . .</b>	<b>613</b>
	Table 91: System Tasks . . . . .	613
<b>Part 11</b>	<b>Appendix</b>	
<b>Chapter 37</b>	<b>Services Overview . . . . .</b>	<b>617</b>
	Table 92: TLB Configuration Limits . . . . .	634
	Table 93: Deterministic Port Block Allocation Commit Constraints . . . . .	655
	Table 94: Comparison of NAT Implementation Methods . . . . .	656
	Table 95: CGNAT Implementation—Feature Comparison by Platform . . . . .	669
	Table 96: CGNAT Translation Types . . . . .	671
	Table 97: ALGs Available by Default . . . . .	672
	Table 98: ALGs Supported by Junos OS . . . . .	683
	Table 99: RealAudio Product Port Usage . . . . .	689
	Table 100: Supported RPC Services . . . . .	690

# About the Documentation

- Documentation and Release Notes on page xxvii
- Documentation Conventions on page xxvii
- Documentation Feedback on page xxix
- Requesting Technical Support on page xxx

## 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 xxviii 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 xxviii defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
<b>Bold text like this</b>	Represents text that you type.	To enter configuration mode, type the <b>configure</b> command:  user@host> <b>configure</b>
Fixed-width text like this	Represents output that appears on the terminal screen.	user@host> <b>show chassis alarms</b>  No alarms currently active
<i>Italic text like this</i>	<ul style="list-style-type: none"> <li>Introduces or emphasizes important new terms.</li> <li>Identifies guide names.</li> <li>Identifies RFC and Internet draft titles.</li> </ul>	<ul style="list-style-type: none"> <li>A policy <i>term</i> is a named structure that defines match conditions and actions.</li> <li><i>Junos OS CLI User Guide</i></li> <li>RFC 1997, <i>BGP Communities Attribute</i></li> </ul>
<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@# <b>set system domain-name</b> <i>domain-name</i>

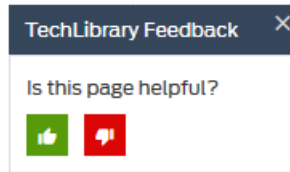
Table 2: Text and Syntax Conventions (continued)

Convention	Description	Examples
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"><li>To configure a stub area, include the <b>stub</b> statement at the <b>[edit protocols ospf area area-id]</b> hierarchy level.</li><li>The console port is labeled <b>CONSOLE</b>.</li></ul>
< > (angle brackets)	Encloses optional keywords or variables.	<b>stub &lt;default-metric <i>metric</i>&gt;;</b>
(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.	<b>broadcast   multicast</b>  <b>(<i>string1</i>   <i>string2</i>   <i>string3</i>)</b>
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	<b>rsvp { # Required for dynamic MPLS only</b>
[ ] (square brackets)	Encloses a variable for which you can substitute one or more values.	<b>community name members [ <i>community-ids</i> ]</b>
Indentation and braces ( { } )	Identifies a level in the configuration hierarchy.	<pre>[edit] routing-options {   static {     route default {       nexthop <i>address</i>;       retain;     }   } }</pre>
;(semicolon)	Identifies a leaf statement at a configuration hierarchy level.	
GUI Conventions		
Bold text like this	Represents graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none"><li>In the Logical Interfaces box, select <b>All Interfaces</b>.</li><li>To cancel the configuration, click <b>Cancel</b>.</li></ul>
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select <b>Protocols&gt;Ospf</b> .

## 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](mailto: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/>.



## PART 1

# Overview

- [Edge Services Director Overview on page 3](#)
- [Getting Started on page 21](#)
- [Tasks Pane on page 31](#)
- [Dashboard on page 45](#)



## CHAPTER 1

# Edge Services Director Overview

- [Understanding the Need for Edge Services Director on page 3](#)
- [Understanding Edge Services Director User Administration on page 5](#)
- [Understanding the Edge Services Director User Interface on page 6](#)
- [Understanding Edge Services Director and the Management Lifecycle Modes on page 15](#)
- [Service Delivery Gateway Overview on page 17](#)
- [Edge Services Director Overview on page 19](#)

## Understanding the Need for Edge Services Director

---

The service delivery gateway (SDG) orchestration of services, by coordination of traffic flows and service interaction, based on policy and subscriber context, immensely simplifies the service deployment. In addition the consolidation of various components necessary to deliver services (such as carrier-grade NAT [CGNAT], stateful firewall, deep-packet inspection, or stateful load balancing) at scale and allows for simplified and reliable services network architecture. An SDG management application is a key component for simplifying and solving major operational challenges of service delivery that leads to service innovation in future.

The SDG management application called *Edge Services Director* is an operations tool that is primarily used by the service provider operations team for comprehensive and centralized service management across different regions, zones and business units catering to different type of customers. For an enterprise, consumer, voice over Long Term Evolution (VoLTE) mobile networks, and others, the SDG management application is implemented on the Junos Space Network Management application and enables service providers the capability to achieve faster IP service rollouts for business needs and reduce overall operating expense (OPEX) costs for managing the service lifecycle. Currently, SDG users need to plan, configure and debug entire SDG network using the CLI interface.

Operators need extensive training in CLI commands and need to be abreast with the latest syntax and format changes in the CLI commands and configuration stanzas. The CLI method of setup and administration is not well-suited for bulk management and requires a longer time to test, deploy, and maintain large networks. Although SNMP may provide some monitoring, it lacks a thorough management capability for service lifecycle management. Service management is a well-understood concept in enterprise and consumer domains. However, service management in operator and service provider

networks is a combination of complex integration of element management systems, operations support systems (OSS), and business support systems (BSS) in the backend. Edge Services Director has a distinct, potent advantage over other management applications available in the market because it is not vendor-specific and provides a cohesive, seamless management capability. The service provider operations personnel use role-based users and workflows to manage services.

The workflow supports deployment specific methods of operations. Edge Services Director, which is the SDG management application, addresses the following business and operational requirements:

- Cost and time to market for new service—Introduction of new value-added services in a faster and streamlined way to meet evolving business needs.
- Reliability of services and network—Proactive monitoring of SDG traffic flow and various components for root cause analysis and faster resolution.
- Flexibility of making changes—Configuration and reconfiguration of services in a much more controlled and isolated environment.
- OPEX savings—Reduction of OPEX cost related to deploying new services and the training required to operate the service
- OPEX savings—Reduction of OPEX cost related to deploying new services and the training required to operate the service

You can employ SDG management in the following network scenarios:

- Simplified and scalable management of hundreds of SDGs
- Configuration and provisioning of SDG services
- Support for large brownfield deployments
- Greenfield deployment support by enabling service planning and configuration templates

Edge Services Director currently supports only brownfield deployments or provisioning and not greenfield deployments. A greenfield deployment refers to the Junos OS base configurations and bootstrapping, core device settings such as routing instances, interfaces and IP addresses, and routing protocols to be available for configuration using the network management application. A brownfield deployment refers to the basic and mandatory device settings already being configured on the devices before they are imported or discovered for additional modifications, such as configuration of services, using the network management application.

- High availability (HA) support for active and standby systems
- Faster and easier issue resolution to isolate and identify problems, and debugging call flows for troubleshooting
- A searchable and sortable inventory of service instances, service components, and the underlying hardware
- Policy and filter management across service instances on the network
- Scalable video traffic monitoring, such as to monitor multicast or unicast video traffic.

- Single pane of view for service monitoring with key performance indicators (KPIs) and threshold values
- Fault and performance management for scalable logging and data collection and correlation of different data sources
- System image and version management, and management of scripts
- Reporting of service usages

**Related Documentation**

- [Understanding Edge Services Director User Administration on page 5](#)
- [Understanding the Edge Services Director User Interface on page 6](#)
- [Understanding Edge Services Director and the Management Lifecycle Modes on page 15](#)
- [Service Delivery Gateway Overview on page 17](#)
- [Edge Services Director Overview on page 19](#)

---

## Understanding Edge Services Director User Administration

Edge Services Director uses the user administration features of the Junos Space platform on which it runs. Using these features, you can add, delete, and edit user accounts and roles and changing user passwords. Refer to the *Junos Space Network Application Platform User Guide* for more information about user administration.

When Edge Services Director is installed, some additional user administration options are available in Junos Space, which are specific to Edge Services Director.

In addition to the Super Administrator role, the following predefined roles are available to Edge Services Director users:

- Edge Services Director - Administrator—Has complete access to all the Edge Services Director modes and user preferences.
- Edge Services Director - Operator—Has access to all modes except the Build mode. Has access to windows and capabilities, such as fault management, performance management, dashboard and monitoring. You can create custom roles to grant users different access rights to the Edge Services Director modes.
- Edge Services Director - Designer—Has access to the Build mode for handling device and service configuration operations such as creation of services and KPI templates.

You can also create custom roles to grant users different access rights to the Connectivity Services Director modes. Edge Services Director modes—Build, Deploy, Monitor, Fault, and Report modes are available to assign to custom user roles in the list of application workspaces and associated tasks.



**NOTE:** The tasks listed under the Edge Services Director modes do not have any effect. Access is controlled at the mode level, so if you grant a role access to a mode, the role has access to all tasks in that mode, regardless of which tasks you select.

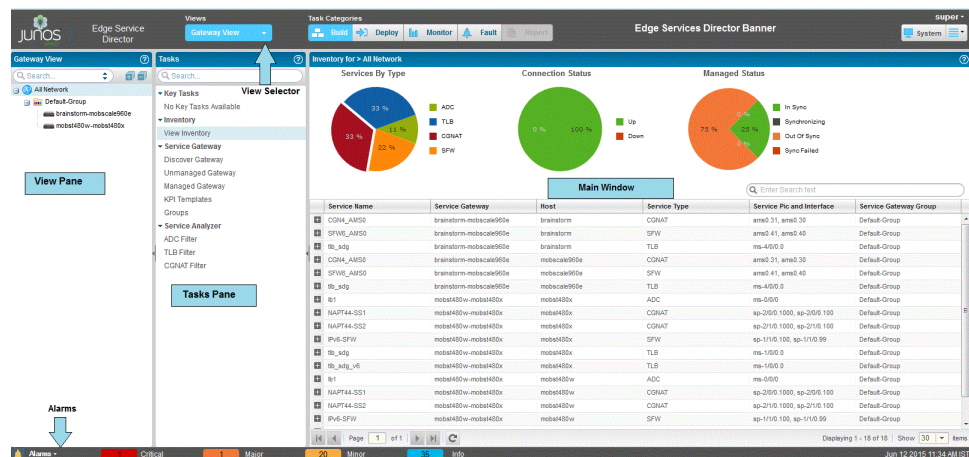
#### Related Documentation

- [Understanding the Need for Edge Services Director on page 3](#)
- [Understanding the Edge Services Director User Interface on page 6](#)
- [Understanding Edge Services Director and the Management Lifecycle Modes on page 15](#)
- [Service Delivery Gateway Overview on page 17](#)
- [Edge Services Director Overview on page 19](#)

## Understanding the Edge Services Director User Interface

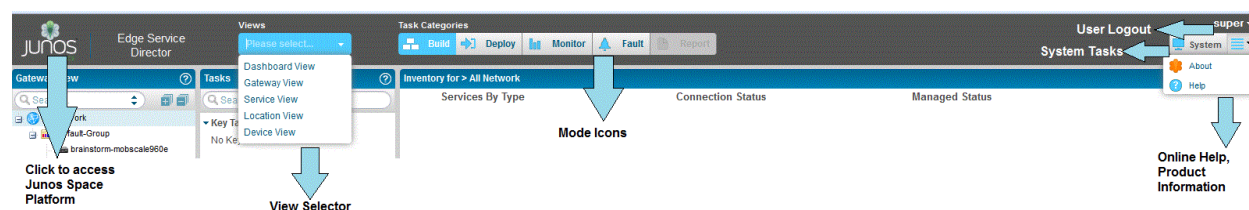
Junos Space Edge Services Director provides a simple-to-use, Web 2.0 user interface that you can access through standard Web browsers. The user interface uses task-based workflows to help you accomplish administrative tasks quickly and efficiently. It provides you with the flexibility to work with single or 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 6 illustrates the main components of the interface.



Use the Edge Services Director banner, shown in [Figure 2 on page 7](#), to select the working mode. You can also use the Edge Services Director banner to perform other global tasks, such as setting up your preferences or accessing Junos Space.

Figure 2: Edge Services Director Banner




The following are the functions of the banner:

- Accessing Junos Space Platform—Click to exit Edge Services Director and open the Junos Space Network Application Platform. You can switch back and forth between Edge Services Director and Junos Space without logging in again.
- View Selector—Select the network view that you want to work in. You can choose from one of the following views:
  - Dashboard View
  - Location View
  - Device View
  - Gateway View
  - Service View
- Mode Icons—Select the mode you want to work in.



**NOTE:** You might not have access to all the Edge Services Director modes. What modes you have access to depends on your assigned user role.

- Login as—Displays the username using which you logged in to Edge Services Director. Click the down arrow next to the username and select the scope of the view, such as global.
- User Log out—Select this icon, which is the rightmost one in the banner, to log out of Edge Services Director and Junos Space.
- Product Information and Online Help —Click the down arrow next to System and select either of the following options:
  - Help—Enables you to 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 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.
- System Tasks and Jobs—Access the system tasks such as viewing audit logs and jobs and collecting troubleshooting logs.

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

## View Pane

In the View pane, Edge Services Director provides you a unified, hierarchal view of your wired, wireless, and data center 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 service delivery gateway (SDG) group in Gateway View, you indicate that the scope for a task is the routers in the SDG group.
- 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 MX240 node in Device View, you indicate that the scope for a task is all MX240 routers in your network.

You can perform the following actions in the View pane:

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

---

### Displaying Devices in Various Network Views

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

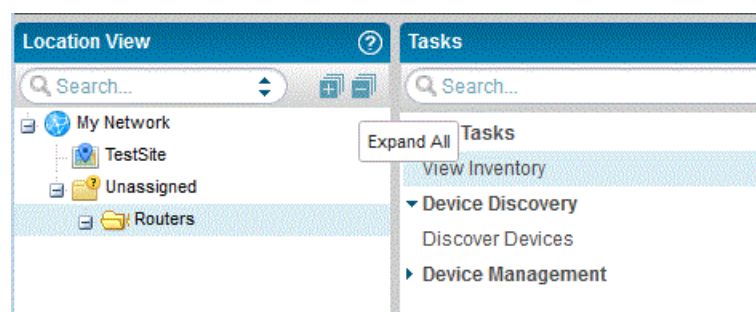
- **Dashboard View**—This is a customizable view that provides information about your network. You can select and add monitoring widgets to the Dashboard View based on your requirements. This is the default view that opens when you log in to Edge Services Director.
- **Location View**—Devices are organized by their physical locations. You build this view by creating sites, buildings, floors, aisles, racks, outdoor areas, and then assigning your routers to these locations.
- **Device View**—Devices are organized by device type, such as routers. Within each device type, devices are organized by device model. For example, all models of MX240 routers are grouped together under one node in the tree.
- **Gateway View**—Service delivery gateways (SDGs) are discovered by SDG discovery workflow. The discovered SDGs are shown in the SDG inventory page. The discovered SDGs can be a part of a high availability (HA) pair or standalone SDGs. You create the network managed by Junos Space Edge Services Director by bringing devices under the administration of the network management application and retrieving the device settings to save in the Edge Services Director database. It provides you with the ability to use device discovery to bring devices under Edge Services Director management,

to customize your view of the devices, to configure devices, and to perform some common device management tasks.

- **Service View**—You can create services, policies, and filters for devices that are managed by Edge Services Director. The service templates and attributes for services, policies, and filters help you classify and control the manner in which packets must be handled by the various services. You can also import objects, which are components or parameters used for creation of services, from the Service Delivery Gateways (SDGs) that are present in the Edge Services Director database or from external XML files.

### 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.

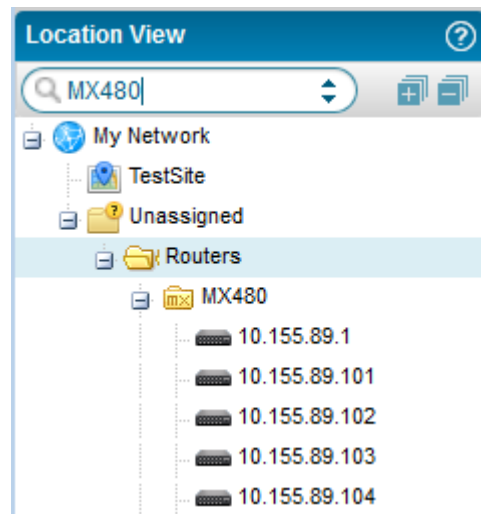
Similarly, to collapse a node in the network tree, select the node and then click the **Collapse All** icon (next to the Expand All icon). The node you selected is collapsed and no nodes under it are shown.

### 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 in the Search box and click the **Search** icon, as shown in [Figure 3 on page 10](#).

Figure 3: Performing Search on the View Pane



Edge Services 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-sensitive: 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. 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 Edge Services Director, you navigate to the task.

## 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 informational 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 bar. You are automatically placed in Fault mode and the Fault mode monitors are displayed.

## Main Window or Workspace

The main window or workspace displays content relevant to the mode, scope, and task you have selected. When you log in to Edge Services Director, the main window displays the dashboard. The dashboard enables you to allow the operators to quickly monitor health and status of the managed devices. The sections or frames on the dashboard allows the operator to understand the device problem or fault at the macro level (comprehensive and widespread network health and status) and the micro level (individual device health and status). The health representation of the devices can be customized based on the monitoring properties defined. You can view all the available

devices that are managed by Edge Services Director from the Device Inventory page. The Device Inventory page is accessible in Device View of Build mode as the default landing page.

## Tables in Edge Services Director

Tables are used throughout Edge Services 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
- Navigating pages
- Displaying the Column Drop-Down Menu
- Sorting on a Column
- Hiding and Exposing Columns
- Searching Table Contents
- Filtering Table Contents

### Moving and Resizing Columns

---

You can reposition and resize columns in a table. To move a column, drag the column head to the new location. Edge Services 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.

### Navigating Pages

---

Paging controls at the bottom of an applicable page allow you to navigate the entries on the pages when the inventory is too large to fit on one page. Using these controls, you can go to a specific page, navigate to the next or previous page, navigate to the first or last page of the inventory, or refresh the inventory view.

### 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 down arrow appears. By clicking the arrow, you display the drop-down menu, as shown in [Figure 4 on page 12](#).

Figure 4: Column Drop-Down Menu

Hostname ↑	IP Address	Serial Number	Platform	OS Version
10.93.213.153		GX0211041838	EX4500-40F	13.1-20130116_cdi_13
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_cdi_13
st-java1021-VC-CORE-1	10.93.202.73	9469818	EX4200-24T	12.3R1.4
st-java1024-J-ACC-1	10.93.202.75	9469695	EX4200-24T	12.3R1.4
st-vent02	10.93.213.193	8520032	EX8216	13.1-20130116_cdi_13
sys-java141	10.93.213.138	BP0208372546	EX4200-48T	13.1-20130116_cdi_13
sys-java20	10.93.213.130	BK0208109492	EX3200-48T	13.1-20130116_cdi_13

### 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.

Edge Services 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 3 on page 12](#).

Table 3: 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 4 on page 12](#). Select the check box beside a column in the drop-down menu to expose it. Clear the check box beside a column to hide it.

As a general rule, Edge Services 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 Edge Services Director to search for the term in all columns in the table. Every table has a predefined default column that the system searches; before it proceeds to search other columns.

You can also choose to search a particular column for a term. Edge Services 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 12.3 Release 4.5, then `os: 12.3R4.5` returns search results, whereas `OS: 12.3R4.5` 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. Edge Services Director uses the AND operator for each search term that you enter. Edge Services 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 MX480 router that is running Junos OS Release 14.1:

- Enter **MX480** as the search term in the text box.
- From the list that appears, select to search the Platform column.

Edge Services Director lists all the MX480 routers in your network.

3. Enter **14.1** as the search term after the comma separator in the text box.
4. From the list, select to search from the OS Version column.

Edge Services Director lists all the MX480 routers in your network that are running Junos OS Release 14.1.

### 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 down 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.

In addition to these functions, Connectivity Services Director displays the date and time in the local time zone on the bottom right corner.

#### Related Documentation

- [Understanding the Need for Edge Services Director on page 3](#)
- [Understanding Edge Services Director User Administration on page 5](#)
- [Understanding Edge Services Director and the Management Lifecycle Modes on page 15](#)
- [Service Delivery Gateway Overview on page 17](#)
- [Edge Services Director Overview on page 19](#)

---

## Understanding Edge Services Director and the Management Lifecycle Modes

---

Junos Space Edge Services Director is a Junos Space application for management of services interfaces of MX Series routers, such as adaptive services interfaces and multiservices interfaces, that provide specific capabilities for manipulating traffic before it is delivered to its destination. Services interfaces enable you to add services to your network incrementally. Providing full network lifecycle management, Edge Services Director simplifies the discovery, configuration, visualization, monitoring, and administration of large networks. Operators 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 Edge Services Director user interface is based on the network management lifecycle. The interface provides five main working modes that are aligned to the network management lifecycle, and a sixth mode for working with Edge Services Director itself. Each mode provides access to different tasks:

- **Build mode**—In Build mode, you can create services, policies, and filters for devices that are managed by Edge Services Director. You can define service templates and attributes of different services. You can also specify policies and filters to classify and control the manner in which packets must be handled by the various services. Configuring a policy has a major impact on the flow of routing information or packets within and through the router.

In Gateway view of Build mode, you create the network managed by Junos Space Edge Services Director by bringing devices under the administration of the network management application and retrieving the device settings to save in the Edge Services Director database. It provides you with the ability to use device discovery to bring devices under Edge Services Director management, to customize your view of the devices, to configure devices, and to perform some common device management tasks. In Device view of Build mode, you can perform software upgrades to devices and perform several device management and configuration file management tasks. You can also back up the Edge Services Director database that contains all the configuration parameters of devices, settings that enable monitoring and management of devices and services, and reports that contain statistics and graphs of the tracked system states. You can restore the data backed up to a different server that runs the Edge Services Director application.

- **Deploy mode**—Deploy mode enables you to deploy configuration changes to devices. You can create a deployment plan for each of the service planning templates, such as the ones defined for ADC or stateful firewall (SFW) services, and the policy or filter templates, such as the packet filter or SFW policy, that you have created. A deployment plan contains details about the settings and configuration parameters that must be propagated and provisioned on the SDGs managed by Edge Services Director. You can also create, update, display, publish, and commission packet filters, stateful firewall policies, and CGNAT policies present on discovered and managed SDGs.
- **Monitor mode**—Monitor mode in Edge Services Director provides visibility into the behavior and performance of your network. Edge Services 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.

- **Fault mode**—Fault mode shows you information about the health of your network and changing conditions of your equipment. Use Fault mode to identify problems with equipment, pinpoint security attacks, or analyze trends and categories of errors. Edge Services Director correlates traps, which describe a condition, into an alarm. Alarms are ranked by their impact on the network.
- **Report mode**—Use Report mode to generate reports from the data that Edge Services Director stores about network performance, status, and activity. In Report mode, 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, and trend and capacity analysis, and generates records that can be archived for compliance requirements.
- **View pane**—On the View pane, Edge Services Director provides you with a unified, hierarchical view of your wired, wireless, and virtual networks in the form of a expand tree that is expandable and collapsible. You can choose from five views, or perspectives, of your network—Dashboard view, Location view, Device view, Gateway view, and Service view. By selecting both a view and a node from the tree, you indicate the scope over which you want an operation or task to occur. The Dashboard view provides a summary, encompassing a pictorial representation of the health and performance of devices and services in your network, which enables you to analyze and troubleshoot the parameters that are causing traffic-handling errors.

System-level tasks include viewing the Edge Services Director user and system audit trail, managing jobs, and gathering logs for troubleshooting. The dashboard enables you to allow the operators to quickly monitor health and status of the managed service delivery gateways (SDGs) through several widgets and monitors. The sections or frames on the dashboard allows the operator to understand the device problem or fault at macro level (comprehensive and widespread network health and status) to micro level (individual SDG health and status). The health representation of the SDGs can be customized based on the monitoring properties defined in SDG templates.

#### Related Documentation

- [Understanding the Need for Edge Services Director on page 3](#)
- [Understanding Edge Services Director User Administration on page 5](#)
- [Understanding the Edge Services Director User Interface on page 6](#)
- [Service Delivery Gateway Overview on page 17](#)
- [Edge Services Director Overview on page 19](#)

---

## Service Delivery Gateway Overview

---

The service delivery gateway (running on the MX Series 3D Universal Edge router) consolidates a variety of best-in-class Gi ("i" for Internet or IP network) network services onto a single platform to reduce cost, increase network resiliency, and increase performance. The Gi interface is the connection between a GGSN and the Internet or destination networks connected to a public land mobile network (PLMN). Costs are reduced by using less rack space, less hardware, reduced power and cooling, less cabling, and simplified network management. Resiliency is increased by leveraging the redundancy features of the MX Series 3D routers and by limiting the number of different boxes and OS types that must be managed. Performance is increased by taking advantage of the ability of the MX Series 3D Universal Edge routers to perform many services at line rate in hardware.

The MX Series routers provide industry-leading packet forwarding performance along with a very compelling set of value added services that include carrier grade NAT, firewall, intrusion prevention service, video optimization, server load balancing, MPLS VPN, IPsec VPN and much more. Many of these services can be performed at line rate by leveraging the Trio chipset on the Packet Forwarding Engines. This makes the MX Series routers an ideal and robust platform upon which to host the service delivery gateway.

The following sections describe some of the services that are required on the service delivery gateways:

- [Carrier-Grade NAT on page 17](#)
- [Firewalls and Intrusion Prevention System on page 18](#)
- [Traffic Direct on page 18](#)
- [Load Balancing and Adaptive Services on page 18](#)

### Carrier-Grade NAT

Carrier-grade network (CGN) is rapidly increasing in importance now that the Internet Assigned Numbering Authority (IANA) has run out of IPv4 addresses. Some operators are already moving to IPv6, but this does not solve the IPv4 exhaust problem because most of the Internet is still only reachable via an IPv4 address. The answer for many mobile operators that are faced with rapid smart phone growth is carrier grade NAT. Juniper Networks provides a complete implementation of CGN on the service delivery gateway. In the mobile operator's domain, the IPv4 address exhaust problem is more severe than in the wireline world because there is exponential growth, and because of the move to always-on connectivity with smartphones. Always-on connectivity indicates that the subscriber has a session and an IP address even when the device is idle. The two approaches that have received the most attention in the mobile world are dual stack and IPv6-only. Dual stack allows the mobile device to access content that is either IPv6 or IPv4 addressable. To make this work in a seamless fashion, the mobile device must have both an IPv4 and an IPv6 session up and active at the same time. This is possible beginning in 3GPP Release 8. The other approach that is receiving a lot of attention is IPv6-only. In this implementation, the mobile device establishes a single IPv6 session, and traffic headed for the IPv4 Internet is translated using NAT64. The drawback is that

the mobile device must use IPv6 native applications. Problems during roaming might also occur, and the device does not work on most Wi-Fi networks, which can be mitigated by using IPv4 when connecting to Wi-Fi.

## Firewalls and Intrusion Prevention System

Firewalls are an essential part of any mobile Gi network that connects to the Internet. In many cases, firewalls are also tightly linked with the CGN function. Some operators require a dedicated security device and Juniper Networks SRX Series Services Gateways provides rich firewall services with industry leading performance and scale. The service delivery gateway, in combination with the SRX Series, allows Juniper Networks to address a wide variety of deployment scenarios. Intrusion prevention system (IPS) takes the firewall concept one step further by analyzing traffic using deep packet inspection (DPI) to identify threat signatures. Juniper's library of threat signatures is constantly upgraded to handle the latest security vulnerabilities. The primary focus of a firewall and IPS function on the Gi network is to prevent attacks from being launched against the mobile network and mobile users from hosts out on the Internet.

## Traffic Direct

An essential part of any mobile packet core design is the method by which data traffic is steered as it moves from the mobile device through to the correct GGSN, and from there to the correct Gi network. Access point nodes (APNs) are the traditional solution to the problem, but they can be administratively complex. Not only must mobile devices be configured with the correct APN, but so must the network infrastructure (SGSNs, DNS, and GGSNs). Juniper Networks has developed Traffic Direct as an alternative to APNs. This is a much simpler solution to the challenge of making sure that users get where they need to go. The Traffic Direct feature sits on a service delivery gateway and can steer traffic from the GGSN to the correct Gi network. There are several instantiations of Traffic Direct, of which the most popular is Static Bypass Traffic Direct. This feature makes use of the service delivery gateway's policy routing feature. The service delivery gateway is capable of routing on any of the elements of the IP header which include the source and destination IP addresses, source and destination port numbers, and protocol type. Forwarding is handled in hardware at line rate by the Juniper Networks Junos Trio chipset. This approach is a simple way of guaranteeing that all users get to the correct Gi network.

## Load Balancing and Adaptive Services

The service delivery gateway services umbrella leverages the Multiservices-Dense Port Concentrator (MS-DPC), in-house Junos services, the Junos Software Development Kit (SDK) and external third-party platforms and applications. Offered services can run in standalone mode or can be consolidated (chaining with next hop routing), as long as the chained combination is meaningful, to concurrently run in the same chassis or blade. Scaling is achieved by adding MS-DPC blades in the chassis. The combination of consolidated services also dictates the number of MS-DPC blades to be used. Needed services that are not directly hosted by the service delivery gateway are collocated with the service delivery gateway within the different service complexes to provide specific value-added services. As an example, such service complexes include the user equipment (UE) DNS service complex and Juniper Networks Mobile Video Optimization service complex.

Some of these service complex functions can be integrated by leveraging Junos SDK capabilities. Service complexes and packet gateways (such as GGSN and PGW) attach to active or standby service delivery gateway in VRRP groups leveraging MC-LAG. The services delivery gateway pair connects to the core routers using LAG. Services delivery gateway can be deployed to act as a CE or a PE router, with BFD enabled. Server load balancing (SLB) towards service complexes is performed using ECMP. RPM probes are configured to provide server status updates in the complex. An event script or an operational script can be leveraged to take appropriate action upon detection of a status change.

Leveraging adaptive delivery controllers (ADC) from MS-DPC is another possible avenue. ADCs for the MX Series 3D Universal Edge Router offers advanced router-integrated ADC functions that enables service providers and enterprises to efficiently scale service capacity and increase service performance.

Configuring load balancing requires an aggregated Multiservices (AMS) system. AMS involves grouping several Multiservices PICs together. An AMS configuration eliminates the need for separate routers within a system. The primary benefit of having an AMS configuration is the ability to support load balancing of traffic across multiple services PICs.

**Related  
Documentation**

- [Understanding the Need for Edge Services Director on page 3](#)
- [Understanding Edge Services Director User Administration on page 5](#)
- [Understanding the Edge Services Director User Interface on page 6](#)
- [Understanding Edge Services Director and the Management Lifecycle Modes on page 15](#)
- [Edge Services Director Overview on page 19](#)

---

## Edge Services Director Overview

Service providers are increasingly using IP Layer 3 through Layer 7 services to differentiate themselves from third-party, external, over-the-top (OTT) providers and provide better customer experience. These IP services are used to provide better end user experience by managing traffic flow per application type, better security, better video quality and other enhanced IP applications. OTT providers are consuming service provider resources and therefore, value-added IP service is the way optimize and offer best returns for network investment.

Juniper Networks service delivery gateway (SDG) is a next-generation services solution framework to address these set of converging, simultaneous challenges that mobile and wireline operators currently face in delivering services. The SDG primarily is a service orchestration solution, which is based on subscriber and service policy contexts to coordinate the traffic flow between services in on-the-box or off-the-box scenarios and also with Juniper Networks devices or third-party devices. In addition, SDG consolidates the common elements necessary to deliver services at scale, such as carrier-grade NAT (CGN), stateful firewall, deep-packet inspection, or stateful load balancing. This mechanism simplifies and accelerates delivery and introduction of new services offering while keeping changes to network and other services at a minimum.

The SDG is equipped to be positioned upstream from a broadband gateway, cable modem termination system (CMTS), or any other aggregation point such as Packet Data Network Gateway (P-GW) on the Gi interface in the network where centralized services can be applied via an IP address or subscriber database. These services on Juniper service cards are configured in a service chain for a specific packet flow to meet a customer's business and network requirement. In addition, policies and filters associated with these services are modified and updated as business evolves and requirement changes. In the future, the service chain of service instances on different service cards becomes a necessity because packet flow within single platform depends on services configured in the platform. Also, tethered services on virtual instances attached to MX Series routers and other Juniper Networks platforms can be complicated.

Over a period of time, as the number of services increase and configuration too increases correspondingly, deployment and management of IP services become dynamic, complex and unmanageable. Configuration and deployment using configuration files and statements through the CLI interface is prone to human errors. Therefore, a robust and a comprehensive GUI-based service management application is required to automate management and monitoring tasks. Amazon Web Services (AWS) enables the simplification of compute, storage, and network management using easy-to-use web application increases customer acquisition and adoption of management apps. SDG reduces operational (OPEX) costs by abstracting and simplifying complex low-level CLI commands and scripts. The services management system called *Edge Services Director* enables faster time to market and better customer experience. Edge Services Director simplifies and enables dynamic SDG Service planning, configuration, and provisioning of settings on MX Series routers so that users can respond to market condition faster with lower OPEX and overhead costs.

**Related  
Documentation**

- [Understanding the Need for Edge Services Director on page 3](#)
- [Understanding Edge Services Director User Administration on page 5](#)
- [Understanding the Edge Services Director User Interface on page 6](#)
- [Understanding Edge Services Director and the Management Lifecycle Modes on page 15](#)
- [Service Delivery Gateway Overview on page 17](#)

## CHAPTER 2

# Getting Started

- [Understanding How to Use the Edge Services Director Interface to View System Information on page 22](#)
- [Getting Started Assistant in Junos Space Platform Overview on page 23](#)
- [Changing Your Password for Edge Services Director on page 24](#)
- [Logging In to Edge Services Director on page 25](#)
- [Logging Out of Edge Services Director on page 27](#)
- [Quickly Accessing Important Monitoring and Troubleshooting Details on page 28](#)

## Understanding How to Use the Edge Services Director Interface to View System Information

---

When you log in to the Edge Services Director application, the initial default page that is displayed is the Dashboard page. The dashboard functionality allows the operators to quickly identify, understand and monitor the health and status of the service delivery gateways (SDGs). The SDG network management application or Edge Services Director tries to simplify the complexity involved in monitoring the health and status of SDGs deployed across networks through following components and visual representation. The dashboard gadget enables you to understand the issue at macro level (overall network health and status) to micro level (an individual SDG health and status). The health representation of the SDGs can be customized based on the monitoring properties defined in SDG templates.

The SDG service Dashboard and Monitoring pages in the Edge Services Director GUI provide a proactive account of the SDG health status and working efficiency of service delivery gateway (SDG) devices in a bird's eye, comprehensive, and intuitive format at the network level, SDG instance, and service levels. A single pane of glass (SPOG) view helps the operator to view various alarms and quickly identify and isolate issues. The dashboard and monitoring feature aggregates and correlates data from different sources such as SNMP and system event logs. The defined key performance indicators (KPIs) and threshold values enable operators to specify monitoring criteria critical for service operations and administration. The performance management view also highlights the top or first three non-confirming SDGs and provides a historical context with time graph and additional data from the logging system.

The following are the salient capabilities and benefits of the dashboard view of SDGs:

- A single pane of glass (SPOG) view of entire SDG deployment
- Configuration of KPI templates to enable the monitoring of health status
- Display of service name, service status, alarm status, and heat map
- A summary of a list of object counts by types
- Chassis view with service status overlay
- Logical view of selected service with component data, such as ingress and egress direction
- High-priority log tickers
- Proactive SNMP traps and alarms, syslogs, and KPI thresholds
- Near real-time CPU and memory usage graph per core and service
- Customized Dashboard view for user roles and profiles
- Performance view with three top non-confirming SDGs CPU, memory, and service status
- Performance view with selected KPI filters
- Comparison graph view between operator-selected SDGs

- The total count for alarm type, status and others
- System health status at SDG instance and component level Hardware details and hardware status within a SDG

#### Related Documentation

- [Getting Started Assistant in Junos Space Platform Overview on page 23](#)
- [Changing Your Password for Edge Services Director on page 24](#)
- [Logging In to Edge Services Director on page 25](#)
- [Logging Out of Edge Services Director on page 27](#)
- [Quickly Accessing Important Monitoring and Troubleshooting Details on page 28](#)

## Getting Started Assistant in Junos Space Platform Overview


In the Junos Space Platform user interface, the Getting Started assistant is a section in the sidebar that shows you how to perform common tasks. The tasks in the Getting Started assistant are workspace specific. The tasks displayed in this section vary according to the workspace. The Getting Started assistant provides instructions on how to perform tasks related to a device, service template, or a policy and filter template configuration.

The Getting Started topics are context-sensitive per application. Getting Started displays all the steps of a task. From a step in a task, you can jump to that point in the user interface to actually complete it. If **Show Getting Started on Startup** checkbox is selected, the Getting Started assistant automatically displays the tasks when you log in. If this checkbox was not selected, click the **Help** icon and click **Getting Started** from the resulting sidebar.

To use a Getting Started assistant:

1. Select an application from the **Applications** list above the task tree.
2. In the sidebar, expand **Getting Started**.

A main Getting Started topic link appears on the sidebar.

If the sidebar is not displayed, select the **Help** (  ) icon at the right side of the Junos Space header. The sidebar appears.

3. Select a main topic.

For example, if you are in the Network Management Platform application user interface, click the **Increase Space Capacity** link. A list of required steps appears in the sidebar. Each step contains a task link and a link to Help.

4. Perform a specific step by clicking the link.

You jump to that point in the user interface. The assistant remains visible on the sidebar to aid navigation to subsequent tasks.

5. Access help for a specific step by clicking the Help icon next to that step.

**Related  
Documentation**

- [Understanding How to Use the Edge Services Director Interface to View System Information on page 22](#)
- [Changing Your Password for Edge Services Director on page 24](#)
- [Logging In to Edge Services Director on page 25](#)
- [Logging Out of Edge Services Director on page 27](#)
- [Quickly Accessing Important Monitoring and Troubleshooting Details on page 28](#)

---

## Changing Your Password for Edge Services Director

---

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

Your username and password are the same in Junos Space and Edge Services Director.

To change your password:

1. From the Edge Services Director user interface, click the Junos Space icon on the Edge Services Director banner.  
The Junos Space Platform user interface is displayed.
2. Click the **User Settings** icon on the Junos Space banner.  
The **Change User Settings** dialog box appears.
3. In the **Old Password** text box, enter your old password.



**NOTE:** Mouse over the information icon (small blue *i*) next to the **New Password** text box to view the rules for password creation. For more information about the password rules, see *Modifying Junos Space Network Management Platform Settings*.

4. In the **New Password** text box, enter your new password. The minimum value for this field is 6 (the default) and the maximum value is 999. The password can include alphanumeric and special characters, but not control characters.
5. In the **Confirm Password** text box, enter your new password again to confirm it.



**NOTE:** The fields on the X.509 Certificate tab are applicable when you want to use certificate-based authentication. If you are using password-based authentication, you can ignore these fields. For more information about certificate-based authentication, see the *Certificate Management Overview* topic in the *Junos Space Network Management Platform Workspaces Feature Guide*.

6. (Optional) Select the **Manage objects from all assigned domains** check box on the **Object Visibility** tab to view and manage objects from all the domains for which you are assigned.

7. Click **OK**.

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

#### Related Documentation

- [Logging In to Edge Services Director on page 25](#)
- [Logging Out of Edge Services Director on page 27](#)

## Logging In to Edge Services Director

You connect to Edge Services Director using your Web browser. The following Web browsers are supported: Internet Explorer 9.0 and 10.0, Mozilla Firefox 3.6 or later, and Google Chrome 17 and later. The minimum screen resolution is 1280 x 1024.

You can connect to Edge Services Director in either of the following ways:

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

```
https://<n.n.n.n>/mainui/?appName=SGD
```

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

Enter the login credentials. After successful login, the Dashboard page of Edge Services Director is displayed.

- 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.

The Junos Space Platform login page is displayed.

To enter the login credentials and open the Junos Space Platform page:

1. From the Junos Space Platform login page, in the **Username** text box, enter your username. For information about how to change your username, consult your system administrator.
2. In the **Password** text box, enter your password. For information about how to change your password, see [“Changing Your Password for Edge Services Director” on page 24](#).
3. (Optional) If the remote authentication server is configured for Challenge/Response, you are presented with the challenge questions. Provide valid responses to the challenge questions you are asked, to log in successfully.
4. Click **Log In**.

The Junos Space home page appears. If the home page is not set, the Junos Space Dashboard page is displayed. If the home page is inaccessible due to role or domain restrictions, a warning message is displayed and the Junos Space Dashboard page is loaded.



**NOTE:** If you are a user with access to more than one domain, then an informational message about switching domains is displayed in a dialog box.

Do one of the following:

- To prevent the informational message from appearing again, ensure that the **Don't show again** check box is selected and click OK. The **Don't show again** check box is selected by default.
  - To allow the informational message to continue appearing, clear the **Don't show again** check box and click OK.
- 

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

The default username and password are the same for both Junos Space and Edge Services Director:

- Username—super
- Password—juniper123

**Related  
Documentation**

- [Changing Your Password for Edge Services Director on page 24](#)
- [Logging Out of Edge Services Director on page 27](#)

---

## Logging Out of Edge Services Director

---

After you finish using Edge Services Director, log out to prevent unauthorized access. You can log out manually or set an automatic logout period for Edge Services Director to automatically log you out.

**Logging out manually**—To log out of Edge Services Director manually, click the down arrow next to the username on the Edge Services Director banner and select Logout from the list.

**Logging out automatically**—Edge Services Director automatically logs you out if you have not performed any action on it, such as by using 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. You can change the setting on the Applications inventory page. Select **Administration > Applications > Network Management Platform > Modify Application Settings** (from the Actions menu) > **User**.

Edge Services Director uses the same automatic logout period as Junos Space.

To change the automatic logout period:

1. Click the System Platform icon on the Edge Services Director banner.  
The logout page appears.
2. Click the **Click here to log in again** link on the logout page to log in to the system again.
3. Navigate to **Administration > Applications**.  
The Applications page is displayed.
4. Right-click **Network Management Platform** and select **Modify Application Settings**.  
The Modify Application Settings page appears.
5. In the Modify Network Management Settings page, select **User**.  
The User page is displayed.
6. In the **Automatic logout after inactivity (minutes)** field, move the slider to modify the automatic logout setting.  
The logout setting is modified.
7. Click **Modify** to save the setting.  
You are returned to the Modify Applications page.

- Related Documentation**
- [Changing Your Password for Edge Services Director on page 24](#)
  - [Logging In to Edge Services Director on page 25](#)

## Quickly Accessing Important Monitoring and Troubleshooting Details

In the task pane, the top part of the navigation tree enables you to select the options corresponding to different activities that you can perform on the devices and services that are managed by Edge Services Director, and configuration settings you can specify. The bottom bar of the GUI enables you to view critical, salient information about the configured devices and services in an intuitive, easily-navigable format. The summarized way in which you can view statistical details enables you to examine the health and operating-efficiency of devices, and the performance of services. It provides a bird's eye, high-level view of parameters that enables effective and simplified troubleshooting and administration. For example, if you find that a particular Service Delivery Gateway (SDG) or and SDG group has recorded a large number of critical or major alarms, you can then navigate to the Monitoring page or the appropriate device settings page to correct and modify the attributes or diagnose the problems that might be generating the alarms.

You can view the following types of details from the quick-access facility that is displayed at the bottom left corner of the Edge Services Director GUI pages in all lifecycle modes and views:

- Service analyzer details
- Status of deployment plans
- Alarms recorded on the devices

To view alarm details:

1. From the bottom left corner of the screen, click the down arrow and select the **Alarms** option from the shortcut menu. By default, the Alarms option is selected.

The Network Alarms table is displayed. A list of alarms classified by severity levels is displayed. The following table describes the alarms displayed:

- 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.

To view the status of deploy plans:

1. From the bottom left corner of the screen, click the down arrow and select the **Deployment** option from the shortcut menu.

The Deployment Plan Status table is displayed. The following describes the deploy plan states that are displayed:

- Number of deploy plans for which approval is pending
- Number of deploy plans in approved state
- Number of deploy plans in rejected state
- Number of deploy plans currently being provisioned on devices
- Number of deploy plans that have been successfully propagated and applied on devices
- Number of deploy plans scheduled for deployment at a future time
- Number of deploy plans for which deployment failed

To view the service or packet analyzer details:

1. From the bottom left corner of the screen, click the down arrow and select the **Service Analyzer** option from the shortcut menu.

The Service Analyzer table is displayed. The following details are displayed:

- Active—Number of service instances that are currently running
- Completed—Number of service instances that have successfully completed
- Stopped—Number of service instances that were halted

#### Related Documentation

- [Understanding How to Use the Edge Services Director Interface to View System Information on page 22](#)
- [Getting Started Assistant in Junos Space Platform Overview on page 23](#)
- [Changing Your Password for Edge Services Director on page 24](#)
- [Logging In to Edge Services Director on page 25](#)
- [Logging Out of Edge Services Director on page 27](#)



## CHAPTER 3

# Tasks Pane

- [Understanding the Build Mode Tasks Pane on page 31](#)
- [Understanding the Deploy Mode Tasks Pane on page 38](#)
- [Understanding the Fault Mode Tasks Pane on page 40](#)
- [Understanding the Monitor Mode Tasks Pane on page 41](#)
- [Understanding the Report Mode Tasks Pane on page 42](#)

### 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 (Location, Device, Gateway, or Service) 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.

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

- **Device Discovery**—Before your devices can be managed by Edge Services Director, you must use device discovery to discover them. As Edge Services Director discovers devices, it adds them to your network view in the View pane. [Table 4 on page 33](#) describes the device discovery tasks.
- **Inventory**—The Device Inventory page lists devices managed by Edge Services Director and provides basic information about the devices, such as IP address and current operating status, and configured services, such as server load balancing (SLB) and carrier grade NAT (CGNAT). The Device Inventory page is available in Build mode. [Table 5 on page 33](#) describes the inventory tasks.
- **Service Gateway**—The service delivery gateway (SDG) devices that are administered, maintained, and monitored from the Edge Services director application are called managed devices. The service delivery gateway (SDG) devices that are not managed and monitored from the Edge Services director application are called unmanaged devices. A service delivery gateway (SDG) device can be combined into a group of devices for easier and streamlined administration. You can create an SDG group for a

particular domain or zone in your network, or for any logical bundling that is needed. Templates contain the KPI parameters that evaluate the health of a SDG device. A system-created default KPI template is available. This system-created KPI template cannot be edited or deleted. [Table 6 on page 33](#) describes the service gateway tasks.

- **Service Analyzer**—You can configure and provision filters for packet analysis, configure filters for CGNAT, ADC, and TLB services. Also, you can start and stop the configured filters. [Table 7 on page 33](#) describes the service analyzer 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, deleting a device, or rebooting a device. [Table 8 on page 34](#) describes the device management tasks.
- **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 9 on page 34](#) describes the location management tasks.
- **Service Template**—You use the service templates to configure the following attributes and settings for the following four types of services: stateful firewall (SFW), carrier-grade network addressing (CGNAT), traffic load balancer (TLB), and application delivery controller (ADC). The service planning functionality enables you to use the Service Designer page to create service templates, which can be used on multiple devices. The Service Designer page lists all service components used to create service templates. According to the business needs, you can configure generic properties in a template and enable the editing of deployment-specific parameters. [Table 10 on page 36](#) describes the connectivity tasks.
- **Service Inventory**—The Services Inventory page lists the services configured in the Edge Services Director database and provides basic information about the configured services, such as adaptive delivery controller (ADC), stateful firewall (SFW), server load balancing (SLB), and carrier grade NAT (CGNAT). The Services Inventory page is available in Build mode and under Service view.
- **Object Builder**—The objects are the constituents or building blocks that are used to create service definitions and policy or filter templates. You can use the Object Builder page to retrieve and transfer the objects or components that have been previously created on the SDGs or devices.. [Table 11 on page 36](#) describes the profile and configuration management tasks.
- **Key Tasks**—Edge Services 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 Edge Services 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 Gateway View of Edge Services Director” on page 65](#) and [“Understanding Build Mode in Service View of Edge Services Director” on page 183](#).

Table 4 on page 33 through Table 11 on page 36 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 4: Device Discovery Tasks**

Task	Description	Scope
Discover Devices	Discovers supported devices, such as routers, in the network and brings them under Edge Services Director management.	Any

**Table 5: Inventory Tasks**

Task	Description	Scope
View Device Inventory	Displays three pie charts that summarize the status of the devices and services in your network environment. You can remove or restore a category (segment) from the pie chart by clicking that segment in the chart	Any
View Service Inventory	Displays the services configured in the Edge Services Director database and provides basic information about the configured services, such as adaptive delivery controller (ADC), stateful firewall (SFW), server load balancing (SLB), and carrier grade NAT (CGNAT). The Services Inventory page is available in Build mode and under Service view.	Any

**Table 6: Service Gateway Tasks**

Task	Description	Scope
Discover Gateway	Discovers and synchronizes physical devices such as MX Series routers that function as service delivery gateways in your network that are managed by Edge Services Director	Any
Unmanaged Gateway	Changes an unmanaged device to a managed device, and modifies managed device and KPI associations.	Any
Managed Gateway	Changes a managed device to an unmanaged device, and modifies managed device and KPI associations.	Any
Groups	Creates and manages a cluster of group of SDGs for easy and effective administration of service and policy definitions.	Any
KPI Templates	Clones, modifies, or deletes KPI templates to be associated with standalone SDGs or a high-availability pair of SDGs.	Any

**Table 7: Service Analyzer Tasks**

Task	Description	Scope
ADC Filter	Configures filters for ADC services. Also, starts and stops the service analyzer filters.	Any
TLB Filter	Configures filters for TLB services. Also, starts and stops the service analyzer filters.	Any

**Table 7: Service Analyzer Tasks (continued)**

Task	Description	Scope
CGNAT Filter	Configures filters for CGNAT services. Also, starts and stops the service analyzer filters.	Any

**Table 8: Device Management Tasks**

Task	Description	Scope
Change Location of Device	Changes where a device is located in Location view.	View: All Object: Individual router
Delete Devices	Deletes a switch or a wireless LAN controller as a managed device from Edge Services Director. If you select a scope that contains more than one switch or controller, you can choose which devices are deleted.	View: All Object: All, except access points
Reboot Devices	Reboots devices. If you select a scope that contains more than one switch or controller, you can choose which devices get deleted.	View: All Object: All
Show Current Configuration	Shows the running configuration on a switch or a wireless LAN controller.	View: All Object: Individual router
SSH to Device	Launches an SSH connection to the selected device.	View: All Object: Individual router
Validate Pending Configuration	Validates configuration changes that have not yet been deployed on devices.	View: All Object: All
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 router
View Physical Inventory	Displays information about the selected device's hardware components.	View: All Object: Individual router

**Table 9: Location Management Tasks**

Task	Description	Scope
Add Building	Creates a new building in the selected site.  <b>NOTE:</b> 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

Table 9: Location Management Tasks (continued)

Task	Description	Scope
Add Floor	Creates a new floor in the selected building.  <b>NOTE:</b> 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.  <b>NOTE:</b> 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
Assign Devices to Building	Assigns routers to a building. You cannot assign access points to a building.	View: Location Object: A building
Assign Devices to Closet	Assigns routers to a closet. You cannot assign access points to a closet.	View: Location Object: A closet
Assign Devices to Floor	Assigns routers to a floor.	View: Location Object: A floor
Assign Devices to Outdoor Area	Assigns routers to an outdoor area.	View: Location Object: An outdoor area
Setup Locations	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.  <b>NOTE:</b> Use this task only to create a site. Do not use it to modify an existing site.	View: Location Object: My Network and any location node within an existing site.

**Table 10: Service Template Tasks**

Task	Description	Scope
Manage ADC Service Templates	Creates, modified, or deletes an ADC service template with attributes and settings for load balancing operations	View: Service Object: Individual SDG or router
Manage CGNAT Service Templates	Creates, modified, or deletes a CGNAT service template with attributes and settings for load balancing operations	View: Service Object: Individual SDG or router
Manage SFW Service Templates	Creates, modified, or deletes a stateful firewall service template with attributes and settings for load balancing operations	View: Service Object: Individual SDG or router
Manage TLB Service Templates	Creates, modified, or deletes a traffic load-balancer (TLB) service template with attributes and settings for load balancing operations	View: Service Object: Individual SDG or router

**Table 11: Object Builder Tasks**

Task	Description	Scope
Import Objects	Retrieves and adds all of the object types that are supported for different services in a single, one-step operation from SDGs or an XML configuration file. You can select an SDG from which you want to import all of the objects contained in it. The supported or applicable objects of CGNAT pools, CGNAT rules, CGNAT rule sets, SFW rules, SFW rule sets, applications, application sets, and real servers can be imported in a bulk manner from a device.	Object: All services
Real Servers	Imports real servers, which are application servers used for traffic or server load balancing. The ADC software monitors the servers in the real-server group and the load-balanced applications running on them.	ADC services
SFW Rules	Imports firewall rules for use in stateful firewall policy creation.	SFW services
SFW Rule Sets	Imports firewall rule sets, which is a collection of rules, for use in stateful firewall policy creation.	SFW services
CGNAT Rules	Imports NAT rules for use in carrier-grade NAT policy creation.	CGNAT services
CGNAT Rule Sets	Imports NAT rule sets, which is a collection of rules for use in carrier-grade NAT policy creation.	CGNAT services
CGNAT Pools	Imports NAT pools for use in carrier-grade NAT policy creation.	CGNAT services

*Table 11: Object Builder Tasks (continued)*

Task	Description	Scope
Applications	Defines application protocols for the stateful firewall and Network Address Translation (NAT) services to use in match condition rules.	SFW and CGNAT services
Application Sets	Imports application sets for use in match conditions or criteria of stateful firewall and NAT rule terms	SFW and CGNAT services

## Understanding the Deploy Mode Tasks Pane

The Tasks pane in Deploy mode lists the available tasks. The Deploy mode tasks that are available depend on 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 12 on page 38](#) describes the configuration deployment tasks.
- **Image Management**—These tasks enable you to manage software images on devices. [Table 13 on page 39](#) describes the image management tasks.
- **Device Management**—These tasks enable you to view the device inventory, resynchronize the configuration of out-of-sync devices, and see extensive configuration settings that are present on a device. [Table 14 on page 39](#) describes the device management tasks.
- **Device Configuration File Management**—These tasks enable you manage configuration files on managed devices. [Table 15 on page 39](#) describes the device configuration file management tasks.
- **Deploy Service**—These tasks enable you to create a deployment plan that contains details about the settings and configuration parameters to be propagated and provisioned on the SDGs managed by Edge Services Director. For each approved deploy plan, a transaction is automatically created by the Edge Services Director application. [Table 16 on page 39](#) describes the service deployment tasks.
- **Service Edit**—This task enables you to view the list of CGNAT, SFW, and packet policy or filter templates as pie charts, whose segments display service policy filters in various states. [Table 17 on page 40](#) describes the task associated with the viewing of service object statistical details.
- **Policy & Filters**—These tasks enable the creation, update, display, publish and commission of packet filters, stateful firewall and NAT policies present on discovered and managed SDGs. [Table 18 on page 40](#) describes the service deployment tasks.
- **Key Tasks**—Edge Services 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 Edge Services 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 12 on page 38](#) through [Table 15 on page 39](#) describe the tasks in each task category.

**Table 12: Configuration Deployment Tasks**

Task	Description
Deploy Configuration Changes	Deploys pending configuration changes to devices.

**Table 12: Configuration Deployment Tasks (continued)**

Task	Description
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 Edge Services Director can collect and manage event and error information from these devices.
View Deployment Jobs	Manages configuration deployment jobs.

**Table 13: 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 14: Device Management Tasks**

Task	Description
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 15: 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 16: Service Deployment Tasks**

Task	Description
Manage Deployment Plans	Enables you to create deployment plans, which contain the configuration settings and attributes of services that must be propagated to SDGs. You can provision the deploy plans to transfer the configuration to devices immediately or schedule the deployment at a later specified time.

**Table 16: Service Deployment Tasks (continued)**

Task	Description
Transactions	Displays all of the transactions generated by the system for approved deploy plans. You can delete a transaction, which causes the transaction to be removed from listing, but does not delete the deployment plan associated with it. In addition, you can view the XML API format of configurations that exist on the device.

**Table 17: Service Edit Tasks**

Task	Description
View Statistics	Displays a set of five pie charts when you select Service Edit from the task pane. The pie charts are displayed for the different policy and service filters, such as ADC, TLB, CGNAT, stateful firewall, and packet filter templates, in various configuration states, such as in-synchronization, out-of-synchronization, and synchronization-in-progress.

**Table 18: Policy and Filter Tasks**

Task	Description
CGNAT	Enables you to create, update, and delete CGNAT policies on selected SDGs.
SFW	Enables you to create, update, and delete stateful firewall policies on selected SDGs.
Packet Filter	Enables you to create, update, and delete packet and service filter policies on selected SDGs.

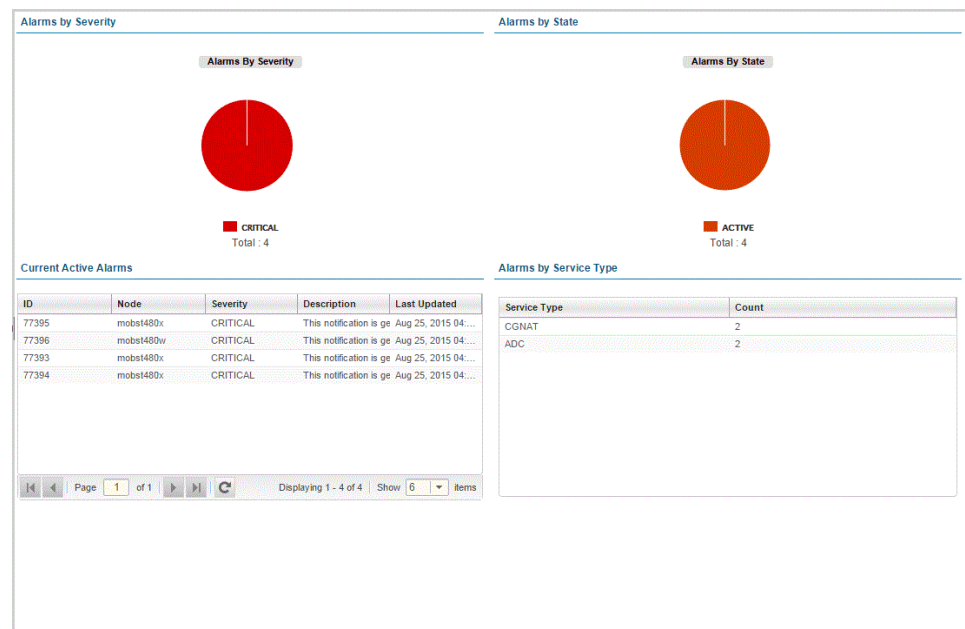
## Understanding the Fault Mode Tasks Pane

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, Edge Services 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 Edge Services 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.

Figure 5: Alarms Page in Fault Mode



## 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 Monitor icon on the Edge Services Director is available or accessible only when you select Gateway View and Service View from the View selector.

Monitor mode in Edge Services Director provides you visibility into your network status and performance. Edge Services 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.

The main purpose and benefit of monitoring functionalities is to allow the operators to quickly monitor the health (working condition), operating efficiency, traffic-handling capacity, and performance status of the managed SDGs and configured services such as ADC, TLB, CGNAT and SFW. The SDG monitoring mechanism is an extensive and ingrained tool; it allows the operator to understand the network health and status by drilling down to all the components of SDG. The SDG status is marked as Green, Red, Orange or Gray, based on the health, availability, performance and other important KPI indicators. Red denotes an emergency condition, which is a system panic or other conditions that cause the routing platform to stop functioning. It also indicates that the device is offline or turned down. Orange denotes an alert, which can be conditions that must be corrected immediately, such as a corrupted system database. Green indicates a notice, which signifies conditions that are not error conditions but are of interest or might warrant special handling. It can also include a severity level equivalent to

informational or debugging messages. Gray signifies an unknown or an unconnected device that is out of synchronization.

The Monitoring page is refreshed automatically every 3 minutes. Static polling occurs to obtain and display data, and asynchronous collection is not used.

The Master and Standby tabs display information about the primary or master, and standby or secondary SDGs in an SDG pair. The Service Wait tab is displayed if the standby device is not fully active after a switchover.

## Understanding the Report Mode Tasks Pane

---

Edge Services 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.

For example, if you want to view inventory reports on only your wireless controllers, you can select the Device view and the Routers > MX960 node in the network tree to provide granular information on just those devices. After selecting the view and node in the network tree, create a report definition. In this definition file you select from a number of preconfigured reports and set the time frame, schedule, and output options.

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 Edge Services Director is running. However, because reports can be large, the report is delivered in a compressed or *zipped* format. and can be stored offline.

- 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 Edge Services 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.



## CHAPTER 4

# Dashboard

- [Understanding the Dashboard on page 45](#)
- [Working with the Dashboard on page 45](#)
- [Using Dashboard Widgets on page 50](#)
- [Alarm Severities and States Overview on page 51](#)
- [Viewing the Detailed Status of KPI Templates Applied to Devices on page 52](#)

### 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 Edge Services Director banner.

#### Related Documentation

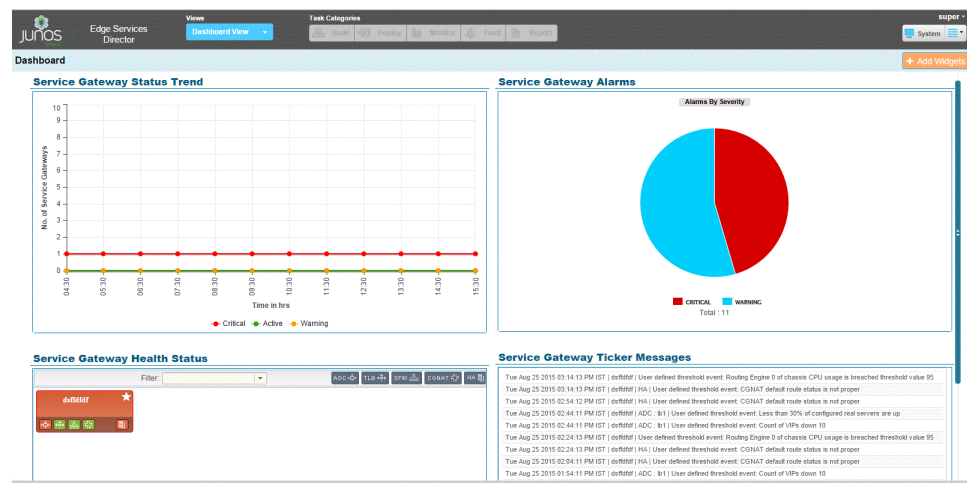
- [Working with the Dashboard on page 45](#)

### Working with the Dashboard

---

When you log in to the Edge Services Director interface, the first page that is displayed is the Dashboard page. After the deployment of the Edge Services Director application, if you have not discovered any SDGs, you are prompted to the next step of discovering devices. The link to the Service Gateways page, which is accessible by clicking the Build icon in the banner, is provided only for users with administrator privileges. The Dashboard page contains several monitors or frames. The following monitors are displayed on the Dashboard page:

Figure 6: Dashboard Page



- [SDG Views on page 46](#)
- [Service Delivery Gateway Alarms on page 47](#)
- [Filters on page 47](#)
- [Specifying KPI Template and Alarm Filters on page 48](#)
- [Service Delivery Gateways Count by Severity on page 49](#)
- [Service Gateway Ticker Updates on page 50](#)
- [Service Delivery Gateway Health Status Trend on page 50](#)

## SDG Views

The dashboard default view is tiled view. In tile view, a high-level, graphical view of the chassis is shown. It indicates the state of the interfaces. When the administrative and operational status of the interface is up, it is displayed in green. If the administrative status is down, the interface is displayed in grey. And, if the administrative status is up and operational status is down, the interface is displayed in amber. The image is a replica of the SDG. If you are connected to a virtual chassis, the image includes all the member switches of the virtual chassis.

The purpose of the view is to try and provide a comprehensive monitoring view of the health and status of deployed SDGs across the network. In this view all the managed SDGs are shown with their appropriate status and health based on the KPI template applied. This view helps the operator to know the health and status across the network, it provides with the operator to quickly see the macro level information, which allows the operator to further analyze the information provided and quickly navigate to individual SDGs and take any further corrective measure required. In tiled view, SDGs are by default sorted based on Red, Orange and Green. Consider a scenario in which an operator deploys n number of SDGs in a network. If the operator finds a difficulty in monitoring the status and health of the entire network, you can organize the dashboard with a tiled view to enable the operator to have a macro-level view on the network health.

You can change the view format to be tile view, group view, or band view. To change the view format, click the **Maximize** icon in the Service Gateway Health Status widget. The

Service Gateway Health Status widget is popped out as a separate dialog box. You can click the **Tile View**, **Group View**, or **Band View** icons in the dialog box to organize the view appropriately in the Dashboard page. SDGs are deployed in the network as zones and the group view, which is a cluster of SDGs typically in a particular zone, helps the operators to quickly reach to a particular zone and find the status and health of the deployed SDG's in that zone. By default it includes all the SDGs across network. Operator can able to view the SDGs based on particular zone. In Tile view, the **Gateways/page** box enables you to customize the number of SDGs displayed per page. You can display 25, 30, 35, or 40 SDGs per page. In Group view, the **Groups/page** box enables you to customize the number of SDG groups displayed per page. You can display 2, 4, 6, or 8 groups per page. The SDG groups are displayed in quadrants on the page.

Click the **Maximum** button at the upper right corner of each of the group boxes to expand and display the particular group in a separate window. SDGs span across the network based on zones. In this group view, all the managed SDG's are grouped based on the zones. The corresponding SDG status and health is shown based on the KPI template applied.

Operators can quickly narrow down network discrepancies and failures based on a particular zone. Assume that in a particular zone, the KPI indication of SDGs in that group is not satisfactory. Using the group view, the operator can identify the particular zone in an efficient, optimal, and faster manner to view the status and health and analyze further. Logical grouping of SDGs, mostly based on zones, avoids the difficulty in monitoring the status and health of the entire network because the specific zone or group can be determined to investigate and drill down to the exact SDG that needs to be rectified.

A band view is also provided. In this view, each of the SDGs in a band are displayed spanning across the frame that shows the views. The bands are displayed one below the other. It provides a graphical display of the chassis in the network based on the bands, which are the different system severities or health conditions, such as red, orange, green, and gray. Double-clicking any of the SDGs in any of the views navigates you to the Monitoring page under Monitor mode for more detailed, in-depth diagnosis and debugging.

## Service Delivery Gateway Alarms

Alarms are displayed below the health status line graph of the SDGs. Critical, major, and minor alarms are displayed in a pie chart with percentage values of each type of alarm. When you move the mouse over the segments of the pie chart, the total number of alarms of each type are displayed.

## Filters

The SDG dashboard filter consists of two types, namely alarms and KPI templates, and favorite SDGs. The filters enable you to quickly and easily sort and segregate the appropriate SDGs that correspond to the KPI templates defined or alarms. The filter capability makes it easier for you to focus on only the SDGs that you are of relevance or interest. The SDGs status and health are colored based on the KPIs set by the operator. The dashboard filter operates in a logical OR format. If either of the filter conditions are satisfied, the display is modified to match the filter condition.

Assume that the CPU threshold value of one of the SDGs is set as 40 percent in a template. If the particular SDG exceeds that threshold, it needs to be displayed as red. When an operator logs into the dashboard and views the tiled pattern of display, all the SDGs are shown across the network. The operator is also viewing the critical SDGs and can decide to filter the SDGs based on the template. Based on the filter chosen, the SDGs are displayed in the dashboard view.

Select **Favorite SDGs** from the Filters box to display only the SDGs that you are interested in or are involved with managing in the entire network.

Select **Alarms | KPIs** from the Filters box to display SDGs that match with the criteria specified in the templates.

Click **Clear Filters** to remove the applied filters. You are prompted to confirm the deletion in such a case. A graphical representation of the components such as ADC and SFW that are defined in the KPI templates are displayed to the right of the Filters box on the Dashboard page.

To configure a filter, see *Specifying KPI Template and Alarm Filters*.

## Specifying KPI Template and Alarm Filters

To specify KPI template and alarm filters for sorting the dashboard devices according to your needs:

1. From the Dashboard page, select **Alarms | KPIs** from the Filter box. If you have already configured a filter, it is applied to determine the match criterion for displaying SDGs. The Filters window is displayed only if you cleared a previously configured filter or are configuring a filter for the first time.
2. Select the **Critical, Major, or Minor** check boxes next to the Alarms field to filter SDGs based on the critical, major, or minor alarms generated for the devices.
3. For the ADC, CGNAT, TLB, SFW, or HA sections, select the check boxes under the R, G, or O columns for the respective fields to cause SDGs that map with the settings defined in the KPI templates to be displayed in the dashboard view.

R refers to SDGs with red status or catastrophic problems, G refers to SDGs with green status or fully functional condition, and O refers to SDGs with orange status or moderately critical problems.

For the ADC section, you can choose the following:

- CPU Status—Working state of the CPU for ADC
- Service Pic Status—Operating status of the services PIC for ADC
- VIP Status—Virtual IP address state used by virtual servers for ADC
- Real Servers—Real servers used by ADC instances

For the CGNAT section, you can choose the following:

- CPU Status—Working state of the CPU
- Service Pic Status—Operating status of the services PIC for CGNAT
- Memory Status—Utilization of memory for CGNAT services
- CPU Utilization—Usage of CPU for CGNAT operations

For the TLB section, you can choose the following:

- CPU Status—Working state of the CPU
- Service Pic Status—Working status of the services PIC for TLB
- Real Server Status—Status of real servers for TLB

For the SFW section, you can choose the following:

- CPU Status—Working state of the CPU
- Service Pic Status—Status of services PIC for stateful firewall

For the HA section, you can choose the following:

- VRRP Status—Status of VRRP
- CGNAT SFW HA—Inter-chassis high availability for CGNAT and firewalls.

4. Click **Apply** to save the settings. Click **Close** to close the Filters window and return to the dashboard.

## Service Delivery Gateways Count by Severity

A set of four boxes are displayed beneath the pane that shows the view of SDGs. These boxes are indicators for the overall set of SDGs that have been deployed. The boxes are colored as orange, red, green, or grey to indicate the health and performance of the SDGs based on the applied KPI templates. Red denotes an emergency condition, which is a system panic or other conditions that cause the routing platform to stop functioning. It also indicates that the device is offline or turned down. Orange denotes an alert, which can be conditions that must be corrected immediately, such as a corrupted system database. Green indicates a notice, which signifies conditions that are not error conditions but are of interest or might warrant special handling. It can also include a severity level equivalent to informational or debugging messages. Gray signifies an unknown or an unconnected device that is out of synchronization.

The configuration state of a device is shown as In Sync when the configuration information in all three repositories match (settings made using the devices CLI, Edge Services Director in Build mode, or Junos Space Network Management Platform). If there is a conflict between the configuration information in one or more of the repositories, the device configuration state is 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 Edge Services Director. You can resynchronize such devices to bring them back to be in synchronization. A number is displayed overlaying each of the boxes to specify the number of SDGs in each of the states or health conditions.

The paging controls that appear at the bottom of the SDG icons that are shown. You can use these controls to browse the SDGs when the inventory of SDGs is too large to fit on one page. The Page box lets you jump to a specific page of managed objects. Type the page number in the Page box and press Enter to jump to that page. The SDGs per page box enables you to customize the number of objects displayed per page. You can also navigate to the specific page of the dashboard view by typing the page number, if the SDGs span across several pages. Otherwise, you can use the first, previous, next, and last page buttons to traverse across pages. The Refresh icon enables you to revise the display and show updated information.

## Service Gateway Ticker Updates

The SDG dashboard ticker constantly updates with the event, messages and logs. This view can be added or removed from the Dashboard View by clicking the Add Widgets button or the cross mark (X) icon respectively. For a displayed syslog message, if you want to further analyze it for troubleshooting and diagnosing the cause of the error, you can use the options under the Monitor mode. At any given point in time, the latest ten messages are displayed and the older messages are flushed out. If you want to view historical messages, you can view the Fault Management, Performance Management, or Monitoring pages. By default, the standard Junos OS format for messages specifies the month, date, hour, minute, and second when the message was logged. Also, the event category and description are displayed.

## Service Delivery Gateway Health Status Trend

The line chart displays the number of SDGs in each of the severity states. The severity is determined by the memory utilized on the routing engine, temperature, CPU load, and fan status. A fan running at normal speed is displayed in green. If the fan is running at maximum speed or not running at all, it is displayed in red. For a virtual chassis the status of the fans for the selected member is displayed. The SDG health status displays the operating condition and working efficiency in a color-coded form based on the configured KPI template for each of the individual SDGs.

A line graph is displayed with the horizontal axis showing the number of SDGs. The vertical axis shows the time in 24-hour clock format at which alarms have been raised or cleared for the SDGs. The green line denotes active SDGs, the orange line denotes warning messages, and the red line denotes critical problems with SDGs. Mouse over the dots on the graph to view details about the number of alarms at a particular point in time. Time is shown in increments of two hours on the horizontal axis, ending with the current time of the local clock on the system. You can expand or collapse the SDG Health Status pane by clicking the double right or double left arrows on the top-right corner of the pane.

**Related Documentation** • [Quickly Accessing Important Monitoring and Troubleshooting Details on page 28](#)

---

## 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 Edge Services 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** • [Working with the Dashboard on page 45](#)

## Alarm Severities and States Overview

By default, the Junos Space Network Management Platform is monitored using a built-in SNMP manager. The Junos Space Network Management Platform node is listed in the node list (Network Monitoring > Node List), and is referred to as the Junos Space Network Management Platform node.

### Alarm Severity

Alarms are ranked by their impact to the network. The following list shows the ranking of alarms in Edge Services 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.

**Indeterminate (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.

## 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.

### Related Documentation

- *Events and Alarms Overview*
- [Understanding Monitor Mode in Edge Services Director on page 559](#)

---

## Viewing the Detailed Status of KPI Templates Applied to Devices

In a network environment, it is essential and important for a network administrator or a supervisor to quickly, easily assess the device performance and operating efficiency to be able to take corrective action and restoration measures for any device alarms, overloaded conditions, or traffic drops observed.

From the [“Dashboard” on page 45](#) page, you can view the KPI templates applied to the SDG devices in an in-depth, granular way before you navigate to the KPI Templates page or the Service Designer page to modify the metrics of KPI settings or service settings respectively. You can view ingrained, extensive information about the KPI settings for SDGs displayed in tile view, group view, or band view.

To view detailed status information of the KPI templates associated with a particular SDG device:

1. From the Dashboard page, right-click an SDG device and select **Status Details**. The Service Gateway KPI Status Details window appears.

The name of the SDG device is displayed at the top of the page. The same color-coding format that is used to display the SDG in the Dashboard page is used in this window.

For example, if the SDG is shown in the tile view or band view in red, the host name is shown in red in the KPI Status Details window.

2. Click the right arrow next to each of the KPI sections or components that are displayed. Only the KPI components applied to the specified SDG are displayed. When you expand each of the KPI components, the attributes or parameters that apply for the KPI are shown. A colored box is displayed for each of the KPI attributes to signify the health and efficiency of the device for the corresponding KPI setting. For example, for ADC, CPU Status, Service Pic Status, VIP Status, and Real Servers might be displayed.
3. Click **Close** after you complete viewing the settings and to return to the Dashboard page.

**Related Documentation** • [Working with the Dashboard on page 45](#)



## PART 2

# System Administration

- [Handling Administrative Tasks on page 57](#)



## CHAPTER 5

# Handling Administrative Tasks

- [Understanding Edge Services Director User Administration on page 57](#)
- [Viewing Audit Logs From Edge Services Director on page 58](#)
- [Managing Jobs on page 59](#)
- [Collecting Logs for Troubleshooting on page 60](#)

### Understanding Edge Services Director User Administration

---

Edge Services Director uses the user administration features of the Junos Space platform on which it runs. Using these features, you can add, delete, and edit user accounts and roles and changing user passwords. Refer to the *Junos Space Network Application Platform User Guide* for more information about user administration.

When Edge Services Director is installed, some additional user administration options are available in Junos Space, which are specific to Edge Services Director.

In addition to the Super Administrator role, the following predefined roles are available to Edge Services Director users:

- Edge Services Director - Administrator—Has complete access to all the Edge Services Director modes and user preferences.
- Edge Services Director - Operator—Has access to all modes except the Build mode. Has access to windows and capabilities, such as fault management, performance management, dashboard and monitoring. You can create custom roles to grant users different access rights to the Edge Services Director modes.
- Edge Services Director - Designer— Has access to the Build mode for handling device and service configuration operations such as creation of services and KPI templates.

You can also create custom roles to grant users different access rights to the Connectivity Services Director modes. Edge Services Director modes—Build, Deploy, Monitor, Fault, and Report modes are available to assign to custom user roles in the list of application workspaces and associated tasks.



**NOTE:** The tasks listed under the Edge Services Director modes do not have any effect. Access is controlled at the mode level, so if you grant a role access to a mode, the role has access to all tasks in that mode, regardless of which tasks you select.

#### Related Documentation

- [Understanding the Need for Edge Services Director on page 3](#)
- [Understanding the Edge Services Director User Interface on page 6](#)
- [Understanding Edge Services Director and the Management Lifecycle Modes on page 15](#)
- [Service Delivery Gateway Overview on page 17](#)
- [Edge Services Director Overview on page 19](#)

## Viewing Audit Logs From Edge Services Director

Audit logs are generated for login activity and tasks that are initiated from the Edge Services 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 Edge Services 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 Edge Services Director banner.
2. Select **View Audit Logs** from the Tasks pane.

The Audit Logs page is displayed with the fields listed in [Table 19 on page 58](#).

**Table 19: 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

*Table 19: Audit Logs Page Fields (continued)*

Field	Description
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"> <li>• Success—Job completed successfully</li> <li>• Failure—Job failed and was terminated</li> <li>• Job Scheduled—Job is scheduled but has not yet started</li> </ul>
Description	A description of the audit log
Job ID	The job ID for any task that runs as a job

## Managing Jobs

Edge Services 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, Edge Services Director enables you to view special pre-filtered versions of this page from various other tasks, such as View Discovery Status or View Image Deployment 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 Edge Services 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 20 on page 60](#). 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 Edge Services Director will be available only from Edge Services Director. These jobs will not be listed in the Job Management pane in Junos Space platform and vice-versa.

*Table 20: 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	The status of the job: <ul style="list-style-type: none"> <li>• Success—Job completed successfully</li> <li>• Failure—Job failed and was terminated</li> <li>• Job Scheduled—Job is scheduled but has not yet started</li> <li>• In progress—Job is has started, but not completed</li> <li>• Cancelled—Job is cancelled</li> </ul>
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.

## Collecting Logs for Troubleshooting

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

Edge Services 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 Edge Services 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 Edge Services Director and Junos Space** link.

Edge Services Director begins collecting the logs and data. It can take a few minutes for Edge Services 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 21 on page 61](#) lists the files included in the **troubleshoot\_yyyy-mm-dd\_hh-mm-ss.zip** file.

**Table 21: 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>
Watchdog log file	<code>/var/log/</code>



## PART 3

# Gateway View of Build Mode

- [About Gateway View of Build Mode on page 65](#)
- [Managing Service Delivery Gateways and Groups on page 85](#)
- [Managing KPI Templates on page 119](#)
- [Viewing the Device Inventory on page 133](#)



## CHAPTER 6

# About Gateway View of Build Mode

- [Understanding Build Mode in Gateway View of Edge Services Director on page 65](#)
- [Understanding Resynchronization of Device Configuration on page 69](#)
- [Importing Devices on page 74](#)
- [Device Discovery Overview on page 77](#)
- [Unmanaged Devices Overview on page 78](#)
- [Working With Managed Devices on page 80](#)
- [Working With Unmanaged Devices on page 80](#)
- [Working With Discovered Devices on page 81](#)
- [Managing Jobs as a System Task on page 81](#)

### Understanding Build Mode in Gateway View of Edge Services Director

---

In Gateway view of Build mode, you create the network managed by Junos Space Edge Services Director by bringing devices under the administration of the network management application and retrieving the device settings to save in the Edge Services Director database. It provides you with the ability to use device discovery to bring devices under Edge Services 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 on page 65](#)
- [Configuring Devices on page 66](#)
- [Viewing the Devices Inventory on page 67](#)
- [Service Delivery Gateway Groups on page 68](#)
- [KPI Templates on page 68](#)

### Discovering Devices

Device discovery finds your network devices and brings them under Edge Services Director management. You provide Edge Services Director with identifying information about the devices you want Edge Services Director to manage—an IP address or hostname, an IP address range, an IP subnetwork, or a CSV file that contains this information. Edge Services Director uses the information to probe the devices by using either ping or SNMP get

requests. If a device probe is successful, Edge Services 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, Edge Services Director adds the device to its database of managed devices. Edge Services 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 Edge Services Director management if they are supported by Edge Services Director.

Besides bringing your devices under Edge Services Director management, device discovery:

- Reads the device configuration and saves it in the Junos Space configuration database. Edge Services 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.
- Imports the device configuration into the Gateway view of Build mode configuration. For more information about importing device configurations, see ["Importing Device Configurations" on page 67](#).

## Configuring Devices

In Gateway view of 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 Gateway view of Build mode configuration in a set of profiles.



**NOTE:** This section does not apply to virtual devices that Edge Services Director manages.

---

### Deploying Device Configurations

---

After you build your device configurations in Gateway view of Build mode, you need to deploy the configurations on the devices. None of the configurations you create in Gateway view of 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 Gateway view of Build mode, the device becomes available in Deploy mode for configuration deployment.

For more information about deploying configuration changes, see *Understanding Deploy Mode in Edge Services Director*.

### Importing Device Configurations

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

As it imports the device configuration, Edge Services Director automatically creates discovery 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, Edge Services Director assigns the existing Port profile to those ports. For the other ports, Edge Services 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 Edge Services 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 Edge Services 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 Gateway view of 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. The Edge Services Director resolves out-of-band configuration changes and synchronizes the Gateway view of Build mode configuration with the device configuration by using the resynchronization of devices functionality.



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

### Viewing the Devices Inventory

This inventory page lists all the SDG hardware and inventory of the chassis components. A graphical representation is provided of the types of services, connection status of the

SDGs or devices, and the configuration status of managed SDGs. A pie chart is displayed to signify these details. Also, you can view full-blown information on the chassis, line cards, and associated hardware components of an SDG and the interface attributes. All of the SDGs that are created are displayed in a tree structure on the left pane of the Inventory page.



**NOTE:** From Service view in Build mode, you can select **View Inventory** from the tasks pane also view comprehensive, consolidated information on each of the services, such as load balancing or CGNAT, from the Inventory page by clicking the plus (+) sign that appears adjacent to each service on the page.

---

## Service Delivery Gateway Groups

Service delivery gateways (SDGs) are discovered by SDG discovery workflow. The discovered SDGs are shown in the SDG inventory page. The discovered SDGs can be a part of a high availability (HA) pair or standalone SDGs. In case of SDG HA pairs, all the actions from SDG management are at the SDG HA pair level; you cannot an action on only the master SDG alone or the standby SDG. SDGs can be grouped as zones or domains. The SDG groups contain one or more SDGs. Each SDG can be part of just one group. User can create, edit, or delete SDG Groups.

## KPI Templates

Templates contain the KPI parameters that evaluate the health of a SDG device. A system-created default KPI template is available. This system-created KPI template cannot be edited or deleted. However, an SDG administrator can clone a new template based on this default template. An administrator-created KPI template can be edited or deleted. During the SDG discovery process, one of the KPI template copy is associated to each and every SDG. The KPI parameters from the KPI template are part of the SDG settings and are used to compute the statuses and condition of services of SDGs as green, orange or red.

**Related Documentation**

- [Working with the Dashboard on page 45](#)

---

## Understanding Resynchronization of Device Configuration

---

In a network managed by Edge Services Director, three separate repositories about device configuration are maintained:

- The configuration information on the devices themselves. Each switch and wireless LAN controller 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 Edge Services Director, Junos Space stores a record of the configuration on that device.

Edge Services 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 Edge Services Director in Build mode. This information takes the form of the profiles assigned to the device, plus the additional configuration, such as LAG and access point configuration, that you can do under device management.

In Edge Services 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, Edge Services 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 Edge Services 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.
- Using RingMaster software.
- 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 Edge Services Director enables you to resynchronize the device configuration state. It covers:

- [The Resynchronize Device Configuration Task on page 70](#)
- [How Resynchronization Works in NSOR Mode on page 70](#)
- [How Resynchronization Works in SSOR Mode on page 71](#)
- [How Edge Services Director Resynchronizes the Build Mode Configuration on page 73](#)

## The Resynchronize Device Configuration Task

Edge Services 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 Edge Services 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 Edge Services Director. Thus when you perform a resynchronization, the Junos Space configuration record and the Edge Services 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 Edge Services Director about the change.
2. Both Junos Space and Edge Services Director set the device configuration state to Out of Sync.
3. 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.
4. If the configuration change does not affect configuration that you can perform in Build mode (for example, routing configuration), Edge Services Director also sets the device configuration state to In Sync after the Junos Space resynchronization completes. All three configuration repositories are now in sync.

If the configuration change affects configuration that you can perform in Build mode, Edge Services Director does not set the device configuration state to In Sync. Instead, it continues to show the device configuration state as Out of Sync because the Build mode configuration does not match the device configuration.

5. To resolve the Out of Sync state in Edge Services Director, use the Resynchronize Device Configuration task in Deploy mode. Edge Services Director updates the Build mode configuration to match the out-of-band changes.
6. Edge Services Director sets the device configuration state to In Sync.



**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 Edge Services 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, Edge Services 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 Edge Services Director about the change.
2. Junos Space sets the device configuration state as Device Changed, and Edge Services Director sets the device configuration state to Out of Sync.

Edge Services 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 Edge Services Director.

3. In Edge Services 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 Edge Services 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 Edge Services Director Build mode configuration remains unchanged.
4. Both Junos Space and Edge Services 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. Edge Services Director shows the device configuration state as In Sync.



**NOTE:** Because the device configuration state is In Sync in Edge Services Director, you can deploy configuration on the device from Edge Services Director at this point. If you do so, the Edge Services 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 Edge Services Director changes the device state to Out of Sync.
3. In Edge Services 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 Edge Services Director, the Build mode configuration is updated to reflect the configuration changes.
4. Edge Services Director sets the device configuration state to In Sync.

If you use Junos Space instead of Edge Services 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 Edge Services 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 Edge Services 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 Edge Services Director. You must use the Resynchronize Device Configuration task to resolve the Out of Sync state.



**NOTE:** When Junos Space is in SSOR mode, we recommend that you do not make out-of-band changes to the cluster configuration on the secondary seeds and member controllers of a mobility domain, such as disabling the cluster on these devices. Use Edge Services Director to modify the cluster configuration on these devices.

## How Edge Services Director Resynchronizes the Build Mode Configuration

A network managed by Edge Services 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 Edge Services Director.

When the configuration contained in all three repositories match, the device configuration state is shown as In Sync in Edge Services 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 Edge Services Director.

- 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 Edge Services 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 Edge Services 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 Edge Services 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 Edge Services Director Build mode configuration.

When a commit operation is performed on a managed device under NSOR, Junos Space Network Management Platform, by default, schedules a resynchronization job to run 20 seconds after the commit operation is received. However, if Junos Space Network Management Platform receives another commit notification within 20 seconds of the previous commit notification, no additional resynchronization jobs are scheduled because Junos Space Network Management Platform resynchronizes both commit operations in one job. This damping feature of automatic resynchronization provides a window of time during which multiple commit operations can be executed on the device, but only one or a few resynchronization jobs are required to resynchronize the Junos Space Network

Management Platform database after multiple configuration changes are executed on the device.

You can change the default value of 20 seconds to any other duration by specifying the value in seconds in the **Administration > Applications > Network Management Platform > Modify Application Settings > Device > Max auto resync waiting time secs** field. For example, if you set the value of this field to 120 seconds, then Junos Space Network Management Platform automatically schedules a resynchronization job to run 120 seconds after the first commit operation is received. If Junos Space Network Management Platform receives any other commit notification within these 120 seconds, it resynchronizes both commit operations in one job.

When Junos Space Network Management Platform receives the device commit notification, the device status is “Out of Sync”. When the resynchronization job begins on the device, the Managed Status for the device displays “Synchronizing” and then “In Sync” after the resynchronization job has completed, unless a pending device commit operation causes the device to display “Out of Sync” while it was synchronizing.

For details about resynchronizing devices, see *Resynchronizing Managed Devices with the Network*.

#### Related Documentation

- [Understanding Build Mode in Gateway View of Edge Services Director on page 65](#)

---

## Importing Devices

You can import device configurations from MX Series devices running Junos OS into the Edge Services Director database.

When importing from a device, the management system connects to the device and imports Data Model (DM) information that contains details of the device configuration. The connection is secured using Secure Server Protocol (SSP), a proprietary encryption method; an always-on connection exists between the management system and the device.

To import a single device, you must have available the following requirements:

- A management interface (fxp0) with the IP address of the device
- A user with full administrative privileges for the NSM administrator
- Device connection information (IP address, connection method) and the device administrator's name and password



**NOTE:** All passwords handled by NSM are case-sensitive.

---

- A physical connection to your network with access to network resources
- Connectivity to the NSM Device Server, which can be with a static IP address
- A Telnet or an SSHv2, and a NETCONF protocol over SSH connection



**NOTE:** After importing a device configuration, log entries from that device begin to appear in the Log Viewer. However, until you update the device from NSM, the following log fields display 0 (or unknown):

- domain
- rulebase
- policy
- rule number
- source zone
- destination zone

After you update the imported device configuration using NSM, the appropriate values are displayed for log entries from the device.

When you import a device configuration, the Log Viewer displays the appropriate values for the device's log entries. This feature eliminates the need to update the device after importing it.

To add devices in a large-scale, bulk manner:

1. From the View selector, select **Gateway View** or **Device View**. The workspaces that are available in this view are displayed. The Gateway view displays the service delivery gateway (SDG) groups and the SDGs that are part of the high availability pair in an SDG group. The Device view displays the SDGs based on the device type, and within the device type, the devices are organized by the device model. For example, all models of MX960 routers are grouped together under one node in the tree.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the View pane, select the All Network item in Gateway view. If you are in Device view, click the plus sign (+) beside the My Network item in the View pane to expand the tree and select the device node you want.
4. Perform either of the following:
  - If you are in Gateway View, select **Services Gateways** from the task pane.  
The Service Gateways page is displayed.
  - If you are in Device View, select **Device Discovery** from the task pane.  
The Service Gateways page is displayed.

5. Under Device Discovery, select the **Discover Devices** option from the task pane. Alternatively, under Services gateways, select the **Discover Gateway** option from the task pane.

The Service Gateways—Discovered Devices view is displayed.

You need not click this button if you are launching the Service Gateways page by navigating from another page or another mode, such as Deploy or Monitor. It is displayed by default. You must click this button only if you are viewing unmanaged or managed SDGs or devices.

6. Click the **Add** icon.

The Discovery Profile window appears.

You can add devices using either the **CSV Upload** button or the Add icon, or both together. This button is available in the IP Details, User Details, and SNMP Details sections of the Discovery Profile window.

7. Click the **CSV Upload** button to add your own CSV files.



**NOTE:** The format of the CSV file that you are uploading should exactly match the format of the sample CSV file.

A dialog box appears.

8. Click **Browse**.

The CSV File Upload dialog box appears.

9. Navigate to the desired CSV file, select it, and then click **Open**.

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

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

**Related  
Documentation**

- [Understanding Build Mode in Gateway View of Edge Services Director on page 65](#)
- [Understanding Resynchronization of Device Configuration on page 69](#)

## Device Discovery Overview

You use device discovery to add devices to Junos Space Edge Services Director application. *Discovery* is the process of finding a device and then synchronizing the device inventory and configuration with the Junos Space Edge Services Director application database. To use device discovery, Junos Space Edge Services Director application must be able to connect to the device.

To discover network devices, Junos Space Edge Services Director application uses the SSH and SNMP protocols. Device authentication initially is handled through administrator login SSH v2 credentials and SNMP v1/v2c or v3 settings, which are part of the device discovery configuration. You can continue to use credentials for these devices thereafter, or you can create and upload RSA keys to devices to allow Junos Space Edge Services Director application to authenticate itself to them automatically during later discoveries.

You can specify a single IP address, a DNS hostname, an IP range, or an IP subnet to discover devices on a network. During discovery, Junos Space Edge Services Director application connects to the physical device and retrieves the running configuration and the status information of the device. To connect with and configure devices, Junos Space Edge Services Director application uses Juniper Network's Device Management Interface (DMI), which is an extension to the NETCONF network configuration protocol.

When discovery succeeds, Junos Space Edge Services Director application creates an object in the Junos Space Edge Services Director application database to represent the physical device and maintains a connection between the object and the physical device so their information is linked.

Junos Space can manage devices in either of the following ways:

- Junos Space initiates and maintains a connection to the device.
- The device initiates and maintains a connection to Junos Space.

By default, Junos Space manages devices by initiating and maintaining a connection to the device. When Junos Space initiates the connection to the device, you can discover and manage devices provided that the management system is behind Network Address Translation (NAT), as Junos Space establishes the SSH tunnel directly to the device. For WW Junos devices, Junos Space uses SSH with an adapter to manage the devices.

If device-initiated connection to Junos Space is enabled, the DMI channel and port 7804 are used and the following (sample) configuration is added on the device to establish the connection to Junos Space:

```
set system services outbound-ssh client 00111D0CEFAC device-id 7CE5FE
set system services outbound-ssh client 00111D0CEFAC secret "$ABC123"
set system services outbound-ssh client 00111D0CEFAC services netconf
set system services outbound-ssh client 00111D0CEFAC 172.22.199.10 port 7804
```

When configuration changes are made in Junos Space Edge Services Director application, for example, when you deploy service orders to activate a service on your network devices, the configuration is pushed to the physical device.

If the network is the system of record (NSOR), when configuration changes are made on the physical device (out-of-band CLI commits and change-request updates), Junos Space Edge Services Director application automatically resynchronizes with the device so that the device inventory information in the Junos Space Edge Services Director application database matches the current device inventory and configuration information. If Junos Space Edge Services Director application is the system of record (SSOR), this resynchronization does not occur and the database is unchanged.

The following device inventory and configuration data is captured and stored in relational tables in the Junos Space Edge Services Director application database:

- Devices—hostname, IP address, credentials
  - Physical Inventory—chassis, FPM board, Power Entry Module (PEM), Routing Engine, Control Board (CB), Flexible PIC Concentrator (FPC), CPU, Physical Interface Card (PIC), transceiver (Xcvr), fan tray
- Junos Space Edge Services Director application displays the model number, part number, serial number, and description for each inventory component, when applicable.
- Logical Inventory—subinterfaces, encapsulation (link-level), type, speed, maximum transmission unit (MTU), VLAN ID
  - License information:
    - License usage summary—license feature name, feature description, licensed count, used count, given count, needed count
    - Licensed feature information—original time allowed, time remaining
    - License SKU information—start date, end date, and time remaining
  - Loopback interface

Other device configuration data is stored in the Junos Space Edge Services Director application database as binary large objects, and is available only to northbound interface (NBI) users.

- Related Documentation**
- [Understanding Build Mode in Gateway View of Edge Services Director on page 65](#)
  - [Understanding Resynchronization of Device Configuration on page 69](#)

---

## Unmanaged Devices Overview

---

Unmanaged devices are non-DMI devices made by vendors other than Juniper Networks, Inc. You can add such devices to Junos Space Edge Services Director manually, or by importing multiple devices simultaneously from a CSV file. You need to provide the IP address or the host name of the unmanaged device, name of the vendor, username of the device, password for the device, SNMP credentials, loopback address details, and key-values that are needed for the device driver to operate. The currently supported SNMP versions are SNMP V1, SNMP V2C, and SNMP V3. You or any other user may need to enter the key and the value for the device driver to operate. You can add multiple key-value pairs for an unmanaged device. The key-value pairs supported are plain-text

string or password. If Junos Space Network Management Platform can communicate with the device using SNMP, the information gathered via SNMP overrides the information that you enter.

When you have added the unmanaged device and installed the appropriate device drivers, the inventory data is fetched from the device. Physical interface and logical interface details can be fetched for an unmanaged device. You can resynchronize the device with the network both in SSOR and NSOR mode. This action will resynchronize the inventory data for the device based on the capabilities supported by the device. Junos Space Network Management Platform does not monitor the connection status of the unmanaged device for which the device driver is installed in Junos Space Network Management Platform. The Device Management table lists NA in the Connection Status column for unmanaged devices.

You need to package the driver code into a JAR file and add it to the Jboss 7 shared module directory. Junos Space Network Management Platform accesses the driver class using module based class loading. You also need to make an entry into the driver registration XML file. Junos Space Network Management Platform reads this XML file when JBOSS starts and populates this XML file into the database. The XML file should include the following parameters:

Parameter	Description
Name	Name of the device driver.
Vendor	Vendor of the device.
DeviceFamily	The family of the device.
Platform	The platform of the device.
DriverClassName	The full class name of the class which extends the driver class from the device platform.
MgtAttrColl	Holds a collection of key-value pairs that are populated in the Advanced Properties section when creating an unmanaged device. This can be omitted if you want to enter the key-value pairs when creating an unmanaged device.
IsDefaultForVendor	When set to true, the driver is used as a default driver for devices of other device family or platform, but from the same vendor.
IsDefaultForFamily	When set to true, the driver is used as a default driver for devices of a different platform, but from the same vendor and device family.

Creating an unmanaged device from a vendor other than Juniper Networks also creates a tag for that vendor (for example, CISCO) and assigns that tag to the device. If you have successfully installed the device driver on Junos Space Network Management Platform, you will be able to fetch the information related to physical and logical interface, manually resynchronize inventory data, and edit the unmanaged device configuration using the Junos Space applications. You can view the changes in the unmanaged device configuration using the View Configuration Change Log action.

- Related Documentation**
- [Importing Devices on page 74](#)
  - [Device Discovery Overview on page 77](#)
  - [Working With Managed Devices on page 80](#)
  - [Working With Unmanaged Devices on page 80](#)
  - [Working With Discovered Devices on page 81](#)

---

## Working With Managed Devices

The service delivery gateway (SDG) devices that are administered, maintained, and monitored from the Edge Services director application are called *managed devices*. For such devices, you can monitor alarms and events, create service templates and assign to the devices, and modify assigned KPI templates. You can perform the following tasks with managed devices from the Service Gateways—Managed Service Gateways page under the Build mode of the GUI interface:

- Remove the managed SDG devices and mark them unmanaged.
- Modify the SDG association with the SDG group and KPI template.
- Modify KPI templates for the SDGs.
- View configuration and compare the configurations of up to four SDGs.
- Export the managed device details to a CSV file.

- Related Documentation**
- [Importing Devices on page 74](#)
  - [Device Discovery Overview on page 77](#)
  - [Unmanaged Devices Overview on page 78](#)
  - [Working With Unmanaged Devices on page 80](#)
  - [Working With Discovered Devices on page 81](#)

---

## Working With Unmanaged Devices

The service delivery gateway (SDG) devices that are not managed and monitored from the Edge Services director application are called *unmanaged devices*. In certain network topologies, you might require certain devices to be configured individually before they are added to the management application. You can perform the following tasks with unmanaged devices from the Service Gateways—Managed Service Gateways page under the Build mode of the GUI interface:

- Bring the unmanaged devices back into the administration of Edge Services Director.
- View discovery log details.
- Searching and filtering unmanaged devices.
- Modify KPI templates for the SDGs.

- Related Documentation**
- [Importing Devices on page 74](#)
  - [Device Discovery Overview on page 77](#)
  - [Unmanaged Devices Overview on page 78](#)
  - [Working With Managed Devices on page 80](#)
  - [Working With Discovered Devices on page 81](#)

---

## Working With Discovered Devices

You can create discovery profiles to specify the parameters to be used for a discovery job, schedule the discovery of devices, modify discovery profiles, and view profile details. You can perform the following tasks with discovered devices from the Service Gateways—Discovered Devices page under the Build mode:

- Create and modify discovery profiles.
- Delete previously created discovery profiles.
- View discovery profile details.
- Schedule the discovery of devices.

- Related Documentation**
- [Importing Devices on page 74](#)
  - [Device Discovery Overview on page 77](#)
  - [Unmanaged Devices Overview on page 78](#)
  - [Working With Managed Devices on page 80](#)
  - [Working With Unmanaged Devices on page 80](#)

---

## Managing Jobs as a System Task

Edge Services 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, Edge Services Director enables you to view special pre-filtered versions of this page from various other tasks, such as View Discovery Status or View Image Deployment 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 the **My Jobs** icon located on the top right of the Edge Services Director banner.

The My Jobs report appears. The My Jobs report displays your 25 most recent jobs. The jobs displayed in the My Jobs report provide information about the status of the

job, percentage completion of the job, the name of the job, and the job ID. The date and time represents the date and time when the job failed (in case the job failed) and the date and time when the job succeeded (in case the job succeeded).

2. You can also click a job in the My Jobs report to view the job on the Job Management page. Clicking the job ID filters the Job Management page to display only that job. To view jobs details, click **Manage Jobs**. The Job Management page appears.
3. Click **Close** to exit the My Jobs page.

To view the job details:

1. in the My Jobs page, select a row and click **Show Details** or double-click a row.
2. 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 20 on page 60](#).

**Table 22: 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	The status of the job: <ul style="list-style-type: none"> <li>• Success—Job completed successfully</li> <li>• Failure—Job failed and was terminated</li> <li>• Job Scheduled—Job is scheduled but has not yet started</li> <li>• In progress—Job is has started, but not completed</li> <li>• Cancelled—Job is cancelled</li> </ul>
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

Table 22: Job Management Page Fields (continued)

Field	Description
Recurrence	The recurrent time when the job will be restarted.

You can clear your jobs from a list of your jobs when these jobs are no longer of interest to you.

To remove the jobs that you have initiated:

1. In the banner of the Junos Space user interface, click the **My Jobs** icon located at the top right.  
  
The My Jobs report appears. The My Jobs report displays your 25 most recent jobs.  
  
The jobs displayed in the My Jobs report provide information about the status of the job, percentage completion of the job, the name of the job, and the job ID. The date and time represents the date and time when the job failed (in case the job failed) and the date and time when the job succeeded (in case the job succeeded).
2. Perform one of the following actions:
  - Click the **Clear Job** icon that appears to the right of the job to remove a job.
  - Click **Clear All My Jobs** at the top of the My Jobs report to clear all your jobs displayed on the My Jobs list.



**NOTE:** Clearing a job from the My Jobs report does not affect the job itself, but only updates the My Jobs view.

3. Click **Close** to exit the My Jobs page.

**Related Documentation**

- [Importing Devices on page 74](#)
- [Device Discovery Overview on page 77](#)
- [Unmanaged Devices Overview on page 78](#)
- [Working With Managed Devices on page 80](#)
- [Working With Unmanaged Devices on page 80](#)
- [Working With Discovered Devices on page 81](#)



## CHAPTER 7

# Managing Service Delivery Gateways and Groups

- [Discovering Devices on page 85](#)
- [Comparing Configuration Settings of Devices on page 92](#)
- [Exporting Managed Device Details to a CSV File on page 94](#)
- [Changing an Unmanaged Device to a Managed Device on page 95](#)
- [Modifying the SDG Group and KPI Templates for a Device on page 97](#)
- [Scheduling the Discovery of Devices on page 98](#)
- [Creating Service Gateway Groups on page 99](#)
- [Managing Service Gateway Groups on page 101](#)
- [Viewing the Service Gateway Details on page 102](#)
- [Searching Unmanaged Devices on page 104](#)
- [Viewing the List of Discovered, Managed, and Unmanaged Devices on page 106](#)
- [Changing a Managed Device to an Unmanaged Device on page 112](#)
- [Modifying Discovery Profiles on page 113](#)
- [Deleting Discovery Profiles on page 114](#)
- [Systems of Record in Junos Space Overview on page 115](#)
- [Resynchronizing Managed SDGs with the Network on page 117](#)

## Discovering Devices

---

You can discover and synchronize physical devices such as MX Series routers that function as service delivery gateways in your network that are managed by Edge Services Director.



**NOTE:** On MX Series routers, Edge Services 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** in the Junos Space Platform page. Select **Network Application Platform** and click **Actions > Modify Application Settings**. Change SSH port for device connection field 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, select **Service Gateways**. The Service Gateways page is displayed. Click **Discover Devices** to create a discovery profile or a job, and to view the previously created discovery profiles

This topic describes:

- [Preparing MX Series Devices for Discovery on page 86](#)
- [Specifying a Discovery Profile and the Target Devices on page 87](#)
- [Specifying SNMP Probes on page 89](#)
- [Specifying Credentials on page 91](#)

## Preparing MX Series Devices for Discovery

Juniper Networks MX Series 3D Universal Edge Routers—MX240, MX480, and MX960—include all standard Ethernet capabilities as well as enhanced mechanisms for service providers to provision and support large numbers of Ethernet services in addition to all Layer 3 services. You can discover these routers and manage them as switching devices from Edge Services Director. However, before discovering these MX devices from Edge Services Director, you must ensure that the Junos OS running on the device is at the required level and that the network service mode is set to LAN.

To prepare an MX Series device for discovery:

1. Log in to the MX Series device by using the CLI.
2. Ensure that the device is running a version of Junos OS that is compatible with Edge Services Director. Use the operational mode command **show version** to determine the Junos OS software release.
3. Commit your changes.

The MX Series device is now discoverable from Edge Services Director.

## Specifying a Discovery Profile and the Target Devices

You can add devices to Edge Services Director for device discovery by using the **Add** icon on the Service Gateways page. A discovery profile is created, which is a discovery job that contains the list of devices and its properties to be retrieved and added to the Edge Services Director database.



**NOTE:** If you want to discover and manage MX Series devices—MX240, MX480, and MX960—from Edge Services Director, you must first make these devices discoverable. For more details see [“Preparing MX Series Devices for Discovery” on page 86](#).

To specify a discovery profile and the target devices that you want Edge Services Director to discover:

1. From the View selector, select **Gateway View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the View selector, select **Gateway View** or **Device View**. The workspaces that are available in this view are displayed. The Gateway view displays the service delivery gateway (SDG) groups and the SDGs that are part of the high availability pair in an SDG group. The Device view displays the SDGs based on the device type, and within the device type, the devices are organized by the device model. For example, all models of MX960 routers are grouped together under one node in the tree.
4. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
5. From the View pane, select the All Network item in Gateway view. If you are in Device view, click the plus sign (+) beside the My Network item in the View pane to expand the tree and select the device node you want.
6. From the task pane in Gateway view, select **Services Gateways**.  
The Service Gateways page is displayed.



**NOTE:** Alternatively, you can select **Device View** from the View selector, click the **Build** icon on the banner, and select **Discover Devices** from the task pane to open the Discovery Profiles window to discover and manage devices.

7. From the task pane, select the **Discover Gateway** option. You need not click this button if you are launching the Service Gateways page by navigating from another page or another mode, such as Deploy or Monitor. It is displayed by default. You must click this button only if you are viewing unmanaged or managed SDGs or devices.
8. Click the **Add** icon. The Discovery Profile window appears.
9. In the **Name** field, enter a name for the device discovery job. No name is shown by default. A job or profile name cannot exceed 128 characters and can contain only letters, numbers, spaces, and some special characters. The special characters allowed are hyphen (-), underscore (\_), period (.), at (@), single quote ('), forward slash (/), and ampersand (&).
10. (Optional) In the **Description** field, type a user-defined description. (a minimum of 2 characters and a maximum limit of 255 characters). The description cannot exceed 256 characters and cannot contain hyphens. The operators who use the profile rely on the description for information on the discovery job.
11. To add individual devices by specifying the IP address credentials, click **Add** in the IP Details table.

The IP Details dialog box appears.

12. Choose one of the following options to specify the target devices:
  - Select the **IP Address** 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.
  - Click **Save** to save the target devices that you specified. When you have added all target devices that you want Edge Services Director to discover, click **Save** in the Discovery Profile window.

The IP Details section displays the addresses of the configured target devices.

13.
  - To edit a target device, select the box that displays with an icon for each added device in the IP Details section 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 the box that displays with an icon for each added device in the IP Details section 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.
14. (Optional) You can proceed to specify the SNMP probes and credentials for the added devices.

## Specifying SNMP Probes

You can specify an SNMP probe to connect to and discover the devices in a network.

To add a probe:

1. From the View selector, select **Gateway View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the task pane, select **Services Gateways**. The Service Gateways page is displayed.
4. Select the **Discover Gateway** option.
5. Click the **Add** icon. The Discovery Profile window appears.
6. Click the **Add** icon in the SNMP Details table. The SNMP Details dialog box is displayed.
7. Select one of the following options and enter the appropriate value in the field provided.
  - Select **SNMP V1/V2C** and specify the community string in the **Community** field.  
The SNMP v1/v2c community string *public* is available by default. The SNMP v1/v2c community string is based on the community string configured on the devices in your network.
  - Select **SNMP V3** and enter the information in the fields provided.
    - a. Enter the SNMP V3 username in the **Username** field.
    - b. Select the privacy protocol (the encryption standard for the SNMP user) from the **Privacy type** list.  
The available options are **AES128**, **DES**, and **None**.
    - c. Enter the password used to generate the key used for encryption in the **Privacy password** field.  
The password must be at least eight characters long. You can include all character classes in a password (alphabetic, numeric, and special characters) except control characters.

- d. Select the authentication type for the SNMP user from the **Privacy type** drop-down list.

The available options are **MD5**, **SHA1**, and **none**.

- e. Enter the password used to generate the key used for authentication in the **Authentication password** field.

The password must be at least eight characters long. You can include all character classes in a password (alphabetic, numeric, and special characters) except control characters.

8. Click **Save** to close the **SNMP Details** dialog box and add the SNMP probe to the **SNMP Settings** list.

The **SNMP Details** section of the Discovery Profile page displays the configured SNMP settings.

You can also click **Cancel** to close the **SNMP Details** dialog box without adding any SNMP probes.

To edit an SNMP probe:

1. Select the SNMP probe that you want to edit and click the **Modify** icon [slanted pencil] to open the **SNMP Details** dialog box.
2. Select one of the following options and enter the appropriate value in the field provided.

You can choose to edit the existing values in the selected SNMP version, or you can select a different SNMP version and enter the desired values.

- Select **SNMP V1/V2C** and specify the community string in the **Community** field. You can enter “public”, “private”, or a predefined string.
- Select **SNMP V3** and enter the information in the fields provided.
  - a. Enter the SNMP version 3 username in the **Username** field.
  - b. Select the privacy protocol—that is, the encryption standard for the SNMP user—from the **Privacy type** list. The available options are **AES128**, **DES**, and **None**.
  - c. Enter the password used to generate the key used for encryption in the **Privacy password** field. The password must be at least eight characters long. You can include all character classes in a password (that is, alphabetic, numeric, and special characters) except control characters.
  - d. Select the authentication type for the SNMP user from the **Privacy type** drop-down list.

The available options are **MD5**, **SHA1**, and **none**.

- e. Enter the password used to generate the key used for authentication in the **Authentication password** field.

The password must be at least eight characters long. You can include all character classes in a password (that is, alphabetic, numeric, and special characters) except control characters.

3. Click **Modify** to save your changes and close the **SNMP Details** dialog box.

The **SNMP Details** section displays the configured SNMP settings.

Alternatively, click **Cancel** to close the dialog box without editing any SNMP probes.

To delete an SNMP probe:

1. Select the SNMP probe that you want to delete in the SNMP Details section and click the **Delete** icon [X].
2. The SNMP probe is removed from the SNMP Details section.

## Specifying Credentials

Optionally, specify an administrator name and password to establish the SSH connection for each target device that you configured. If you are using key-based authentication, you do not need to do this step. To specify the credentials:



**NOTE:** Alternatively, you can select **Device View** from the View selector, click the **Build** icon on the banner, and select **Discover Devices** from the task pane to open the Discovery Profiles window to discover and manage devices.

1. From the View selector, select **Gateway View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.

The functionalities that you can configure in this mode are displayed in the task pane.

3. From the task pane, select **Services Gateways**.

The Service Gateways page is displayed.

4. Select the **Discover Gateway** option.
5. Click the **Add** icon. The Discovery Profile window appears.

6. Click the **Add** icon in the User Details table. The User Details dialog box is displayed.
7. Specify the administrator username and password, and confirm the password. The name and password must match the name and password configured on the device.  
Save the user name and password that you specified by selecting **Save**.  
The User Details section of the Discovery Profile window displays the administrator user names that you configured.

**Related Documentation**

- [Importing Devices on page 74](#)
- [Device Discovery Overview on page 77](#)
- [Unmanaged Devices Overview on page 78](#)
- [Working With Managed Devices on page 80](#)
- [Working With Unmanaged Devices on page 80](#)
- [Working With Discovered Devices on page 81](#)

---

## Comparing Configuration Settings of Devices

---

You can compare the configuration of a SDG with any other SDG that is discovered. You can contrast and view the configuration settings of a master device with another master device in two SDG high availability pairs, of a standby device with another standby device, or of a master and a standby device in the same SDG pair. You can compare the settings of up to four devices simultaneously. After you select the desired device and initiate the comparison operation, you can also add more devices until the maximum limit of four devices is reached for comparison.

The services are displayed and the objects or components of each service instance are also shown. You can filter and view a specified service component or the components of all services. A red minus mark denotes that the particular parameter or element is not available or configured on the specified device. A green tick mark denotes that the particular parameter is available on the corresponding device. For attributes that can contain values, the associated values are shown for the appropriate device. Otherwise, for attributes that can either be disabled or enabled, the red minus icon or green tick icon provides a graphical indication. To compare the configuration settings of two or more devices

To compare the configuration settings of two or more devices

1. From the View selector, select **Gateway View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.

3. From the task pane, select **Services Gateways**.

The Service Gateways page is displayed.

4. From the task pane, select the **Managed Gateway** option.

The list of managed SDGs appears.

5. Perform the following from the View pane:

- a. Click the **All Network** item. The list of discovered SDG groups that contain devices are displayed.
- b. Click the + sign to expand the tree beside the SDG groups. At least, two discovered devices must be present for which you want to compare the services configured. You must select the pair of devices within the SDG group.
- c. Select the SDG device pair. The list of managed devices are displayed.



**NOTE:** You must select a minimum of two devices, if the SDG is not a high availability pair of devices.

6. Select the **Compare Configuration** option from the task pane that displays the configuration settings contrasted between the two devices you selected. The Compare Configuration View page is displayed.

The list of service types are displayed in the leftmost column of the table.

**Figure 7: Compare Configuration View Page**

	mobst480w	mobst480x
ADC	✓	✓
b1	✓	✓
TLS	✓	✓
tlb_sdq	✓	✓
tlb_sdq_v6	✓	✓
SFW	✓	✓
IPv6-SFW	✓	✓
SFW Rules	✓	✓
IPv6-SFW	✓	✓
Match Direction	input	input
Term	✓	✓
ACCEPT	✓	✓
REJECT	✓	✓
Next-hop Service Inside Interface	✓	✗
sp-1/1/0.99	✓	✓
Next-hop Service Outside Interface	✓	✓
CGNAT	✓	✓
NAPT44-SS1	✓	✓
NAPT44-SS2	✓	✓

7. From the View drop-down list, select one of the following options:

- **All Configs**—Causes all of the services configurations to be displayed.
  - **ADC**—Causes the application delivery service components to be displayed.
  - **TLB**—Causes the traffic load balancer service components to be displayed.
  - **SFW**—Causes the stateful firewall service components to be displayed.
  - **CGNAT**—Causes the carrier-grade NAT service components to be displayed.
8. From the **Select Devices** list, select more devices up to the maximum limit of four devices. Click the cross mark beside each selected device if you want to remove it from the list, and select a different device instead.
  9. Click the **Compare Configuration** icon adjacent to the Select Devices drop-down list. The devices you select cause their configurations to be displayed on the page.
  10. Click **Refresh** to refresh the displayed configuration. This action reads the device configuration of the selected device, perform the comparison, and updates the display to highlight the differences between the devices.
  11. Click **Close** after you finish viewing the configuration comparison and to return to the page that lists the devices.

**Related Documentation**

- [Discovering Devices on page 85](#)
- [Exporting Managed Device Details to a CSV File on page 94](#)
- [Changing an Unmanaged Device to a Managed Device on page 95](#)
- [Modifying the SDG Group and KPI Templates for a Device on page 97](#)
- [Scheduling the Discovery of Devices on page 98](#)

---

## Exporting Managed Device Details to a CSV File

The Service Gateways—Managed Service Gateways page lists all of the devices that are currently being controlled and provisioned by the Edge Services Director application. You can export Service Now device data to CSV and Excel file formats. A CSV file is a plaintext file that stores each data record separated by a comma. 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 devices listing and delineates the fields with commas so that it imports into popular spreadsheet programs

To export the device data in CSV format:

1. From the View selector, select **Gateway View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.

The functionalities that you can configure in this mode are displayed in the task pane.

3. From the View pane, select the All Network item in Gateway view.

4. Select **Services Gateways** from the task pane.

The Service Gateways page is displayed.

5. Select the **Manage Service Gateways** option.

The Service Gateways—Managed Service Gateways page appears with the list of managed SDGs. If the SDGs are configured in a high availability pair, the details of both the devices in the pair are exported.

6. Select the check box next to the managed SDG or SDG pair that you want to export to a CSV file.

7. Click the **Export Service Gateway Details** icon.

The Export SDG dialog box is displayed.

8. Export the device inventory information to the CSV file. You can export information about selected devices or export information about all of the devices managed by Junos Space. Click either the **Export Selected** button or the **Export All** button to begin creating the CSV file.

9. Download the resulting CSV file. Now that you have the CSV report, you can import that CSV file into other applications such as those you use for asset management.

**Related  
Documentation**

- [Discovering Devices on page 85](#)
- [Comparing Configuration Settings of Devices on page 92](#)
- [Changing an Unmanaged Device to a Managed Device on page 95](#)
- [Modifying the SDG Group and KPI Templates for a Device on page 97](#)
- [Scheduling the Discovery of Devices on page 98](#)

---

## Changing an Unmanaged Device to a Managed Device

The Service Gateways—Unmanaged Devices page lists all of the devices that are currently not being controlled and provisioned by the Edge Services Director application. You can enable management of such devices from Edge Services Director.

To convert an unmanaged device to a managed device:

1. From the View selector, select **Gateway View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the View pane, select the All Network item in Gateway view.
4. Select **Services Gateways** from the task pane.  
The Service Gateways page is displayed.
5. Select the **Unmanaged Gateway** option.  
The list of unmanaged SDGs appears.
6. Select the check box next to the unmanaged SDG that you want to bring under Edge Services Director administration.
7. Click the **Manage Service Gateways** icon.  
The Manage Service Gateway dialog box is displayed, with the SDG name and description of the SDG displayed in the respective fields.
8. From the Service Gateway Group list, select the SDG group with which you want to associate the SDG to be managed. Alternatively, click the green plus sign (+) beside the list to create a new SDG group. For more information, see [“Creating Service Gateway Groups” on page 99](#).
9. From the KPI list, select the KPI template to be associated with the selected SDG. Alternatively, click the green plus sign (+) beside the list to create a new KPI template by modeling it on an existing, system-defined KPI template. For more information, see [“Cloning a KPI Template” on page 121](#).
10. Click the **Apply** button to save the settings.  
An informational message is displayed stating that the settings are successfully applied to the selected device.
11. Click the **Manage** button to classify the device as a managed device.  
The Manage Status dialog box is displayed with the name of the SDG device and the status.

The device becomes a managed devices and is removed from the unmanaged devices listing. The device is added to the list of managed devices.

12. Select the **Auto Refresh** check box at the bottom of the Service Gateways -- Unmanaged devices page to indicate that the page needs to be refreshed automatically. The default value is three seconds. When you deselect this check box, the page is not refreshed periodically by itself; instead you can click the **Refresh** icon to update the page contents for viewing. Also, when you deselect the auto-refresh functionality, a message is displayed to denote that auto-refresh is turned off and of the number of updates that have not been viewed since the last refresh operation. The date and time at which the page was last updated is shown.

#### Related Documentation

- [Discovering Devices on page 85](#)
- [Comparing Configuration Settings of Devices on page 92](#)
- [Exporting Managed Device Details to a CSV File on page 94](#)
- [Modifying the SDG Group and KPI Templates for a Device on page 97](#)
- [Scheduling the Discovery of Devices on page 98](#)

## Modifying the SDG Group and KPI Templates for a Device

The Service Gateways—Managed Service Gateways page lists all of the devices that are currently being controlled and provisioned by the Edge Services Director application. You can modify the KPI template and the SDG group that are mapped to the standalone SDG or the high availability pair of SDGs.

To modify the SDG details:

1. From the View selector, select **Gateway View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the View pane, select the All Network item in Gateway view. You can also click the plus sign (+) beside the All Network item in the View pane to expand the tree and select the device node you want.
4. From the task pane, select **Services Gateways**.  
The Service Gateways page is displayed.
5. Select the **Manage Service Gateways** option.

The list of managed SDGs appears. If the SDGs are configured in a high availability pair, the details of both the devices in the pair are exported.

6. Select the check box next to the managed SDG or SDG pair that you want to modify.

The SDG is available for modification.

7. Click the down arrow in the **Modify** button at the top of the table of managed SDGs that are listed, and select **Service Gateway**.

The Modify Service Gateway dialog box is displayed.

8. In the **Description** field, edit the user-defined comment or description as needed.

9. From the **SDG Group** and **KPI Template** lists, select the SDG group and KPI template you want to associate the SDGs with.

10. Click **Modify** to save the edited settings.

#### Related Documentation

- [Discovering Devices on page 85](#)
- [Comparing Configuration Settings of Devices on page 92](#)
- [Exporting Managed Device Details to a CSV File on page 94](#)
- [Changing an Unmanaged Device to a Managed Device on page 95](#)
- [Scheduling the Discovery of Devices on page 98](#)

---

## Scheduling the Discovery of Devices

The Discovery Profiles page displays the discovery jobs that you have previously created. After you specify a discovery profile that contains the list of devices or SDG hosts that need to be retrieved and added to the Edge Services Director database to facilitate easier administration, you must configure the discovery operation. You can choose to discover the devices immediately or to plan for the discovery to happen at a specified future time.

To specify the scheduling details for discovering devices:

1. From the View selector, select **Gateway View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the View pane, select the All Network item in Gateway view.

4. From the task pane, select **Services Gateways**.

The Service Gateways page is displayed.

5. Select the **Discover Gateway** option.

The list of discovery profiles appears.

6. Select the check box next to the discovery profile that you want to schedule for discovering devices.

7. Click **Discover Device(s) Now** above the table of displayed profiles if you want to discover the devices immediately. A dialog box confirming that discovery has been initiated is displayed.

8. Alternatively, click **Discover Device(s) Later** 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. The calendar picker and the drop-down list for selection of time are displayed beside the **Discover Device(s) Later** button.

Select a date from the calendar and the time from the list. The time is shown in increments of 15 minutes from 12:00 AM - 11:45 PM.



**NOTE:** The selected time in the scheduler corresponds to Junos Space server time but is mapped to the local time zone of the client computer.

9. Click the **Schedule** button. The discovery process occurs at the designated time on the specified day.

After you have configured the device discovery options, you can view the device discovery status from the Service Gateways page with the Discover Devices view.

#### Related Documentation

- [Discovering Devices on page 85](#)
- [Comparing Configuration Settings of Devices on page 92](#)
- [Exporting Managed Device Details to a CSV File on page 94](#)
- [Changing an Unmanaged Device to a Managed Device on page 95](#)
- [Modifying the SDG Group and KPI Templates for a Device on page 97](#)

---

## Creating Service Gateway Groups

A service delivery gateway (SDG) device can be combined into a group of devices for easier and streamlined administration. You can create an SDG group for a particular domain or zone in your network, or for any logical bundling that is needed.

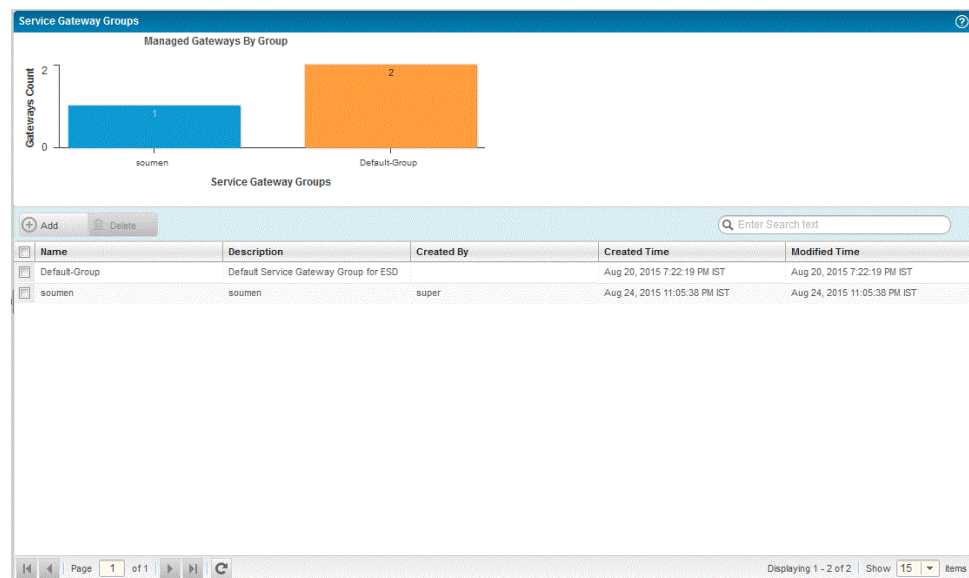
The Service Gateway Groups page displays all of the created SDG groups.

To create a SDG group:

1. From the View selector, select **Gateway View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the View pane, select the All Network item in Gateway view.
4. Select **Services Gateway > Groups** from the task pane.

The Service Gateway Groups page is displayed.

*Figure 8: Service Gateway Groups Page*



5. Click the **Add** icon.  
The Create SDG Group dialog box appears.
6. In the **Name** field, enter a unique name for the template (limit of 63 alphanumeric characters without spaces).
7. (Optional) Enter a description of the template in the **Description** field (limit of 255 characters).
8. Click **Create** to save the SDG group and return to the page that displays all the configured groups.

- Related Documentation**
- [Managing Service Gateway Groups on page 101](#)
  - [Searching Unmanaged Devices on page 104](#)
  - [Viewing the List of Discovered, Managed, and Unmanaged Devices on page 106](#)
  - [Changing a Managed Device to an Unmanaged Device on page 112](#)
  - [Modifying Discovery Profiles on page 113](#)
  - [Deleting Discovery Profiles on page 114](#)

---

## Managing Service Gateway Groups

A service delivery gateway (SDG) device can be combined into a group of devices for easier and streamlined administration. You can create an SDG group for a particular domain or zone in your network, or for any logical bundling that is needed. When you modify the details of a managed device, you can change the SDG group that is assigned to it. The listing of SDG groups provides you with an agglomerative view of all the SDGs present in a particular group at a point in time.

The Service Gateway Groups page displays all of the created SDG groups. You can perform the following tasks on this page:

- Create SDG groups
- Delete SDG groups (You cannot delete the default group named Default-Group.)
- Search SDG groups

To view the configured SDG groups:

1. From the View selector, select **Gateway View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the View pane, select the All Network item in Gateway view.
4. Select **Services Gateway > Groups** from the task pane.

The Service Gateway Groups page is displayed.

The top half of the page displays a bar chart. The SDG names are displayed on the horizontal axis and the count of SDGs are displayed on the vertical axis. A color-coding format is used to represent the bars on the chart. Mouse over each bar in the chart to highlight and display the SDG name

The following fields are displayed in the lower half of the page:

Field	Description
Name	Unique name of the SDG group.
Description	User-defined description of the SDG group.
Created By	Name of the user that created the DG group.
Created Time	Date and time at which the SDG group was created.
Modified Time	Date and time at which the SDG group was last updated.

#### Related Documentation

- [Creating Service Gateway Groups on page 99](#)
- [Searching Unmanaged Devices on page 104](#)
- [Viewing the List of Discovered, Managed, and Unmanaged Devices on page 106](#)
- [Changing a Managed Device to an Unmanaged Device on page 112](#)
- [Modifying Discovery Profiles on page 113](#)
- [Deleting Discovery Profiles on page 114](#)

## Viewing the Service Gateway Details

---

You can view the details of managed SDGs, such as the Junos OS version running on the device and the model number of the device, and the names and types of different services, such as ADC, TLB, stateful firewall, and CGNAT, on the managed SDGs. The high-level view you can obtain enables you to examine the existing system configuration and services on a device and modify them according to your topology needs by navigating to the gateway and service workspaces.

To view the details of managed service gateways:

1. From the View selector, select **Gateway View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the View pane, select the All Network item. Expand the tree to select the SDG in an SDG group.
4. From the task pane, select **Services Gateways**.

The Service Gateways page is displayed.

- From the task pane, select the **Managed Gateway** option.

The list of managed SDGs appears.

- In the main window, click the plus sign (+) next to the SDG pairs to expand the tree and view the pair of devices in the SDG group or pair. Select the check box next to the individual SDG for which you want to view the device details.

The Service Gateway Details page is displayed.

**Figure 9: Service Gateway Details Page**

**Service Gateway Details**

Name : mobst480w-mobst480x  
 Description :  
 Service Gateway Group : Default-Group  
 KPI Template : SystemDefault\_12\_1

**Master**

Host Name : mobst480w  
 Version : 12.1X43.11  
 Platform : MX480  
 IP Address : 10.213.0.1  
 Connection Status : up

**Services**

Name	Type	ServicePic
lib1	ADC	[ms-0/0/0]
IPv6-SFW	SFW	[sp-1/1/0.100, sp-1/1/0.99]
tlb_sdg	TLB	[ms-1/0/0.0]
tlb_sdg_v6	TLB	[ms-1/0/0.0]
NAPT44-SS1	CGNAT	[sp-2/0/0.1000, sp-2/0/0.100]
NAPT44-SS2	CGNAT	[sp-2/1/0.1000, sp-2/1/0.100]

Close

Table 23 on page 103 describes the fields displayed on the Service Gateway Details page.

**Table 23: Fields on the Service Gateway Details Page**

Field	Description
Name	Hostnames of the SDGs in the SDG group.
Description	User-defined description of the SDG group.
Service Gateway Group	Name of the SDG group.
KPI Template	Name of the KPI template associated with the SDG.

*Table 23: Fields on the Service Gateway Details Page (continued)*

Field	Description
Host Name	Hostname of the device.
Version	Software version the device is running.
IP Address	IP address configured for the device.
Platform	Model number of the device.
Connection Status	Device's state: <ul style="list-style-type: none"> <li>• UP—Edge Services Director can communicate with the device.</li> <li>• DOWN—Edge Services Director cannot communicate with the device.</li> </ul>
Services	Displays the details of configured services.
Name	Name of the service configured on the SDG.
Type	Type of the service configured on the SDG, such as ADC, TLB, stateful firewall, or CGNAT.
Service Pic	Services PIC and interface details, such as multiservices PIC or adaptive services PIC with the FPC slot, PIC, and port attributes

7. Click **Close** after you finish viewing the gateway details.

#### Related Documentation

- [Discovering Devices on page 85](#)
- [Exporting Managed Device Details to a CSV File on page 94](#)
- [Changing an Unmanaged Device to a Managed Device on page 95](#)
- [Modifying the SDG Group and KPI Templates for a Device on page 97](#)
- [Scheduling the Discovery of Devices on page 98](#)

## Searching Unmanaged Devices

The Service Gateways—Unmanaged Devices page lists all of the devices that are not currently being controlled and provisioned by the Edge Services Director application. Use the search mechanism to filter and isolate information about a specific device. Use this facility to specify complex sorting and filtering criteria for all devices.

The search functionality is an effective tool that helps to search a discovered device from the available list, based on various criteria such as Node Name, IP Address, Service available, HA State. The search utility also supports wildcards. For example, you can search and sort devices of the same platform type or OS version.

To search and filter discovered devices:

1. From the View selector, select **Gateway View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the View pane, select the All Network item in Gateway view.
4. Select **Services Gateways** from the task pane.  
The Service Gateways page is displayed.
5. Click the **Unmanaged Devices** button.  
The list of unmanaged SDGs that have been discovered appears.
6. Click the double right arrow icon to the far right of the screen in the line of toolbar icons.  
The Search drop-down list and Name field are displayed.
7. Enter the search criteria by selecting the parameter that you want to use to filter in the **Search** list. You can select one of the following values:
  - Host Name
  - IP Address
  - Platform Type
  - Version
  - HA State

Enter the value for the search parameter you selected in the text field adjacent to the drop-down list, and click the magnifying glass icon.

The page refreshes to display the devices that match the specified criterion.

8. To save the search criterion you specified for future purposes, enter a name for the search in the **Name** drop-down list and click the Save icon (floppy drive icon) to save the search filter.

You can also edit or delete the search filters by selecting them from the **Name** drop-down list and clicking the **Edit** or **Delete** icons respectively.

**Related  
Documentation**

- [Creating Service Gateway Groups on page 99](#)
- [Managing Service Gateway Groups on page 101](#)

- [Viewing the List of Discovered, Managed, and Unmanaged Devices on page 106](#)
- [Changing a Managed Device to an Unmanaged Device on page 112](#)
- [Modifying Discovery Profiles on page 113](#)
- [Deleting Discovery Profiles on page 114](#)

---

## Viewing the List of Discovered, Managed, and Unmanaged Devices

There are three types of views displayed on the Service Gateways page, depending on whether you select the **Discover Devices**, **Unmanaged Devices**, or **Manage Service Gateways** button. These views enable you to examine the discovery profiles, devices managed by Edge Services Director, and SDGs that are not currently managed.

To view the list of discovered, managed, and unmanaged devices:

1. From the View selector, select **Gateway View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the View pane, select the All Network item in Gateway view.
4. Select **Services Gateways** from the task pane. The Service Gateways page is displayed.
5. Select the **Discover Gateway** option.

The Service Gateways—Discovered Devices view is displayed. The following table describes the fields in this view:

Field	Description
Discovery Profile	Unique name of the discovery profile.
Description	User-defined description of the profile.
Created By	Name of the user that created the profile.
Created Time	Date and time at which the profile was first created.
Modified Time	Date and time at which the profile was last updated.

Field	Description
Last Execution Status	<p>Status of the last discovery process performed for a profile. This column indicates whether the discovery was completed or aborted.</p> <p>Click the link in this column to view extensive details about a discovery profile. The <i>DiscoveryProfileName: Last Execution Status</i> window appears. The Discovery Profile -- Last Execution Status window is divided into two panes. The top half of the page displays a bar chart that denotes the number of devices in each of the states during their discovery and process of addition to the Edge Services Manager database. The x-axis displays the different states of the devices during the discovery process and the y-axis denotes the number of devices corresponding to each state. Mouse over the different segments of the bar chart to highlight and view the total number of devices in each state. Click any of the states in the color-coding legend box to display only the details pertaining to that state to be shown beneath the graph in the table. The following color-coding legend denotes the devices in the different states:</p> <ul style="list-style-type: none"> <li>• Dark green—Denotes the devices that are already added and managed by Edge Services Director</li> <li>• Light green—Denotes the devices for which discovery succeeded</li> <li>• Yellow—Denotes the devices for which discovery is in progress</li> <li>• Red—Denotes the devices for which discovery failed</li> <li>• Dark orange—Denotes the devices for which synchronization failed</li> <li>• Light orange—Denotes the devices for which a timeout has occurred in the connection from Edge Services Director</li> <li>• Pink—Denotes the devices for which discovery is skipped</li> </ul> <p>See <a href="#">Table 24 on page 107</a> for a description of the fields shown in the table of this dialog box.</p>
In Progress	Indicates whether a discovery job is currently running.
Scheduled	Indicates whether a discovery is scheduled for a future time.

**Table 24: Fields in the Last Execution Status Dialog Box**

Field	Description
Host Name	Device name
IP Address	IP address
Description	User-defined description of the profile

*Table 24: Fields in the Last Execution Status Dialog Box (continued)*

Field	Description
Status	<p>Indicates whether the device's configuration is in sync with Edge Services Director's version:</p> <ul style="list-style-type: none"> <li>• <b>Already Added</b>—Denotes that the device has been discovered and are being currently managed.</li> <li>• <b>Succeeded</b>—Denotes the devices for which discovery succeeded.</li> <li>• <b>Discovered</b>—Denotes the devices that have been discovered and retrieved.</li> <li>• <b>Failed</b>—Denotes the devices for which discovery failed.</li> <li>• <b>Sync Failed</b>—Denotes devices for which synchronization of the device configuration with the Edge Services Director database failed.</li> <li>• <b>Timedout</b>—Denotes devices for which the establishment of connection from Edge Services Director to the devices failed because a timeout occurred.</li> <li>• <b>Skipped</b>—Denotes devices that were not discovered and were skipped from being brought into Edge Services Director again because no configuration setting changes were observed on such devices</li> </ul>

6. Select the **Unmanaged Devices** option beneath Service Gateways in the task pane. The Service Gateways—Unmanaged devices view is displayed. The following fields are displayed in this view:

Field	Description
Host Name	Host name of the device.
Version	Software version the device is running.
IP Address	IP address configured for the device.
Device Family	Device family to which the device belongs. Hardware family such as Junos OS is displayed.
Platform	Model number of the device.
Current HA Status	<p>Present high availability status that indicates if the SDGs are configured in a redundancy pair. It denotes whether the device is a master or a standby device.</p> <p><b>NOTE:</b> Not applicable for devices that are not in a high availability pair.</p>
Deployment HA Status	<p>High availability status after deployment that indicates if the SDGs are configured in a redundancy pair. It denotes whether the device is a master or a standby device.</p> <p><b>NOTE:</b> Not applicable for devices that are not in a high availability pair.</p>
Connection Status	<p>Device's state:</p> <ul style="list-style-type: none"> <li>• <b>UP</b>—Edge Services Director can communicate with the device.</li> <li>• <b>DOWN</b>—Edge Services Director cannot communicate with the device.</li> </ul>

Field	Description
Peer IP	IP address of the peer device in a redundancy group.
ADC	Whether the ADC service is configured. A tick mark indicates the service is configured, and a gray minus sign indicates the service is unavailable.
TLB	Whether the TLB service is configured. A tick mark indicates the service is configured, and a gray minus sign indicates the service is unavailable.
CGNAT	Whether the CGNAT service is configured. A tick mark indicates the service is configured, and a gray minus sign indicates the service is unavailable.
SFW	Whether the stateful firewall service is configured. A tick mark indicates the service is configured, and a gray minus sign indicates the service is unavailable.
Routing Instance	Whether the routing instance is configured. A tick mark indicates the service is configured, and a gray minus sign indicates the service is unavailable.
Managed Status	<p>Indicates whether the device's configuration is in sync with Edge Services Director's version:</p> <ul style="list-style-type: none"> <li>• Connecting—The device is being contacted to establish a connection.</li> <li>• In Sync—The configuration on the device is in sync with the Edge Services Director configuration for the device.</li> <li>• Out Of Sync—The configuration on the device does not match the Edge Services Director configuration for the device. This state is usually the result of the device configuration being altered outside of Edge Services Director.</li> </ul> <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"> <li>• Synchronizing—The device configuration is in the process of being resynchronized.</li> <li>• Sync failed—An attempt to resynchronize an Out Of Sync device failed.</li> </ul>

Select the **Auto Refresh** check box at the bottom of the Service Gateways -- Unmanaged devices page to indicate that the page needs to be refreshed automatically. The default value is three seconds. When you deselect this check box, the page is not refreshed periodically by itself; instead you can click the **Refresh** icon to update the page contents for viewing. Also, when you deselect the auto-refresh functionality, a message is displayed to denote that auto-refresh is turned off and of the number of updates that have not been viewed since the last refresh operation. The date and time at which the page was last updated is shown.

7. Select the **Managed Service Gateways** option beneath Service Gateways from the task pane.

The Service Gateways—Managed Service Gateways view is displayed. The following fields are displayed in this view:

Field	Description
Name	Host names of the devices in an SDG high availability pair or the standalone SDG device.
Service Gateway Group	Name of the SDG group associated with the device.
KPI Template	Name of the KPI template associated with the SDG.
Host Name	Host name of the device.
Version	Software version the device is running.
IP Address	IP address configured for the device.
Platform	Model number of the device.
Current HA Status	High availability status that indicates if the SDGs are configured in a redundancy pair. It denotes whether the device is a master or a standby device.  <b>NOTE:</b> Not applicable for devices that are not in a high availability pair.
PM Status	Whether the performance management utility is running and statistical counters are being computed. If this field indicates Managed, it denotes that performance management is successfully operating on the device.  <b>NOTE:</b> If the PM Status field denotes a value that is other than Managed, you might need to examine the device and services settings to take the required corrective action for performance management to work properly. For example, you might need to change the device to unmanaged and perform troubleshooting.
PM Job Status	Whether the jobs to retrieve counters and values from devices to display performance management statistics are running. If this field indicates a value other than Running, you might need to stop the PM collection utility, rectify the settings, and restart the PM collection utility.
ADC	Whether the ADC service is configured. A tick mark indicates the service is configured, and a gray minus sign indicates the service is unavailable.
TLB	Whether the TLB service is configured. A tick mark indicates the service is configured, and a gray minus sign indicates the service is unavailable.
CGNAT	Whether the CGNAT service is configured. A tick mark indicates the service is configured, and a gray minus sign indicates the service is unavailable.
SFW	Whether the stateful firewall service is configured. A tick mark indicates the service is configured, and a gray minus sign indicates the service is unavailable.

Field	Description
Connection Status	<p>Device's state:</p> <ul style="list-style-type: none"> <li>• UP—Edge Services Director can communicate with the device.</li> <li>• DOWN—Edge Services Director cannot communicate with the device.</li> </ul>
Managed Status	<p>Indicates whether the device's configuration is in sync with Edge Services Director's version:</p> <ul style="list-style-type: none"> <li>• Connecting—The device is being contacted to establish a connection.</li> <li>• In Sync—The configuration on the device is in sync with the Edge Services Director configuration for the device.</li> <li>• Out Of Sync—The configuration on the device does not match the Edge Services Director configuration for the device. This state is usually the result of the device configuration being altered outside of Edge Services Director. 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.</li> <li>• Synchronizing—The device configuration is in the process of being resynchronized.</li> <li>• Sync failed—An attempt to resynchronize an Out Of Sync device failed.</li> </ul>



**NOTE:** After you bring a device under the control and provisioning of Edge Services Director, you might require to define services, policy filters, and modify certain configuration settings of the managed devices before you want to start the collection of performance management (PM) statistics and counters. You can select a device or the high availability pair of SDGs, and select **Start PM Collection**. An information message is displayed to indicate whether a successful start of the retrieval of monitoring details has occurred. Else, an error message denotes a failure in the attempt to start collection of monitoring information. By default, retrieval and computation of statistics is enabled. You can terminate the collection of PM statistics at any time by selecting a device or pair of devices, and selecting **Stop PM Collection**. The option to start and stop collection of monitoring statistics is a toggle button. Alternatively, to change a managed device to be unmanaged and remove it from the administration and monitoring of Edge Services Director, you can also right-click a device or a high availability pair of SDGs in the list of devices that are displayed in the Managed Service Gateways page, and select **Stop Managing**.

#### Related Documentation

- [Creating Service Gateway Groups on page 99](#)
- [Managing Service Gateway Groups on page 101](#)
- [Searching Unmanaged Devices on page 104](#)
- [Changing a Managed Device to an Unmanaged Device on page 112](#)
- [Modifying Discovery Profiles on page 113](#)

- [Deleting Discovery Profiles on page 114](#)

## Changing a Managed Device to an Unmanaged Device

---

The Service Gateways—Managed Service Gateways page lists all of the devices that are currently being controlled and provisioned by the Edge Services Director application. You can remove the management of such devices from Edge Services Director. For example, in a certain deployment, you might require certain device characteristics to be separately configured without a bulk application of settings. In such a case, you can mark the device as unmanaged, perform the configurations manually using the device CLI interface, and later decide to add it to the managed devices.

To convert a managed device to an unmanaged device:

1. From the View selector, select **Gateway View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the View pane, select the All Network item in Gateway view.
4. Click the **Manage Service Gateways** button.

The list of managed SDGs appears. If the SDGs are configured in a high availability pair, an attempt to unmanage such SDGs causes both the master and standby devices in the redundancy group to become unmanaged.



**NOTE:** After you bring a device under the control and provisioning of Edge Services Director, you might require to define services, policy filters, and modify certain configuration settings of the managed devices before you want to start the collection of performance management (PM) statistics and counters. You can right-click a device or the high availability pair of SDGs, and select **Start PM Collection**. An information message is displayed to indicate whether a successful start of the retrieval of monitoring details has occurred. Else, an error message denotes a failure in the attempt to start collection of monitoring information. By default, retrieval and computation of statistics is enabled. You can terminate the collection of PM statistics at any time by right-clicking a device or pair of devices, and selecting **Stop PM Collection**. The option to start and stop collection of monitoring statistics is a toggle button. Alternatively, to change a managed device to be unmanaged and remove it from the administration and monitoring of Edge Services Director, you can also right-click a device or a high availability pair of SDGs in the list of devices that are displayed in the Managed Service Gateways page, and select **Stop Managing**.

5. Select the check box next to the managed SDG or SDG pair that you want to remove from Edge Services Director administration.

6. Click the **Unmanage Service Gateway** icon.

The device becomes an unmanaged device and is removed from the managed devices listing. The device is added to the list of unmanaged devices.

#### Related Documentation

- [Creating Service Gateway Groups on page 99](#)
- [Managing Service Gateway Groups on page 101](#)
- [Searching Unmanaged Devices on page 104](#)
- [Viewing the List of Discovered, Managed, and Unmanaged Devices on page 106](#)
- [Modifying Discovery Profiles on page 113](#)
- [Deleting Discovery Profiles on page 114](#)

## Modifying Discovery Profiles

The Discovery Profiles page displays the discovery jobs that you have previously created. You can edit the properties of a discovery profile, such as adding more devices into a job or updating the SNMP settings.

To modify a configured discovery profile:

1. From the View selector, select **Gateway View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the View selector, select **Gateway View** or **Device View**. The workspaces that are available in this view are displayed. The Gateway view displays the service delivery gateway (SDG) groups and the SDGs that are part of the high availability pair in an SDG group. The Device view displays the SDGs based on the device type, and within the device type, the devices are organized by the device model. For example, all models of MX960 routers are grouped together under one node in the tree.
4. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
5. From the View pane, select the All Network item in Gateway view. If you are in Device view, click the plus sign (+) beside the My Network item in the View pane to expand the tree and select the device node you want.

6. From the task pane in Gateway view, select **Services Gateways**.

The Service Gateways page is displayed.



**NOTE:** Alternatively, you can select **Device View** from the View selector, click the **Build** icon on the banner, and select **Discover Devices** from the task pane to open the Discovery Profiles window to discover and manage devices.

7. Select the **Discover Gateway** option in Gateway view. Alternatively, in Device view, select the **Discover Devices** option from the task pane.
8. Select the check box next to the discovery profile that you want to modify.
9. Click the pencil icon above the table of discovery profiles to modify the selected profile.  
The Discovery Profile window appears.
10. Modify or add the discovery profile properties by clicking the plus sign or the pencil icon in the IP Details, SNMP Details, and User Details tables.
11. After you finish modifying all the necessary settings, click **Save** to save the modified profile in the database.

#### Related Documentation

- [Creating Service Gateway Groups on page 99](#)
- [Managing Service Gateway Groups on page 101](#)
- [Searching Unmanaged Devices on page 104](#)
- [Viewing the List of Discovered, Managed, and Unmanaged Devices on page 106](#)
- [Changing a Managed Device to an Unmanaged Device on page 112](#)
- [Deleting Discovery Profiles on page 114](#)

---

## Deleting Discovery Profiles

The Discovery Profiles page displays the discovery jobs that you have previously created. You can delete a discovery profile if you do not need it for discovering devices.

To delete a configured discovery profile:

1. From the View selector, select **Gateway View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.

The functionalities that you can configure in this mode are displayed in the task pane.

3. Select **Services Gateways** from the task pane in Gateway view. Alternatively, select **Device Discovery** from the task pane in Device view.

The Service Gateways page is displayed.

4. Select the **Discover Gateway** option in Gateway view. Else, select the **Discover Devices** option in Device view.

The list of discovery profiles is displayed.

5. Select the check box next to the discovery profile that you want to delete.
6. Click the red minus (-) icon above the table of listed templates. You are prompted to confirm the deletion.
7. Click **OK** to confirm the deletion. The corresponding profile is removed.

#### Related Documentation

- [Creating Service Gateway Groups on page 99](#)
- [Managing Service Gateway Groups on page 101](#)
- [Searching Unmanaged Devices on page 104](#)
- [Viewing the List of Discovered, Managed, and Unmanaged Devices on page 106](#)
- [Changing a Managed Device to an Unmanaged Device on page 112](#)
- [Modifying Discovery Profiles on page 113](#)

---

## Systems of Record in Junos Space Overview

Although by default the Junos Space network you are administering is the system of record (SOR)—each device defines its own official state—you may prefer to have the Junos Space Network Management Platform database contain the official state of the network, enabling you to restore that official state if unwanted out-of-band changes are made to a device. This feature enables you to designate Junos Space Network Management Platform as the SOR if you prefer.

- [Systems of Record on page 115](#)
- [Implications on device management on page 116](#)

### Systems of Record

A network managed by Junos Space Network Management Platform contains two repositories of information about the devices in the network: the devices themselves (each device defines and reports its official state) and the Junos Space Network Management Platform database (which contains information that is reported by the device during device discovery). One of these repositories must have precedence over

the other as the accepted desirable state. By default, the network itself is the system of record (NSOR).

In NSOR, when a local user commits a change in the configuration of a network device, the commit operation triggers a report via system log to Junos Space Network Management Platform. The values in the Junos Space Network Management Platform database are automatically changed to match the new device values, and the timestamps are synchronized. Thus the devices control the contents of the database.

As of version 12.2, you can designate the Junos Space Network Management Platform database values as having precedence over any values configured locally at a device. In this scenario, Junos Space Network Management Platform (database) is the system of record (SSOR). It contains the configurations that the Junos Space administrator considers best for the network devices. If an out-of-band commit operation is executed on a network device, Junos Space Network Management Platform receives a system log message, but the values in the Junos Space Network Management Platform database are not automatically changed or synchronized. Instead, the administrator can choose whether or not to overwrite the device's local changes by pushing the accepted configuration to the device from the Junos Space Network Management Platform database.

The choice of pushing the Junos Space Network Management Platform configuration is left to the administrator because the local device changes may, for example, be part of a temporary test that the administrator would not want to interrupt. However, if the tester forgets to reset the configuration at the end of the test, the administrator might then push the SSOR configuration to the device.

## Implications on device management

The basic difference between NSOR and SSOR lies in whether or not the Junos Space Network Management Platform database is automatically synchronized when changes are made to a network device, and which set of values has precedence.

Setting the Junos Space Network Management Platform database as the system of record does not protect your network from local changes. The device notifies Junos Space Network Management Platform via system log when the changes occur, and it does not resynchronize, so you still have the previous configuration and you can reset the remote device quickly if you need to do so. In an NSOR scenario, Junos Space Network Management Platform is also notified via system log. You can still push a more desirable configuration to the device, but this process is less efficient.

In the NSOR scenario, you can disable automatic resynchronization. When autoresynchronization is turned off, the server continues to receive notifications and goes into the out-of-sync state; however, autoresynchronization does not run on the device. You can manually resynchronize a device in such a case.

NSOR with automatic resynchronization disabled is not equivalent to SSOR: manually resynchronizing under NSOR updates the values in the Junos Space Network Management Platform database to reflect those on the device. This never happens under SSOR, where the Junos Space Network Management Platform database values have precedence over the device values, and synchronizing them involves pushing the database values to the device, effectively resetting the device's out-of-band changes.

**Related Documentation** • [Resynchronizing Managed SDGs with the Network on page 117](#)

## Resynchronizing Managed SDGs with the Network

If the network is the system of record, you can resynchronize a managed device at any time. For example, when a managed device is updated by a device administrator from the device's native GUI or CLI, you can resynchronize the device configuration in the Junos Space Network Management Platform database with the physical device. (If Junos Space Network Management Platform is the system of record, this capability is not available.)

To resynchronize a device:

1. From the View selector, select **Gateway View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. Do either of the following:
  - Select **Services Gateways > Unmanaged Gateway** from the task pane. The Unmanaged Devices page is displayed.
  - Select **Services Gateways > Managed Gateway** from the task pane. The Managed Devices page is displayed.
4. Select the devices you want to resynchronize and click the **Re-synch Hosts** button above the table of listed service delivery gateways (SDGs) or SDG pairs.  
The Resynchronize Devices pop-up window is displayed.
5. Click **Confirm**.

When a resynchronization job is scheduled to run but another resynchronization job on the same device is in progress, Junos Space Network Management Platform delays the scheduled resynchronization job. The time delay is determined by the damper interval that you set from the application workspace. By default the time delay is 20 seconds. The scheduled job is delayed as long as the other resynchronization job to the same device is in progress. When the job that is currently running finishes, the scheduled resynchronization job starts.



**NOTE:** You can check whether a managed device was resynchronized with the network, from the Job Details page. To go to the Job Details page, double-click the ID of the resynchronization job on the Job Management page. The Description column on this page specifies whether the managed device was resynchronized with the network. If the managed device was not resynchronized with the network, the column lists the reason for failure.

- Related Documentation**
- [Viewing the Device Inventory Page in Device View of Edge Services Director on page 176](#)
  - [Resynchronizing Device Configuration on page 342](#)
  - [Systems of Record in Junos Space Overview on page 115](#)

## CHAPTER 8

# Managing KPI Templates

- [Understanding Measurement Points, Key Performance Indicators, and Baseline Values on page 119](#)
- [Cloning a KPI Template on page 121](#)
- [Deleting KPI Templates on page 128](#)
- [Managing KPI Templates on page 129](#)
- [Viewing KPI Templates on page 130](#)
- [Modifying a KPI Template Associated with a Service Gateway on page 131](#)

## Understanding Measurement Points, Key Performance Indicators, and Baseline Values

This chapter topic provides guidelines for monitoring the service quality of an IP network. It describes how service providers and network administrators can use information provided by Juniper Networks routers to monitor network performance and capacity. You should have a thorough understanding of the SNMP and the associated MIB supported by Junos OS.



**NOTE:** For a good introduction to the process of monitoring an IP network, see RFC 2330, *Framework for IP Performance Metrics*.

This topic contains the following sections:

- [Measurement Points on page 119](#)
- [Basic Key Performance Indicators on page 120](#)
- [Setting Baselines on page 120](#)

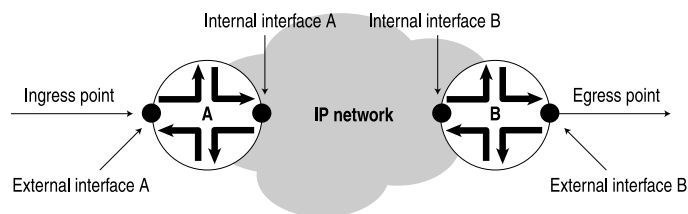
### Measurement Points

Defining the measurement points where metrics are measured is equally as important as defining the metrics themselves. This section describes measurement points within the context of this chapter and helps identify where measurements can be taken from a service provider network. It is important to understand exactly where a measurement point is. Measurement points are vital to understanding the implication of what the actual measurement means.

An IP network consists of a collection of routers connected by physical links that are all running the Internet Protocol. You can view the network as a collection of routers with an ingress (entry) point and an egress (exit) point. See [Figure 10 on page 120](#).

- Network-centric measurements are taken at measurement points that most closely map to the ingress and egress points for the network itself. For example, to measure delay across the provider network from Site A to Site B, the measurement points should be the ingress point to the provider network at Site A and the egress point at Site B.
- Router-centric measurements are taken directly from the routers themselves, but be careful to ensure that the correct router subcomponents have been identified in advance.

**Figure 10: Network Entry Points**



**NOTE:** [Figure 10 on page 120](#) does not show the client networks at customer premises, but they would be located on either side of the ingress and egress points. Although this chapter does not discuss how to measure network services as perceived by these client networks, you can use measurements taken for the service provider network as input into such calculations.

## Basic Key Performance Indicators

For example, you could monitor a service provider network for three basic key performance indicators (KPIs):

- *Availability* measures the “reachability” of one measurement point from another measurement point at the network layer (for example, using ICMP ping). The underlying routing and transport infrastructure of the provider network will support the availability measurements, with failures highlighted as unavailability.
- *Health* measures the number and type of errors that are occurring on the provider network, and can consist of both router-centric and network-centric measurements, such as hardware failures or packet loss.
- *Performance* of the provider network measures how well it can support IP services (for example, in terms of delay or utilization).

## Setting Baselines

How well is the provider network performing? We recommend an initial three-month period of monitoring to identify a network’s normal operational parameters. With this information, you can recognize exceptions and identify abnormal behavior. You should

continue baseline monitoring for the lifetime of each measured metric. Over time, you must be able to recognize performance trends and growth patterns.

Within the context of this chapter, many of the metrics identified do not have an allowable operational range associated with them. In most cases, you cannot identify the allowable operational range until you have determined a baseline for the actual variable on a specific network.

**Related  
Documentation**

- [Cloning a KPI Template on page 121](#)
- [Deleting KPI Templates on page 128](#)
- [Managing KPI Templates on page 129](#)
- [Viewing KPI Templates on page 130](#)

---

## Cloning a KPI Template

You clone a template definition to quickly create a new template definition with a new name but same properties. To modify a template definition without disabling templates based upon that definition, first clone the definition, then modify the clone.

1. From the View selector, select **Gateway View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the View pane, select the All Network item in Gateway view.
4. From the task pane, select **KPI Templates**.

The KPI Templates page is displayed.



.....  
**NOTE:** You can search for a specific template by entering the search criteria in the search field, in the table above the list of displayed templates. You can also search the templates based on the SDG, host, or SDG group.  
.....

5. Select the template you want to clone.

The Clone KPI Template dialog box is displayed.

Figure 11: Clone KPI Template Window

**Clone KPI Template - 14111 (v14.1)**

TLB | CGNAT | SFW | Chassis | HA

☒ TLB KPIs

☒ RI Composite Next Hop Status ■ When Next Hop is Not Available ■ When Next Hop is Available

☒ Real Servers up(%) ■ 0-4 ■ 5-30 ■ 31-100

☒ Service PIC Status ■ When Down ■ When Up

☒ Server Group Status ■ Server Group Down Status > 0 ■ Server Group Down Status = 0

☒ VIP Status ■ VIP Down is > 0 ■ VIP Down is = 0

☒ VIP Client Packet Drop Status ■ Packet Drop > 0 ■ Packet Drop = 0

Save Cancel

6. In the **Name** field, type a user-defined template definition name. A template definition name cannot exceed 128 characters and can contain only letters, numbers, spaces, and some special characters. The special characters allowed are hyphen (-), underscore (\_), period (.), at (@), single quote ('), forward slash (/), and ampersand (&).
7. (Optional) In the **Description** field, type a user-defined description. (limit of 255 characters). The description cannot exceed 256 characters. The operators who use the template definition to create templates rely on the description for information on the template definition.
8. Click **Save** to save the template.

The dialog box closes and the KPI Template window appears.

9. On the **ADC** tab, define the KPI settings for the adaptive delivery controller (ADC) service. Fill in the fields under this tab as indicated in the following table:

**Table 25: ADC Tab**

Field	Description
ADC KPIs	Select to enable configuration of KPIs for the adaptive delivery controller (ADC).
VIP Status	Select to configure virtual IP address status. Red is displayed for any failure in the virtual IP status, and green is displayed if there is no virtual IP failure.
Real Servers Up(%)	Select to configure real servers. Red denotes 0-30 percent of real servers are up and active, yellow denotes 31-60 percent of real servers are up, and green denotes 61-100 percent of total real servers are up.
Connection Table Count(K)	Select to configure connection table utilization. The connection table contains the online information on the current open connections that are handled by the ADC software. Red denotes 81-100 kilobytes of the table are utilized, yellow denotes 61-80 kilobytes of utilization, and green denote 0-60 kilobytes of utilization.
CPU Status Control	Select to configure CPU utilization for control packets as a status indicator. Red denotes CPU utilization of 81-100 percent, yellow is for 41-80 percent, and green is for 0-40 percent.
CPU Status (Data)	Select to configure CPU utilization for data packets as a status indicator. Red denotes CPU utilization of 96-100 percent, yellow is for 70-95 percent, and green is for 0-69 percent.
CPU Status (DataCores)	Select to configure CPU utilization for data cores as a status indicator. Red denotes CPU utilization of 1-5 cores, yellow is for 6-9 cores, and green is for 10-21 cores.  <b>NOTE:</b> A core represents a single CPU unit and multiple cores on a chip often share a memory bus or I/O bus. Virtual processor is a way to optimize the use of a core by permitting more threads to execute on the same core, while one thread is awaiting a memory or bus operation. Cores are an accurate metric of actual performance, as the VP's optimization is not constant, but depends on the workload.
NPU Allocation Failure	Select to configure network processing unit (NPU) allocation failures as a health status indicator. Red represents a failure, and green represents no failures.
DP Allocation Failure	Select to configure allocation failures in data plane as a health status marker in the template. Red represents a failure, and green represents no failures.
Service PIC Status	Select to configure the status of services PICs, such as adaptive or multiservices PICs. Red represents a failure, and green represents no failures.
Egress Interface Status	Select to configure the status of egress interfaces for services applications. Red represents a failure, and green represents no failures.

10. On the **TLB** tab, define the KPI settings for the adaptive delivery controller (ADC) service. Fill in the fields under this tab as indicated in the following table:

**Table 26: TLB Tab**

Field	Description
TLB KPIs	Select to enable configuration of KPIs for traffic load balancer (TLB).
RI Composite Next Hop Status	Select to configure the TLB routing instance composite next-hop status. Red is displayed if next-hop is not available, and green is displayed if next-hop is available.
Real Servers Status	Select to configure status of real servers. Red denotes 31-100 percent of real servers are used, yellow denotes 5-30 percent of real servers are utilized, and green denotes 0-4 percent of total real servers are utilized. Servers are configured to be available for hash-based, next-hop session distribution.
CPU Status(%)	Select to configure CPU status for as a status indicator. Red denotes CPU status of 91-100 percent, yellow is for 60-90 percent, and green is for 0-59 percent.
Service PIC Status	Select to configure the status of services PICs, such as adaptive or multiservices PICs. Red represents a failure, and green represents no failures.
Egress Interface Status	Select to configure the status of egress interfaces for services applications. Red represents a failure, and green represents no failures.

11. On the **CGNAT** tab, define the KPI settings for the adaptive delivery controller (ADC) service. Fill in the fields under this tab as indicated in the following table:

**Table 27: CGNAT Tab**

Field	Description
CGNAT KPIs	Select to enable configuration of KPIs for carrier-grade NAT (CGNAT) services.
CPU Status(%)	Select to configure CPU status for as a status indicator. Red denotes CPU status of 81-100 percent, yellow is for 60-80 percent, and green is for 0-59 percent.
Packet Drop Status)	Select to configure packet drop probability. Red denotes one or more packets are dropped, and green denotes no packet drops.
Memory Status(%)	Select to configure the working status of memory. Red denotes 0-30 percent of efficiency, yellow denotes 31-60 percent of efficiency, and green denotes 61-100 percent of efficiency.
NAT Pool Status(%)	Select to configure utilization of NAT address pools utilization. Red denotes 96-100 percent of utilization, yellow denotes 85-95 percent of utilization , and green denotes 0-84 percent of utilization .

**Table 27: CGNAT Tab (continued)**

Field	Description
Service PIC Status	Select to configure the status of services PICs, such as adaptive or multiservices PICs. Red represents a failure, and green represents no failures.
CPU Utilization(%)	Select to configure CPU utilization for control packets as a status indicator. Red denotes CPU utilization of 30-00 percent, yellow is for 20-29 percent, and green is for 0-10 percent.

12. On the **SFW** tab, define the KPI settings for the adaptive delivery controller (ADC) service. Fill in the fields under this tab as indicated in the following table:

**Table 28: SFW Tab**

Field	Description
SFW KPIs	Select to enable configuration of KPIs for stateful firewall services.
CPU Status(%)	Select to configure CPU status for as a status indicator. Red denotes CPU status of 81-100 percent, yellow is for 60-80 percent, and green is for 0-59 percent.
Packet Drop Status	Select to configure packet drop probability. Red denotes one or more packets are dropped, and green denotes no packet drops.
Memory Status(%)	Select to configure the working status of memory. Red denotes 0-30 percent of efficiency, yellow denotes 31-60 percent of efficiency, and green denotes 61-100 percent of efficiency.
NAT Pool Status(%)	Select to configure utilization of NAT address pools utilization. Red denotes 96-100 percent of utilization, yellow denotes 85-95 percent of utilization , and green denotes 0-84 percent of utilization .
Service PIC Status	Select to configure the status of services PICs, such as adaptive or multiservices PICs. Red represents a failure, and green represents no failures.

13. On the **Chassis** tab, define the KPI settings for the adaptive delivery controller (ADC) service. Fill in the fields under this tab as indicated in the following table:

**Table 29: Chassis Tab**

Field	Description
Chassis KPIs	Select to enable configuration of KPIs for chassis operations.
Chassis Level KPIs	Select to enable configuration of KPIs for chassis-level processes and parameters.

**Table 29: Chassis Tab (continued)**

Field	Description
Power Supply Failure	Select to configure failure of power supplies as a status marker. Red denotes a failure, and green denotes no failure.
Fan Failure	Select to configure failure of fans as a status marker. Red denotes a failure, and green denotes no failure.
FRU Failure	Select to configure failure of FRUs as a status marker. Red denotes a failure, and green denotes no failure.
Over Temperature	Select to configure temperature of chassis as a metric. Red denotes 95-100 percent of over temperature condition, yellow denotes 80-94 percent, and green denotes 0-70 percent.
CPU Utilization(%)	Select to configure CPU utilization of the chassis as a metric. Red denotes 95-100 percent of CPU usage, yellow denotes 80-94 percent, and green denotes 0-70 percent.
Memory Status(%)	Select to configure the percentage of memory used by the entire chassis. If this number exceeds 80 percent, you might experience a software problem (memory leak). Red denotes 95-100 percent of usage, yellow denotes 80-94 percent, and green denotes 0-70 percent.
Slot Level KPIs	Select to enable configuration of KPIs for slot-based processes and parameters.
<b>RE Tab</b>	
RE	Select to configure Routing Engine characteristics.
Temperature	Select to configure temperature of Routing Engines as a metric. Red denotes 95-100 percent of over temperature condition, yellow denotes 80-94 percent, and green denotes 0-70 percent.
CPU Utilization(%)	Select to configure CPU utilization of Routing Engines as a metric. Red denotes 95-100 percent of CPU usage, yellow denotes 80-94 percent, and green denotes 0-70 percent.
Memory(Heap)	Select to configure the percentage of heap space (dynamic memory) being used by the Routing Engine. If this number exceeds 80 percent, you might experience a software problem (memory leak). Red denotes 95-100 percent of usage, yellow denotes 80-94 percent, and green denotes 0-70 percent.
Buffer	Select to configure the percentage of buffer memory space being used by the Routing Engine processor for buffering internal messages. Red denotes 95-100 percent of usage, yellow denotes 80-94 percent, and green denotes 0-70 percent.
<b>Service PIC Tab</b>	
Service PIC (MS-DPC)	Select to configure service PIC characteristics.

Table 29: Chassis Tab (continued)

Field	Description
Temperature	Select to configure temperature of service PIC as a metric. Red denotes 95-100 percent of over temperature condition, yellow denotes 80-94 percent, and green denotes 0-70 percent.
CPU Utilization(%)	Select to configure CPU utilization of services PICs as a metric. Red denotes 95-100 percent of CPU usage, yellow denotes 80-94 percent, and green denotes 0-70 percent.
Memory(Heap)	Select to configure the percentage of heap space (dynamic memory) being used by the services PICs. If this number exceeds 80 percent, you might experience a software problem (memory leak). Red denotes 95-100 percent of usage, yellow denotes 80-94 percent, and green denotes 0-70 percent.
Buffer	Select to configure the percentage of buffer memory space being used by the service PICs for buffering internal messages. Red denotes 95-100 percent of usage, yellow denotes 80-94 percent, and green denotes 0-70 percent.
Packet Drop Status)	Select to configure packet drop probability. Red denotes one or more packets are dropped, and green denotes no packet drops.
<b>FPC Tab</b>	
FPC	Select to configure FPC characteristics.
Temperature	Select to configure temperature of the FPC as a metric. Red denotes 95-100 percent of over temperature condition, yellow denotes 80-94 percent, and green denotes 0-70 percent.
CPU Utilization(%)	Select to configure CPU utilization of FPCs as a metric. Red denotes 95-100 percent of CPU usage, yellow denotes 80-94 percent, and green denotes 0-70 percent.
Memory(Heap)	Select to configure the percentage of heap space (dynamic memory) being used by the FPCs. If this number exceeds 80 percent, you might experience a software problem (memory leak). Red denotes 95-100 percent of usage, yellow denotes 80-94 percent, and green denotes 0-70 percent.
Buffer	Select to configure the percentage of buffer memory space being used by the FPCs for buffering internal messages. Red denotes 95-100 percent of usage, yellow denotes 80-94 percent, and green denotes 0-70 percent.

14. On the **HA** tab, define the KPI settings for the high availability service. Fill in the fields under this tab as indicated in the following table:

*Table 30: HA Tab*

Field	Description
HA KPIs	Select to enable configuration of KPIs for high availability services.
SDG Status	Select to specify the status of SDGs. Red denotes a failure, and green indicates no failure.
BGP Advertising	Select to configure packet drops during BGP advertisements. Red denotes a packet drop, and green denotes no packet drops.
VRRP Status	Select to configure the status of Virtual Router Redundancy Protocol (VRRP). Red denotes a failure, and green indicates no failure.
CGNAT SFW HA Status	Select to configure inter-chassis high availability status of carrier-grade NAT. Red denotes a failure, and green indicates no failure.
CGNAT Route Status	Select to configure the status of routes in CGNAT. Red denotes a failure, and green indicates no failure.
ADC VIP Route Status	Select to configure the status of routes in the ADC virtual server. Red denotes a failure, and green indicates no failure.
TLB RI Route Status	Select to configure the route status in a TLB routing instance. Red denotes a failure, and green indicates no failure.

15. Click **Save** to save the template definition for the different service types.

**Related  
Documentation**

- [Understanding Measurement Points, Key Performance Indicators, and Baseline Values on page 119](#)

## Deleting KPI Templates

The KPI Templates page displays the definitions you created to configure KPIs for the different services in your service delivery gateway (SDG) devices. The templates are sorted by name. You can delete a KPI template that is not referenced by an SDG.

To delete the configured KPI templates:

1. From the View selector, select **Gateway View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the task pane, select **KPI Templates**.

The KPI Templates page is displayed.

4. Select the template you want to delete.
5. Click the red minus (-) icon above the table of listed templates. You are prompted to confirm whether you want to delete the selected template. Click **OK** to confirm the deletion.

The corresponding template is removed from the database.

- Related Documentation**
- [Understanding Measurement Points, Key Performance Indicators, and Baseline Values on page 119](#)
  - [Cloning a KPI Template on page 121](#)
  - [Managing KPI Templates on page 129](#)
  - [Viewing KPI Templates on page 130](#)

---

## Managing KPI Templates

---

The KPI Templates page displays the definitions you created to configure KPIs for the different services in your service delivery gateway (SDG) devices. The templates are sorted by name.

The KPI Templates page provides the metrics that are used in evaluating the health and operating efficiency of an SDG. A preconfigured, system-supplied KPI template is available and it is not editable. You can create a copy of the predefined, system template and edit it for your needs. During a discovery of an SDG, all the KPI templates in the system are displayed and you can associate the appropriate KPI template with the SDG. During KPI template association, a copy of the selected KPI template is associated with the SDG. This behavior indicates that the base KPI template has no link to the KPI details on the SDG after association. As a result, any changes performed to the base template are not propagated to the SDG. Instead, you must modify the KPI template for each SDG. The KPI Templates page enables you to edit the KPI details of a selected SDG.

You can perform the following tasks from this page under Build mode:

- View KPI templates
- Clone a KPI template
- Delete a KPI template
- Modify a KPI Template

- Related Documentation**
- [Understanding Measurement Points, Key Performance Indicators, and Baseline Values on page 119](#)
  - [Deleting KPI Templates on page 128](#)
  - [Viewing KPI Templates on page 130](#)

## Viewing KPI Templates

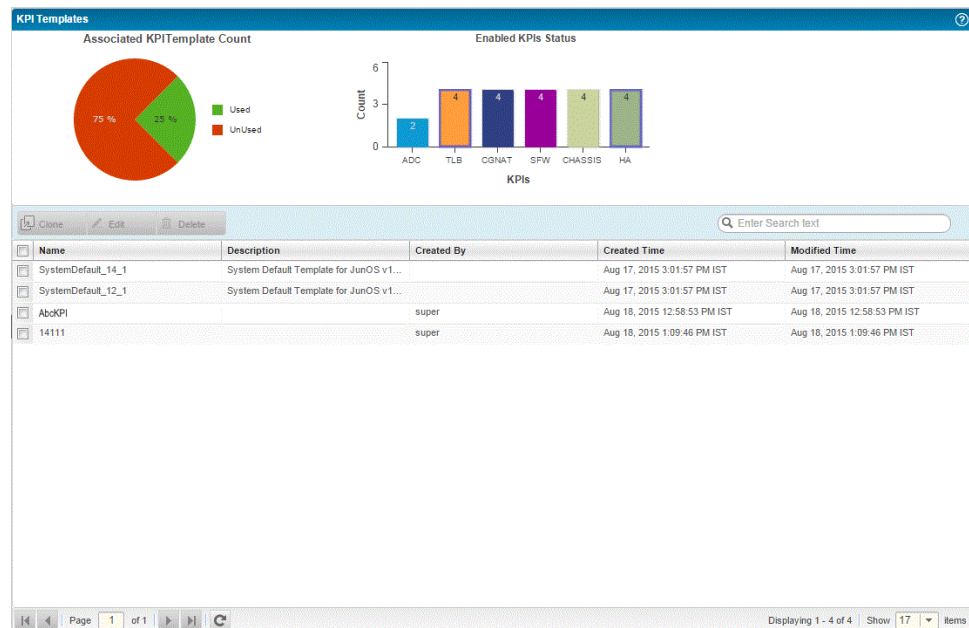
The KPI Templates page displays the definitions you created to configure KPIs for the different services in your service delivery gateway (SDG) devices. The templates are sorted by name.

To view the configured KPI templates:

1. From the View selector, select **Gateway View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the View pane, select the All Network item in Gateway view.
4. From the task pane, select **KPI Templates**.

The KPI Templates page is displayed. The page is divided into two halves. The top part of the page displays two graphs, while the bottom part of the page displays all the configured KPI templates.

Figure 12: KPI Templates Page



Of the two graphs, one of them is the Associated KPI Template Count graph. This pie graph illustrates the percentage of templates, out of the total number of templates, that are associated with SDGs. The Used type denotes the templates associated with SDGs, while the Unused type denotes the templates that are not mapped to any SDG. The other

graph is a bar chart. The Enabled KPIs Status bar chart shows the different service types on the x-axis and the number of KPIs that are defined for each service type on the y-axis. For example, a count of 5 for the TLB service type signifies that five KPI templates contain TLB service attributes.

The lower half of the Templates page displays the following fields in a tabular view:

You clone a template definition to quickly create a new template definition with a new name but same properties. To modify a template definition without disabling templates based upon that definition, first clone the definition, then modify the clone.

**Table 31: KPI Templates View**

Field	Description
Name	Name of the KPI template.
Description	User-defined description of the template.
Created By	Name of the user that created the template
Created Time	Time and date when the KPI template was created. The displayed timezone depends on the server timezone.
Modified Time	Time and date when the KPI template was last updated. The displayed timezone depends on the server timezone.

**Related Documentation**

- [Understanding Measurement Points, Key Performance Indicators, and Baseline Values on page 119](#)
- [Cloning a KPI Template on page 121](#)
- [Deleting KPI Templates on page 128](#)
- [Managing KPI Templates on page 129](#)

## Modifying a KPI Template Associated with a Service Gateway

You can modify the KPI template characteristics for the different services. You can select a managed device with which a KPI template is associated and change the KPI settings to specify monitoring criteria critical for service operations and administration. On the dashboard, the SDGs are colored as orange, red, green, or gray to indicate the health and performance of the SDGs based on the applied KPI templates.

To modify the KPI template associated with a managed service gateway:

1. From the View selector, select **Gateway View**.

The workspaces that are applicable to this view are displayed.

2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.

The functionalities that you can configure in this mode are displayed in the task pane.

3. From the View pane, select the All Network item in Gateway view. You can also click the plus sign (+) beside the All Network item in the View pane to expand the tree and select the device node you want.

4. From the task pane, select **Services Gateways**.

The Service Gateways page is displayed.

5. Select the **Manage Service Gateways** option.

The list of managed SDGs appears. If the SDGs are configured in a high availability pair, the details of both the devices in the pair are exported.

6. In the main window, click the plus sign (+) next to the SDG pairs to expand the tree and view the pair of devices in the SDG group or pair. Select the check box next to the individual SDG for which you want to view the device details.

The SDG for which you want to modify KPI template details is selected.

7. Click the down arrow in the **Modify** button at the top of the table of managed SDGs that are listed, and select **KPI Details**.

The Edit KPI Details for Service Gateway dialog box is displayed.

8. Modify the settings as described in [“Cloning a KPI Template” on page 121](#).

9. Click **Save** to save the template definition for the different service types.

You are returned to the Managed Service Gateways page.

#### Related Documentation

- [Discovering Devices on page 85](#)
- [Exporting Managed Device Details to a CSV File on page 94](#)
- [Changing an Unmanaged Device to a Managed Device on page 95](#)
- [Modifying the SDG Group and KPI Templates for a Device on page 97](#)
- [Scheduling the Discovery of Devices on page 98](#)

## CHAPTER 9

# Viewing the Device Inventory

- [Viewing the Device Inventory Page on page 133](#)
- [Viewing Device Statistics on page 141](#)
- [Viewing Configuration Details of Services on Devices on page 144](#)
- [Viewing Discovery Logs on page 146](#)
- [Viewing Discovery Profiles on page 147](#)

## Viewing the Device Inventory Page

---

The Device Inventory page lists devices managed by Edge Services Director and provides basic information about the devices, such as IP address and current operating status, and configured services, such as server load balancing (SLB) and carrier grade NAT (CGNAT). The Device Inventory page is available in Build mode.

Hardware inventory information shows the slots that are available for a device and provides information about power supplies, chassis cards, fans, part numbers, and so forth. Edge Services Director displays hardware inventory by device name, based on data retrieved both from the device during discovery and resynchronization operations, and from the data stored in the hardware catalog. For each managed device, the hardware catalog provides descriptions for field replaceable units (FRUs), part numbers, model numbers, and the pluggable locations from which empty slots are determined.

The Device Inventory page provides two pie charts that summarize the status of the devices and services in your network environment. You can remove or restore a category (segment) from the pie chart by clicking that segment in the chart.

To view the device inventory page:

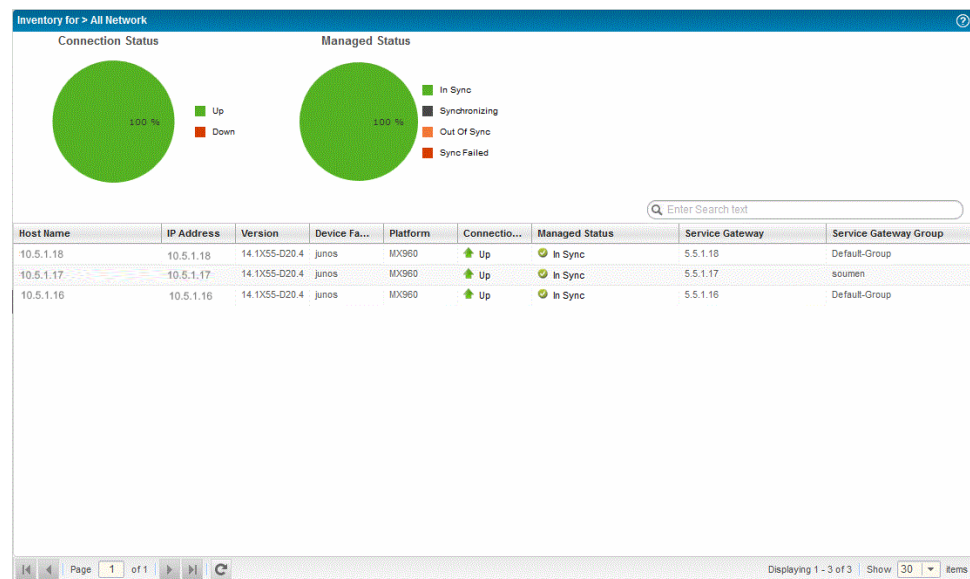
1. From the View selector, select **Gateway View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.

- From the View pane, select the All Network item in Gateway view. You can also click the plus sign (+) beside the All Network item in the View pane to expand the tree and select the device node you want.

- From the task pane in Gateway view, select **Inventory > View Inventory**.

The Device Inventory page for the entire network, the SDG group, or the particular SDG is displayed, depending on the view or perspective you selected.

**Figure 13: Device Inventory Page**



The following charts are displayed in the top half of the page:

- Connection State—Shows the proportion of devices that are up or down. In this chart, Virtual Chassis count as one device. The possible connection states are:
  - UP—Device is connected to Edge Services Director.
  - DOWN—Device is not connected to Edge Services Director.
  - N/A—Device connection state is not available.
- Managed Status—Shows the proportion of devices in each configuration state. See [Table 32 on page 134](#) for definitions of the configuration states.

**Table 32: Managed Status Pie Chart**

Field	Description
In Sync	The configuration on the device is in sync with the Edge Services Director configuration for the device.

*Table 32: Managed Status Pie Chart (continued)*

Field	Description
Out Of Sync	The configuration on the device does not match the Edge Services Director configuration for the device. This state is usually the result of the device configuration being altered outside of Edge Services Director.  You cannot deploy configuration on a device from Edge Services 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 or is an access point.

Mouse over a pie segment to view the actual number of devices and the percentage represented by that pie segment.

In the left pane, which displays all the configured SDGs in a tree format, the **All Service Gateways** option is selected by default. The right pane that displays the pie graphs and the corresponding tabular details of services, hardware, and interfaces in the bottom part of the right pane are corresponding with the **All Service Gateways** selection. You can select a specific SDG from the left pane and the right pane details are appropriately shown.

The lower half of the Inventory page contains three tabs—Gateway, Hardware, and Interface. These tabs are displayed only if you expand the All Network item in the View pane and select a device node or SDG. Otherwise, only the fields under the Gateway tab are displayed.

[Table 23 on page 103](#) describes the fields under the Gateway tab of the Device Inventory table.

*Table 33: Fields Under the Gateway Tab*

Field	Description
Host Name	Hostname of the device.
Version	Software version the device is running.
IP Address	IP address configured for the device.
Device Family	Device family to which the device belongs. Hardware family such as Junos OS is displayed.
Platform	Model number of the device.

**Table 33: Fields Under the Gateway Tab (continued)**

Field	Description
Connection Status	<p>Device's state:</p> <ul style="list-style-type: none"> <li>UP—Edge Services Director can communicate with the device.</li> <li>DOWN—Edge Services Director cannot communicate with the device.</li> </ul>
Managed Status	<p>Indicates whether the device's configuration is in sync with Edge Services Director's version:</p> <ul style="list-style-type: none"> <li>Connecting—The device is being contacted to establish a connection.</li> <li>In Sync—The configuration on the device is in sync with the Edge Services Director configuration for the device.</li> <li>Out Of Sync—The configuration on the device does not match the Edge Services Director configuration for the device. This state is usually the result of the device configuration being altered outside of Edge Services Director. 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.</li> <li>Synchronizing—The device configuration is in the process of being resynchronized.</li> <li>Sync failed—An attempt to resynchronize an Out Of Sync device failed.</li> </ul>
Service Gateway	Name of the service delivery gateway.
Service Gateway Group	Name of the group to which the SDG is assigned.

[Table 34 on page 136](#) describes the fields under the Hardware tab of the Device Inventory table.

**Table 34: Fields Under the Hardware Tab**

Field	Description
Module	Name of the SDG and the platform type, such as MX240 or MX480. Click the plus sign (+) to expand the tree to display the components of the device, such as chassis, PIC, CPU, and PIC parameters. Information about the chassis, midplane, craft interface (FPM), power midplane (PMP), Power Supply Modules (PSMs), Power Distribution Modules (PDMs), Routing Engines, Control Boards (CBs) and Switch Processor Mezzanine Boards (SPMBs), Switch Fabric Boards (SFBs), Flexible PIC Concentrators (FPCs), PICs, adapter cards (ADCs) and fan trays is displayed.
Model Number	Model number of the FRU hardware component.
Model	Model of the FRU component.
Part Number	Part number of the chassis component.

*Table 34: Fields Under the Hardware Tab (continued)*

Field	Description
Serial Number	Serial number of the chassis component. The serial number of the backplane is also the serial number of the router chassis. Use this serial number when you need to contact Juniper Networks Customer Support about the router or switch chassis.

*Table 34: Fields Under the Hardware Tab (continued)*

Field	Description
Description	

Table 34: Fields Under the Hardware Tab (continued)

Field	Description
	<p>Brief description of the hardware item:</p> <ul style="list-style-type: none"> <li>• Type of power supply.</li> <li>• Type of PIC. If the PIC type is not supported on the current software release, the output states <b>Hardware Not Supported</b>.</li> <li>• Type of FPC: <b>FPC Type 1</b>, <b>FPC Type 2</b>, <b>FPC Type 3</b>, <b>FPC Type 4</b>, or <b>FPC TypeOC192</b>. On EX Series switches, a brief description of the FPC. On the J Series routers, the FPC type corresponds to the Physical Interface Module (PIM). The following list shows the PIM abbreviation in the output and the corresponding PIM name. <ul style="list-style-type: none"> <li>• <b>2x FE</b>—Either two built-in Fast Ethernet interfaces (fixed PIM) or dual-port Fast Ethernet PIM</li> <li>• <b>4x FE</b>—4-port Fast Ethernet ePIM</li> <li>• <b>1x GE Copper</b>—Copper Gigabit Ethernet ePIM (one 10-Mbps, 100-Mbps, or 1000-Mbps port)</li> <li>• <b>1x GE SFP</b>—SFP Gigabit Ethernet ePIM (one fiber port)</li> <li>• <b>4x GE Base PIC</b>—Four built-in Gigabit Ethernet ports on a J4350 or J6350 chassis (fixed PIM)</li> <li>• <b>2x Serial</b>—Dual-port serial PIM</li> <li>• <b>2x T1</b>—Dual-port T1 PIM</li> <li>• <b>2x E1</b>—Dual-port E1 PIM</li> <li>• <b>2x CTIE1</b>—Dual-port channelized T1/E1 PIM</li> <li>• <b>1x T3</b>—T3 PIM (one port)</li> <li>• <b>1x E3</b>—E3 PIM (one port)</li> <li>• <b>4x BRI S/T</b>—4-port ISDN BRI S/T PIM</li> <li>• <b>4x BRI U</b>—4-port ISDN BRI U PIM</li> <li>• <b>1x ADSL Annex A</b>—ADSL 2/2+ Annex A PIM (one port, for POTS)</li> <li>• <b>1x ADSL Annex B</b>—ADSL 2/2+ Annex B PIM (one port, for ISDN)</li> <li>• <b>2x SHDSL (ATM)</b>—G SHDSL PIM (2-port two-wire module or 1-port four-wire module)</li> <li>• <b>1x TGM550</b>—TGM550 Telephony Gateway Module (Avaya VoIP gateway module with one console port, two analog <b>LINE</b> ports, and two analog <b>TRUNK</b> ports)</li> <li>• <b>1x DS1 TIM510</b>—TIM510 E1/T1 Telephony Interface Module (Avaya VoIP media module with one E1 or T1 trunk termination port and ISDN PRI backup)</li> <li>• <b>4x FXS, 4x FXO, TIM514</b>—TIM514 Analog Telephony Interface Module (Avaya VoIP media module with four analog <b>LINE</b> ports and four analog <b>TRUNK</b> ports)</li> <li>• <b>4x BRI TIM521</b>—TIM521 BRI Telephony Interface Module (Avaya VoIP media module with four ISDN BRI ports)</li> <li>• <b>Crypto Accelerator Module</b>—For enhanced performance of cryptographic algorithms used in IP Security (IPsec) services</li> <li>• <b>MPC M 16x 10GE</b>—16-port 10-Gigabit Module Port Concentrator that supports SFP+ optical transceivers. (Not on EX Series switches.)</li> </ul> </li> <li>• For hosts, the Routing Engine type.</li> <li>• For small form-factor pluggable transceiver (SFP) modules, the type of fiber: <b>LX</b>, <b>SX</b>, <b>LH</b>, or <b>T</b>.</li> <li>• LCD description for EX Series switches (except EX2200 switches).</li> </ul>

**Table 34: Fields Under the Hardware Tab (continued)**

Field	Description
	<ul style="list-style-type: none"> <li>• <b>MPC2</b>—1-port MPC2 that supports two separate slots for MICs.</li> <li>• <b>MPC3E</b>—1-port MPC3E that supports two separate slots for MICs (MIC-3D-1X100GE-CFP and MIC-3D-20GE-SFP) on MX960, MX480, and MX240 routers. The MPC3E maps one MIC to one PIC (1 MIC, 1 PIC), which differs from the mapping of legacy MPCs.</li> <li>• 100GBASE-LR4, pluggable CFP optics</li> <li>• Supports the Enhanced MX Switch Control Board with fabric redundancy and existing SCBs without fabric redundancy.</li> <li>• Interoperates with existing MX Series line cards, including Flexible Port Concentrators (FPC), Dense Port Concentrators (DPCs), and Modular Port Concentrators (MPCs).</li> <li>• <b>MPC4E</b>—Fixed configuration MPC4E that is available in two flavors: MPC4E-3D-32XGE-SFPP and MPC4E-3D-2CGE-8XGE on MX2020, MX960, MX480, and MX240 routers.</li> <li>• LCD description for MX Series routers</li> </ul>

Table 35 on page 140 describes the fields under the Interface tab of the Device Inventory table.

**Table 35: Fields Under the Interface Tab**

Field	Description
Host Name	Hostname of the SDG.
Physical Interface Name	Name of the physical interface.
IP Address	IP address configured on the interface.
MAC Address	MAC address configured on the interface
Operation Status	Operational status of the physical interface: Up, Down.
Admin Status	Administrative state of the interface: Enabled or Disabled. If the interface is disabled, it can provide network connectivity, but it cannot provide power to connected devices.
Link Level Type	Encapsulation type configured on the interface.
Link Type	Data transmission type.
Speed	Speed at which the interface is running.
MTU	Maximum transmission unit size on the physical interface.
Description	Configured textual description of the interface.

To view configuration and run-time information for devices:

1. Sort the table by mousing over the column header for the data you want to sort by and clicking the down arrow. Select **Sort Ascending** or **Sort Descending**.
2. Show columns not in the default table view, or hide columns, as follows:
  1. Mouse over any column header and click the down arrow.
  2. Select **Columns** from the menu.
  3. Select the check boxes for columns that you want to view. Clear the check boxes for columns that you want to hide.
3. View information about devices as follows:
  - To restrict the display of devices, enter a search criterion of one or more characters in the Search bar and press Enter.

All devices that match the search criterion are shown in the main display area.

**Related  
Documentation**

- [Viewing Device Statistics on page 141](#)
- [Viewing Configuration Details of Services on Devices on page 144](#)
- [Viewing Discovery Logs on page 146](#)
- [Viewing Discovery Profiles on page 147](#)

---

## Viewing Device Statistics

The Devices statistics page provides three types of data for managed devices:

- Device Count by Platform—The number of Juniper Networks devices organized by type
- Device Status—The connection status of managed devices on the network
- Device Count by OS—The number of devices running a particular Junos OS release

To view device statistics, from the Junos Space Network Management Platform user interface, select **Devices**. The Devices landing page is displayed. This page displays the charts related to devices.

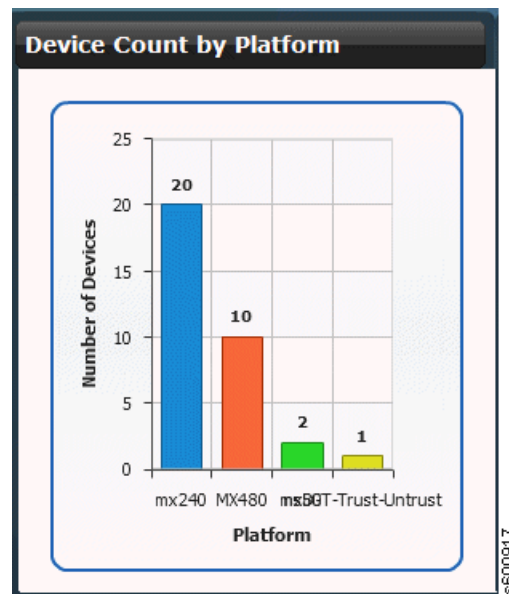
This topic includes the following tasks:

- [Viewing the Number of Devices by Platform on page 142](#)
- [Viewing Connection Status for Devices on page 142](#)
- [Viewing Devices by Junos OS Release on page 143](#)

## Viewing the Number of Devices by Platform

Figure 14 on page 142 shows the Device Count by Platform report. The bar chart shows the number of Juniper Networks devices on the y-axis discovered by platform type on the x-axis. Each vertical bar in the chart displays the number of managed devices for a platform.

Figure 14: Device Count by Platform Report



To view more detailed information about devices per platform:

- Click a bar in the bar graph. The Device Management inventory page appears filtered by the device type you selected. See *Viewing Managed Devices*.

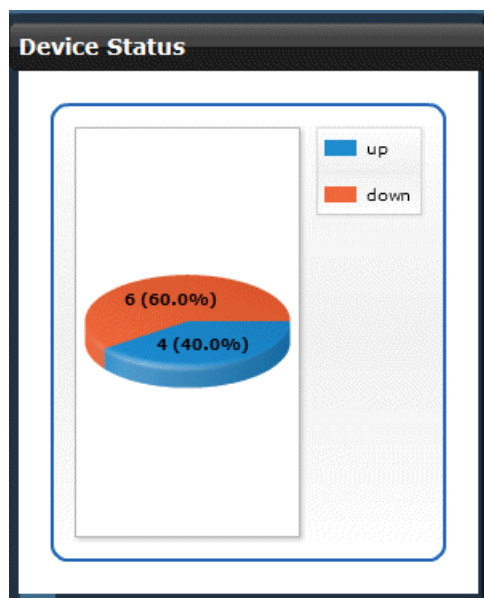
To save the bar chart as an image or to print for presentations or reporting:

- Right-click the bar chart and use the menu to save or print the image.

## Viewing Connection Status for Devices

Figure 15 on page 143 shows the Device Status report. The pie chart displays the percentage and number of devices that are connected and disconnected on the network. The up or down status is expressed as a percentage of the total number of devices.

Figure 15: Device Status Report



To view more detailed device status information:

- Click a slice in the pie chart. The Device Management inventory page appears filtered by the devices that are up or down. See *Viewing Managed Devices*.

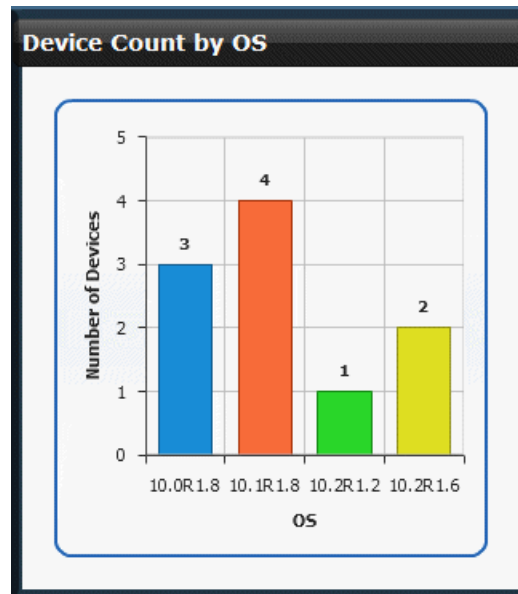
To save the pie chart as an image or to print for presentations or reporting:

- Right-click the bar chart and use the menu to save or print the image.

## Viewing Devices by Junos OS Release

Figure 16 on page 144 shows the Device Count by OS report. The bar chart shows the number of Juniper Networks devices on the network (the y-axis) categorized by running a certain Junos OS release (the x-axis).

Figure 16: Device Count by OS Report



To view more detailed information about devices running a particular Junos OS release:

- Click a bar in the chart. The Device Management inventory page appears. See *Viewing Managed Devices* in the *Junos Space Network Application Platform User Guide* for details.

To save the pie chart as an image or to print for presentations or reporting:

- Right-click the bar chart and use the menu to save or print the image.

#### Related Documentation

- [Viewing the Device Inventory Page on page 133](#)
- [Viewing Configuration Details of Services on Devices on page 144](#)
- [Viewing Discovery Logs on page 146](#)
- [Viewing Discovery Profiles on page 147](#)

## Viewing Configuration Details of Services on Devices

The view configuration capability enables you to see extensive, comprehensive configuration settings of a device. The settings of each service configured, such as ADC or TLB, and the options or substatements for each service, such as service PICs and client-facing interfaces, are displayed. For each service instance configured on a device, you can view granular information on each attribute of a service instance.

The configuration details are displayed in property view and configuration view. The *property view* is useful if you want a GUI, tree-based structure of display. In this view, you can drill-down the tree and view data about each of the service attributes. Property view is simple view of configuration as key value pair. The dynamic fields in form view are defined using parameters. The *configuration view* is beneficial if you are familiar with the CLI interface structure and want to view service attributes in the form of configuration stanzas and hierarchy levels. This display is similar to the **show** command that you can use at a certain **[edit]** hierarchy level to view the defined settings. Each level in the hierarchy is indented to indicate each statement's relative position in the hierarchy.

To view the configuration details of services on a device:

1. From the View selector, select **Gateway View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the View pane, select the All Network item in Gateway view. You can also click the plus sign (+) beside the All Network item in the View pane to expand the tree and select the device node you want.
4. Select **Services Gateways** from the task pane. The Service Gateways page is displayed.
5. Select the check box next to the discovered device for which you want to view the services configured.  
The device for which you want to view the configuration of services is selected.
6. Click the **View Configuration** icon above the table of displayed devices. The View Service page is displayed.  
The page is divided into three panes. The left pane displays a tree of all configured services. Click the plus sign (+) for each service to expand and view the service instances contained in a service. The middle pane displays the components or attributes of the selected service instance. The rightmost pane displays the attributes of the service instance in property or config view.
7. Click the **Service Details** option from the left pane. The list of services corresponding to the selected SDG pair is displayed.
8. Mouse over the middle pane that lists the service instance components to highlight the component and view its name.
9. Select the service instance from the left pane. Drill down until you locate the instance you need. The graphical representation of the components of the service instance are

shown in the middle pane. The categories and the elements of the components are shown.

When you click a different component or attribute of the service instance, the property or config view is refreshed accordingly.

10. Select the **Property View** tab if you want to view the parameters in a tree-based, key value pair structure. Select the **Config View** tab if you want to view the parameters in a CLI structure.

#### Related Documentation

- [Viewing the Device Inventory Page on page 133](#)
- [Viewing Device Statistics on page 141](#)
- [Viewing Discovery Logs on page 146](#)
- [Viewing Discovery Profiles on page 147](#)

---

## Viewing Discovery Logs

Discovery logs are prepared and stored on each stage of the SDG discovery. All these discovery logs are stored in the discovery audit log database. When you initiate the request to view the system event logging messages, the logs for the selected device are retrieved and displayed with the timestamp of the recording.

To view the details of a discovery log:

1. From the View selector, select **Gateway View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the View pane, select the All Network item in Gateway view. You can also click the plus sign (+) beside the All Network item in the View pane to expand the tree and select the device node you want.
4. Select **Services Gateways** from the task pane. The Service Gateways page is displayed.
5. Select the **Unmanaged Devices** option. The list of discovered devices are displayed.
6. Select the check box next to the discovered device for which you want to view the logs.
7. Click the **View Discovery Logs** icon above the table of displayed devices. The Discovery Log window is displayed.

The timestamp is the UTC time in database that is mapped to the local time zone of client computer. The description of the log is displayed in color-coded format. Red indicates error severity, orange indicates severity level of warning, and black indicates an informational message.

8. In the Service Delivery Host Discovery Log window, you can sort and view the log messages that pertain to a severity level to quickly, effectively identify and separate only the logs that are of relevance to you. To filter the logs based on a severity level, select the check boxes next to the severity levels, such as Error, Warning, or Info and click the search icon to display based on the match criterion. Click the red cross (x) icon to clear the applied filter and display the logs corresponding to all severity levels. If you rediscover a device, the logs display a cumulative, consolidated list of all of the messages generated during all of the discovery attempts.
9. Click **Refresh** if you want updated snapshots of the logs to be displayed. The refresh process causes a request to be sent to the device and retrieval of the latest logs occurs.
10. After you finish viewing the profile settings, click **Close** to return to the page that displays all the discovery profiles.

#### Related Documentation

- [Viewing the Device Inventory Page on page 133](#)
- [Viewing Device Statistics on page 141](#)
- [Viewing Configuration Details of Services on Devices on page 144](#)
- [Viewing Discovery Profiles on page 147](#)

## Viewing Discovery Profiles

The Discovery Profiles page displays the discovery jobs that you have previously created. You can examine the parameters contained in a discovery profile before you modify or want to create a fresh profile. All of the discovery details, such as the device attributes and credentials for connecting to it, are shown.

To view the details of a configured discovery profile:

1. From the View selector, select **Gateway View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the View pane, select the All Network item in Gateway view. You can also click the plus sign (+) beside the All Network item in the View pane to expand the tree and select the device node you want.

4. Select **Services Gateways** from the task pane. The Service Gateways page is displayed.
5. Select the **Discover Devices** option.
6. Select the check box next to the discovery profile that you want to view.
7. Click the **View Discovery Profile Details** icon above the table of displayed profiles. The View Discovery Profile Details window is displayed.

The following fields are displayed in this window:

Field	Description
Name	Unique name of the discovery profile
Description	User-defined description of the profile
Discover Targets	
Type	Indicates whether hostname or IP address of the devices in the profile are configured.
Value	Hostname or IP address of the devices in the profile
Credentials	
UserName	Name of the user to connect to the devices in the discovery profile.
Password	Password for authentication to the devices in the profile displayed as a set of asterisks
Protocol Details	
SNMP Version	Version of SNMP, such as SNMPv1, v2C, or v3.
Details	Parameters such as the privacy and authentication details for SNMP.

8. After you finish viewing the profile settings, click **Close** to return to the page that displays all the discovery profiles.

**Related  
Documentation**

- [Viewing the Device Inventory Page on page 133](#)
- [Viewing Device Statistics on page 141](#)
- [Viewing Configuration Details of Services on Devices on page 144](#)
- [Viewing Discovery Logs on page 146](#)

## PART 4

# Location and Device Views of Build Mode

- [Location View Configuration on page 151](#)
- [Device Management on page 173](#)



## CHAPTER 10

# Location View Configuration

- [Understanding Build Mode in Location and Device Views of Edge Services Director on page 151](#)
- [Understanding the Location View on page 154](#)
- [Assigning and Unassigning Devices to a Location on page 155](#)
- [Changing the Location of a Device on page 157](#)
- [Configuring Buildings on page 159](#)
- [Configuring Floors on page 160](#)
- [Configuring Outdoor Areas on page 162](#)
- [Creating a Site on page 163](#)
- [Deleting Sites, Buildings, Floors, Wiring Closets, and Devices on page 164](#)
- [Setting Up Closets on page 166](#)
- [Setting Up the Location View on page 168](#)

## Understanding Build Mode in Location and Device Views of Edge Services Director

In Build mode, you build the network managed by Junos Space Edge Services Director. It provides you with the ability to use device discovery to bring devices under Edge Services 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 on page 151](#)
- [Building the Location and Custom Views on page 152](#)
- [Configuring Devices on page 153](#)
- [Managing Devices on page 154](#)

## Discovering Devices

Device discovery finds your network devices and brings them under Edge Services Director management. You provide Edge Services Director with identifying information about the devices you want Edge Services Director to manage—an IP address or hostname, an IP address range, an IP subnetwork, or a CSV file that contains this information. Edge Services Director uses the information to probe the devices by using either ping or SNMP get

requests. If a device probe is successful, Edge Services 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, Edge Services Director adds the device to its database of managed devices. Edge Services 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 Edge Services Director management if they are supported by Edge Services Director.

Besides bringing your devices under Edge Services Director management, device discovery:

- Reads the device configuration and saves it in the Junos Space configuration database. Edge Services 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 69](#).
- Imports the device configuration into the Build mode configuration. For more information about importing device configurations, see [“Importing Device Configurations” on page 67](#).

## Building the 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 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.

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 Edge Services 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 Edge Services 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, radios, logical entities. For example, you can create a class-of-service (CoS) profile that contains settings that are appropriate for point-to-point, Layer 3 VPN, and VPLS services that you can manage, provision, and monitor in Service View of Edge Services Director.

In addition to creating configuration profiles, in Build mode you can configure Link Aggregation Groups (LAGs) on routers.

### 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.

### Importing Device Configurations

---

As part of device discovery, Edge Services Director analyzes the configuration of a newly discovered device and automatically imports the configuration into the Build mode configuration for that device. For example, as part of the discovery of a wireless LAN controller, Edge Services Director imports the configurations of wireless access points from the controller and makes them available for viewing and modification under the Manage Access Point task for that controller.

As it imports the device configuration, Edge Services 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, Edge Services Director assigns the existing Port profile to those ports. For the other ports, Edge Services 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 Edge Services 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 Edge Services 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 Edge Services 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 69](#).



**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

- [Working with the Dashboard on page 45](#)

## Understanding the Location View

---

The Location View is one of the perspectives that Edge Services 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 or data center; 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. Access points can be assigned only to outdoor areas and floors. 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

- [Assigning and Unassigning Devices to a Location on page 155](#)
- [Changing the Location of a Device on page 157](#)
- [Configuring Buildings on page 159](#)
- [Configuring Floors on page 160](#)
- [Configuring Outdoor Areas on page 162](#)
- [Creating a Site on page 163](#)
- [Deleting Sites, Buildings, Floors, Wiring Closets, and Devices on page 164](#)
- [Setting Up Closets on page 166](#)
- [Setting Up the Location View on page 168](#)

## Assigning and Unassigning Devices to a Location

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 example, you cannot assign access points to buildings or closets. However, you can assign access points to floors or outdoor areas. 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 168](#).

This topic describes:

- [How to Assign or Unassign Devices on page 156](#)
- [Assigning Devices on page 157](#)

## 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 Edge Services 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 standalone service delivery gateway (SDG) or a high-availability pair of SDGs that contains a master and standby SDG, or MX Virtual Chassis devices to buildings, floors, aisles, and closets.

While assigning MX Series 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:** Edge Services Director displays the Virtual Chassis systems in the Location view network tree only if the following conditions are met:

- MX Series Virtual Chassis 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 systems are not displayed in the network tree.

### Related Documentation

- [Changing the Location of a Device on page 157](#)
- [Configuring Buildings on page 159](#)
- [Configuring Floors on page 160](#)
- [Configuring Outdoor Areas on page 162](#)
- [Creating a Site on page 163](#)
- [Deleting Sites, Buildings, Floors, Wiring Closets, and Devices on page 164](#)
- [Setting Up Closets on page 166](#)
- [Setting Up the Location View on page 168](#)

## Changing the Location of a Device

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. You can move an access point to another floor or to another outdoor area. However, you cannot move an access point to a building or wiring closet. 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 on page 158](#)
- [Changing the Location of a Device on page 158](#)

## 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. You can move an access point, only to another floor or outdoor area.
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 36 on page 158](#).

**Table 36: Contents of Selected Device Details**

Field	Description
Device Name	Hostname
Device IP	Device Address
Device Family	Hardware family of products, for example, Junos-MX.
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 on page 154](#)
- [Assigning and Unassigning Devices to a Location on page 155](#)
- [Configuring Buildings on page 159](#)
- [Configuring Floors on page 160](#)
- [Configuring Outdoor Areas on page 162](#)
- [Creating a Site on page 163](#)

## Configuring Buildings

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 on page 159](#)
- [Adding or Editing a Building for a Location on page 159](#)

### 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 Edge Services 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 37 on page 159](#) describes the fields needed to establish a building.

*Table 37: Add or Edit Building Fields*

Field	Description
Building Name	Type a representative name for the building. The Building Name is a required field.

*Table 37: Add or Edit Building Fields (continued)*

Field	Description
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 on page 154](#)
  - [Assigning and Unassigning Devices to a Location on page 155](#)
  - [Changing the Location of a Device on page 157](#)
  - [Configuring Floors on page 160](#)
  - [Configuring Outdoor Areas on page 162](#)
  - [Creating a Site on page 163](#)

---

## Configuring Floors

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 on page 160](#)
- [Adding or Editing a Building Floor for a Location on page 161](#)

### 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 Edge Services 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 38 on page 161](#).

**Table 38: 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 <b>Device Inventory</b> page in the default view.
Cancel	Discards any configuration changes.

- Related Documentation**
- [Understanding the Location View on page 154](#)
  - [Assigning and Unassigning Devices to a Location on page 155](#)
  - [Changing the Location of a Device on page 157](#)
  - [Configuring Buildings on page 159](#)
  - [Configuring Outdoor Areas on page 162](#)
  - [Creating a Site on page 163](#)

---

## Configuring Outdoor Areas

You can associate an outdoor area to a site or a building for wireless coverage 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 on page 162](#)
- [Configuring an Outdoor Area on page 162](#)

### 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 Edge Services 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 39 on page 162](#) describes the fields and buttons necessary to create or change an outdoor area.

*Table 39: Outdoor Area Fields*

Field	Description
Outdoor Area Name	Type the name of the outdoor area. Edge Services 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 on page 154](#)
  - [Assigning and Unassigning Devices to a Location on page 155](#)
  - [Changing the Location of a Device on page 157](#)
  - [Configuring Buildings on page 159](#)
  - [Configuring Floors on page 160](#)
  - [Creating a Site on page 163](#)

## Creating a Site

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 Edge Services Director, the first step is to create a site.

This topic describes:

- [How to Add or Edit a Location Site on page 163](#)
- [Creating or Editing a Site on page 163](#)

### How to Add or Edit a Location Site

1. Click the Build Mode icon in the Edge Services 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 40 on page 163](#).

*Table 40: 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 40: 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 because it sets the regulatory country code for wireless devices. Edge Services Director validates the country code against the country codes in the network's controllers and access points. If the codes do not match, a warning message is sent.</p>
---------	--

**Related Documentation**

- [Understanding the Location View on page 154](#)
- [Assigning and Unassigning Devices to a Location on page 155](#)
- [Changing the Location of a Device on page 157](#)
- [Configuring Buildings on page 159](#)
- [Configuring Floors on page 160](#)
- [Configuring Outdoor Areas on page 162](#)

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

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 on page 164](#)
- [Deleting Sites on page 165](#)
- [Deleting Buildings on page 165](#)
- [Deleting Floors on page 165](#)
- [Deleting Closets on page 165](#)
- [Deleting Devices on page 165](#)

### How to Delete a Location Object

1. Ensure you are in the Build mode and Location view. Click **Build** in the Edge Services 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 on page 154](#)
- [Assigning and Unassigning Devices to a Location on page 155](#)
- [Changing the Location of a Device on page 157](#)
- [Configuring Buildings on page 159](#)
- [Configuring Floors on page 160](#)
- [Configuring Outdoor Areas on page 162](#)
- [Creating a Site on page 163](#)

---

## Setting Up Closets

---

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 on page 166](#)
- [Adding or Editing a Wiring Closet on page 167](#)

### How to Add or Edit a Closet

To add or change a wiring closet:

1. Click the Build Mode icon in the Edge Services 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 on page 154](#)
- [Assigning and Unassigning Devices to a Location on page 155](#)
- [Changing the Location of a Device on page 157](#)
- [Configuring Buildings on page 159](#)
- [Configuring Floors on page 160](#)
- [Configuring Outdoor Areas on page 162](#)
- [Creating a Site on page 163](#)

## 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 41 on page 168](#) describes the devices that you can assign to each of the location component.

**Table 41: Devices that can be Assigned to each Location Component**

Component	Devices that can be assigned
Site	None
Building	MX Series routers
Floor	MX Series routers
Closet	MX Series routers
Aisle	None
Rack	MX Series routers
Outdoor Area	MX Series routers

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 Edge Services 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.

Edge Services 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 163](#).

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

Edge Services 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. Edge Services 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 for wireless coverage 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**.

Edge Services 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 159](#).

- 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 162](#).
  - 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 160](#).
  - 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.

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

**Related  
Documentation**

- [Understanding the Location View on page 154](#)
- [Assigning and Unassigning Devices to a Location on page 155](#)
- [Changing the Location of a Device on page 157](#)
- [Configuring Buildings on page 159](#)
- [Configuring Floors on page 160](#)
- [Configuring Outdoor Areas on page 162](#)
- [Creating a Site on page 163](#)



## CHAPTER 11

# Device Management

- [Accessing a Device's CLI from Edge Services Director on page 173](#)
- [Deleting Devices from Edge Services Director on page 174](#)
- [Rebooting Devices from Edge Services Director on page 175](#)
- [Viewing the Device Inventory Page in Device View of Edge Services Director on page 176](#)
- [Viewing the Physical Inventory of Devices on page 178](#)
- [Viewing a Device's Current Configuration from Edge Services Director on page 179](#)

### Accessing a Device's CLI from Edge Services Director

Edge Services Director enables you to connect to the CLI for devices in your network, using SSH.

This topic describes the steps to connect to a router 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 router or controller opens in a separate browser tab or window depending on your browser settings. Refer to the [MX Series documentation](#) for more information about using the CLI for MX Series routers.



**NOTE:** Any configuration changes you make to a device, using the CLI qualify as out-of-band changes in Edge Services 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 the Edge Services Director User Interface on page 6](#)

---

## Deleting Devices from Edge Services Director

---

You can delete devices that are no longer used from Edge Services 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 Edge Services 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, Edge Services 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 router family in Device View and click Delete Devices, only routers that belong to that router family are displayed.

To delete devices, complete the following tasks:

1. Select the check box adjacent to the router that you want to delete.
2. Click **Done**.

Edge Services 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 Edge Services Director User Interface on page 6](#)

## Rebooting Devices from Edge Services Director

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 Edge Services Director User Interface on page 6](#)

## Viewing the Device Inventory Page in Device View of Edge Services Director

The Device Inventory page lists devices managed by Edge Services 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.

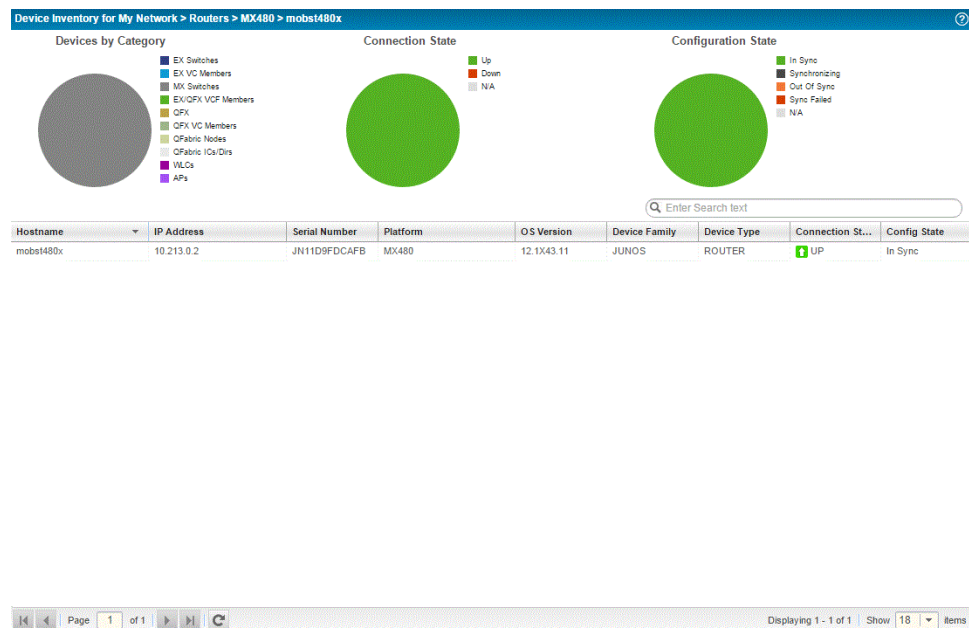
The scope you have selected in the View pane and the network view that you have selected from the View selector determines which devices are listed in the Device Inventory page. For example:

- If you are in the Device View and select My Network, all devices managed by Edge Services Director are listed.
- 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 Category—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 32 on page 134](#) for definitions of the configuration states.

Figure 17: Device Inventory Page



Mouse over a pie segment to view the actual number of devices and the percentage represented by that pie segment.

Table 32 on page 134 describes the fields in the Device Inventory table.

**Table 42: 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.
Device Family	Device family of the device, such as JUNOS for MX Series routers.
Device Type	Type of the device: <ul style="list-style-type: none"> <li>• ROUTER—MX Series routers</li> <li>• AP—Wireless LAN access point</li> <li>• Fabric Member—QFabric member switch</li> <li>• QFabric—QFabric system</li> <li>• Switch—Standalone switch</li> <li>• VC—Virtual Chassis master</li> <li>• VC Member—Virtual Chassis member switch</li> </ul>
Connection State	Connection status of the device in Edge Services Director: <ul style="list-style-type: none"> <li>• UP—Device is connected to Edge Services Director.</li> <li>• DOWN—Device is not connected to Edge Services Director.</li> <li>• N/A—Access point state is unavailable to Edge Services Director.</li> </ul>
Config State	Displays the configuration status of the device: <ul style="list-style-type: none"> <li>• In Sync—The configuration on the device is in sync with the Edge Services Director configuration for the device.</li> <li>• Out Of Sync—The configuration on the device does not match the Edge Services Director configuration for the device. This state is usually the result of the device configuration being altered outside of Edge Services Director. You cannot deploy configuration on a device from Edge Services Director when the device is Out Of Sync. To resolve this state, use the Resynchronize Device Configuration task in Deploy mode.</li> <li>• Sync failed—An attempt to resynchronize an Out Of Sync device failed.</li> <li>• Synchronizing—The device configuration is in the process of being resynchronized.</li> <li>• N/A—The device is down or is an access point.</li> </ul>

*Table 42: Fields in the Device Inventory Table (continued)*

Field	Description
Manageability State	Displays if the device is directly manageable or not.  This is a hidden field. To display the Manageability State field, click any column, click the down arrow to expand the list, select <b>Columns</b> from the list, and then enable <b>Manageability State</b> .

## 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. Edge Services 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 an MX Series router 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 MX Series routers, 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 43 on page 178](#).

*Table 43: 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, Edge Services 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	Model number of the FRU hardware component.
Part Number	Part number of the MX Series router chassis component.
Serial Number	The hardware serial number of the device.
Description	The description about the component.

- Related Documentation**
- [Understanding the Edge Services Director User Interface on page 6](#)

## Viewing a Device's Current Configuration from Edge Services Director

---

You can view a device's current configuration from Edge Services Director. This is a convenient way to view device configurations without leaving Edge Services Director.

To view a device's current configuration:

1. Click **Build** or **Deploy** in the Edge Services 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.



**NOTE:** Juniper Networks devices require a license to activate the feature. To understand more about Edge Services Director Licenses, see, [Licenses for Network Management](#). Please refer to the Licensing Guide for general information about License Management. Please refer to the product Data Sheets for further details, or contact your Juniper Account Team or Juniper Partner.

- Related Documentation**
- [Understanding the Edge Services Director User Interface on page 6](#)



## PART 5

# Service View of Build Mode

- [About Build Mode in Service View on page 183](#)
- [Using the Service Designer on page 185](#)
- [Using the Object Builder on page 305](#)
- [Managing Packet Analyzers on page 321](#)



## CHAPTER 12

# About Build Mode in Service View

- [Understanding Build Mode in Service View of Edge Services Director on page 183](#)

## Understanding Build Mode in Service View of Edge Services Director

---

In Build mode, you can create services, policies, and filters for devices that are managed by Edge Services Director. The service templates and attributes for services, policies, and filters help you classify and control the manner in which packets must be handled by the various services.

Configuring a policy has a major impact on the flow of routing information or packets within and through the router. For example, you can configure a routing policy that does not allow routes associated with a particular customer to be placed in the routing table. As a result of this routing policy, the customer routes are not used to forward data packets to various destinations and the routes are not advertised by the routing protocol to neighbors.

You can also import objects, which are components or parameters used for creation of services, from the Service Delivery Gateways (SDGs) that are present in the Edge Services Director database or from external XML files.

This topic contains the following sections that describe the different workspaces or utilities that you can access from Build mode:

- [Service Designer on page 183](#)
- [Services Inventory on page 184](#)
- [Object Builder on page 184](#)

### Service Designer

The service planning functionality enables you to use the Service Designer page to create service templates, which can be used on multiple devices. The Service Designer page lists all service components used to create service templates. According to the business needs, you can configure generic properties in a template and enable the editing of deployment-specific parameters. The operator can then easily and quickly configure the service on a large number of devices. You can use the Service Designer page to define and manage stateful firewall (SFW), carrier-grade NAT (CGNAT), application delivery controller (ADC), and traffic load balancing (TLB) services.



**NOTE:** Edge Services Director currently supports only brownfield deployments and not greenfield deployments. A greenfield deployment refers to the Junos OS base configurations and bootstrapping, core device settings such as routing instances, interfaces and IP addresses, and routing protocols to be available for configuration using the network management application. A brownfield deployment refers to the basic and mandatory device settings already being configured on the devices before they are imported or discovered for additional modifications, such as configuration of services, using the network management application.

As a designer, you can also modify service parameters and definitions by using the View Service page that you can open from the Service Gateways -- Unmanaged devices page in the Build mode without using service templates for updating services details. All the service components are listed on the Service Designer page so that the designer can use choose components and create the new service template.

---

## Services Inventory

The Services Inventory page lists the services configured in the Edge Services Director database and provides basic information about the configured services, such as adaptive delivery controller (ADC), stateful firewall (SFW), server load balancing (SLB), and carrier grade NAT (CGNAT). The Services Inventory page is available in Build mode and under Service view.

## Object Builder

Objects are constituents or building blocks that are used to create service definitions and policy or filter templates. You can use the Object Builder page to retrieve and transfer the objects or components that have been previously created on the SDGs or devices. The objects might reside on the managed SDGs or SDG groups if the objects were defined using the appropriate configuration statements and parameters in the Junos CLI interface of the respective SDGs. This mechanism of importing object settings enables you to easily, quickly, and optimally use the object definitions when you create service and policy templates.

### Related Documentation

- [Understanding Edge Services Director and the Management Lifecycle Modes on page 15](#)

## CHAPTER 13

# Using the Service Designer

- [Object Builder Overview on page 185](#)
- [Planning and Deployment of Service Templates Overview on page 187](#)
- [Service Templates Overview on page 189](#)
- [Filtering Service Templates on page 189](#)
- [Restoring Service Template Configurations on page 190](#)
- [Viewing Service Templates on page 192](#)
- [Viewing the Services Inventory Page on page 193](#)
- [Using the Actions Menu on the Service Template and Service Edit Pages on page 195](#)
- [Viewing a Graphical Statistic of Service Templates on page 202](#)
- [Creating and Managing ADC Service Templates on page 203](#)
- [Creating and Managing CGNAT Service Templates on page 234](#)
- [Creating and Managing SFW Service Templates on page 252](#)
- [Creating and Managing TLB Service Templates on page 275](#)
- [Modifying Individual Service Instances and Deploying to Devices on page 300](#)

### Object Builder Overview

---

You can use the Object Builder workspace in Edge Services Director to create objects to be used by firewall policies, VPNs, and NAT policies. These objects are stored in the Junos Space database. You can reuse these objects with multiple security policies, VPNs, and NAT policies.

You can use the Object Builder workspace to create, modify, clone, and delete the following objects:

- Addresses and address groups
- Services and service groups
- Variables

You cannot delete any of the objects that you created in Object Builder (except Template definition and Templates) if they are already used in a firewall policy, NAT policy, or any other service definition.

Object Builder supports concurrent editing of its objects, with a save as option to save your changes with a different name.

Concurrent editing is supported for the following objects:

- Addresses and address groups
- Application signatures
- Stateful firewall rules
- Stateful firewall rule sets
- CGNAT pools
- CGNAT rule sets
- CGNAT rules
- Real server settings

If you attempt to save your changes to an object that has been modified since you began editing, you receive an error message.

**Related  
Documentation**

- [Understanding the Object Builder on page 305](#)
- [Importing All Types of Objects on page 306](#)
- [Importing SFW Rule Sets on page 308](#)
- [Importing SFW Rules on page 310](#)
- [Importing Real Server Settings on page 312](#)
- [Importing CGNAT Rule Sets on page 313](#)
- [Importing CGNAT Rules on page 315](#)
- [Importing CGNAT Pools on page 316](#)
- [Importing Applications on page 318](#)
- [Importing Application Sets on page 319](#)

---

## Planning and Deployment of Service Templates Overview

---

The service planning functionality of Edge Services Director enables you to create service templates and deploy the same service template configuration to multiple devices. As a designer, when you create a service template, you can configure generic properties and modify it to suit your network deployment needs, thereby enabling streamlined and simplified administration of services (such as stateful firewall [SFW], carrier-grade NAT [CGNAT], application delivery controller [ADC], and traffic load balancing [TLB]) on service delivery gateways (SDGs) in your topology.

This topic contains the following sections that describe the sequence of operations performed for planning and deploying service templates:

- [Planning Workflow for Service Templates on page 187](#)
- [Deployment Workflow for Service Templates on page 187](#)

### Planning Workflow for Service Templates

The fundamental workflow of planning templates is derived from the existing devices inventory and framework:

- The designer creates the service template by using the available inventory service components and structure model.
- The designer can import the discovered service data while creating the service template for the existing service data values of the device.
- While creating the service template, the designer can add or modify service parameter values and restrict the access level for each service parameter for the operator. The designer can set the following access levels for each service parameters to operator in the planning template:
  - Read-only (The configuration parameter is read-only for operator as part of provisioning.)
  - Editable (The configuration parameter is editable as part of provisioning.)
  - Mandatory (The configuration parameter is part of provisioning but operator must provide the values.)
  - Device-Specific (The configuration parameter value needs to be entered by the operator for each device during deployment.)

The designer must publish the service templates to the operator to use in the creation of deployment plans.

An operator can create the service deployment plan using the planning template so that one deployment plan can be applied on multiple devices. This method of deploy reduces the scope for human errors that can occur with the CLI interface.

### Deployment Workflow for Service Templates

The following workflow is used the deployment process:

1. An operator uses only published planning templates to create deployment plans for a single SDG service or multiple SDG devices.
2. The operator modifies or adds data in the allowed service specific parameters according to the access permissions specified by the designer and associate the deployment plan with a single SDG device or multiple SDG devices.
3. An operator publishes the deployment plans with the device association for the designer to review and approve.
4. The administrator must approve the configuration changes for each device for each service deployment plan.
5. The operator has a copy of the service planning template while creating the deployment plan. After creating a deployment plan, there is no association between the deployment plan and planning template. Changes made by the operator to the deployment template are maintained in their own copy and are not reflected in the original planning template and vice-versa.
6. The deployment plan is assigned to multiple devices and sent for approval. After a deployment plan is associated with a device, the device contains its own copy of the deployment plan. For example, if one deployment plan was created and associated with four devices, you see four deployment plans separately on each device in the service deployment plan. The operator can edit the deployment plan for each device if needed.

The status of a deployment plan determines the kinds of tasks that a user can perform:

- Add – Create a deployment plan; the status that immediately follows this status is the Unpublish state.
- Update – Update a deployment plan.
- Delete – Delete a deployment plan. Only plans that are in the Unpublish state can be deleted.
- Publish – Publish the deployment plan. In this state, the operator waits for an approval from the designer before the plan can be deployed to a device.
- Unpublish – Unpublish the published deployment plan to make more changes.
- Approve – The administrator or designer approves the published deployment plan.
- Reject – The administrator or designer rejects the published deployment plan.

A deployment plan can obtain any one of the following status:

- Discovered – This is the default state for filter discovered and stored in the inventory.
- Unpublished – New, updated, and deleted filters are saved in draft or Unpublished state initially.
- Published – After all the changes are done, the filter in draft status is ready for admin or designer approval and is published.
- Rejected – An administrator or designer can reject the published filter to disapprove updates.

- **Approved** – An administrator or designer can approve the published filter to concur changes.
- **Commissioned** – An administrator or designer can commission filters to push to devices.
- **Commission Failed** – This state is assigned to a filter if commissioning of the filter fails.

**Related Documentation** • [Service Templates Overview on page 189](#)

## Service Templates Overview

You use the service templates to configure the following attributes and settings for the following four types of services:

- Stateful firewall (SFW)
- Carrier-Grade network addressing (CGNAT)
- Traffic load balancer (TLB)
- Application delivery controller (ADC)

After you create and publish the service templates, you can use these templates to create service deployment plans. The deployment plans are defined with the SDGs on which services need to be deployed. After the deployment plans are published and approved, the services can be deployed to become effective on the relevant SDGs.

**Related Documentation** • [Planning and Deployment of Service Templates Overview on page 187](#)

## Filtering Service Templates

You can filter service templates to sort and identify the service definitions that are of interest or necessary for your network needs. To filter service templates based on their states:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the task pane, select **Service Template > Manage Service Templates**. The Service Templates page is displayed.
4. From the View pane, perform one of the following tasks:

- Click the **ADC** button.

The list of ADC service templates is displayed. You need not click this button if you are launching the Service Designer page for the first time or are navigating to this page from another mode or a different page. You need to click this button only if you are viewing the other service templates, such as CGNAT or TLB.

- Click the **SFW** button.

The list of SFW templates is displayed.

- Click the **TLB** button.

The list of TLB templates is displayed.

- Click the **CGNAT** button.

The list of CGNAT templates is displayed.

5. To filter and sort the display of service templates, enter the name of the template as a match criterion in the Search box and click the Search icon.

The page refreshes to display the service templates that match with the specified criterion. You can use the paging controls to navigate across multiple pages of objects as necessary.

#### Related Documentation

- [Planning and Deployment of Service Templates Overview on page 187](#)

---

## Restoring Service Template Configurations

Restoring a configuration file means either merging the contents of a configuration file with the existing configuration file on the device, or overriding the device's running configuration file with a candidate configuration file.

When you restore a configuration file, an audit log entry is automatically generated.

To restore a service template configuration file from Edge Services Director to a device:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the task pane, select **Service Edit**.

On the right pane, pie charts corresponding to the configured services and policy filters are displayed if you view the page without drilling-down the tree in the task pane to select a particular service or policy.

- Select **Service Edit > Packet Filter** to display the Packet Filter page

- Select **Service Edit > SFW Policy and Filters** to display the SFW Policy and Filters page
  - Select **Service Edit > CGNAT Policy and Filter** to display the CGNAT Policy and Filter page
4. On the appropriate Filters page, such as a packet filter or stateful firewall filter, select the SDG host or the device whose configuration you want to restore. (To restore all of them, select the check box in the column header next to the first column header.)
  5. From the Actions menu, select **Restore**.  
The **Restore** dialog box appears, displaying the name of the filter, name of the selected file, the name of the device, and the version that is to be restored to the device. By default, the latest version is merged with the existing configuration on the device.
  6. In the **Version** column, click next to the version number displayed and select the appropriate version from the drop-down list that appears.  
The timestamp that is displayed adjacent to the version number indicates the time at which the version of the configuration was commissioned.
  7. In the **Associated Devices** column, select the devices to which you want to restore the selected configuration. You can override the device's running configuration file with a candidate configuration file.
  8. To initiate the restoration of the configuration, click **Restore**.  
The **Restore Configuration Files** dialog box that appears indicates the successful scheduling of the restoration.  
The word Success in the Status column on the Job Management page indicates that the restoration is successful. If a device cannot be accessed, the device is skipped, and the job status indicates a failure.
  9. Click **OK** to close the Restore Configuration Files dialog box.
  10. Alternatively, to abort the restoration, click **Cancel** to close the Restore dialog box.

**Related  
Documentation**

- [Service Templates Overview on page 189](#)
- [Filtering Service Templates on page 189](#)
- [Viewing Service Templates on page 192](#)
- [Using the Actions Menu on the Service Template and Service Edit Pages on page 195](#)

## Viewing Service Templates

---

On the Service Designer page, you can view the collection of service templates defined for several applications, such as stateful firewall or CGNAT.

To view the list of service templates, such as ADC, SFW, CGNAT, or TLB templates:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.

3. From the task pane, select **Service Template > View Statistics**.

The Service Designer page displays a bar graph in the top pane of the page. The total number of service templates of each type is displayed on the vertical axis and the service type is shown on the horizontal axis. A color-coding format is used to represent the bars on the graph. Published service templates are shown in olive green color and unpublished service templates are shown in blue color. Mouse over each bar in the chart to highlight and display the number of templates published or unpublished for each type of service.

4. From the View pane, perform one of the following tasks:
  - Click the All Services item to view all of the service types, such as ADC, TLB, SFW, and CGNAT.
  - Click the **ADC** button.  
The list of ADC service templates is displayed. You need not click this button if you are launching the Service Designer page for the first time or are navigating to this page from another mode or a different page. You need to click this button only if you are viewing the other service templates, such as CGNAT or TLB.
  - Click the **SFW** button.  
The list of SFW templates is displayed.
  - Click the **TLB** button.  
The list of TLB templates is displayed.
  - Click the **CGNAT** button.  
The list of CGNAT templates is displayed.

[Table 44 on page 193](#) describes the fields displayed on the Service Designer page:

*Table 44: Service Designer View*

Field	Description
Name	Name of the service template.
Description	User-defined description of the template.
Created By	Name of the user who created the template
Created Time	Time and date when the template was created. The server time zone determines the time zone displayed on this page.
Modified Time	Time and date when the template was last updated. The server time zone determines the time zone displayed on this page.

5. Click the **Add** icon above the list of templates to create a new template.
6. Click the **Edit** icon above the list of templates to modify an existing template.
7. Click the **Delete** icon above the list of templates to delete an existing template.

#### Related Documentation

- [Service Templates Overview on page 189](#)
- [Filtering Service Templates on page 189](#)
- [Using the Actions Menu on the Service Template and Service Edit Pages on page 195](#)

## Viewing the Services Inventory Page

The Services Inventory page lists the services configured in the Edge Services Director database and provides basic information about the configured services, such as adaptive delivery controller (ADC), stateful firewall (SFW), server load balancing (SLB), and carrier grade NAT (CGNAT). The Services Inventory page is available in Build mode and under Service view.

To view the services inventory page:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the task pane, select **Inventory > View Inventory**. The Services Inventory page is displayed.

Figure 18: Services Inventory Page

Inventory for > All Services

Enter Search text

Service Name	Service Gateway	Host	Service Type	Service Pic and Interface	Service Gateway Group
b1	dsfctd	mobs1480x	ADC	ms-0/0/0	dsdsd
NAPT44-SS1	dsfctd	mobs1480x	CGNAT	sp-2/0/0 1000, sp-2/0/0 100	dsdsd
NAPT44-SS2	dsfctd	mobs1480x	CGNAT	sp-2/1/0 1000, sp-2/1/0 100	dsdsd
IPv6-SFW	dsfctd	mobs1480x	SFW	sp-1/1/0 100, sp-1/1/0 99	dsdsd
lib_sdg	dsfctd	mobs1480x	TLB	ms-1/0/0 0	dsdsd
lib_sdg_v6	dsfctd	mobs1480x	TLB	ms-1/0/0 0	dsdsd
b1	dsfctd	mobs1480w	ADC	ms-0/0/0	dsdsd
NAPT44-SS1	dsfctd	mobs1480w	CGNAT	sp-2/0/0 1000, sp-2/0/0 100	dsdsd
NAPT44-SS2	dsfctd	mobs1480w	CGNAT	sp-2/1/0 1000, sp-2/1/0 100	dsdsd
IPv6-SFW	dsfctd	mobs1480w	SFW	sp-1/1/0 100, sp-1/1/0 99	dsdsd
lib_sdg	dsfctd	mobs1480w	TLB	ms-1/0/0 0	dsdsd
lib_sdg_v6	dsfctd	mobs1480w	TLB	ms-1/0/0 0	dsdsd

Page 1 of 1

Displaying 1 - 12 of 12 | Show 30 items

4. From the View pane, do one of the following:

- Select **ADC** to open the Inventory > ADC page on the right pane.
- Select **TLB** to open the Inventory > TLB page on the right pane.
- Select **CGNAT** to open the Inventory > CGNAT page on the right pane.
- Select **SFW** to open the Inventory > SFW page on the right pane.

Table 45 on page 194 describes the fields on the Services page.

Table 45: Fields on the Services Page

Field	Description
Service Name	Name of the configured service, such as stateful firewall or CGNAT. Click the plus sign (+) beside each service to view extensive information on attributes configured for the service.
Service Gateway	Name of the service delivery gateway.
Host	Hostname of the device.
Service Type	Type of the service, such as ADC, SFW, CGNAT, or TLB.
Service Pic and Interface	Services PIC and interface details, such as multiservices PIC or adaptive services PIC with the FPC slot, PIC, and port attributes.
SDG Group	Name of the group to which the SDG is assigned.

To view configuration and run-time information for services:

1. Sort the table by mousing over the column header for the data you want to sort by and clicking the down arrow. Select **Sort Ascending** or **Sort Descending**.
2. Show columns not in the default table view, or hide columns, as follows:
  1. Mouse over any column header and click the down arrow.
  2. Select **Columns** from the menu.
  3. Select the check boxes for columns that you want to view. Clear the check boxes for columns that you want to hide.
3. View information about devices as follows:
  - To restrict the display of devices, enter a search criterion of one or more characters in the Search bar and press Enter.All devices that match the search criterion are shown in the main display area.

**Related  
Documentation**

- [Viewing Device Statistics on page 141](#)
- [Viewing Configuration Details of Services on Devices on page 144](#)
- [Viewing Discovery Logs on page 146](#)
- [Viewing Discovery Profiles on page 147](#)

---

## Using the Actions Menu on the Service Template and Service Edit Pages

You can use the Actions menu on the Service Template and Service Edit pages for ADC, TLB, CGNAT, and SFW service instances to publish, unpublish, and clone the defined service instances. You can also create a deployment plan for the service or disregard the changes done to the service.

- [Publishing a Service Template on page 196](#)
- [Unpublishing a Service Template on page 197](#)
- [Exporting a Service to a CSV File on page 198](#)
- [Cloning a Service Template on page 199](#)
- [Creating a Deploy Plan and Provisioning Services Immediately on page 200](#)

## Publishing a Service Template

You need to publish a service template definition when you want to make it available to create device templates from the template definition.

To publish a service template:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.

The functionalities that you can configure in this mode are displayed in the task pane.

3. From the task pane, select **Service Template > Manage Service Templates**.

The Service Templates page is displayed in the right pane, listing all the previously defined service instances.

4. From the View pane, perform one of the following tasks:

- Click the **ADC** button.

The list of ADC service instances is displayed. You need not click this button if you are launching the Service Designer page for the first time or are navigating to this page from another mode or a different page. You need to click this button only if you are viewing the other service instances, such as CGNAT or TLB.

- Click the **SFW** button.

The list of SFW templates is displayed.

- Click the **TLB** button.

The list of TLB templates is displayed.

- Click the **CGNAT** button.

The list of CGNAT templates is displayed.

5. Select a template, and click the **Publish** button.

The filter status changes from “Draft” to “Published”. The Publish option is available only if all selected filters are assigned the Draft status.

## Unpublishing a Service Template

To make a template definition unavailable to operators, you must unpublish it. You must also unpublish a definition before you can modify or delete it. If you unpublish a definition that is already being used as the basis for templates, all templates based on that definition are disabled. Republishing the definition alone is not enough to reenable the templates. The templates must be reviewed before they can be reenabled.

To unpublish a service template:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.

3. From the task pane, select **Service Template > Manage Service Templates**.

The Service Templates page is displayed in the right pane, listing all the previously defined service instances.

4. From the View pane, perform one of the following tasks:

- Click the **ADC** button.

The list of ADC service instances is displayed. You need not click this button if you are launching the Service Designer page for the first time or are navigating to this page from another mode or a different page. You need to click this button only if you are viewing the other service instances, such as CGNAT or TLB.

- Click the **SFW** button.

The list of SFW templates is displayed.

- Click the **TLB** button.

The list of TLB templates is displayed.

- Click the **CGNAT** button.

The list of CGNAT templates is displayed.

5. Select a template, and click the **Unpublish** button above the table of listed templates.

The filter status changes from “Published” to “Draft”. The Unpublish option is available only if all selected filters are assigned the Published status.

## Exporting a Service to a CSV File

You can export the service template settings and parameters to a comma-separated value (.csv) file to open it by using a spreadsheet or any other business application on your client computer.

To export a service template to CSV file:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.

3. From the task pane, select **Service Template > Manage Service Templates**.

The Service Templates page is displayed in the right pane, listing all the previously defined service instances.

4. From the View pane, perform one of the following tasks:

- Click the **ADC** button.

The list of ADC service instances is displayed. You need not click this button if you are launching the Service Designer page for the first time or are navigating to this page from another mode or a different page. You need to click this button only if you are viewing the other service instances, such as CGNAT or TLB.

- Click the **SFW** button.

The list of SFW templates is displayed.

- Click the **TLB** button.

The list of TLB templates is displayed.

- Click the **CGNAT** button.

The list of CGNAT templates is displayed.

5. Click the **Actions** menu, and select **Export to CSV** from the drop-down menu.

The Export dialog box appears.

6. Export the policy information to the CSV file. You can export information about selected devices or export information about all of the devices managed by Junos Space.
  - Select the check box next to the managed SDG or SDG pair that you want to export to a CSV file, and click the **Export Selected** button to export the policy information about selected devices and begin creating the CSV file.
  - Click the **Export All** button to export the policy information for all the devices and begin creating the CSV file.

A progress bar is displayed to indicate the percentage of completion of the export job. After the export job is completed, a download link is displayed that you can click to download the CSV file.

## Cloning a Service Template

You clone a template definition to quickly create a new template definition with a new name but same properties. To modify a template definition without disabling templates based upon that definition, first clone the definition, then modify the clone.

To clone a service template:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.

3. From the task pane, select **Service Template > Manage Service Templates**.

The Service Templates page is displayed in the right pane, listing all the previously defined service instances.

4. From the View pane, perform one of the following tasks:

- Click the **ADC** button.

The list of ADC service instances is displayed. You need not click this button if you are launching the Service Designer page for the first time or are navigating to this page from another mode or a different page. You need to click this button only if you are viewing the other service instances, such as CGNAT or TLB.

- Click the **SFW** button.

The list of SFW templates is displayed.

- Click the **TLB** button.

The list of TLB templates is displayed.

- Click the **CGNAT** button.

The list of CGNAT templates is displayed.

5. Select the template you want to clone
6. Click the **Clone** button above the table of displayed templates.
7. In the Name field, type a user-defined template definition name. A template definition name cannot exceed 128 characters and can contain only letters, numbers, spaces, and some special characters. The special characters allowed are hyphen (-),

underscore (\_), period (.), at (@), single quote ('), forward slash (/), and ampersand (&).

8. (Optional) In the Description field, type a user-defined description. (limit of 255 characters). The description cannot exceed 256 characters. The operators who use the template definition to create templates rely on the description for information on the template definition.
9. Click Save to save the template. The dialog box closes and the Manage Service Templates window appears.

## Creating a Deploy Plan and Provisioning Services Immediately

To deploy a deployment plan and policies immediately:

1. From the View selector, select **Gateway View** or **Service View**. The workspaces that are applicable to this view are displayed. In Gateway view, the devices in the entire network are displayed, organized by the device types and the device models within each device type. In Service View, the different types of services are displayed in the View pane.
2. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the View pane, select the All Network item in Gateway view. Click the plus sign (+) beside the All Network item in the View pane to expand the tree and select the device node you want.
4. From the task pane, select **Service Edit**. The Service Instances page is displayed.
5. In the Service Edit page, from the tree that lists the SDGs, select **All Service Gateways**, or the SDG or SDG pair for which you want to view the previously configured policy or filter templates. This step is applicable only if you selected Gateway View.  
  
The page is divided into two panes. The list of SDGs are displayed on the left pane. You can drill-down to the SDG or pair of SDGs for which you want to process policies or filters. The policy and filter rules are displayed in the right pane.
6. Alternatively, if you are in Service view, from the View pane, perform one of the following tasks:

- Click the **ADC** button.

The list of ADC service instances is displayed. You need not click this button if you are launching the Service Designer page for the first time or are navigating to this page from another mode or a different page. You need to click this button only if you are viewing the other service instances, such as CGNAT or TLB.

- Click the **SFW** button.

The list of SFW templates is displayed.

- Click the **TLB** button.

The list of TLB templates is displayed.

- Click the **CGNAT** button.

The list of CGNAT templates is displayed.

7. From the task pane, select **Service Edit**. The Service Instances page is displayed.

8. In the Service Instances page, from the tree that lists the SDGs, select **All Service Gateways**, or the SDG or SDG pair for which you want to view the previously configured policy or filter templates.

You can drill-down to the SDG or pair of SDGs for which you want to process policies or filters. The service instances associated with each SDG in an SDG pair are displayed.

9. In the Service Instances page, select a service instance and click the **Lock** icon.

The corresponding service instance is locked and is available for modifications.

10. Click the **Send for Deployment** button.

- If you create a deployment plan from Gateway view of Deploy mode, the Deployment Plan Summary dialog box appears, with the service name, type, and status listed.

Click **Send** to create a deployment plan.

- If you create a deployment plan from Service view of Deploy mode, the Edit Service Instance page is displayed. You can modify the SDGs associated with the service instance and also modify the service instance attributes as necessary by either clicking the buttons corresponding to the various settings at the top of the wizard page to directly traverse to the page you want to modify or clicking the navigation buttons at the bottom of the wizard page to go to the different pages of the wizard. Click **Finish** to create a deployment plan.

The configuration deployment job runs. To view the status or results of the deployment job, you can view the Deployment Plans page. In the Deployment Plans page, the Provision Status and Message columns are updated indicating the progress of commission. If the deploy is successful, the status denotes Commissioned. If the deploy fails, the status changes to Commission Failed.

Alternatively, you can select **Discard changes** from the **Actions** menu to ignore the modifications done to a policy or filter template.

- Related Documentation**
- [Service Templates Overview on page 189](#)
  - [Filtering Service Templates on page 189](#)
  - [Viewing Service Templates on page 192](#)

## Viewing a Graphical Statistic of Service Templates

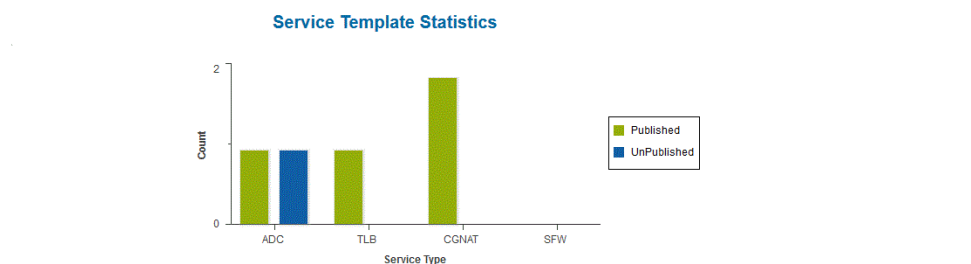
---

To view the total number of service templates that are previously configured in the Edge Services Director database and are in the published or unpublished states:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the task pane, select **Service Template > View Statistics**.

The Service Template Statistics page is displayed.

*Figure 19: Service Template Statistics Page*



The page displays a bar graph in the top pane of the page. The count of service templates of each type is displayed on the vertical axis and the service type is shown on the horizontal axis. A color-coding format is used to represent the bars on the graph. Published service templates are shown in olive green color and unpublished service templates are shown in blue color. Mouse over each bar in the chart to highlight and display the number of templates published or unpublished for each type of service.

- Related Documentation**
- [Creating and Managing ADC Service Templates on page 203](#)
  - [Creating and Managing CGNAT Service Templates on page 234](#)
  - [Creating and Managing SFW Service Templates on page 252](#)
  - [Creating and Managing TLB Service Templates on page 275](#)

---

## Creating and Managing ADC Service Templates

You can configure the adaptive delivery controller (ADC) software within your router to balance user session traffic among a group of available servers that provide shared services. The ADC software uses Junos OS firewall filters, Junos OS routing instances of type forwarding-instance, and Junos OS logical interfaces and interface address families (units and addresses) defined on the Multiservices-DPCs running the ADC software.

You can perform the following tasks with the Service Designer page for ADC:

- Create an ADC service template with attributes and settings for load balancing operations.
- Modify an existing ADC template to meet the network needs and deployment scenarios.
- Delete an existing template.
- [Creating an ADC Service Template on page 204](#)
- [Importing an ADC Service Template on page 207](#)
- [Creating a Deployment Plan on page 209](#)
- [Creating a Real Server on page 210](#)
- [Creating a Group for Real Servers on page 212](#)
- [Load-Balancing Methods for Real-Server Groups on page 214](#)
- [Creating a Client-Facing Interface and Routing Instance on page 216](#)
- [Creating a Server-Facing Interface and Routing Instance on page 218](#)
- [Creating a Services PIC for an ADC Service Template on page 219](#)
- [Creating a Health Check for an ADC Service Template on page 221](#)
- [Creating a Custom Health Check for an ADC Instance on page 222](#)
- [Creating a Virtual Service for an ADC Service Template on page 225](#)
- [Creating a Virtual Server for an ADC Service Template on page 228](#)
- [Creating a Firewall Rule for an ADC Service Template on page 229](#)
- [Modifying ADC Service Templates on page 232](#)

## Creating an ADC Service Template

To configure a new ADC service template:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.

The functionalities that you can configure in this mode are displayed in the task pane.

3. From the task pane, select **Service Templates**.

The Manage Service Templates page is displayed.

4. Click the **ADC** button.

The list of ADC service templates is displayed. You need not click this button if you are launching the Service Designer page for the first time or are navigating to this page from another mode or a different page. You need to click this button only if you are viewing the other service templates, such as CGNAT or TLB.

The Service Designer page displays a bar graph in the top pane of the page. The total number of service templates of each type is displayed on the vertical axis and the service type is shown on the horizontal axis. A color-coding format is used to represent the bars on the graph. Published service templates are shown in olive green color and unpublished service templates are shown in blue color. Mouse over each bar in the chart to highlight and display the number of templates published or unpublished for each type of service.

5. Click the **Add** icon.

The Create an ADC Planning Template window appears.

Figure 20: Create ADC Service Template Window

6. In the Name field, enter a name for the service template or profile (limit of 63 alphanumeric characters without spaces).
7. In the ADC Instance Name field, enter a name for the service instance (limit of 63 alphanumeric characters without spaces). Each service instance that you define can be applied to a single SDG or multiple SDGs.

8. (Optional) Alternatively, instead of creating a new template entirely, click the **Import** button to clone an existing template by importing it. You can import the parameters defined for a previous ADC service instance and customize only the settings that are necessary.

Imported templates are created without any device assigned to them. To use these templates, you must associate a device with the policy.

The Import Services dialog box is displayed. See *Importing an ADC Service Template* for step-wise details on importing an ADC service template.

9. The Create an ADC Planning Template window displays the individual elements or components of the service with a graphical icon for each of the service elements and the corresponding names in separate boxes. You can add, edit, or delete these service elements in a template.

The Property View tab and the Config View tab are displayed on the right pane of the template window. The Property View tab provides a tree-based structure of the parameters defined in a service template. You can expand the tree and view details of each component. A key value pair representation is shown. Each of the components can be treated as categories of the service template shown in the property view.

The Config View tab displays the elements or components specified for a service template in the form of configuration stanzas and hierarchy levels. This display is similar to the show command that you can use at a certain [edit] hierarchy level to view the defined settings. Each level in the hierarchy is indented to indicate each statement's relative position in the hierarchy. Each level is generally set off with braces, with an open brace ( { ) at the beginning of each hierarchy level and a closing brace ( } ) at the end. If the statement at a hierarchy level is empty, the braces are not displayed. Each leaf statement ends with a semicolon ( ; ), as does the last statement in the hierarchy.

- a. Click the green tick mark (✓) displayed at the top-right corner of each of the service element boxes to create a new element. If the green tick mark is not shown, it indicates that the user role does not have the permission to create an element.
- b. Click the red cross mark (x) displayed at the top-right corner of the icons of each element if you want to delete the existing configuration. The user with designer role has permissions to remove or edit elements.
- c. If the red cross mark is not displayed beside a particular icon, it signifies that the element cannot be deleted.
- d. The diamond icon that contains an orange tick mark within it at the top-right corner of the service component name denotes that the particular element can be modified. The absence of this icon denotes that the user does not have permissions to modify the attributes of the service component.
- e. Double-click each icon pertaining to a service element to view or edit its settings. If you do not possess the permission to modify the element, a view-only dialog box with the attributes of the selected element is shown. Otherwise, an editable dialog box enables you to modify the settings.
- f. Click the Maximize icon displayed at the top-right corner of the rectangle or box that shows all of the values or entities of a particular component of a service template. The specified component or attribute is displayed as a separate dialog box, listing all of the values of the particular component. You can add, modify, or delete the listed values.
- g. While creating the new service template, the designer can add or modify service parameter values and also restrict the access level for each service parameter for the operator. The designer can set following access levels for each service parameters to operator in planning template. Click the new icon (cascading files icon) displayed at the top-left corner of each of the element boxes to open the shortcut menu. You can click one of the following radio buttons:
  - Read-only (the configuration parameter is read-only for operator as part of provisioning)
  - Editable (the configuration parameter is editable as part of provisioning)
  - Device-Specific (the configuration parameter value needs to be entered by the operator for each device during deployment)
- h. In the ADC Configuration Parameters box, do the following:

- Select the **Failed Server Loyalty** check box to enable failed server protection. If any server in a server group fails, the remaining servers continue to provide access to vital applications and data. The failed server can be brought back up without interrupting access to services.
  - Select the **Clear on Tcp Reset** check box to clear the adaptive load-balancing mechanism when a Reset flag is received in a TCP packet.
- i. Click **Save** to save the service template configuration. Else, click **Close** to discard the changes to the template.
  - j. Click **Save & Publish** to save and publish the service template configuration. The designer must publish the service templates to the operator to use in the creation of deployment plans. After a filter or policy is published, it goes for peer review and approval. After approval, the filter or policy is deployed to device.

## Importing an ADC Service Template

To create a clone of an existing ADC template by importing it:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the task pane, select **Service Templates**. The Manage Service Templates page is displayed.
4. Click the **ADC** button. The list of ADC service templates is displayed. You need not click this button if you are launching the Service Designer page for the first time or are navigating to this page from another mode or a different page. You need to click this button only if you are viewing the other service templates, such as CGNAT or TLB.
5. Click the **Add** icon. The Create an ADC Planning Template window appears.
6. Enter the name of the template and the service instance in the respective fields.
7. Click the **Import** button. The Import Services dialog box appears.

You can import the service templates assigned to SDGs or choose from a list of all of the predefined templates in the database. Also, you can either import all of the components of a service or specific components.

8. Perform one of the following for the Import section:
  - Select the **From Existing Service Gateway** radio button if you want to import the CGNAT rule from SDGs that are present in the Edge Services Director database.

- Select the **From XML** radio button if you want to import the CGNAT rule from an XML configuration file on an external system.
9. If you selected the option to import the object from SDGs, do the following:
- Click the **Normal View** tab to view the list of SDGs. You can search for specific SDGs by entering a search item and clicking the **Search** icon.  
  
Alternatively, click the **Group View** tab to view the list of SDG groups. You can search for specific SDG groups by entering a search item and clicking the **Search** icon.
  - Click the plus sign (+) next to the All Service Gateways item to expand the tree structure that displays the list of SDGs or SDG groups. If the SDG pair is configured, you can select one of the devices, master or standby, from which you want to import the object.  
  
Alternatively, if you selected the **Group View** tab, you can select an SDG from the groups displayed from which you want to import the object.
  - Click **Import**. The object is added to the database and can be used during configuration of services or policies.
10. If you selected the option to import from an XML file, do the following:
- Click **Browse** beside the File Name field to navigate to the path where an XML file is available to be imported.
  - Click **Upload**. The service template is added to the database and can be used during configuration of services or policies.
11. Do one of the following to import all components of a selected template or only a particular component of a template. For the components that are not imported, you need to specify the definitions of the components afresh.
- Select the check boxes next to all of the service instances that are displayed for the selected SDG or SDG group, or for the XML file that you uploaded. In such a case, all of the elements or parameters of the selected template or instance are imported.
  - Alternatively, select the check box next to a particular or group of service instances to import only a specific component of the selected template

For example, if the service instance you are importing contains **Routing Interface Details** from the list of individual service components being retrieved to the service template you are creating, you can import the client-facing and server-facing interface and routing instances. The interface and routing instance where client packets are received from the list of all the items that belong to the devices in the inventory form the client-facing set. The interface and routing instance through which packets traverse to servers from the list of all the items that belong to the devices in the inventory form the server-facing set.



**NOTE:** Client-facing interfaces—The device interfaces where client traffic is received. Traffic arriving on these interfaces is handled by the ADC software and destined to be routed to the virtual IP addresses and filter destination addresses configured in the instance. At least one client-facing interface must be specified for each adc-instance. A client-facing interface can be shared between instances.

Server-facing interfaces—The device interfaces where servers are connected, usually through switches or routers. Traffic to the servers is routed to these interfaces. At least one server-facing interface must be specified for each load-balancing instance; a server-facing interface can be shared between instances. The same device interface can be used as a client-facing interface in one (or more) adcinstances, and as a server-facing interface in other instances.

12. Similarly, you can select other components and import them to the template. Save the imported components to add them to the template you are creating by using the imported template as a base.

## Creating a Deployment Plan

You must have previously defined service templates and policy or filter templates before you can create a deployment plan.

To create a deployment plan and assigning devices to it:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the task pane, select **Service Edit**. The Manage Service Templates page is displayed.
4. Click the **ADC** button. The list of ADC service templates is displayed. You need not click this button if you are launching the Service Designer page for the first time or are navigating to this page from another mode or a different page. You need to click this button only if you are viewing the other service templates, such as CGNAT or TLB.
5. Select the check boxes next to the SDGs or SDG groups that you want to assign to the plan. Based on your selection of a service or a policy template, the components or attributes are shown for the corresponding device.

6. From the boxes that show the components of a service template, you can edit, delete, or add elements to it. If you do not have permissions to update a template, the corresponding icons are not shown.
7. Click the down arrow in the **Actions** menu and select **Send for Deployment** to create a deployment plan for the particular service template and save the plan.
  - If you create a deployment plan from Gateway view of Deploy mode, the Deployment Plan Summary dialog box appears, with the service name, type, and status listed.  
Click **Send** to create a deployment plan.
  - If you create a deployment plan from Service view of Deploy mode, the Edit Service Instance page is displayed. You can modify the SDGs associated with the service instance and also modify the service instance attributes as necessary by either clicking the buttons corresponding to the various settings at the top of the wizard page to directly traverse to the page you want to modify or clicking the navigation buttons at the bottom of the wizard page to go to the different pages of the wizard. Click **Finish** to create a deployment plan.

A deploy plan is created for the service template with the devices that are assigned to it when you view the Deployment Plans page.
8. Alternatively, you can select **Discard changes** from the **Actions** menu to ignore the modifications done to a policy or filter template.
9. From the Deployment plans page, you can select **Reject** or **Approve** from the Actions drop-down list to reject or approve the deployment plan and make it available for commissioning to the devices.

## Creating a Real Server

To create a real server as a component for the ADC template:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the task pane, select **Service Templates**. The Manage Service Templates page is displayed.
4. Click the **ADC** button. The list of ADC service templates is displayed.
5. Click the **Add** icon. The Create an ADC Planning Template window appears.

6. Enter the name of the template and the service instance in the respective fields.
7. Click the green plus sign in the Real Servers box. The Addition of Real Server dialog box appears.



**NOTE:** For the service elements that you can configure using the Object Builder workspace, such as applications and rules, when you click the green plus sign (+) at the top-right corner of each of the service element boxes, the shortcut menu is displayed. Click the **Create New** radio button to create the service component afresh. Alternatively, click the **Import from Object Builder** radio button to open a dialog box that enables you to select from the list of service elements that are present in the database of Edge Services Director and import them into the service template.

If a green tick mark is shown beside a field in the dialog box, it denotes that you can add attributes for that component. A red cross mark shows that you can delete that particular attribute for that component.

8. In the Name field, enter the name to identify the real server. Make sure the servers are connected via a router interface that is defined as a server-facing interface for the adc-instance. For each real server, you must assign a real-server name and specify its actual IP address.
9. In the Address Family field, select **IPv4** to specify an IPv4 address, or select **IPv6** to enter the IPv6 address of the real server.
10. In the IP Address field, specify the IP address of the real server.
11. In the Health check section, select the check box and specify the following:
  - In the Interval field, specify the amount of time, in seconds, between polls of the real server by the router.



**NOTE:** The ADC software monitors the servers in the real-server group and the load-balanced applications running on them. If a router detects that a server or application has failed, it does not direct any new connection requests to that server. When a service fails, the ADC software can remove the individual service from the load-balancing algorithm without affecting other services provided by that server. By default, the router checks the status of each service on each real server every five (5) seconds. Sometimes, the real server can be too busy processing connections to respond to health checks. If a service does not respond to four consecutive health checks, the router, by default, declares the service unavailable. You can modify both the health check interval and the number of retries.

- In the Failure-retries field, specify the number of times the router attempts its check on the real server before marking the server as unavailable. In the Recovery-retries field, specify the number of times the router attempts to recover the real-server connection.
  - In the Recovery Retries field, set the number of recovery retries to attempt to determine server recovery. The range is from 1 through 63.
12. In the Listing Ports section, click the plus sign to add as many ports as needed for the real server. Enter the port number in the Port field. For example, you might require ports for the common application ports and the applications that use them, such as 8080 for HTTP and 443 for HTTPS.
  13. In the Content String section, click the plus sign to add as many content strings as needed to be added for the real server. Enter the string for matching traffic to be sent to the real server in the String field. ADC software supports two content-string methods (URL hashing and URL pattern matching) and all Layer 4 load-balancing methods. If you do not add a defined string (or add the defined string any), the server handles any request. Content string handling applies to the DNS, RTSP, HTTP services, and to filters.

You can assign one or more content strings to each real server. When more than one URL string is assigned to a real server, requests matching any string are redirected to that real server. There is also a special string known as "any" that matches all content.
  14. Click **Save** to save the service template configuration. Else, click **Close** to discard the changes to the template.

## Creating a Group for Real Servers

Define the group and assign real servers to it. The real servers in any given group must have an IP address accessible to the module that performs the SLB functions. This IP routing is most easily accomplished by placing the servers on a network local to the router. Routing to the server can be used as long as it does not violate the topology rules outlined.

A group is a collection of multiple servers with the same content, so that client requests can be load-balanced between them.

To create a group of real servers:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.

The functionalities that you can configure in this mode are displayed in the task pane.
3. From the task pane, select **Service Templates**.

The Manage Service Templates page is displayed.

4. Click the **ADC** button.

The list of ADC service templates is displayed.

5. Click the **Add** icon.

The Create an ADC Planning Template window appears.

6. In the Template Name field, enter a name for the service template or profile (limit of 63 alphanumeric characters without spaces).
7. In the Instance Name field, enter a meaningful, easily-identifiable name for the service instance (limit of 255 characters). Each service instance you define can be applied to a single or multiple SDGs.
8. Click the green plus sign in the Server Groups box. The Addition of Group dialog box appears.



**NOTE:** For the service elements that you can configure using the Object Builder workspace, such as applications and rules, when you click the green plus sign (+) at the top-right corner of each of the service element boxes, the shortcut menu is displayed. Click the **Create New** radio button to create the service component afresh. Alternatively, click the **Import from Object Builder** radio button to open a dialog box that enables you to select from the list of service elements that are present in the database of Edge Services Director and import them into the service template.

If a green tick mark is shown beside a field in the dialog box, it denotes that you can add attributes for that component. A red cross mark shows that you can delete that particular attribute for that component.

9. In the Name field, enter the name for the real servers group.
10. In the Group Unit field, specify the unit on a group. In general, the unit is used when the traffic is going out from the ADC software to the server. To support virtual routers on the server side, each server is assigned a unit. When the traffic is going out from the ADC software to this server, the traffic goes out from the matching Multiservices-DPC NPU IFL (ms-x/y/z.#, where # is the unit). This allows you to attach the relevant IFL to a virtual router and attach the server to this virtual router. If the unit is not configured on the server, the unit is taken from the group configuration. If the unit is not configured in the group, the unit is taken from the adc-instance configuration. If no unit is configured, the ADC software uses the default unit (unit 0).

For example, if you specify the unit as 40, it sets all servers inside this group to use unit 40, unless a unit is configured on a specific server inside the group.

11. From the Load Balance Method list, select the method of load balancing for the real servers group. Load-balancing methods are used for selecting which real-server in a group receives the next client connection. The available metrics include hash, least connections, round-robin, response (response time), and bandwidth.
12. In the Real Servers section, assign the real servers to be part of the group. Select the real servers from the Available column and click the right arrow to move the server to the Selected column.

13. In the Health Check section, do one of the following:

The ADC software monitors the servers in the real-server group and the load-balanced applications running on them. If a router detects that a server or application has failed, it does not direct any new connection requests to that server. When a service fails, the ADC software can remove the individual service from the load-balancing algorithm without affecting other services provided by that server. By default, the router checks the status of each service on each real server every five (5) seconds. Sometimes, the real server can be too busy processing connections to respond to health checks. If a service does not respond to four consecutive health checks, the router, by default, declares the service unavailable. You can modify both the health check interval and the number of retries.

- Select the **DNS** radio button to configure DNS health checking. Enter the hostname for which health verification needs to be performed..
  - Select the **HTTP** radio button to configure HTTP-based health check. HTTP-based health checks can include the hostname for Host headers. The Host header and health check URL are constructed from the Virtual server hostname, domain name, and the server group health check field. Enter the URL for which health check is needed and the HTTP header method, such as GET, PUT, POST, DELETE , and. PATCH. Select the **Use Head Method** that causes the HTTP Head method to retrieve HTTP headers only.
  - Select the **PING** radio button to configure ping-based health checking. Ping health checks verify if the real server is alive.
  - Select the **SSLHELLO** radio button to sets Secure Sockets Layer (SSL) hello health-check parameters. SSL version 2 (SSLv2) is used for the SSL health check
  - Select the **SCRIPT** radio button to create a custom-based health check. From the Custom Health Check field, specify **tcp** or **udp** as the protocol for the script to use in a custom health check. A script is made up of one or more TCP or UDP command containers. A script can contain any number of these containers, up to the allowable number of characters that a script supports.
14. Click **Save** to save the service template configuration. Else, click **Close** to discard the changes to the template.

## Load-Balancing Methods for Real-Server Groups

The following methods for real server groups are supported:

- **Hash**—The hash load-balancing method uses IP address information in the client request to select a server. For virtual-services, the client source IP address is used. All requests from a specific client are sent to the same server. This is useful for applications where client information must be retained between sessions. When selecting a server, a mathematical hash of the relevant IP address information is used as an index into the list of currently available servers. Any given IP address information always has the same hash result, providing natural persistence, as long as the server list is stable. When a configured server becomes unavailable, clients bound to operational servers continue to be bound to the same servers for future sessions and clients bound to unavailable servers are rehashed to select an operational server. Some services allow you to hash using the client-ip and port. This is done using the source-port-inhash parameter. There are more hash options in filters, that are set using the load-balancing-hash parameter.
- **Least Connections**—With the least-connections load-balancing method, the number of connections currently open on each real server is measured in real time. The server with the fewest current connections is considered to be the best choice for the next client connection request. This option is the most self-regulating, with the fastest servers typically getting the most connections over time.
- **Round-Robin**—With the round-robin load-balancing method, new connections are issued to each server in turn; that is, the first real server in the group gets the first connection, the second real server gets the next connection, followed by the third real server, and so on. When all the real servers in this group have received at least one connection, the issuing process starts over with the first real server.
- **Response Time**—The response-time load-balancing method uses real-server response time to assign sessions to servers. The response time between the servers and the load-balancing module is used as the weighting factor. The router monitors and records the amount of time it takes for each real server to reply to a health check to adjust the real-server weights. The weights are adjusted so they are inversely proportional to a moving average of response time. In such a scenario, a server with half the response time as another server receives a weight twice as large. Note: The effects of the response-time or bandwidth weighting apply directly to the real servers and are not necessarily confined to the group. When response-time or bandwidth-metered real servers are also used in other groups that use the least connections, round-robin, or hash methods, the response-time or bandwidth weights are applied on top of the method calculations for the affected real servers. Since the response-time or bandwidth weight changes dynamically, this can produce fluctuations in traffic distribution for the groups that use the least-connections, round-robin, or hash load-balancing methods.
- **Bandwidth** The bandwidth load-balancing method uses real-server octet counts to assign sessions to a server. The load-balancing module monitors the number of octets sent between the server and the module. Then, the real-server weights are adjusted so they are inversely proportional to the number of octets that the real server processes during the last interval. Servers that process more octets are considered to have less available bandwidth than servers that have processed fewer octets. For example, the server that processes half the amount of octets over the last interval receives twice the weight of the other servers. The higher the bandwidth used, the smaller the weight assigned to the server. Based on this weighting, the subsequent requests go to the

server with the highest amount of free bandwidth. These weights are automatically assigned.



**NOTE:** The effects of the response-time or bandwidth weighting apply directly to the real servers and are not necessarily confined to the group. When response-time or bandwidth-metered real servers are also used in other groups that use the leastconnections, round-robin, or hash methods, the response-time or bandwidth weights are applied on top of the method calculations for the affected real servers. Since the response-time or bandwidth weight changes dynamically, this can produce fluctuations in traffic distribution for the groups that use the least-connections, round-robin, or hash load-balancing methods.

## Creating a Client-Facing Interface and Routing Instance

Clients and servers can be connected through the same router port. Each port in use on the router can be configured to process client requests, server traffic, or both:

Client-facing interfaces—Router ports through which client requests to the virtual server are received.

Server-facing interfaces—Router ports to which servers are connected (directly or through routing). Responses to clients are received on the router through these ports.

To assign a client-facing instance and interface to an ADC template:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the task pane, select **Service Templates**.  
The Manage Service Templates page is displayed.
4. Click the **ADC** button.  
The list of ADC service templates is displayed.
5. Click the **Add** icon.  
The Create an ADC Planning Template window appears.
6. In the Template Name field, enter a name for the service template or profile (limit of 63 alphanumeric characters without spaces).

7. In the Instance Name field, enter a meaningful, easily-identifiable name for the service instance (limit of 255 characters). Each service instance you define can be applied to a single or multiple SDGs.
8. Click the green plus sign in the Client-Facing box. The Client facing dialog box appears.



**NOTE:** For the service elements that you can configure using the Object Builder workspace, such as applications and rules, when you click the green plus sign (+) at the top-right corner of each of the service element boxes, the shortcut menu is displayed. Click the **Create New** radio button to create the service component afresh. Alternatively, click the **Import from Object Builder** radio button to open a dialog box that enables you to select from the list of service elements that are present in the database of Edge Services Director and import them into the service template.

If a green tick mark is shown beside a field in the dialog box, it denotes that you can add attributes for that component. A red cross mark shows that you can delete that particular attribute for that component.

9. From the Service Gateway Name field, select the SDG group with which the service element must be associated.
10. From the Host Name field, select the SDG in the SDG high-availability pair of active and standby SDGs.
11. In the Device Inventory Routing Instances section, select the check box next to the routing instance of the SDG that must be used for packets arriving from clients or users. All the routing instances from the inventory of devices are listed.
12. In the Device Inventory Interfaces section, select the check box next to the interface instance of the SDG that must be used for packets arriving from clients or users. All of the interfaces from the inventory of devices are listed.
13. Click **OK** to save the settings. Else, click **Cancel** to discard the configuration.

## Creating a Server-Facing Interface and Routing Instance

Clients and servers can be connected through the same router port. Each port in use on the router can be configured to process client requests, server traffic, or both:

Client-facing interfaces—Router ports through which client requests to the virtual server are received.

Server-facing interfaces—Router ports to which servers are connected (directly or through routing). Responses to clients are received on the router through these ports.

To assign a server-facing instance and interface to an ADC template:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the task pane, select **Service Templates**.  
The Manage Service Templates page is displayed.
4. Click the **ADC** button.  
The list of ADC service templates is displayed.
5. Click the **Add** icon.  
The Create an ADC Planning Template window appears.
6. In the Template Name field, enter a name for the service template or profile (limit of 63 alphanumeric characters without spaces).
7. In the Instance Name field, enter a meaningful, easily-identifiable name for the service instance (limit of 255 characters). Each service instance you define can be applied to a single or multiple SDGs.
8. Click the green plus sign in the Client-Facing box. The Client facing dialog box appears.



**NOTE:** For the service elements that you can configure using the Object Builder workspace, such as applications and rules, when you click the green plus sign (+) at the top-right corner of each of the service element boxes, the shortcut menu is displayed. Click the **Create New** radio button to create the service component afresh. Alternatively, click the **Import from Object Builder** radio button to open a dialog box that enables you to select from the list of service elements that are present in the database of Edge Services Director and import them into the service template.

If a green tick mark is shown beside a field in the dialog box, it denotes that you can add attributes for that component. A red cross mark shows that you can delete that particular attribute for that component.

9. From the Service Gateway Name field, select the SDG group with which the service element must be associated.
10. From the Host Name field, select the SDG in the SDG high-availability pair of active and standby SDGs.
11. In the Device Inventory Routing Instances section, select the check box next to the routing instance of the SDG that must be used for packets traversing to the servers. All the routing instances from the inventory of devices are listed.
12. In the Device Inventory Interfaces section, select the check box next to the interface instance of the SDG that must be used for packets to be sent to the servers. All of the interfaces from the inventory of devices are listed.
13. Click **OK** to save the settings. Else, click **Cancel** to discard the configuration.

## Creating a Services PIC for an ADC Service Template

Multiservices (ms-) interfaces are the physical multiservices interfaces of a device that are used to run the load-balancing instance application. The more multiservices interfaces used for a loadbalancing instance, the more capacity and processing power the instance has. At least one MS interface must be specified for each adc-instance, up to eight interfaces can run the same instance. A multiservices interface is associated exclusively to a single load-balancing instance (it cannot be shared between instances).

To assign a services interface to an ADC template:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.

The functionalities that you can configure in this mode are displayed in the task pane.

3. From the task pane, select **Service Templates**.

The Manage Service Templates page is displayed.

4. Click the **ADC** button.

The list of ADC service templates is displayed.

5. Click the **Add** icon.

The Create an ADC Planning Template window appears.

6. Enter the name of the template and the service instance in the respective fields.

7. Click the green plus sign in the Service Pic box. The Service Pic dialog box appears.



**NOTE:** For the service elements that you can configure using the Object Builder workspace, such as applications and rules, when you click the green plus sign (+) at the top-right corner of each of the service element boxes, the shortcut menu is displayed. Click the **Create New** radio button to create the service component afresh. Alternatively, click the **Import from Object Builder** radio button to open a dialog box that enables you to select from the list of service elements that are present in the database of Edge Services Director and import them into the service template.

If a green tick mark is shown beside a field in the dialog box, it denotes that you can add attributes for that component. A red cross mark shows that you can delete that particular attribute for that component.

8. From the Service Gateway Name field, select the SDG group with which the service element must be associated.
9. From the Host Name field, select the SDG in the SDG high-availability pair of active and standby SDGs.
10. Select the check box next to the ms- interface of an SDG that must be assigned to the ADC template.
11. Click **OK** to save the settings. Else, click **Cancel** to discard the configuration.

## Creating a Health Check for an ADC Service Template

The ADC software does health checking on each defined server (see Health Checking, page 183). In order for the traffic to get from the ADC software to the server, a source IP with the same subunit as the server must be defined. Usually all subunits that are in use in a certain adc-instance must have a matching IP address with the same subunit defined in the instance.

The health check itself is defined at the group parameter. Select a health check based on the application running on the real server in question. If the real server is an LDAP server, for example, use the LDAP health check method. It is important to make sure that the server can answer connections from the IP address configured. This source IP address must be “routable” back to the router. Each server in the load-balancing instance has a sub-unit attached to it. Before the ADC software sends a health check to a server, it checks the sub-unit attached to the server, then chooses the source IP address to use for this server health check according to the address configured under the same unit in the health-check-source configuration. As a result, each sub-unit attached to a server must have a matching address in the healthcheck- source configuration. This way the ADC software can send health checks to servers using this sub-unit. When no health check address is defined for the unit, all servers with this unit are in a failed status. Family inet is the only supported family under the health-check-source configuration.

To configure a health check source for an ADC template:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the task pane, select **Service Templates**.  
The Manage Service Templates page is displayed.
4. Click the **ADC** button.  
The list of ADC service templates is displayed.
5. Click the **Add** icon.  
The Create an ADC Planning Template window appears.
6. In the Template Name field, enter a name for the service template or profile (limit of 63 alphanumeric characters without spaces).
7. In the Instance Name field, enter a meaningful, easily-identifiable name for the service instance (limit of 255 characters). Each service instance you define can be applied to a single or multiple SDGs.

8. Click the green plus sign in the Health Check box. The Addition of Health Check dialog box appears.



**NOTE:** For the service elements that you can configure using the Object Builder workspace, such as applications and rules, when you click the green plus sign (+) at the top-right corner of each of the service element boxes, the shortcut menu is displayed. Click the **Create New** radio button to create the service component afresh. Alternatively, click the **Import from Object Builder** radio button to open a dialog box that enables you to select from the list of service elements that are present in the database of Edge Services Director and import them into the service template.

If a green tick mark is shown beside a field in the dialog box, it denotes that you can add attributes for that component. A red cross mark shows that you can delete that particular attribute for that component.

9. Specify the unit of the health check source in the Unit field. as part of the auto-configuration, the ADC software defines IFLs and IFAs (units and addresses) on the Multiservices-DPC. These IFLs require a unique unit number that is used later in auto-configured filters to direct traffic. By default, the units used by the ADC software for automatic configuration are in the range of 10,000 to 11,032.
10. Select the **IPv4 Family** check box to specify IPv4 as the address protocol family.
11. Specify the IPv4 address of the source for health verification in the IP Address field.
12. Select the **IPv6 Family** check box to specify IPv6 as the address protocol family.
13. Specify the IPv6 address prefix of the source for health verification in the IP Address field.
14. Click **Save** to save the settings. Else, click **Cancel** to discard the configuration.

## Creating a Custom Health Check for an ADC Instance

You can configure the ADC software to send a series of health-check requests to real servers or real-server groups and monitor the responses. Health checks are supported for TCP and UDP protocols, using either binary or ASCII content.

Health check scripts dynamically verify application and content availability by executing a sequence of tests based on send and expect commands. You can configure the ADC software to send a series of health check requests to real servers or realserver groups and monitor the responses. Both ASCII and binary-based scripts, for TCP and UDP protocols, can be used to verify application and content availability.

To configure a custom health-check script for an ADC template:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the task pane, select **Service Templates**.  
The Manage Service Templates page is displayed.
4. Click the **ADC** button.  
The list of ADC service templates is displayed.
5. Click the **Add** icon.  
The Create an ADC Planning Template window appears.
6. In the Template Name field, enter a name for the service template or profile (limit of 63 alphanumeric characters without spaces).
7. In the Instance Name field, enter a meaningful, easily-identifiable name for the service instance (limit of 255 characters). Each service instance you define can be applied to a single or multiple SDGs.
8. Click the green plus sign in the Custom Health Check box. The Addition of Custom Health Check dialog box appears.



**NOTE:** For the service elements that you can configure using the Object Builder workspace, such as applications and rules, when you click the green plus sign (+) at the top-right corner of each of the service element boxes, the shortcut menu is displayed. Click the Create New radio button to create the service component afresh. Alternatively, click the Import from Object Builder radio button to open a dialog box that enables you to select from the list of service elements that are present in the database of Edge Services Director and import them into the service template.

If a green tick mark is shown beside a field in the dialog box, it denotes that you can add attributes for that component. A red cross mark shows that you can delete that particular attribute for that component.

9. Specify the name of the script to be used for health-check in the Script Name field.  
A script is made up of one or more TCP or UDP command containers. A script can

contain any number of these containers, up to the allowable number of characters that a script supports.

10. Select the type of protocol for custom health-check from the **Command Type** list. You can select either **TCP** or **UDP**. Commands exist to open a connection to a specific TCP or UDP port, send a request to the server, and expect an ASCII string or binary pattern. Only one protocol can be configured per script.

11. Specify the name of the command for custom health-check in the Command Name field.

The name of the TCP or UDP command for script-based health-check is a container for one or more commands.

12. Click the **Add** icon to create a health-check command. The Health Check Command dialog box is displayed.

You can also select the check boxes beside existing commands from the list of previously configured commands from the Custom Health Check dialog box if you want to assign them to the health-check script. Click **Save** to save the settings.

13. In the Health Check Command dialog box, enter the unique identifier for the command to be used for diagnosing and monitoring the health of servers or URLs using script-based checking in the Command ID field.

14. Select the type of command for script-based health monitoring from the **Command Type** list.

The following are the currently available commands for building a script-based health check:

- **open**—Specifies which destination real-server UDP port to use; for example, OPEN 9201. After entering the destination port, you is prompted to specify a protocol; choose **udp**.
- **send**—Specifies the send content in raw hexadecimal format.
- **binary-send (for binary content only)**—Used to specify binary content (in hexadecimal format) for the request packet.
- **expect**—Specify the expected content in raw hexadecimal format.
- **binary-expect (for binary content only)**—Used to specify the binary content (in hex format) to be expected from the server response packet.
- **offset (for binary content only)**—Specifies the offset from the beginning of the binary data area to start matching the content specified in the binary-expect command. The offset command is supported for both UDP and TCP-based health checks. Specify the offset command after a binary-expect command if an offset is desired. If this command is not present, an offset of zero is assumed.

- **depth** (for binary content only)—Specifies the number of bytes in the IP packet that should be examined. If no offset value is specified, depth is specified from the beginning of the packet. When depth is not specified, it is the length of the content. This means that the content is expected exactly at the offset specified (or 0 when the offset is not specified).
  - **wait**—Specifies a wait interval before the expected response is returned. The wait window begins when the send string is sent from the ADC. If the expected response is received within the window, the wait step passes. Otherwise, the health check fails. The wait window is in units of milliseconds. When the wait value is not specified the script waits according to the realserver configured interval.
15. Enter a value corresponding to the command type selected in the Value field. You can enter one of the following types of values based on the command type:
    - **binary-expect and binary-send hexadecimal-value**—Specifies the content to expect from the server response packet using hexadecimal format.
    - **depth number**—Specifies the number of bytes in the IP packet that should be examined. If no offset value is specified, depth is specified from the beginning of the packet. Default: The default value is the length of the content.
    - **offset number**—Specifies the offset from the beginning of the binary data area to start matching the content specified in the binary-expect command. The offset command is supported for both UDP-based and TCP-based health checks. If you require an offset, specify the offset command after a binary-expect command. Default: 0
    - **binary-expect, binary-send, and expect wait interval**—Specifies a wait interval before the expected response is returned. The wait interval begins when the send string is sent from the ADC software. If the expected response is received within the interval, the wait step passes. Otherwise, the health check fails. The wait interval is expressed in units of milliseconds. When the wait interval is not specified, the script waits according to the real server configured interval. Range: 0 through 65535
    - **send text**—Specifies the send content in raw hexadecimal format.
    - **open port**—Specifies which destination real-server UDP port to use; for example, open 9201.
  16. Click **Save** to save your settings in the Health Check Command dialog box. You are returned to the Custom Health Check dialog box and the newly configured command is added to the list shown.
  17. Click **OK** to save the settings in the Custom Health Check dialog box. Else, click **Cancel** to discard the configuration.

## Creating a Virtual Service for an ADC Service Template

A virtual service is a service that is being load-balanced across the servers in the group; for example, dns-virtual-service. The service belongs to a virtual server, that defines the IP address through which the service is accessible to the client. The service is accessed through one or more predefined application ports (TCP or UDP). The virtual server defines

the IP address to which client requests are sent. The virtual service defines a destination port within the virtual-server IP address. The virtual service configuration includes parameters relevant to the processing of client requests to this service. The service is actually provided by the real servers in the group defined in the virtual service.

To configure a virtual service for an ADC template:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the task pane, select **Service Templates**.  
The Manage Service Templates page is displayed.
4. Click the **ADC** button.  
The list of ADC service templates is displayed.
5. Click the **Add** icon.  
The Create an ADC Planning Template window appears.
6. In the Template Name field, enter a name for the service template or profile (limit of 63 alphanumeric characters without spaces).
7. In the Instance Name field, enter a meaningful, easily-identifiable name for the service instance (limit of 255 characters). Each service instance you define can be applied to a single or multiple SDGs.
8. Click the green plus sign in the Virtual Service box. The Addition of Virtual Service dialog box appears.



**NOTE:** If a green tick mark is shown beside a field in the dialog box, it denotes that you can add attributes for that component. A red cross mark shows that you can delete that particular attribute for that component.

9. In the Name field, specify the name of the virtual service (limit of 128 characters).
10. In the Address field, specify the IP address of the virtual server.
11. From the Service Type list, select **DNS** to set up the DNS service for the virtual server. You can also select other service types such as plain, HTTP, or SSL.

IP server load balancing allows you to configure your ADC software for server load balancing based on the client's IP address only. Typically, the client IP address is used with the client port number to produce a session identifier. When the Layer 3 option is enabled, the ADC software uses only the client IP address as the session identifier.

12. In the Server Listening Port field, specify the port number the server uses to listen or receive connection requests. The range is from 0 through 65,534. You can change the destination port of traffic to a specific port by using this field setting.
13. From the Protocol list, select **TCP** or **UDP** to specify the application type of virtual service.
14. From the Group list, select the name of a real server group configured to be used for this virtual service.
15. In the Service Timeout field, configure the service-timeout parameter to the amount of time that idle connections should remain in the connection table before being removed, in minutes (0 to 32768). The default, when the parameter is not set, is to use the timeout configured for the real server, typically 10 minutes.
16. Select the **Fast Load Balancing** check box to specify the connection table needs to be used for requests only.

Traffic to virtual services is managed using the connection table. Each connection is recorded in the table. Usually, the connection table is used both for the request processing and for reply processing. In request processing, the ADC software looks for a corresponding entry to check persistency information, finds the appropriate real-server address and listening port, and uses it to send the request to the server. In reply processing, the ADC software looks for a corresponding entry to know how to change the source address from a real-server address and listening port back to the virtualserver address and service port. In some cases, faster traffic processing can be achieved by not checking the connection table for the response path, but by using another, more efficient, mechanism for the address and port translation.

17. Select the **Send Traffic to VIP** check box to redirect the packets to the virtual IP address configured for the virtual server associated with the virtual service. When a certain VIP is available, the route to this VIP exists in the routing-instance. This allows the dynamic protocol to publish the VIP as owned by the router. When the virtual IP address is not available (i.e., all the servers for this VIP are down), the route is redrawn using the routing-instance. This causes the routing protocol to redraw the route to this IP from its publications. In turn, traffic to this VIP is no longer be routed to this specific router.
18. Click **Save** to save the settings. Else, click **Cancel** to discard the configuration.

## Creating a Virtual Server for an ADC Service Template

Each virtual server can be configured to support up to 8 service ports and is limited to a total of 1023 services per router. If more than eight service ports are required for a virtual address, you can define multiple virtual servers with the same address. The protocol setting specifies whether this virtual service is a TCP or UDP application. The port setting specifies the application port for this application.

To configure a virtual server for an ADC template:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the task pane, select **Service Templates**.  
The Manage Service Templates page is displayed.
4. Click the **ADC** button.  
The list of ADC service templates is displayed.
5. Click the **Add** icon.  
The Create an ADC Planning Template window appears.
6. In the Template Name field, enter a name for the service template or profile (limit of 63 alphanumeric characters without spaces).
7. In the Instance Name field, enter a meaningful, easily-identifiable name for the service instance (limit of 255 characters). Each service instance you define can be applied to a single or multiple SDGs.
8. Click the green plus sign in the Virtual Server box. The Addition of Virtual Server dialog box appears.



**NOTE:** For the service elements that you can configure using the Object Builder workspace, such as applications and rules, when you click the green plus sign (+) at the top-right corner of each of the service element boxes, the shortcut menu is displayed. Click the **Create New** radio button to create the service component afresh. Alternatively, click the **Import from Object Builder** radio button to open a dialog box that enables you to select from the list of service elements that are present in the database of Edge Services Director and import them into the service template.

If a green tick mark is shown beside a field in the dialog box, it denotes that you can add attributes for that component. A red cross mark shows that you can delete that particular attribute for that component.

9. In the **Name** field, enter the name of the virtual server. The virtual server defines the IP address to which client requests are sent.
10. In the **Address** field, specify the IP address of the virtual server.
11. From the **Type** list, select **DNS** to set up the DNS service for the virtual server. You can also select other service types such as LDAP, HTTP, or SNMP.
12. In the **Virtual Services** section, select a virtual service from the **Available** column and click the right arrow to move the service to the **Selected** column,
13. Click **Save** to save the settings. Else, click **Cancel** to discard the configuration.

## Creating a Firewall Rule for an ADC Service Template

ADC filter terms are an ordered list of terms. Each filter term is composed from a match clause (ADC Filter Terms—"from" Clause) that defines the match criteria, and a then clause (ADC Filter Terms—"then" Clause) that defines the action and behavior with traffic that matches the term. An ADC filter term name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose the entire name in quotation marks (" "). Each term name must be unique within a filter. You can specify multiple terms in the ADC filter, effectively chaining together a series of match action operations to apply to the packets. You can also use the go-to action so that, when a match condition is met, the evaluation continues from the go-to term, rather than terminating. ADC filter terms are evaluated in the order in which you specify them in the configuration. To reorder terms, use the configuration mode insert command. For example, the command `insert term up before term start` places the term up before the term start. Up to 2048 filter terms can be configured on the module. Descriptive names can be used to define filter terms. Each filter can be set to perform from or then actions, based on any combination of the filter options.

### ADC Filter Terms—"from" Clause

In the from statement in the ADC filter term, you specify conditions that the packet must match for the action in the then statement to be taken. All conditions in the from statement must match for the action to be taken. The order in which you specify match conditions is not important, because a packet must match all the conditions in a term for a match to occur. If you specify no match conditions in a term, that term matches all packets. In the from clause you can indicate Layer 4 information to match traffic:

- source-address—Source IP address or range.
- destination-address—Destination IP address or range (dip and dmask).
- protocol tcp | udp—Match using either TCP or UDP protocol. By default, both are matched.
- source-port—TCP/UDP application or source port or source port range (such as 31000 to 33000). The service number specified on the module must match the service specified on the server.
- destination-port—TCP/UDP application or destination port or destination port range (such as 31000 to 33000).



**NOTE:** Advanced filtering options such as TCP flags are available. Using these filter criteria, you could create a single filter that blocks external Telnet traffic to your main server except from a trusted IP address. Another filter could warn you if FTP access is attempted from a specific IP address. Another filter could redirect all incoming e-mail traffic to a server where it can be analyzed for spam. The options are nearly endless

---

### ADC Filter Terms—"then" Clause

A filter term then statement instructs the filter what to do once the filtering criteria are matched. These actions are defined in the then clause of the filter term. You can specify one of the following filter actions:

- accept—Allows the frame to pass (by default). It is processed according to its destination: either handled by ADC virtual services or by the router and sent to its destination.
- discard—Discards frames that fit this filter's profile. They are not processed further.
- go-to term—Match to the specified term and continue classification from there. Note: The target term must appear further down the list than the currently evaluated term.
- http-redirect—Allows you to specify a target term name that the filter search should jump to when a match occurs. The http-redirect causes filter processing to jump to a designated filter, effectively skipping over a block of filter terms. Filter searching then continues from the designated filter term. To specify the new filter, use the http-redirect command.
- load-balance—Redirects frames that fit this filter's profile, such as for web cache redirection. In addition, Layer 4 processing must be used.

- **content-term**—Traffic is further matched against content strings, when matched. The content term then clause is effective. When the content-term is not matched there is no further filter term matching.
- **log**—Generates system log messages when the filter term is hit. This option can be used in conjunction with other term actions.
- **per-packet-load-balancing**—To improve efficiency, by default, filter processing is performed only on the first frame in each session. Subsequent frames in the session are assumed to match the same criteria and are automatically treated in the same way as the initial frame. Sessions that match a filter term are logged in the connection table for immediate processing of subsequent frames, rather than a full search to find a matching term. Some types of filtering (such as TCP flag) require each frame in the session to be filtered separately. To set this behavior, set **per-packet-load-balancing** for the relevant filters.

To configure a virtual server for an ADC template:

1. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
2. From the task pane, select **Service Templates**.  
The Manage Service Templates page is displayed.
3. Click the **ADC** button.  
The list of ADC service templates is displayed.
4. Click the **Add** icon.  
The Create an ADC Planning Template window appears.
5. In the Template Name field, enter a name for the service template or profile (limit of 63 alphanumeric characters without spaces).
6. In the Instance Name field, enter a meaningful, easily-identifiable name for the service instance (limit of 255 characters). Each service instance you define can be applied to a single or multiple SDGs.
7. Click the green plus sign in the Firewall Rules box. The Addition of Firewall Rule dialog box appears.



**NOTE:** If a green tick mark is shown beside a field in the dialog box, it denotes that you can add attributes for that component. A red cross mark shows that you can delete that particular attribute for that component.

8. Select the element for the from clause that specifies the match criterion or filter condition.
9. Select the element for the then clause that specifies the action modifier to be performed.
10. Click **Save** to save the settings. Else, click **Cancel** to discard the configuration.

## Modifying ADC Service Templates

On the Service Designer page, you can view the collection of service templates defined for several applications, such as stateful firewall or CGNAT.

To modify service template instances, such as ADC, SFW, CGNAT, or TLB templates:

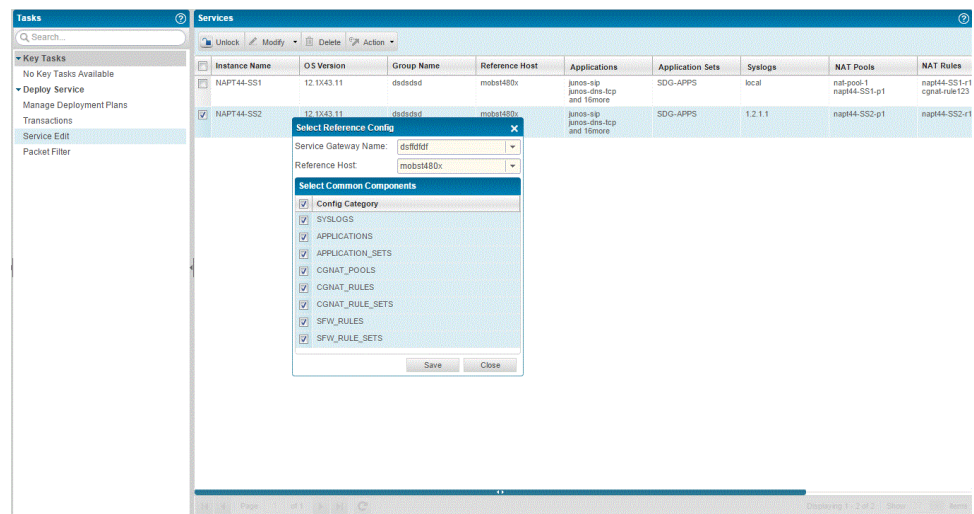
1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the task pane, select **Deploy Service > Service Edit**.  
The Service Instances page is displayed in the right pane, listing all the previously defined service templates.
4. From the View pane, perform one of the following tasks:
  - Click the **ADC** button.  
The list of ADC service templates is displayed. You need not click this button if you are launching the Service Designer page for the first time or are navigating to this page from another mode or a different page. You need to click this button only if you are viewing the other service templates, such as CGNAT or TLB.
  - Click the **SFW** button.  
The list of SFW templates is displayed.
  - Click the **TLB** button.  
The list of TLB templates is displayed.
  - Click the **CGNAT** button.  
The list of CGNAT templates is displayed.
5. In the main window, click the plus sign (+) next to the SDG pairs to expand the tree and view the pair of devices in the SDG group or pair. Select the check box next to the SDG pair or individual SDG for which you want to modify settings. In an SDG pair, you can select a single SDG or both the SDGs in the in the redundancy pair of devices.



**NOTE:** Alternatively, you can also modify service templates from Service View in Build Mode by selecting the Service Templates > Manage Service Templates from the task pane, selecting a service instance, and clicking the Modify button. You can also modify ADC and TLB service templates from Gateway View in Deploy mode by selecting the SDG pair or SDG from the View pane, selecting Service Edit from the task pane, and selecting the TLB service from the main window that displays all the previously configured template instances to lock and modify it.

- Click the **Lock** icon above the table of listed packet filters. The Select Reference Config dialog box is displayed.

Figure 21: Select Reference Config Dialog Box



- From the Service Gateway Name drop-down list, select the SDG group to which the packet filter must be applied.
- From the Host Name drop-down list, select the hostname of the SDG.
- In the Select Common Components section, select the check boxes beside the service modules or components, such as packet filters, SFW rules, or CGNAT rules, that are displayed. The displayed components depend on the attributes that are previously defined for that selected packet filter. For example, if the service policy is for stateful firewall, SFW rules and SFW rule sets are shown. Select the check box beside Config Category to select all the service components.
- Click **Save** to save the modified association.
- Select the check box beside the template you want to modify.

12. Open the **Modify** menu above the list of templates to modify an existing template, and select the component or service attribute, such as application or rule, that you want to edit.
13. Perform one of the following from the drop-down menu displayed for each component:
  - To retrieve the service component and import into the database of Edge Services Director, select **Import Object**. The Import Services dialog box appears. You can import the service templates assigned to SDGs or choose from a list of all of the predefined templates in the database. Also, you can either import all of the components of a service or specific components.
  - To create the component afresh, select **Create New**. The Create page corresponding to the service component appears. You can define the attributes for the service component in the same manner as you define the elements during the creation of a service template.

**Related Documentation**

- [Service Templates Overview on page 189](#)
- [Filtering Service Templates on page 189](#)
- [Viewing Service Templates on page 192](#)
- [Using the Actions Menu on the Service Template and Service Edit Pages on page 195](#)

---

## Creating and Managing CGNAT Service Templates

---

Each carrier-grade NAT rule consists of a set of terms, similar to a service filter. A term consists of the following:

from statement—Specifies the match conditions and applications that are included and excluded. The from statement is optional in NAT rules.

then statement—Specifies the actions and action modifiers to be performed by the router software. The then statement is mandatory in NAT rules.

You can perform the following tasks with the Service Designer page for CGNAT:

- Create a CGNAT service template with attributes and settings for NAT operations.
- Modify an existing CGNAT template to meet the network needs and deployment scenarios.
- Delete an existing template.
- [Creating a CGNAT Service Template on page 235](#)
- [Modifying CGNAT Service Templates on page 238](#)
- [Creating a Deployment Plan on page 240](#)
- [Importing a CGNAT Service Template on page 241](#)
- [Creating a Service Set on page 243](#)

- [Creating a Syslog on page 247](#)
- [Creating a Rule on page 249](#)
- [Creating a Rule Set on page 250](#)
- [Creating a Pool on page 251](#)

## Creating a CGNAT Service Template

To configure a new CGNAT service template:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the task pane, select **Service Templates**.  
The Manage Service Templates page is displayed.
4. Click the **CGNAT** button. The list of CGNAT service templates is displayed.

The Service Designer page displays a bar graph in the top pane of the page. The count of service templates of each type is displayed on the vertical axis and the service type is shown on the horizontal axis. A color-coding format is used to represent the bars on the graph. Published service templates are shown in olive green color and unpublished service templates are shown in blue color. Mouse over each bar in the chart to highlight and display the number of templates published or unpublished for each type of service.

5. Click the **Add** icon. The Select Version dialog box appears.
6. Select **Junos 12.1** if you want to create a template based on the Junos OS Release 12.1. Alternatively, select **Junos 14.1** if you want to create a template based on the Junos OS Release 14.1.



**NOTE:** All the service template components described in this section can be created for templates that are based on both the Junos OS Releases 12.1 and 14.1. The service elements or components that are additionally available for configuration when you select the Junos OS 14.1 version are explicitly mentioned in the relevant steps of the procedure.

The Create a CGNAT Planning Template window appears.

Figure 22: Create CGNAT Service Template Window

7. In the Name field, enter a name for the service template or profile (limit of 63 alphanumeric characters without spaces).
8. In the Description field, enter a meaningful, easily-identifiable name for the service instance (limit of 255 alphanumeric characters). Each service instance you define can be applied to a single or multiple SDGs.
9. Instead of creating a new template entirely, you can import the parameters defined for a previous CGNAT service instance and customize only the settings that are necessary. Imported templates are created without any device assigned to them. To use these templates, you must associate a device with the policy. To clone an existing template by importing it, click the **Import** button.

The Import Services dialog box is displayed. See *Importing a CGNAT Service Template* for step-wise details on importing a CGNAT service template.

10. The Create a CGNAT Planning Template window displays the individual elements or components of the service with a graphical icon for each of the service elements and the corresponding names in separate boxes. You can add, edit, or delete these service elements in a template.

The Property View tab and the Config View tab are displayed on the right pane of the template window. The Property View tab provides a tree-based structure of the parameters defined in a service template. You can expand the tree and view details of each component. A key value pair representation is shown. Each of the components can be treated as categories of the service template shown in the property view.

The Config View tab displays the elements or components specified for a service template in the form of configuration stanzas and hierarchy levels. This display is

similar to the show command that you can use at a certain [edit] hierarchy level to view the defined settings. Each level in the hierarchy is indented to indicate each statement's relative position in the hierarchy. Each level is generally set off with braces, with an open brace ( { ) at the beginning of each hierarchy level and a closing brace ( } ) at the end. If the statement at a hierarchy level is empty, the braces are not displayed. Each leaf statement ends with a semicolon ( ; ), as does the last statement in the hierarchy.

- a. Click the green tick mark (✓) displayed at the top-right corner of each of the service element boxes to create a new element. If the green tick mark is not shown, it indicates that the user role does not have the permission to create an element.
- b. Click the red cross mark (x) displayed at the top-right corner of the icons of each element if you want to delete the existing configuration. The user with designer role has permissions to remove or edit elements.
- c. If the red cross mark is not displayed beside a particular icon, it signifies that the element cannot be deleted.
- d. The diamond icon that contains an orange tick mark within it at the top-right corner of the service component name denotes that the particular element can be modified. The absence of this icon denotes that the user does not have permissions to modify the attributes of the service component.
- e. Double-click each icon pertaining to a service element to view or edit its settings. If you do not possess the permission to modify the element, a view-only dialog box with the attributes of the selected element is shown. Otherwise, an editable dialog box enables you to modify the settings.
- f. Click **Save** to save the service template configuration. Else, click **Close** to discard the changes to the template.
- g. Click the Maximize icon displayed at the top-right corner of the rectangle or box that shows all of the values or entities of a particular component of a service template. The specified component or attribute is displayed as a separate dialog box, listing all of the values of the particular component. You can add, modify, or delete the listed values.
- h. While creating the new service template, the designer can add or modify service parameter values and also restrict the access level for each service parameter for the operator. The designer can set following access levels for each service parameters to operator in planning template. Click the new icon (cascading files icon) displayed at the top-left corner of each of the element boxes to open the shortcut menu. You can click one of the following radio buttons:
  - Read-only (the configuration parameter is read-only for operator as part of provisioning)
  - Editable (the configuration parameter is editable as part of provisioning)
  - Device-Specific (the configuration parameter value needs to be entered by the operator for each device during deployment)
- i. Click **Save & Publish** to save and publish the service template configuration. The designer must publish the service templates to the operator to use in the creation

of deployment plans. After a filter or policy is published, it goes for peer review and approval. After approval, the filter or policy is deployed to device.

## Modifying CGNAT Service Templates

On the Service Designer page, you can view the collection of service templates defined for several applications, such as stateful firewall or CGNAT.

To modify service template instances, such as ADC, SFW, CGNAT, or TLB templates:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.

3. From the task pane, select **Deploy Service > Service Edit**.

The Service Instances page is displayed in the right pane, listing all the previously defined service templates.

4. From the View pane, perform one of the following tasks:

- Click the **ADC** button.

The list of ADC service templates is displayed. You need not click this button if you are launching the Service Designer page for the first time or are navigating to this page from another mode or a different page. You need to click this button only if you are viewing the other service templates, such as CGNAT or TLB.

- Click the **SFW** button.

The list of SFW templates is displayed.

- Click the **TLB** button.

The list of TLB templates is displayed.

- Click the **CGNAT** button.

The list of CGNAT templates is displayed.

5. In the main window, click the plus sign (+) next to the SDG pairs to expand the tree and view the pair of devices in the SDG group or pair. Select the check box next to the SDG pair or individual SDG for which you want to modify settings. In an SDG pair, you can select a single SDG or both the SDGs in the in the redundancy pair of devices.

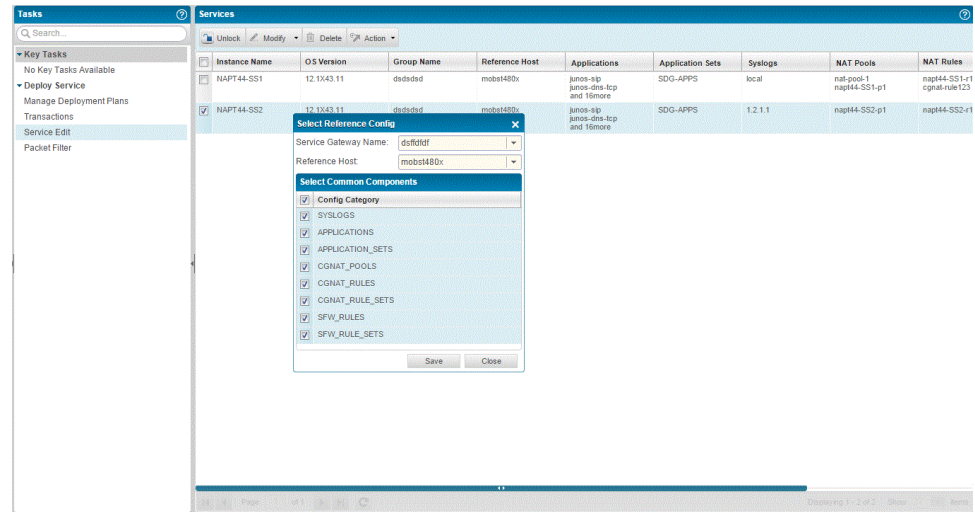


**NOTE:** Alternatively, you can also modify service templates from Service View in Build Mode by selecting the **Service Templates > Manage Service Templates** from the task pane, selecting a service instance, and clicking the **Modify** button.

---

- Click the **Lock** icon above the table of listed packet filters. The Select Reference Config dialog box is displayed.

Figure 23: Select Reference Config Dialog Box



- From the Service Gateway Name drop-down list, select the SDG group to which the packet filter must be applied.
- From the Host Name drop-down list, select the hostname of the SDG.
- In the Select Common Components section, select the check boxes beside the service modules or components, such as packet filters, SFW rules, or CGNAT rules, that are displayed. The displayed components depend on the attributes that are previously defined for that selected packet filter. For example, if the service policy is for stateful firewall, SFW rules and SFW rule sets are shown. Select the check box beside Config Category to select all the service components.
- Click **Save** to save the modified association.
- Select the check box beside the template you want to modify.
- Open the **Modify** menu above the list of templates to modify an existing template, and select the component or service attribute, such as application or rule, that you want to edit.
- Perform one of the following from the drop-down menu displayed for each component:
  - To retrieve the service component and import into the database of Edge Services Director, select **Import Object**. The Import Services dialog box appears. You can import the service templates assigned to SDGs or choose from a list of all of the

predefined templates in the database. Also, you can either import all of the components of a service or specific components.

- To create the component afresh, select **Create New**. The Create page corresponding to the service component appears. You can define the attributes for the service component in the same manner as you define the elements during the creation of a service template.

## Creating a Deployment Plan

You must have previously defined service templates and policy or filter templates before you can create a deployment plan.

To create a deployment plan and assigning devices to it:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the task pane, select **Service Edit**.  
The Manage Service Templates page is displayed.
4. Click the **CGNAT** button.  
The list of CGNAT service templates is displayed.
5. Select the check boxes next to the SDGs or SDG groups that you want to assign to the plan. Based on your selection of a service or a policy template, the components or attributes are shown for the corresponding device.
6. From the boxes that show the components of a service template, you can edit, delete, or add elements to it. If you do not have permissions to update a template, the corresponding icons are not shown.
7. Click the down arrow in the **Actions** menu and select **Send for Deployment** to create a deployment plan for the particular service template and save the plan.

If you create a deployment plan from Service view of Deploy mode, the Edit Service Instance page is displayed. You can modify the SDGs associated with the service instance and also modify the service instance attributes as necessary by either clicking the buttons corresponding to the various settings at the top of the wizard page to directly traverse to the page you want to modify or clicking the navigation buttons at the bottom of the wizard page to go to the different pages of the wizard. Click **Finish** to create a deployment plan.

A deploy plan is created for the service template with the devices that are assigned to it when you view the Deployment Plans page.

8. Alternatively, you can select **Discard changes** from the **Actions** menu to ignore the modifications done to a policy or filter template.
9. From the Deployment plans page, you can select **Reject** or **Approve** from the Actions drop-down list to reject or approve the deployment plan and make it available for commissioning to the devices.

## Importing a CGNAT Service Template

To create a clone of an existing CGNAT template by importing it:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.

3. From the task pane, select **Service Templates**.

The Manage Service Templates page is displayed.

4. Click the **CGNAT** button.

The list of CGNAT service templates is displayed. You need not click this button if you are launching the Service Designer page for the first time or are navigating to this page from another mode or a different page. You need to click this button only if you are viewing the other service templates, such as CGNAT or CGNAT.

5. Click the **Add** icon.

The Create a CGNAT Planning Template window appears.

6. In the Name field, enter a name for the service template or profile (limit of 63 alphanumeric characters without spaces).
7. In the Description field, enter a meaningful, easily-identifiable name for the service instance (limit of 255 alphanumeric characters). Each service instance you define can be applied to a single or multiple SDGs.
8. Click the **Import** button.

The Import Services dialog box appears.

You can import the service templates assigned to SDGs or choose from a list of all of the predefined templates in the database. Also, you can either import all of the components of a service or specific components.

9. Do one of the following for the Import section:

- Select the **From Existing Service Gateway** radio button if you want to import the CGNAT rule from SDGs that are present in the Edge Services Director database.
- Select the **From XML** radio button if you want to import the CGNAT rule from an XML configuration file on an external system.

10. If you selected the option to import the object from SDGs, do the following:

- Click the **Normal View** tab to view the list of SDGs. You can search for specific SDGs by entering a search item and clicking the **Search** icon.

Alternatively, click the **Group View** tab to view the list of SDG groups. You can search for specific SDG groups by entering a search item and clicking the **Search** icon.

- Click the plus sign (+) next to the All Service Gateways item to expand the tree structure that displays the list of SDGs or SDG groups. If the SDG pair is configured, you can select one of the devices, master or standby, from which you want to import the object.

Alternatively, if you selected the **Group View** tab, you can select an SDG from the groups displayed from which you want to import the object.

- Click **Import**. The object is added to the database and can be used during configuration of services or policies.

11. If you selected the option to import from an XML file, do the following:

- Click **Browse** beside the File Name field to navigate to the path where an XML file is available to be imported.
- Click **Upload**. The service template is added to the database and can be used during configuration of services or policies.

12. Do one of the following to import all components of a selected template or only a particular component of a template. For the components that are not imported, you need to specify the definitions of the components afresh.

- Select the check boxes next to all of the service instances that are displayed for the selected SDG or SDG group, or for the XML file that you uploaded. In such a case, all of the elements or parameters of the selected template or instance are imported.
- Alternatively, select the check box next to a particular or group of service instances to import only a specific component of the selected template

13. Similarly, you can select other components and import them to the template. Save the imported components to add them to the template you are creating by using the imported template as a base.

## Creating a Service Set

A service set is a collection of services to be performed by an Adaptive Services (AS) or Multiservices PIC. To create a service set as a component for the CGNAT template:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the task pane, select **Service Templates**.  
The Manage Service Templates page is displayed.
4. Click the **CGNAT** button.  
The list of CGNAT service templates is displayed.
5. Click the **Add** icon.  
The Create a CGNAT Planning Template window appears.
6. In the Name field, enter a name for the service template or profile (limit of 63 alphanumeric characters without spaces).
7. In the Description field, enter a meaningful, easily-identifiable name for the service instance (limit of 255 alphanumeric characters). Each service instance you define can be applied to a single or multiple SDGs.
8. Click the green plus sign in the Service Set box.  
The Addition of Service Set dialog box appears.



**NOTE:** For the service elements that you can configure using the Object Builder workspace, such as applications and rules, when you click the green plus sign (+) at the top-right corner of each of the service element boxes, the shortcut menu is displayed. Click the Create New radio button to create the service component afresh. Alternatively, click the Import from Object Builder radio button to open a dialog box that enables you to select from the list of service elements that are present in the database of Edge Services Director and import them into the service template.

If a green tick mark is shown beside a field in the dialog box, it denotes that you can add attributes for that component. A red cross mark shows that you can delete that particular attribute for that component.

9. In the Name field, enter the name to identify the service set. Rules are combined into rule sets, and are associated with a service set for each application such as firewall or CGNAT.

10. In the Sampling Service Choices section, do one of the following:

- Click **Interface Services** to configure an interface-style service set. An interface service set is used as an action modifier across an entire interface
- In the Service Interfaces field, specify the name for the adaptive services interface associated with an interface-wide service set.

When you have defined and grouped the service rules by configuring the service-set definition, you can apply services to one or more interfaces installed on the router. When you apply the service set to an interface, it automatically ensures that packets are directed to the PIC.

- From the **Load Balancing Options** section, configure the high availability (HA) options.

The following hash keys can be configured in the egress direction: **destination-ip** (Use the destination IP address of the flow to compute the hash used in load balancing.) and **source-ip** (Use the source IP address of the flow to compute the hash used in load balancing.)

- Click the green tick mark beside the Egress Key element to configure the hash keys to be used in the egress flow direction. The configuration is mandatory if you are using AMS for Network Address Translation (NAT). This configuration is not mandatory if you are using AMS for stateful firewall; if the hash keys are not configured, then the defaults are chosen.
- Click the green tick mark beside the Ingress Key element to configure the hash keys to be used in the ingress flow direction. The configuration is mandatory if you are using AMS for Network Address Translation (NAT). This configuration is not mandatory if you are using AMS for stateful firewall; if the hash keys are not configured, then the defaults are chosen.

Configure the hash keys used for load balancing in aggregated multiservices (AMS) for service applications (Network Address Translation [NAT], stateful firewall, application-level gateway [ALG], HTTP header enrichment, and mobility). The hash keys supported in the ingress and egress direction are the source IP address and destination IP address.

Hash keys are used to define the load-balancing behavior among the various members in the AMS group. For example, if **hash-keys** is configured as **source-ip**, then the hashing would be performed based on the source IP address of the packet. Therefore, all packets with the same source IP address land on the same member. Hash keys must be configured with respect to the traffic direction: ingress or egress. For example, if **hash-keys** is configured as **source-ip** in the ingress direction, then it should be configured as **destination-ip** in the egress direction. This is required to ensure that the packets of the same flow reach the same member of the AMS group.

The configuration of the ingress and egress hash keys is mandatory if you are using AMS for NAT. This configuration is not mandatory if you are using AMS for stateful firewall; if the hash keys are not configured, then the defaults are chosen. Refer to [Table 46 on page 245](#) for the supported hash keys.

The resource-triggered option enables anchor session PICs to use the load or resource information from the anchor services PICs to select the AMS member will anchor the services for the subscriber for load balancing among AMS members. In addition, for mobile subscriber-aware services (such as HTTP header enrichment), you must configure the **resource-triggered** statement, which means that the load balancing is not done using the ingress and egress keys.

**Table 46: Hash Keys Supported for AMS for Service Applications**

Service Set at Ingress Interface			Service Set at Egress Interface	
Hash Keys for NAT				
NAT Type	Ingress hash key	Egress hash key	Ingress hash key	Egress hash key
source static	Destination IP address	Source IP address	Source IP address	Destination IP address
source dynamic	Source IP address	Destination IP address	Destination IP address	Source IP address
Network Address Port Translation (NAPT)	Source IP address	Destination IP address	Destination IP address	Source IP address
destination static	Source IP address	Destination IP address	Destination IP address	Source IP address
Hash Keys for Stateful Firewall				
Stateful Firewall	Destination IP address	Source IP address	Destination IP address	Source IP address
Stateful Firewall	Source IP address	Destination IP address	Source IP address	Destination IP address



**NOTE:** If NAT is used in the service set (along with stateful firewall and ALG), then the hash keys should be based on the NAT type; otherwise, the hash keys of the stateful firewall should be used.

- Click **Next Hop Services** to configure a next-hop style service set. A next-hop service set is a route-based method of applying a particular service. Only packets destined for a specific next hop are serviced by the creation of explicit static routes.

- In the **Inside Interface** list, specify the interface type of the service interface associated with the service set applied inside the network. For inline IP reassembly, set the interface type to local. Also, specify the name and logical unit number of the service interface associated with the service set applied inside the network.

When a next-hop service is configured, the AS or Multiservices PIC is considered to be a two-legged module with one leg configured to be the inside interface (inside the network) and the other configured as the outside interface (outside the network).

- In the **Outside Interface** list, specify the interface type of the service interface associated with the service set applied outside the network. For inline IP reassembly, set the interface type to local. Also, specify the name and logical unit number of the service interface associated with the service set applied outside the network.
- In the **Service Interface Pool** list, select the name of the pool of logical interfaces configured at the [edit services service-interface-pools pool pool-name] hierarchy level. You can configure a service interface pool only if the service set has a PGCP rule configured. The service set cannot contain any other type of rule.

- 

- Click **Sampling Services** to configure a sampling service set.
  - In the Service Interface field, specify the service interface, which is the interface the sampling is taken from. In the case of a sampling service set, the service interface must be a Multiservices PIC interface with a subunit number of 0 (zero). The subunit number defaults to 0. The reverse-flow statement is not mandatory. All sampled traffic is considered to be forward traffic. If you set the reverse-flow statement, it is ignored.
- Select the **Replication Service** check box to configure the services replication options for inter-chassis high availability on MS-MIC and MS-MPC. This field is available only if you selected the Junos OS 12.1 version.
  - In the Replication Threshold field, specify the number of seconds for the replication threshold. When a flow has been active for more than the number of seconds specified as a threshold, flow state information is replicated to the backup device. Make sure that the replication-threshold value is than the open-timeout value (the timeout period for establishing a TCP connection). The default value of the replication threshold is 180 seconds. This value is also the minimum.
  - Select the **Stateful Firewall** check box to replicate stateful firewall state information.
  - Select the **NAT** check box to replicate NAT44 information.

11. Select the **Service Set Options** check box to specify the service set options to apply to a service set. This field is available only if you selected the Junos OS 14.1 version.

12. In the Redundancy Set ID field, specify a unique identifier in the range of 1 through 100 for the redundancy set. The redundancy group IDs that the service redundancy daemon (srd) uses are associated with those configured for the ICCP daemon (iccpd) through

the existing ICCP configuration hierarchy by using the same redundancy group ID in the configuration of the services redundancy group. This field is available only if you selected the Junos OS 14.1 version.

The actions to be performed when configured redundancy events occur are defined in redundancy policies. Redundancy policies are associated with redundancy sets; they are analogous to rules associated with service sets. Redundancy sets are associated to redundancy groups by redundancy group IDs. Redundancy group details are defined by the underlying ICCPd configuration. Finally, service sets and redundancy sets are associated through the **redundancy-sets** statement in service sets configuration.

13. In the CGNAT Rule Sets section, select the rule set you want to associate with the service set from the Available column and click the right arrow to move to the Selected column.
14. In the CGNAT Rules section, select the rule you want to associate with the service set from the Available column and click the right arrow to move to the Selected column.
15. In the CGNAT Syslogs section, select the syslog you want to associate with the service set from the Available column and click the right arrow to move to the Selected column.
16. Click **Save** to save the service template configuration. Else, click **Close** to discard the changes to the template.

## Creating a Syslog

You can enable system logging. The system log information from the Adaptive Services or Multiservices PIC is passed to the kernel for logging in the /var/log directory. This setting overrides any syslog statement setting included in the service set or interface default configuration.

To create a syslog for the CGNAT template:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the task pane, select **Service Templates**.  
The Manage Service Templates page is displayed.
4. Click the **CGNAT** button.

The list of CGNAT service templates is displayed.

5. Click the **Add** icon.

The Create a CGNAT Planning Template window appears.

6. In the Name field, enter a name for the service template or profile (limit of 63 alphanumeric characters without spaces).
7. In the Description field, enter a meaningful, easily-identifiable name for the service instance (limit of 255 alphanumeric characters). Each service instance you define can be applied to a single or multiple SDGs.
8. Click the green plus sign in the Server Groups box.

The Addition of Group dialog box appears.



**NOTE:** For the service elements that you can configure using the Object Builder workspace, such as applications and rules, when you click the green plus sign (+) at the top-right corner of each of the service element boxes, the shortcut menu is displayed. Click the **Create New** radio button to create the service component afresh. Alternatively, click the **Import from Object Builder** radio button to open a dialog box that enables you to select from the list of service elements that are present in the database of Edge Services Director and import them into the service template.

If a green tick mark is shown beside a field in the dialog box, it denotes that you can add attributes for that component. A red cross mark shows that you can delete that particular attribute for that component.

9. In the Name field, enter the name for the syslog component. Specify the fully qualified domain name or IP address for the syslog server.
10. In the Services list, specify the system logging severity level. It assigns a severity level to the facility. Valid entries include:
  - **alert**—Conditions that should be corrected immediately.
  - **any**—Matches any level.
  - **critical**—Critical conditions.
  - **emergency**—Panic conditions.
  - **error**—Error conditions.
  - **info**—Informational messages.

- **notice**—Conditions that require special handling.
  - **warning**—Warning messages.
11. From the Facility Override list, select the override for the default facility for system log reporting. Valid values include:
    - authorization**
    - daemon**
    - ftp**
    - kernel**
    - local0** through **local7**
    - user**
  12. In the Log Prefix field, set the system logging prefix value for all logging to the system log host.
  13. In the Port field, specify the port number to be used for connection with the remote syslog server.
  14. In the Class section, set the class of applications to be logged to the system log.
    - **alg-logs**—Log application-level gateway events.
    - **ids-logs**—Log intrusion detection system events.
    - **nat-logs**—Log Network Address Translation events.
    - **packet-logs**—Log general packet-related events.
    - **session-logs**—Log session open and close events.
    - **session-logs open**—Log session open events only.
    - **session-logs close**—Log session close events.
    - **stateful-firewall-logs**—Log stateful firewall events.
  15. Click **Save** to save the service template configuration. Else, click **Close** to discard the changes to the template.

## Creating a Rule

To create a rule for the CGNAT template:

1. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
2. From the task pane, select **Service Templates**.  
The Manage Service Templates page is displayed.

3. Click the **CGNAT** button.

The list of CGNAT service templates is displayed.

4. Click the **Add** icon.

The Create a CGNAT Planning Template window appears.

5. Enter the name of the template and the service instance in the respective fields.

6. Click the green plus sign in the Server Groups box. The Addition of Group dialog box appears.



**NOTE:** If a green tick mark is shown beside a field in the dialog box, it denotes that you can add attributes for that component. A red cross mark shows that you can delete that particular attribute for that component.

7. From the **Rule** list, select one of the previously configured rules.

The rules that you configured in the Service Templates workspace for CGNAT, packet filter, or CGNAT are displayed.

8. Click **Save** to save the service template configuration. Else, click **Close** to discard the changes to the template.

## Creating a Rule Set

The rule-set statement defines a collection of stateful firewall rules that determine what actions the router software performs on packets in the data stream. You define each rule by specifying a rule name and configuring terms. Then, you specify the order of the rules by including the **rule-set** statement at the **[edit services stateful-firewall]** hierarchy level with a rule statement for each rule.

The router software processes the rules in the order in which you specify them in the configuration. If a term in a rule matches the packet, the router performs the corresponding action and the rule processing stops. If no term in a rule matches the packet, processing continues to the next rule in the rule set. If none of the rules matches the packet, the packet is dropped by default.

To create a rule set for the CGNAT template:

1. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.

The functionalities that you can configure in this mode are displayed in the task pane.

2. From the task pane, select **Service Templates**.

The Manage Service Templates page is displayed.

3. Click the **CGNAT** button.

The list of CGNAT service templates is displayed.

4. Click the **Add** icon.

The Create a CGNAT Planning Template window appears.

5. Enter the name of the template and the service instance in the respective fields.

6. Click the green plus sign in the Rule Sets box.

The Addition of Rule Sets dialog box appears.



**NOTE:** For the service elements that you can configure using the Object Builder workspace, such as applications and rules, when you click the green plus sign (+) at the top-right corner of each of the service element boxes, the shortcut menu is displayed. Click the **Create New** radio button to create the service component afresh. Alternatively, click the **Import from Object Builder** radio button to open a dialog box that enables you to select from the list of service elements that are present in the database of Edge Services Director and import them into the service template.

If a green tick mark is shown beside a field in the dialog box, it denotes that you can add attributes for that component. A red cross mark shows that you can delete that particular attribute for that component.

7. Specify the rule set name the router uses when applying this service.
8. Select the rules that you want to group into a rule set from the Available column and click the right arrow to move the rules to the Selected column.
9. Click **Save** to save the service template configuration. Else, click **Close** to discard the changes to the template.

## Creating a Pool

To create an address pool for the CGNAT template:

1. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
2. From the task pane, select **Service Templates**.  
The Manage Service Templates page is displayed.

3. Click the **CGNAT** button.

The list of CGNAT service templates is displayed.

4. Click the **Add** icon.

The Create a CGNAT Planning Template window appears.

5. Enter the name of the template and the service instance in the respective fields.

6. Click the green plus sign in the NAT Pools box. The Addition of NAT Pool dialog box appears.



**NOTE:** For the service elements that you can configure using the Object Builder workspace, such as applications and rules, when you click the green plus sign (+) at the top-right corner of each of the service element boxes, the shortcut menu is displayed. Click the **Create New** radio button to create the service component afresh. Alternatively, click the **Import from Object Builder** radio button to open a dialog box that enables you to select from the list of service elements that are present in the database of Edge Services Director and import them into the service template.

If a green tick mark is shown beside a field in the dialog box, it denotes that you can add attributes for that component. A red cross mark shows that you can delete that particular attribute for that component.

7. From the **Pool** list, select one of the previously configured pools. The pools that you configured in the Service Templates workspace for CGNAT are displayed.
8. Click **Save** to save the service template configuration. Else, click **Close** to discard the changes to the template.

#### Related Documentation

- [Service Templates Overview on page 189](#)
- [Filtering Service Templates on page 189](#)
- [Viewing Service Templates on page 192](#)
- [Using the Actions Menu on the Service Template and Service Edit Pages on page 195](#)

---

## Creating and Managing SFW Service Templates

Each stateful firewall rule consists of a set of terms, similar to a service filter. A term consists of the following:

from statement—Specifies the match conditions and applications that are included and excluded. The from statement is optional in stateful firewall rules.

then statement—Specifies the actions and action modifiers to be performed by the router software. The then statement is mandatory in stateful firewall rules.

You can perform the following tasks with the Service Designer page for SFW:

- Create an SFW service template with attributes and settings for stateful firewall operations.
- Modify an existing SFW template to meet the network needs and deployment scenarios.
- Delete an existing template.
- [Creating an SFW Service Template on page 253](#)
- [Modifying SFW Service Templates on page 256](#)
- [Creating a Deployment Plan on page 258](#)
- [Importing an SFW Service Template on page 259](#)
- [Creating a Service Set on page 261](#)
- [Creating an Application on page 265](#)
- [Creating an Application Set on page 268](#)
- [Creating a Syslog on page 269](#)
- [Creating a Rule on page 271](#)
- [Creating a Rule Set on page 272](#)
- [Creating a Services PIC for an SFW Service Template on page 274](#)

## Creating an SFW Service Template

To configure a new SFW service template:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the task pane, select **Service Templates**.  
The Manage Service Templates page is displayed.
4. Click the **SFW** button.

The list of SFW service templates is displayed.

The Service Designer page displays a bar graph in the top pane of the page. The count of service templates of each type is displayed on the vertical axis and the service type is shown on the horizontal axis. A color-coding format is used to represent the bars on the graph. Published service templates are shown in olive green color and unpublished service templates are shown in blue color. Mouse over each bar in the

chart to highlight and display the number of templates published or unpublished for each type of service.

5. Click the **Add** icon. The Select Version dialog box appears.
6. Select **Junos 12.1** if you want to create a template based on the Junos OS Release 12.1. Alternatively, select **Junos 14.1** if you want to create a template based on the Junos OS Release 14.1.



**NOTE:** All the service template components described in this section can be created for templates that are based on both the Junos OS Releases 12.1 and 14.1. The service elements or components that are additionally available for configuration when you select the Junos OS 14.1 version are explicitly mentioned in the relevant steps of the procedure.

The Create an SFW Planning Template window appears.

**Figure 24: Create SFW Service Template Window**

7. In the Template Name field, enter a name for the service template or profile (limit of 63 alphanumeric characters without spaces).
8. In the Description field, enter a meaningful, easily-identifiable name for the service instance (limit of 255 characters). Each service instance you define can be applied to a single or multiple SDGs.
9. Instead of creating a new template entirely, you can import the parameters defined for a previous SFW service instance and customize only the settings that are necessary. Imported templates are created without any device assigned to them. To use these

templates, you must associate a device with the policy. To clone an existing template by importing it, click the **Import** button.

The Import Services dialog box is displayed. See *Importing an SFW Service Template* for step-wise details on importing an SFW service template.

10. The Create an SFW Planning Template window displays the individual elements or components of the service with a graphical icon for each of the service elements and the corresponding names in separate boxes. You can add, edit, or delete these service elements in a template.



**NOTE:** The Property View tab and the Config View tab are displayed on the right pane of the template window. The Property View tab provides a tree-based structure of the parameters defined in a service template. You can expand the tree and view details of each component. A key value pair representation is shown. Each of the components can be treated as categories of the service template shown in the property view.

The Config View tab displays the elements or components specified for a service template in the form of configuration stanzas and hierarchy levels. This display is similar to the show command that you can use at a certain [edit] hierarchy level to view the defined settings. Each level in the hierarchy is indented to indicate each statement's relative position in the hierarchy. Each level is generally set off with braces, with an open brace ( { ) at the beginning of each hierarchy level and a closing brace ( } ) at the end. If the statement at a hierarchy level is empty, the braces are not displayed. Each leaf statement ends with a semicolon ( ; ), as does the last statement in the hierarchy.

- a. Click the green tick mark (✓) displayed at the top-right corner of each of the service element boxes to create a new element. If the green tick mark is not shown, it indicates that the user role does not have the permission to create an element.
- b. Click the red cross mark (x) displayed at the top-right corner of the icons of each element if you want to delete the existing configuration. The user with designer role has permissions to remove or edit elements.
- c. If the red cross mark is not displayed beside a particular icon, it signifies that the element cannot be deleted.
- d. The diamond icon that contains an orange tick mark within it at the top-right corner of the service component name denotes that the particular element can be modified. The absence of this icon denotes that the user does not have permissions to modify the attributes of the service component.
- e. Double-click each icon pertaining to a service element to view or edit its settings. If you do not possess the permission to modify the element, a view-only dialog box with the attributes of the selected element is shown. Otherwise, an editable dialog box enables you to modify the settings.

- f. Click **Save** to save the service template configuration. Else, click **Close** to discard the changes to the template.
- g. Click the Maximize icon displayed at the top-right corner of the rectangle or box that shows all of the values or entities of a particular component of a service template. The specified component or attribute is displayed as a separate dialog box, listing all of the values of the particular component. You can add, modify, or delete the listed values.
- h. While creating the new service template, the designer can add or modify service parameter values and also restrict the access level for each service parameter for the operator. The designer can set following access levels for each service parameters to operator in planning template. Click the new icon (cascading files icon) displayed at the top-left corner of each of the element boxes to open the shortcut menu. You can click one of the following radio buttons:
  - Read-only (the configuration parameter is read-only for operator as part of provisioning)
  - Editable (the configuration parameter is editable as part of provisioning)
  - Device-Specific (the configuration parameter value needs to be entered by the operator for each device during deployment)
- i. Click **Save & Publish** to save and publish the service template configuration. The designer must publish the service templates to the operator to use in the creation of deployment plans. After a filter or policy is published, it goes for peer review and approval. After approval, the filter or policy is deployed to device.

## Modifying SFW Service Templates

On the Service Designer page, you can view the collection of service templates defined for several applications, such as stateful firewall or CGNAT.

To modify service template instances, such as ADC, SFW, CGNAT, or TLB templates:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the task pane, select **Deploy Service > Service Edit**.  
The Service Instances page is displayed in the right pane, listing all the previously defined service templates.
4. From the View pane, perform one of the following tasks:

- Click the **ADC** button.

The list of ADC service templates is displayed. You need not click this button if you are launching the Service Designer page for the first time or are navigating to this page from another mode or a different page. You need to click this button only if you are viewing the other service templates, such as CGNAT or TLB.

- Click the **SFW** button.

The list of SFW templates is displayed.

- Click the **TLB** button.

The list of TLB templates is displayed.

- Click the **CGNAT** button.

The list of CGNAT templates is displayed.

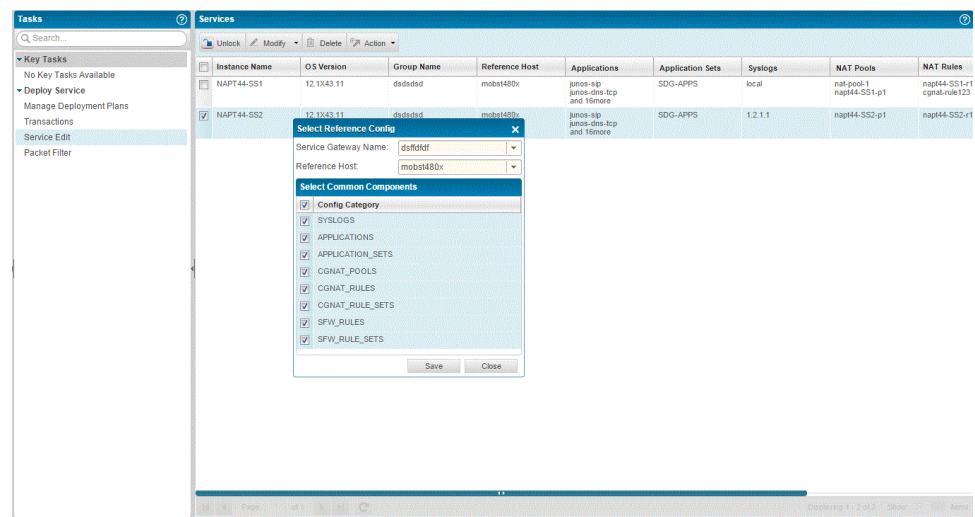
- In the main window, click the plus sign (+) next to the SDG pairs to expand the tree and view the pair of devices in the SDG group or pair. Select the check box next to the SDG pair or individual SDG for which you want to modify settings. In an SDG pair, you can select a single SDG or both the SDGs in the in the redundancy pair of devices.



**NOTE:** Alternatively, you can also modify service templates from Service View in Build Mode by selecting the Service Templates > Manage Service Templates from the task pane, selecting a service instance, and clicking the Modify button.

- Click the **Lock** icon above the table of listed packet filters. The Select Reference Config dialog box is displayed.

Figure 25: Select Reference Config Dialog Box



7. From the Service Gateway Name drop-down list, select the SDG group to which the packet filter must be applied.
8. From the Host Name drop-down list, select the hostname of the SDG.
9. In the Select Common Components section, select the check boxes beside the service modules or components, such as packet filters, SFW rules, or CGNAT rules, that are displayed. The displayed components depend on the attributes that are previously defined for that selected packet filter. For example, if the service policy is for stateful firewall, SFW rules and SFW rule sets are shown. Select the check box beside Config Category to select all the service components.
10. Click **Save** to save the modified association.
11. Select the check box beside the template you want to modify.
12. Open the **Modify** menu above the list of templates to modify an existing template, and select the component or service attribute, such as application or rule, that you want to edit.
13. Perform one of the following from the drop-down menu displayed for each component:
  - To retrieve the service component and import into the database of Edge Services Director, select **Import Object**. The Import Services dialog box appears. You can import the service templates assigned to SDGs or choose from a list of all of the predefined templates in the database. Also, you can either import all of the components of a service or specific components.
  - To create the component afresh, select **Create New**. The Create page corresponding to the service component appears. You can define the attributes for the service component in the same manner as you define the elements during the creation of a service template.

## Creating a Deployment Plan

You must have previously defined service templates and policy or filter templates before you can create a deployment plan.

To create a deployment plan and assigning devices to it:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the task pane, select **Service Edit**.

The Manage Service Templates page is displayed.

4. Click the **SFW** button.

The list of SFW service templates is displayed.

5. Select the check boxes next to the SDGs or SDG groups that you want to assign to the plan. Based on your selection of a service or a policy template, the components or attributes are shown for the corresponding device.
6. From the boxes that show the components of a service template, you can edit, delete, or add elements to it. If you do not have permissions to update a template, the corresponding icons are not shown.
7. Click the down arrow in the **Actions** menu and select **Send for Deployment** to create a deployment plan for the particular service template and save the plan.

If you create a deployment plan from Service view of Deploy mode, the Edit Service Instance page is displayed. You can modify the SDGs associated with the service instance and also modify the service instance attributes as necessary by either clicking the buttons corresponding to the various settings at the top of the wizard page to directly traverse to the page you want to modify or clicking the navigation buttons at the bottom of the wizard page to go to the different pages of the wizard. Click **Finish** to create a deployment plan.

A deploy plan is created for the service template with the devices that are assigned to it when you view the Deployment Plans page.

8. Alternatively, you can select **Discard changes** from the **Actions** menu to ignore the modifications done to a policy or filter template.
9. From the Deployment plans page, you can select **Reject** or **Approve** from the Actions drop-down list to reject or approve the deployment plan and make it available for commissioning to the devices.

## Importing an SFW Service Template

To create a clone of an existing SFW template by importing it:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the task pane, select **Service Templates**.

The Manage Service Templates page is displayed.

4. Click the **SFW** button.

The list of SFW service templates is displayed.

You need not click this button if you are launching the Service Designer page for the first time or are navigating to this page from another mode or a different page. You need to click this button only if you are viewing the other service templates, such as CGNAT or SFW.

5. Click the **Add** icon.

The Create an SFW Planning Template window appears.

6. In the Template Name field, enter a name for the service template or profile (limit of 63 alphanumeric characters without spaces).
7. In the Description field, enter a meaningful, easily-identifiable name for the service instance (limit of 255 characters). Each service instance you define can be applied to a single or multiple SDGs.
8. Click the **Import** button.

The Import Services dialog box appears.

You can import the service templates assigned to SDGs or choose from a list of all of the predefined templates in the database. Also, you can either import all of the components of a service or specific components.

9. Do one of the following for the Import section:
  - Select the **From Existing Service Gateway** radio button if you want to import the CGNAT rule from SDGs that are present in the Edge Services Director database.
  - Select the **From XML** radio button if you want to import the CGNAT rule from an XML configuration file on an external system.
10. If you selected the option to import the object from SDGs, do the following:
  - Click the **Normal View** tab to view the list of SDGs. You can search for specific SDGs by entering a search item and clicking the **Search** icon.

Alternatively, click the **Group View** tab to view the list of SDG groups. You can search for specific SDG groups by entering a search item and clicking the **Search** icon.
  - Click the plus sign (+) next to the All Service Gateways item to expand the tree structure that displays the list of SDGs or SDG groups. If the SDG pair is configured, you can select one of the devices, master or standby, from which you want to import the object.

Alternatively, if you selected the **Group View** tab, you can select an SDG from the groups displayed from which you want to import the object.

- Click **Import**. The object is added to the database and can be used during configuration of services or policies.
11. If you selected the option to import from an XML file, do the following:
    - Click **Browse** beside the File Name field to navigate to the path where an XML file is available to be imported.
    - Click **Upload**. The service template is added to the database and can be used during configuration of services or policies.
  12. Do one of the following to import all components of a selected template or only a particular component of a template. For the components that are not imported, you need to specify the definitions of the components afresh.
    - Select the check boxes next to all of the service instances that are displayed for the selected SDG or SDG group, or for the XML file that you uploaded. In such a case, all of the elements or parameters of the selected template or instance are imported.
    - Alternatively, select the check box next to a particular or group of service instances to import only a specific component of the selected template
  13. Similarly, you can select other components and them to the template. Save the imported components to add them to the template you are creating by using the imported template as a base.

## Creating a Service Set

A service set is a collection of services to be performed by an Adaptive Services (AS) or Multiservices PIC. To create a service set as a component for the SFW template:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the task pane, select **Service Templates**.  
The Manage Service Templates page is displayed.
4. Click the **SFW** button.  
The list of SFW service templates is displayed.
5. Click the **Add** icon.  
The Create an SFW Planning Template window appears.
6. Enter the name of the template and the service instance in the respective fields.

7. Click the green plus sign in the Service Set box.

The Addition of Service Set dialog box appears.



**NOTE:** For the service elements that you can configure using the Object Builder workspace, such as applications and rules, when you click the green plus sign (+) at the top-right corner of each of the service element boxes, the shortcut menu is displayed. Click the **Create New** radio button to create the service component afresh. Alternatively, click the **Import from Object Builder** radio button to open a dialog box that enables you to select from the list of service elements that are present in the database of Edge Services Director and import them into the service template.

If a green tick mark is shown beside a field in the dialog box, it denotes that you can add attributes for that component. A red cross mark shows that you can delete that particular attribute for that component.

8. In the Name field, enter the name to identify the service set. Rules are combined into rule sets, and are associated with a service set for each application such as firewall or CGNAT.

9. In the Sampling Service Choices section, do one of the following:

- Click **Interface Services** to configure an interface-style service set. An interface service set is used as an action modifier across an entire interface
  - In the Service Interfaces field, specify the name for the adaptive services interface associated with an interface-wide service set.

When you have defined and grouped the service rules by configuring the service-set definition, you can apply services to one or more interfaces installed on the router. When you apply the service set to an interface, it automatically ensures that packets are directed to the PIC.

- From the **Load Balancing Options** section, configure the high availability (HA) options.

The following hash keys can be configured in the egress direction: **destination-ip** (Use the destination IP address of the flow to compute the hash used in load balancing.) and **source-ip** (Use the source IP address of the flow to compute the hash used in load balancing.)

- Click the green tick mark beside the Egress Key element to configure the hash keys to be used in the egress flow direction. The configuration is mandatory if you are using AMS for Network Address Translation (NAT). This configuration is not mandatory if you are using AMS for stateful firewall; if the hash keys are not configured, then the defaults are chosen.
- Click the green tick mark beside the Ingress Key element to configure the hash keys to be used in the ingress flow direction. The configuration is mandatory if you are using AMS for Network Address Translation (NAT). This configuration is

not mandatory if you are using AMS for stateful firewall; if the hash keys are not configured, then the defaults are chosen.

Configure the hash keys used for load balancing in aggregated multiservices (AMS) for service applications (Network Address Translation [NAT], stateful firewall, application-level gateway [ALG], HTTP header enrichment, and mobility). The hash keys supported in the ingress and egress direction are the source IP address and destination IP address.

Hash keys are used to define the load-balancing behavior among the various members in the AMS group. For example, if **hash-keys** is configured as **source-ip**, then the hashing would be performed based on the source IP address of the packet. Therefore, all packets with the same source IP address land on the same member. Hash keys must be configured with respect to the traffic direction: ingress or egress. For example, if **hash-keys** is configured as **source-ip** in the ingress direction, then it should be configured as **destination-ip** in the egress direction. This is required to ensure that the packets of the same flow reach the same member of the AMS group.

The configuration of the ingress and egress hash keys is mandatory if you are using AMS for NAT. This configuration is not mandatory if you are using AMS for stateful firewall; if the hash keys are not configured, then the defaults are chosen. Refer to [Table 46 on page 245](#) for the supported hash keys.

The resource-triggered option enables anchor session PICs to use the load or resource information from the anchor services PICs to select the AMS member will anchor the services for the subscriber for load balancing among AMS members. In addition, for mobile subscriber-aware services (such as HTTP header enrichment), you must configure the **resource-triggered** statement, which means that the load balancing is not done using the ingress and egress keys.

**Table 47: Hash Keys Supported for AMS for Service Applications**

Service Set at Ingress Interface			Service Set at Egress Interface	
Hash Keys for NAT				
NAT Type	Ingress hash key	Egress hash key	Ingress hash key	Egress hash key
source static	Destination IP address	Source IP address	Source IP address	Destination IP address
source dynamic	Source IP address	Destination IP address	Destination IP address	Source IP address
Network Address Port Translation (NAPT)	Source IP address	Destination IP address	Destination IP address	Source IP address
destination static	Source IP address	Destination IP address	Destination IP address	Source IP address
Hash Keys for Stateful Firewall				
Stateful Firewall	Destination IP address	Source IP address	Destination IP address	Source IP address
Stateful Firewall	Source IP address	Destination IP address	Source IP address	Destination IP address



**NOTE:** If NAT is used in the service set (along with stateful firewall and ALG), then the hash keys should be based on the NAT type; otherwise, the hash keys of the stateful firewall should be used.

- Click **Next Hop Services** to configure a next-hop style service set. A next-hop service set is a route-based method of applying a particular service. Only packets destined for a specific next hop are serviced by the creation of explicit static routes.
- In the **Inside Interface** list, specify the interface type of the service interface associated with the service set applied inside the network. For inline IP reassembly, set the interface type to local. Also, specify the name and logical unit number of the service interface associated with the service set applied inside the network.

When a next-hop service is configured, the AS or Multiservices PIC is considered to be a two-legged module with one leg configured to be the inside interface (inside the network) and the other configured as the outside interface (outside the network).

- In the **Outside Interface** list, specify the interface type of the service interface associated with the service set applied outside the network. For inline IP reassembly, set the interface type to local. Also, specify the name and logical unit number of the service interface associated with the service set applied outside the network.
- In the **Service Interface Pool** list, select the name of the pool of logical interfaces configured at the [edit services service-interface-pools pool pool-name] hierarchy level. You can configure a service interface pool only if the service set has a PGCP rule configured. The service set cannot contain any other type of rule.
- 
- Click **Sampling Services** to configure a sampling service set.
  - In the Service Interface field, specify the service interface, which is the interface the sampling is taken from. In the case of a sampling service set, the service interface must be a Multiservices PIC interface with a subunit number of 0 (zero). The subunit number defaults to 0. The reverse-flow statement is not mandatory. All sampled traffic is considered to be forward traffic. If you set the reverse-flow statement, it is ignored.
- Select the **Replication Service** check box to configure the services replication options for inter-chassis high availability on MS-MIC and MS-MPC.
  - In the Replication Threshold field, specify the number of seconds for the replication threshold. When a flow has been active for more than the number of seconds specified as a threshold, flow state information is replicated to the backup device. Make sure that the replication-threshold value is than the open-timeout value (the timeout period for establishing a TCP connection). The default value of the replication threshold is 180 seconds. This value is also the minimum.
- Select the **Stateful Firewall** check box to replicate stateful firewall state information.

- Select the **NAT** check box to replicate NAT44 information.
10. In the SFW Rule Sets section, select the rule set you want to associate with the service set from the Available column and click the right arrow to move to the Selected column.
  11. In the SFW Rules section, select the rule you want to associate with the service set from the Available column and click the right arrow to move to the Selected column.
  12. In the SFW Syslogs section, select the syslog you want to associate with the service set from the Available column and click the right arrow to move to the Selected column.
  13. Click **Save** to save the service template configuration. Else, click **Close** to discard the changes to the template.

## Creating an Application

You can define application protocols for the stateful firewall and Network Address Translation (NAT) services to use in match condition rules. An application protocol, or application layer gateway (ALG), defines application parameters using information from network Layer 3 and above. Examples of such applications are FTP and H.323.

To create an application for an SFW rule term:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the task pane, select **Service Templates**.  
The Manage Service Templates page is displayed.
4. Click the **SFW** button.  
The list of SFW service templates is displayed.
5. Click the **Add** icon.  
The Create a SFW Planning Template window appears.
6. In the Name field, enter a name for the service template or profile (limit of 63 alphanumeric characters without spaces).

7. In the Description field, enter a meaningful, easily-identifiable name for the service instance (limit of 255 alphanumeric characters). Each service instance you define can be applied to a single or multiple SDGs.

8. Click the green plus sign in the Applications box.

The Create an Application dialog box appears.



**NOTE:** For the service elements that you can configure using the Object Builder workspace, such as applications and rules, when you click the green plus sign (+) at the top-right corner of each of the service element boxes, the shortcut menu is displayed. Click the Create New radio button to create the service component afresh. Alternatively, click the Import from Object Builder radio button to open a dialog box that enables you to select from the list of service elements that are present in the database of Edge Services Director and import them into the service template.

If a green tick mark is shown beside a field in the dialog box, it denotes that you can add attributes for that component. A red cross mark shows that you can delete that particular attribute for that component.

9. In the Name field, enter the name to identify the application.
10. From the Protocol drop-down list, specify the networking protocol type or number to match in an application definition. The following text values are supported: **TCP**, **UDP**, **ICMP**, and **GRE**. Based on the selection, the dialog box refreshes to display additional fields applicable for the protocol.
11. From the Application Protocol drop-down list, specify the application protocol name. Application protocols are also called application layer gateways (ALGs). The application-protocol setting allows you to specify which of the supported application protocols (ALGs) to configure and include in an application set for service processing. Valid entries include the following:

**dns**—Domain Name Service

**icmp**—ICMP

**rtsp**—Real Time Streaming Protocol

**tftp**—Trivial File Transfer Protocol

Based on the selection, the dialog box refreshes to display additional fields applicable for the application protocol.

12. In the **Inactivity Timeout (secs)** field, specify the length of time, in seconds, for which the application is inactive before it times out. The default is 30 seconds.

13. In the **ICMP Type** field, specify the Internet Control Message Protocol (ICMP) code type. The ICMP code and type provide additional specification, in conjunction with the network protocol, for packet matching in an application definition. Normally, you specify this match in conjunction with the protocol match statement to determine which protocol is being used on the port. The only value available in this field is **ECHO\_REQUEST**.



**NOTE:** From the Junos OS CLI, to configure ICMP settings, include the `icmp-code` and `icmp-type` statements at the `[edit applications application application-name]` hierarchy level:

In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed): `echo-reply` (0), `echo-request` (8), `info-reply` (16), `info-request` (15), `mask-request` (17), `mask-reply` (18), `parameter-problem` (12), `redirect` (5), `router-advertisement` (9), `router-solicit` (10), `source-quench` (4), `time-exceeded` (11), `timestamp` (13), `timestamp-reply` (14), or `unreachable` (3).

14. From the Source Port type list, do one of the following:

- Select **RANGE** to configure a range of source ports for the application, and enter the upper limit and lower limit of the range of ports in the Start Value and End Value fields. You can specify a value in the range of 1 through 65,535.
- Select **SINGLE** to configure a single port number as the source port, and enter the number in the Port Value field.
- Select **NA** if you do not want to specify a port number.

The TCP or UDP source and destination port provide additional specification, in conjunction with the network protocol, for packet matching in an application definition. To configure ports, you must define one source or destination port. Normally, you specify this match in conjunction with the protocol match statement to determine which protocol is being used on the port.

15. From the Destination Port type list, do one of the following:

- Select **RANGE** to configure a range of destination ports for the application, and enter the upper limit and lower limit of the range of ports in the Start Value and End Value fields. You can specify a value in the range of 1 through 65,535.



**NOTE:** If you specify a value of 0 as a destination port or beginning of a destination report range, you will receive the following error: `application application-name' TCP Destination Port 0 Invalid error: configuration check-out failed`

- Select **SINGLE** to configure a single port number as the destination port, and enter the number in the Port Value field.
- Select **NA** if you do not want to specify a port number.

16. Click **Save** to save the application.

## Creating an Application Set

You can group the applications you have defined into a named object as an application set.

To create an application set for an SFW rule term:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the task pane, select **Service Templates**.  
The Manage Service Templates page is displayed.
4. Click the **SFW** button.  
The list of SFW service templates is displayed.
5. Click the **Add** icon.  
The Create a SFW Planning Template window appears.
6. In the Name field, enter a name for the service template or profile (limit of 63 alphanumeric characters without spaces).
7. In the Description field, enter a meaningful, easily-identifiable name for the service instance (limit of 255 alphanumeric characters). Each service instance you define can be applied to a single or multiple SDGs.
8. Click the green plus sign in the Applications box.  
The Create an Application dialog box appears.



**NOTE:** For the service elements that you can configure using the Object Builder workspace, such as applications and rules, when you click the green plus sign (+) at the top-right corner of each of the service element boxes, the shortcut menu is displayed. Click the **Create New** radio button to create the service component afresh. Alternatively, click the **Import from Object Builder** radio button to open a dialog box that enables you to select from the list of service elements that are present in the database of Edge Services Director and import them into the service template.

If a green tick mark is shown beside a field in the dialog box, it denotes that you can add attributes for that component. A red cross mark shows that you can delete that particular attribute for that component.

9. In the **Name** field, enter the name to identify the application set.
10. In the **Application** section, the application set selector dialog box is displayed. Select the applications or application sets that need to be added to the rule term in the from the Available column and click the right arrow to move these applications or application sets to the Selected column.
11. Click **Save** to save the application set.

## Creating a Syslog

You can enable system logging. The system log information from the Adaptive Services or Multiservices PIC is passed to the kernel for logging in the /var/log directory. This setting overrides any syslog statement setting included in the service set or interface default configuration.

To create a syslog for the SFW template:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the task pane, select **Service Templates**.  
The Manage Service Templates page is displayed.
4. Click the **SFW** button.  
The list of SFW service templates is displayed.

5. Click the **Add** icon.

The Create an SFW Planning Template window appears.

6. In the Template Name field, enter a name for the service template or profile (limit of 63 alphanumeric characters without spaces).
7. In the Description field, enter a meaningful, easily-identifiable name for the service instance (limit of 255 characters). Each service instance you define can be applied to a single or multiple SDGs.
8. Click the green plus sign in the Server Groups box.

The Addition of Group dialog box appears.



**NOTE:** If a green tick mark is shown beside a field in the dialog box, it denotes that you can add attributes for that component. A red cross mark shows that you can delete that particular attribute for that component.

9. In the Name field, enter the name for the syslog component. Specify the fully qualified domain name or IP address for the syslog server.
10. In the Services list, specify the system logging severity level. It assigns a severity level to the facility. Valid entries include:
  - **alert**—Conditions that should be corrected immediately.
  - **any**—Matches any level.
  - **critical**—Critical conditions.
  - **emergency**—Panic conditions.
  - **error**—Error conditions.
  - **info**—Informational messages.
  - **notice**—Conditions that require special handling.
  - **warning**—Warning messages.
11. From the Facility Override list, select the override for the default facility for system log reporting. Valid values include:

**authorization**

**daemon**

**ftp**

**kernel**

**local0** through **local7**

**user**

12. In the Log Prefix field, set the system logging prefix value for all logging to the system log host.
13. In the Port field, specify the port number to be used for connection with the remote syslog server.
14. In the Class section, set the class of applications to be logged to the system log.
  - **alg-logs**—Log application-level gateway events.
  - **ids-logs**—Log intrusion detection system events.
  - **nat-logs**—Log Network Address Translation events.
  - **packet-logs**—Log general packet-related events.
  - **session-logs**—Log session open and close events.
  - **session-logs open**—Log session open events only.
  - **session-logs close**—Log session close events.
  - **stateful-firewall-logs**—Log stateful firewall events.
15. Click **Save** to save the service template configuration. Else, click **Close** to discard the changes to the template.

## Creating a Rule

To create a rule for the SFW template:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the task pane, select **Service Templates**.  
The Manage Service Templates page is displayed.
4. Click the **SFW** button.  
The list of SFW service templates is displayed.
5. Click the **Add** icon.  
The Create an SFW Planning Template window appears.

6. In the Template Name field, enter a name for the service template or profile (limit of 63 alphanumeric characters without spaces).
7. In the Description field, enter a meaningful, easily-identifiable name for the service instance (limit of 255 characters). Each service instance you define can be applied to a single or multiple SDGs.
8. Click the green plus sign in the Server Groups box. The Addition of Group dialog box appears.



**NOTE:** If a green tick mark is shown beside a field in the dialog box, it denotes that you can add attributes for that component. A red cross mark shows that you can delete that particular attribute for that component.

9. From the **Rule** list, select one of the previously configured rules. The rules that you configured in the Service Templates workspace for SFW, packet filter, or CGNAT are displayed.
10. Click **Save** to save the service template configuration. Else, click **Close** to discard the changes to the template.

## Creating a Rule Set

The rule-set statement defines a collection of stateful firewall rules that determine what actions the router software performs on packets in the data stream. You define each rule by specifying a rule name and configuring terms. Then, you specify the order of the rules by including the **rule-set** statement at the **[edit services stateful-firewall]** hierarchy level with a rule statement for each rule.

The router software processes the rules in the order in which you specify them in the configuration. If a term in a rule matches the packet, the router performs the corresponding action and the rule processing stops. If no term in a rule matches the packet, processing continues to the next rule in the rule set. If none of the rules matches the packet, the packet is dropped by default.

To create a rule set for the SFW template:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the task pane, select **Service Templates**.

The Manage Service Templates page is displayed.

4. Click the **SFW** button.

The list of SFW service templates is displayed.

5. Click the **Add** icon.

The Create an SFW Planning Template window appears.

6. In the Template Name field, enter a name for the service template or profile (limit of 63 alphanumeric characters without spaces).
7. In the Description field, enter a meaningful, easily-identifiable name for the service instance (limit of 255 characters). Each service instance you define can be applied to a single or multiple SDGs.
8. Click the green plus sign in the Rule Sets box.

The Addition of Rule Sets dialog box appears.



**NOTE:** For the service elements that you can configure using the Object Builder workspace, such as applications and rules, when you click the green plus sign (+) at the top-right corner of each of the service element boxes, the shortcut menu is displayed. Click the **Create New** radio button to create the service component afresh. Alternatively, click the **Import from Object Builder** radio button to open a dialog box that enables you to select from the list of service elements that are present in the database of Edge Services Director and import them into the service template.

If a green tick mark is shown beside a field in the dialog box, it denotes that you can add attributes for that component. A red cross mark shows that you can delete that particular attribute for that component.

9. Specify the rule set name the router uses when applying this service.
10. Select the rules that you want to group into a rule set from the Available column and click the right arrow to move the rules to the Selected column.
11. Click **Save** to save the service template configuration. Else, click **Close** to discard the changes to the template.

## Creating a Services PIC for an SFW Service Template

Multiservices (ms-) interfaces are the physical multiservices interfaces of a device that are used to run the load-balancing instance application. The more multiservices interfaces used for a loadbalancing instance, the more capacity and processing power the instance has. At least one MS interface must be specified for each adc-instance, up to eight interfaces can run the same instance. A multiservices interface is associated exclusively to a single load-balancing instance (it cannot be shared between instances).

To assign a services interface to an SFW template:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the task pane, select **Service Templates**.  
The Manage Service Templates page is displayed.
4. Click the **SFW** button.  
The list of SFW service templates is displayed.
5. Click the **Add** icon.  
The Create an SFW Planning Template window appears.
6. In the Template Name field, enter a name for the service template or profile (limit of 63 alphanumeric characters without spaces).
7. In the Description field, enter a meaningful, easily-identifiable name for the service instance (limit of 255 characters). Each service instance you define can be applied to a single or multiple SDGs.
8. Click the green plus sign in the Service Pic box.

The Service Pic dialog box appears.



**NOTE:** If a green tick mark is shown beside a field in the dialog box, it denotes that you can add attributes for that component. A red cross mark shows that you can delete that particular attribute for that component.

9. Select the check box next to the ms- interface of an SDG that must be assigned to the SFW template.
10. Click **OK** to save the settings. Else, click **Cancel** to discard the configuration.

**Related Documentation**

- [Service Templates Overview on page 189](#)
- [Filtering Service Templates on page 189](#)
- [Viewing Service Templates on page 192](#)
- [Using the Actions Menu on the Service Template and Service Edit Pages on page 195](#)

## Creating and Managing TLB Service Templates

Before you configure the traffic load balancer (TLB) software, install the TLB application package on the services PIC used for the server health monitoring function. Once you have installed the application package, you can configure or re-configure TLB as needed. To create a complete application, you must also define interfaces and routing information. You can optionally define firewall filters and policy options in order to differentiate TLB traffic.

You can perform the following tasks with the Service Designer page for TLB:

- Create a TLB service template with attributes and settings for load balancing operations.
- Modify an existing TLB template to meet the network needs and deployment scenarios.
- Delete an existing template.
- [Creating a TLB Service Template on page 276](#)
- [Creating a Deployment Plan on page 279](#)
- [Modifying TLB Service Templates on page 280](#)
- [Importing a TLB Service Template on page 282](#)
- [Creating a Real Server on page 284](#)
- [Creating a Group for Real Servers on page 285](#)
- [Creating a Services PIC for a TLB Service Template on page 288](#)
- [Creating a Network Monitor Profile for a TLB Service Template on page 289](#)
- [Creating a Command for Script-Based Health Checks on page 291](#)
- [Creating a Server Bypass Filter on page 292](#)
- [Creating a Virtual Service for a TLB Service Template on page 293](#)
- [Creating a Client-Facing Interface and Routing Instance on page 296](#)
- [Creating a Server-Facing Interface and Routing Instance on page 298](#)

## Creating a TLB Service Template

To configure a new TLB service template:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the task pane, select **Service Templates**. The Manage Service Templates page is displayed.
4. Click the **TLB** button.

The list of TLB service templates is displayed.

The Service Designer page displays a bar graph in the top pane of the page. The count of service templates of each type is displayed on the vertical axis and the service type is shown on the horizontal axis. A color-coding format is used to represent the bars on the graph. Published service templates are shown in olive green color and unpublished service templates are shown in blue color. Mouse over each bar in the chart to highlight and display the number of templates published or unpublished for each type of service.

5. Click the **Add** icon. The Select Version dialog box is displayed.
6. Select **Junos 12.1** if you want to create a template based on the Junos OS Release 12.1. Alternatively, select **Junos 14.1** if you want to create a template based on the Junos OS Release 14.1.



**NOTE:** All the service template components described in this section can be created for templates that are based on both the Junos OS Releases 12.1 and 14.1. The service elements or components that are additionally available for configuration when you select the Junos OS 14.1 version are explicitly mentioned in the relevant steps of the procedure.

---

The Create a TLB Planning Template window appears.

Figure 26: Create TLB Service Template Window

7. In the Template Name field, enter a name for the service template or profile (limit of 63 alphanumeric characters without spaces).
8. In the Instance Name field, enter a meaningful, easily-identifiable name for the service instance (limit of 255 characters). Each service instance you define can be applied to a single or multiple SDGs.
9. Instead of creating a new template entirely, you can import the parameters defined for a previous TLB service instance and customize only the settings that are necessary. Imported templates are created without any device assigned to them. To use these templates, you must associate a device with the policy. To clone an existing template by importing it, click the **Import** button.

The Import Services dialog box is displayed. See *Importing a TLB Service Template* for step-wise details on importing a TLB service template.

10. The Create a TLB Planning Template window displays the individual elements or components of the service with a graphical icon for each of the service elements and the corresponding names in separate boxes. You can add, edit, or delete these service elements in a template.

The Property View tab and the Config View tab are displayed on the right pane of the template window. The Property View tab provides a tree-based structure of the parameters defined in a service template. You can expand the tree and view details of each component. A key value pair representation is shown. Each of the components can be treated as categories of the service template shown in the property view.

The Config View tab displays the elements or components specified for a service template in the form of configuration stanzas and hierarchy levels. This display is

similar to the show command that you can use at a certain [edit] hierarchy level to view the defined settings. Each level in the hierarchy is indented to indicate each statement's relative position in the hierarchy. Each level is generally set off with braces, with an open brace ( { ) at the beginning of each hierarchy level and a closing brace ( } ) at the end. If the statement at a hierarchy level is empty, the braces are not displayed. Each leaf statement ends with a semicolon ( ; ), as does the last statement in the hierarchy.

- a. Click the green tick mark (✓) displayed at the top-right corner of each of the service element boxes to create a new element. If the green tick mark is not shown, it indicates that the user role does not have the permission to create an element.
- b. Click the red cross mark (x) displayed at the top-right corner of the icons of each element if you want to delete the existing configuration. The user with designer role has permissions to remove or edit elements.
- c. If the red cross mark is not displayed beside a particular icon, it signifies that the element cannot be deleted.
- d. The diamond icon that contains an orange tick mark within it at the top-right corner of the service component name denotes that the particular element can be modified. The absence of this icon denotes that the user does not have permissions to modify the attributes of the service component.
- e. Double-click each icon pertaining to a service element to view or edit its settings. If you do not possess the permission to modify the element, a view-only dialog box with the attributes of the selected element is shown. Otherwise, an editable dialog box enables you to modify the settings.
- f. Click **Save** to save the service template configuration. Else, click **Close** to discard the changes to the template.
- g. Click the Maximize icon displayed at the top-right corner of the rectangle or box that shows all of the values or entities of a particular component of a service template. The specified component or attribute is displayed as a separate dialog box, listing all of the values of the particular component. You can add, modify, or delete the listed values.
- h. While creating the new service template, the designer can add or modify service parameter values and also restrict the access level for each service parameter for the operator. The designer can set following access levels for each service parameters to operator in planning template. Click the new icon (cascading files icon) displayed at the top-left corner of each of the element boxes to open the shortcut menu. You can click one of the following radio buttons:
  - Read-only (the configuration parameter is read-only for operator as part of provisioning)
  - Editable (the configuration parameter is editable as part of provisioning)
  - Device-Specific (the configuration parameter value needs to be entered by the operator for each device during deployment)
- i. Click **Save & Publish** to save and publish the service template configuration. The designer must publish the service templates to the operator to use in the creation

of deployment plans. After a filter or policy is published, it goes for peer review and approval. After approval, the filter or policy is deployed to device.

## Creating a Deployment Plan

You must have previously defined service templates and policy or filter templates before you can create a deployment plan.

To create a deployment plan and assigning devices to it:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the task pane, select **Service Edit**.  
The Manage Service Templates page is displayed.
4. Click the **TLB** button.  
The list of TLB service templates is displayed.
5. Select the check boxes next to the SDGs or SDG groups that you want to assign to the plan. Based on your selection of a service or a policy template, the components or attributes are shown for the corresponding device.
6. From the boxes that show the components of a service template, you can edit, delete, or add elements to it. If you do not have permissions to update a template, the corresponding icons are not shown.
7. Click the down arrow in the **Actions** menu and select **Send for Deployment** to create a deployment plan for the particular service template and save the plan.
  - If you create a deployment plan from Gateway view of Deploy mode, the Deployment Plan Summary dialog box appears, with the service name, type, and status listed.  
Click **Send** to create a deployment plan.
  - If you create a deployment plan from Service view of Deploy mode, the Edit Service Instance page is displayed. You can modify the SDGs associated with the service instance and also modify the service instance attributes as necessary by either clicking the buttons corresponding to the various settings at the top of the wizard page to directly traverse to the page you want to modify or clicking the navigation buttons at the bottom of the wizard page to go to the different pages of the wizard. Click **Finish** to create a deployment plan.

A deploy plan is created for the service template with the devices that are assigned to it when you view the Deployment Plans page.

8. Alternatively, you can select **Discard changes** from the **Actions** menu to ignore the modifications done to a policy or filter template.
9. From the Deployment plans page, you can select **Reject** or **Approve** from the Actions drop-down list to reject or approve the deployment plan and make it available for commissioning to the devices.

## Modifying TLB Service Templates

On the Service Designer page, you can view the collection of service templates defined for several applications, such as stateful firewall or CGNAT.

To modify service template instances, such as ADC, SFW, CGNAT, or TLB templates:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the task pane, select **Deploy Service > Service Edit**.  
The Service Instances page is displayed in the right pane, listing all the previously defined service templates.
4. From the View pane, perform one of the following tasks:
  - Click the **ADC** button.  
The list of ADC service templates is displayed. You need not click this button if you are launching the Service Designer page for the first time or are navigating to this page from another mode or a different page. You need to click this button only if you are viewing the other service templates, such as CGNAT or TLB.
  - Click the **SFW** button.  
The list of SFW templates is displayed.
  - Click the **TLB** button.  
The list of TLB templates is displayed.
  - Click the **CGNAT** button.  
The list of CGNAT templates is displayed.
5. In the main window, click the plus sign (+) next to the SDG pairs to expand the tree and view the pair of devices in the SDG group or pair. Select the check box next to the

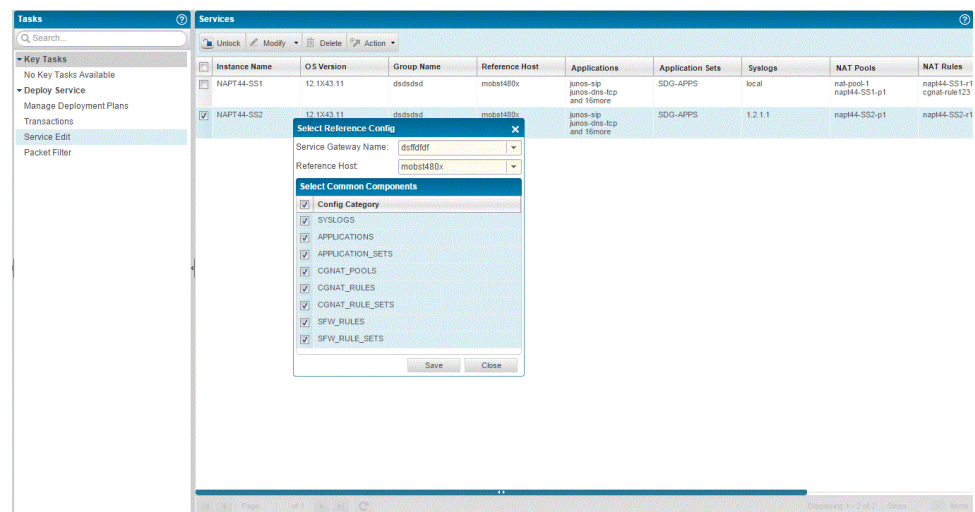
SDG pair or individual SDG for which you want to modify settings. In an SDG pair, you can select a single SDG or both the SDGs in the in the redundancy pair of devices.



**NOTE:** Alternatively, you can also modify service templates from Service View in Build Mode by selecting the **Service Templates > Manage Service Templates** from the task pane, selecting a service instance, and clicking the **Modify** button. You can also modify ADC and TLB service templates from Gateway View in Deploy mode by selecting the SDG pair or SDG from the View pane, selecting **Service Edit** from the task pane, and selecting the TLB service from the main window that displays all the previously configured template instances to lock and modify it.

- Click the **Lock** icon above the table of listed packet filters. The **Select Reference Config** dialog box is displayed.

*Figure 27: Select Reference Config Dialog Box*



- From the Service Gateway Name drop-down list, select the SDG group to which the packet filter must be applied.
- From the Host Name drop-down list, select the hostname of the SDG.
- In the Select Common Components section, select the check boxes beside the service modules or components, such as packet filters, SFW rules, or CGNAT rules, that are displayed. The displayed components depend on the attributes that are previously defined for that selected packet filter. For example, if the service policy is for stateful firewall, SFW rules and SFW rule sets are shown. Select the check box beside Config Category to select all the service components.
- Click **Save** to save the modified association.

11. Select the check box beside the template you want to modify.
12. Open the **Modify** menu above the list of templates to modify an existing template, and select the component or service attribute, such as application or rule, that you want to edit.
13. Perform one of the following from the drop-down menu displayed for each component:
  - To retrieve the service component and import into the database of Edge Services Director, select **Import Object**. The Import Services dialog box appears. You can import the service templates assigned to SDGs or choose from a list of all of the predefined templates in the database. Also, you can either import all of the components of a service or specific components.
  - To create the component afresh, select **Create New**. The Create page corresponding to the service component appears. You can define the attributes for the service component in the same manner as you define the elements during the creation of a service template.

## Importing a TLB Service Template

To create a clone of an existing TLB template by importing it:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.

The functionalities that you can configure in this mode are displayed in the task pane.
3. From the task pane, select **Service Templates**.

The Manage Service Templates page is displayed.
4. Click the **TLB** button.

The list of TLB service templates is displayed. You need not click this button if you are launching the Service Designer page for the first time or are navigating to this page from another mode or a different page. You need to click this button only if you are viewing the other service templates, such as CGNAT or TLB.
5. Click the **Add** icon.

The Create a TLB Planning Template window appears.
6. In the Template Name field, enter a name for the service template or profile (limit of 63 alphanumeric characters without spaces).

7. In the Instance Name field, enter a meaningful, easily-identifiable name for the service instance (limit of 255 characters). Each service instance you define can be applied to a single or multiple SDGs.

8. Click the **Import** button. The Import Services dialog box appears.

You can import the service templates assigned to SDGs or choose from a list of all of the predefined templates in the database. Also, you can either import all of the components of a service or specific components.

9. Do one of the following for the Import section:

- Select the **From Existing Service Gateway** radio button if you want to import the CGNAT rule from SDGs that are present in the Edge Services Director database.
- Select the **From XML** radio button if you want to import the CGNAT rule from an XML configuration file on an external system.

10. If you selected the option to import the object from SDGs, do the following:

- Click the **Normal View** tab to view the list of SDGs. You can search for specific SDGs by entering a search item and clicking the **Search** icon.

Alternatively, click the **Group View** tab to view the list of SDG groups. You can search for specific SDG groups by entering a search item and clicking the **Search** icon.

- Click the plus sign (+) next to the All Service Gateways item to expand the tree structure that displays the list of SDGs or SDG groups. If the SDG pair is configured, you can select one of the devices, master or standby, from which you want to import the object.

Alternatively, if you selected the **Group View** tab, you can select an SDG from the groups displayed from which you want to import the object.

- Click **Import**. The object is added to the database and can be used during configuration of services or policies.

11. If you selected the option to import from an XML file, do the following:

- Click **Browse** beside the File Name field to navigate to the path where an XML file is available to be imported.
- Click **Upload**. The service template is added to the database and can be used during configuration of services or policies.

12. Do one of the following to import all components of a selected template or only a particular component of a template. For the components that are not imported, you need to specify the definitions of the components afresh.

- Select the check boxes next to all of the service instances that are displayed for the selected SDG or SDG group, or for the XML file that you uploaded. In such a case, all of the elements or parameters of the selected template or instance are imported.

- Alternatively, select the check box next to a particular or group of service instances to import only a specific component of the selected template
13. Similarly, you can select other components and import them to the template. Save the imported components to add them to the template you are creating by using the imported template as a base.

## Creating a Real Server

To create a real server as a component for the TLB template:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the task pane, select **Service Templates**.  
The Manage Service Templates page is displayed.
4. Click the **TLB** button.  
The list of TLB service templates is displayed.
5. Click the **Add** icon.  
The Create a TLB Planning Template window appears.
6. In the Template Name field, enter a name for the service template or profile (limit of 63 alphanumeric characters without spaces).
7. In the Instance Name field, enter a meaningful, easily-identifiable name for the service instance (limit of 255 characters). Each service instance you define can be applied to a single or multiple SDGs.
8. Click the green plus sign in the Real Servers box. The Addition of Real Server dialog box appears.



**NOTE:** For the service elements that you can configure using the Object Builder workspace, such as applications and rules, when you click the green plus sign (+) at the top-right corner of each of the service element boxes, the shortcut menu is displayed. Click the **Create New** radio button to create the service component afresh. Alternatively, click the **Import from Object Builder** radio button to open a dialog box that enables you to select from the list of service elements that are present in the database of Edge Services Director and import them into the service template.

If a green tick mark is shown beside a field in the dialog box, it denotes that you can add attributes for that component. A red cross mark shows that you can delete that particular attribute for that component.

9. In the Name field, enter the name to identify the real server. Make sure the servers are connected via a router interface that is defined as a server-facing interface for the adc-instance. For each real server, you must assign a real-server name and specify its actual IP address.
10. In the Address Family field, select **IPv4** to specify an IPv4 address, or select **IPv6** to enter the IPv6 address of the real server.
11. In the IP Address field, specify the IP address of the real server.
12. Click **Save** to save the service template configuration. Else, click **Close** to discard the changes to the template.

## Creating a Group for Real Servers

Define the group and assign real servers to it. The real servers in any given group must have an IP address accessible to the module that performs the SLB functions. This IP routing is most easily accomplished by placing the servers on a network local to the router. Routing to the server can be used as long as it does not violate the topology rules outlined.

A group is a collection of multiple servers with the same content, so that client requests can be load-balanced between them.

To create a group of real servers:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.

The functionalities that you can configure in this mode are displayed in the task pane.

3. From the task pane, select **Service Templates**.

The Manage Service Templates page is displayed.

4. Click the **TLB** button.

The list of TLB service templates is displayed.

5. Click the **Add** icon.

The Create a TLB Planning Template window appears.

6. In the Template Name field, enter a name for the service template or profile (limit of 63 alphanumeric characters without spaces).

7. In the Instance Name field, enter a meaningful, easily-identifiable name for the service instance (limit of 255 characters). Each service instance you define can be applied to a single or multiple SDGs.

8. Click the green plus sign in the Server Groups box. The Addition of Group dialog box appears.



**NOTE:** For the service elements that you can configure using the Object Builder workspace, such as applications and rules, when you click the green plus sign (+) at the top-right corner of each of the service element boxes, the shortcut menu is displayed. Click the **Create New** radio button to create the service component afresh. Alternatively, click the **Import from Object Builder** radio button to open a dialog box that enables you to select from the list of service elements that are present in the database of Edge Services Director and import them into the service template.

If a green tick mark is shown beside a field in the dialog box, it denotes that you can add attributes for that component. A red cross mark shows that you can delete that particular attribute for that component.

---

9. In the Name field, enter the name for the real servers group.

10. Do the following in the Routing Instance section:
  - a. Select the **Routing Instance Selection** check box to configure a routing instance for TLB to steer traffic.
  - b. Click the green plus sign next to the Routing Instance field. The Routing Instances dialog box appears.
  - c. From the Service Gateway Name field, select the SDG group with which the service element must be associated.
  - d. From the Host Name field, select the SDG in the SDG high-availability pair of active and standby SDGs.
  - e. In the MS Interfaces section, select the check box next to the routing instance of the SDG that must be used for packets arriving from clients or users. All the routing instances from the inventory of devices are listed.
11. Select the **Real service rejoin options** check box to allow a server to rejoin the group automatically when it comes up. When a previously down server is returned to service, all flows belonging to that server based on hashing return to it, impacting performance for the returned flows. For this reason, the automatic rejoining of a server to an active group can be disabled.
12. From the Health Check Interface Sub Unit list, select the subunit to be used for health monitoring. Select the number of the unit to edit. A health-check source address must be set for each unit on which real servers are configured, in order to allow sending health checks to the servers. This field is applicable only for Junos OS 14.1 version.
13. From the Real Server IP Type field, select **IPv4** or **IPv6** to configure an IPv4 or IPv6 addresses for real servers.
14. In the Real Servers section, assign the real servers to be part of the group. Select the real servers from the Available column and click the right arrow to move the server to the Selected column.
15. In the Network Monitoring Profiles section, select the profile from the Available column and click the right arrow to move the profile to the Selected column.
16. Click **Save** to save the service template configuration. Else, click **Close** to discard the changes to the template.

## Creating a Services PIC for a TLB Service Template

Multiservices (ms-) interfaces are the physical multiservices interfaces of a device that are used to run the load-balancing instance application. The more multiservices interfaces used for a loadbalancing instance, the more capacity and processing power the instance has. At least one MS interface must be specified for each adc-instance, up to eight interfaces can run the same instance. A multiservices interface is associated exclusively to a single load-balancing instance (it cannot be shared between instances).

To assign a services interface to a TLB template:

1. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.

2. From the task pane, select **Service Templates**.  
The Manage Service Templates page is displayed.

3. Click the **TLB** button.  
The list of TLB service templates is displayed.

4. Click the **Add** icon.  
The Create a TLB Planning Template window appears.

5. In the Template Name field, enter a name for the service template or profile (limit of 63 alphanumeric characters without spaces).
6. In the Instance Name field, enter a meaningful, easily-identifiable name for the service instance (limit of 255 characters). Each service instance you define can be applied to a single or multiple SDGs.

7. Click the green plus sign in the Service Pics box.  
The Service Pic dialog box appears.



**NOTE:** If a green tick mark is shown beside a field in the dialog box, it denotes that you can add attributes for that component. A red cross mark shows that you can delete that particular attribute for that component.

8. From the Service Gateway Name field, select the SDG group with which the service element must be associated.
9. From the Host Name field, select the SDG in the SDG high-availability pair of active and standby SDGs.

10. Select the check box next to the ms- interface of an SDG that must be assigned to the TLB template.
11. Click **OK** to save the settings. Else, click **Cancel** to discard the configuration.

## Creating a Network Monitor Profile for a TLB Service Template

To configure a network monitor profile to perform health and welfare validation of servers for a TLB template:

1. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
2. From the task pane, select **Service Templates**.  
The Manage Service Templates page is displayed.
3. Click the **TLB** button.  
The list of TLB service templates is displayed.
4. Click the **Add** icon.  
The Create a TLB Planning Template window appears.
5. In the Template Name field, enter a name for the service template or profile (limit of 63 alphanumeric characters without spaces).
6. In the Instance Name field, enter a meaningful, easily-identifiable name for the service instance (limit of 255 characters). Each service instance you define can be applied to a single or multiple SDGs.
7. Click the green plus sign in the Network Monitor Profile box. The Addition of Network Monitor Profile dialog box appears.



**NOTE:** If a green tick mark is shown beside a field in the dialog box, it denotes that you can add attributes for that component. A red cross mark shows that you can delete that particular attribute for that component.

8. In the Name field, enter the name of the network monitor profile used to monitor the health of servers in the group.
9. In the Probe Interval field, specify the amount of time, in seconds, between polls of the real server by the router.



**NOTE:** The ADC software monitors the servers in the real-server group and the load-balanced applications running on them. If a router detects that a server or application has failed, it will not direct any new connection requests to that server. When a service fails, the ADC software can remove the individual service from the load-balancing algorithm without affecting other services provided by that server. By default, the router checks the status of each service on each real server every five (5) seconds. Sometimes, the real server can be too busy processing connections to respond to health checks. If a service does not respond to four consecutive health checks, the router, by default, declares the service unavailable. You can modify both the health check interval and the number of retries.

10. In the Failure Retries field, specify the number of times the router will attempt its check on the real server before marking the server as unavailable.
11. In the Recover Retries field, specify the number of times the router will attempt to recover the real-server connection.
12. In the **TCP Choices** drop-down list, select one of the supported health checking protocols, such as TCP, HTTP, or ICMP.
13. In the TCP Choices section, do one of the following:
  - a. Select the **HTTP** radio button to select HTTP for health checks. Specify the name of the host, HTTP method such as PUT, GET, OPTIONS, or POST, the URL for which health check needs to be performed, and the port to be used for server health monitoring.
  - b. Select the **ICMP** radio button to select ICMP for health check probes.
  - c. Select the **TCP** radio button to select TCP for health check probes. Specify the port number to be used for monitoring the health and welfare of the server or URL using the SSL-based health probes in the Port field. You can specify this value only if you create the TLB service template based on the Junos OS 14.1 version.
  - d. Select the **SSLHELLO** radio button to sets Secure Sockets Layer (SSL) hello health-check parameters. SSL version 2 (SSLv2) is used for the SSL health check. Specify the port number to be used for monitoring the health and welfare of the server or URL using the SSL-based health probes in the Port field. You can specify this the SSL-hello health check setting only if you create the TLB service template based on the Junos OS 14.1 version.
  - e. Select the **CUSTOM** radio button to create a custom-based health check. From the Protocol field, specify **tcp** or **udp** as the protocol for the script to use in a custom

health check. A script is made up of one or more TCP or UDP command containers. A script can contain any number of these containers, up to the allowable number of characters that a script supports.

In the Command ID field, specify the command ID for the commands to be used. Multiple command lines are usually required in order to specify a full script.

In the Port field, specify the port number to be used for custom-based health check mechanism.

Health check scripts dynamically verify application and content availability by executing a sequence of tests based on send and expect commands. See the *Creating a Command for Script-Based Health Checks* section for detailed information.

14. Click **Save** to save the settings. Else, click **Cancel** to discard the configuration.

## Creating a Command for Script-Based Health Checks

You can create commands for building a script-based health check. You can configure this service element only if you create a TLB service template using the Junos OS 14.1 version.

To create a custom network monitoring profile command for script-based health checks.

1. In the Create Networking Profile dialog box, select the check box next to the SEND or EXPECT row under the Command column of the table.
2. Click the pencil icon to specify the command attributes. The Create Custom Networking Profile Command dialog box appears.
3. If you selected the SEND type, it is displayed in the Command Type field.
4. In the Send Type list, perform either of the following:
  - Select **BINARY** to specify binary content (in hexadecimal format) for the request packet.
  - Select **ASCII** to specify ASCII content (in hexadecimal format) for the request packet.
5. In the Value field, specify the content to be sent in raw hexadecimal format or the binary content to send using raw hexadecimal format for the request packet.
6. If you selected the EXPECT type, it is displayed in the Command Type field.
7. In the Send Type list, perform either of the following:
  - Select **BINARY** to specify binary content (in hexadecimal format) to be expected from the server response packet.
  - Select **ASCII** to specify ASCII content (in hexadecimal format) to be expected from the server response packet.

8. In the Value field, specify the content to be returned in the server response packet in raw hexadecimal format or the binary content to receive using raw hexadecimal format for the response packet.
9. For binary content only, in the Offset field, specify the offset from the beginning of the binary data area to start matching the content specified in the binary-expect command. The offset command is supported for both UDP and TCP-based health checks. If this value is not present, an offset of zero is assumed.
10. For binary content only, in the Length field, specify the number of bytes in the IP packet that should be examined. If no offset value is specified, depth is specified from the beginning of the packet. When depth is not specified, it is the length of the content. This means that the content is expected exactly at the offset specified (or 0 when the offset is not specified).
11. Click Save to save the custom network monitor profile configuration. Else, click Close to discard the changes to the custom health check profile.

## Creating a Server Bypass Filter

You can configure this service element only if you create a TLB service template using the Junos OS 14.1 version.

To configure a virtual service for a TLB template:

1. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
2. From the task pane, select **Service Templates**.  
The Manage Service Templates page is displayed.
3. Click the **TLB** button.  
The list of TLB service templates is displayed.
4. Click the **Add** icon.  
The Create a TLB Planning Template window appears.
5. In the Template Name field, enter a name for the service template or profile (limit of 63 alphanumeric characters without spaces).
6. In the Instance Name field, enter a meaningful, easily-identifiable name for the service instance (limit of 255 characters). Each service instance you define can be applied to a single or multiple SDGs.

7. Click the green plus sign in the Server Bypass Filters box. The Create Server Bypass Filter dialog box appears.



**NOTE:** If a green tick mark is shown beside a field in the dialog box, it denotes that you can add attributes for that component. A red cross mark shows that you can delete that particular attribute for that component.

8. From the Service Gateway Name field, select the SDG group with which the service element must be associated.
9. From the Host Name field, select the SDG in the SDG high-availability pair of active and standby SDGs.
10. From the table, select the check boxes beside the filters to specify the filters used to bypass rephrase as health-check traffic from real servers.
11. Click **Save** to save the settings. Else, click **Cancel** to discard the configuration.

## Creating a Virtual Service for a TLB Service Template

The virtual service provides an address that is associated with a the group of servers to which traffic is directed as determined by hash-based session distribution and server health monitoring. You may optionally specify filters and routing instances to steer traffic for TLB.

The virtual service configuration identifies:

- The group of servers to which sessions are distributed
- The session distribution hashing method

TLB doesn't require a specific virtual IP. VIPs 0.0.0.0 or 0::0 are acceptable.

To configure a virtual service for a TLB template:

1. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
2. From the task pane, select **Service Templates**.  
The Manage Service Templates page is displayed.
3. Click the **TLB** button.  
The list of TLB service templates is displayed.
4. Click the **Add** icon.

The Create a TLB Planning Template window appears.

5. In the Template Name field, enter a name for the service template or profile (limit of 63 alphanumeric characters without spaces).
6. In the Instance Name field, enter a meaningful, easily-identifiable name for the service instance (limit of 255 characters). Each service instance you define can be applied to a single or multiple SDGs.
7. Click the green plus sign in the Virtual Service box. The Addition of Virtual Service dialog box appears.



**NOTE:** If a green tick mark is shown beside a field in the dialog box, it denotes that you can add attributes for that component. A red cross mark shows that you can delete that particular attribute for that component.

8. In the Name field, specify the name of the virtual service.
9. In the Address field, specify a non-zero address for the virtual service.
10. From the Mode field, select one of the following:
  - **translated**—In complex network topologies, the TLB software functions can be managed using a client Network Address Translation (NAT) IP address on the server-facing interfaces traffic. When the client requests services from the TLB software virtual server, the client sends its own IP address for use as a return address. If a NAT IP address is configured for the Multiservices-DPC NPU, the TLB software replaces the client's source IP address with the TLB software NAT IP address before sending the request to the real server. This process is called client NAT. The real server uses the NAT IP address as the destination address for any response. Load-balancing traffic is forced to return through the TLB software and through the same Multiservices-DPC NPU, regardless of alternate paths. Once the TLB software receives the translated IP address, it puts the original client IP address into the destination address and sends the packet to the client. This process is transparent to the client.
  - **direct-server-return**—Direct Server Return health checks are used to verify the existence of a server-provided service where the server replies directly back to the client without responding through the virtual-server IP address. In this configuration, the server is configured with a real-server IP address and virtualserver IP address. The virtual-server IP address is configured to be the same address as your virtual-server IP address. When Direct Server Return health checks are used, the specified health check is sent originating from the configured health check address. It is destined for the virtualserver IP address with the MAC address that was acquired from the real-server IP Address Resolution Protocol (ARP) entry. Direct Server Return is configured at the group level. If a group is configured with

“direct-server-return” the health check performed is sent to the virtual IP and not to the actual server IPs. The TLB software lets you to perform health checks for Direct Server Return configurations (for more information, see Direct Server Return). The router is able to verify that the server correctly responds to requests made to the virtual-server IP address, as required in Direct Server Return configurations. To perform this function, the real-server IP address is replaced with the virtualserver IP address in the health check packets that are forwarded to the real servers for health checking. With this feature enabled, the health check will fail if the real server is not properly configured with the virtual-server IP address.

- **layer2-direct-server-return**—Use transparent mode processing with Layer 2 direct server return (DSR). Some clients may need the Direct Server Return (DSR) feature, which allows the server to respond directly to the client. This capability is useful for sites where large amounts of data flows from servers to clients, such as with content providers or portal sites that typically have asymmetric traffic patterns. DSR and content-intelligent Layer 7 routing cannot be performed at the same time because content intelligent routing requires that all frames return to the router for connection splicing. DSR requires that the server be set up to receive frames that have a destination IP address that is equal to the virtual-server IP address.
11. From the Group list, select the name of a real server group configured to be used for this virtual service.
  12. Select the **Routing Instance Selection** check box to specify a routing instance to be used for this application type of virtual service.
  13. Do the following in the Routing Instance section:
    - a. Click the green plus sign next to the Routing Instance field. The Routing Instances dialog box appears.
    - b. From the Service Gateway Name field, select the SDG group with which the service element must be associated.
    - c. From the Host Name field, select the SDG in the SDG high-availability pair of active and standby SDGs.
    - d. In the MS Interfaces section, select the check box next to the routing instance of the SDG that must be used for packets arriving from clients or users. All the routing instances from the inventory of devices are listed.
  14. In the Rebalance threshold field, specify the limit for rebalancing of traffic. This field is applicable only for Junos OS 12.1 version.
  15. In the Route metric field, specify a routing metric for the virtual service. This field is applicable only for Junos OS 12.1 version.

16. In the Server Protocol section, do the following. This section and the associated fields are applicable only for Junos OS 14.1 version.
  - In the Name field, specify a service name to denote the translated mode details for the specified service. Packets destined to this virtual ip-address + virtual-port + protocol are load balanced to the appropriate server. The destination IP address and port are replaced by the real services IP address and the server-listening-port (configured here).
  - In the Virtual Port field, specify the virtual port number for the virtual service.
  - In the Server Listening Port field, specify the port number the server uses to listen or receive connection requests. The range is from 0 through 65,534. You can change the destination port of traffic to a specific port by using this field setting.
  - From the Protocol list, select **TCP** or **UDP** to specify the protocol type of virtual service.
17. From the Server Interfaces section, select the interfaces from the Available column and click the right arrow to move the hash method to the Selected column.
18. From the **Load balance method** list, select the hash method used for enhanced ECMP load balancing from the Available column and click the right arrow to move the hash method to the Selected column. You can specify **source-ip**, **destination-ip**, or **protocol**
  - destination-ip—Hash based on destination IP address.
  - protocol—Hash based on protocol.
  - source-ip—Hash based on source IP address.
19. Click **Save** to save the settings. Else, click **Cancel** to discard the configuration.

## Creating a Client-Facing Interface and Routing Instance

You can configure this service element only if you create a TLB service template using the Junos OS 14.1 version.

Clients and servers can be connected through the same router port. Each port in use on the router can be configured to process client requests, server traffic, or both:

Client-facing interfaces—Router ports through which client requests to the virtual server are received.

Server-facing interfaces—Router ports to which servers are connected (directly or through routing). Responses to clients are received on the router through these ports.

To assign a client-facing instance and interface to an ADC template:



**NOTE:** Starting with Edge Services Director Release 1.1, you can specify multiple client-facing and server-facing virtual routing and forwarding (VRF) instances when you create or modify a TLB service template that is based on Junos OS Release 14.1. You can select the check boxes beside multiple routing instances in the Create Client Facing and Server Facing dialog boxes that you can open from the Create TLB Service Template window in Gateway View of Build mode. You can also associate multiple client-facing and server-facing VRF instances from the enhanced service edit mode (which you can access from Service View of Deploy mode, with TLB selected in View pane and Service Edit selected in the Tasks pane, and selecting the check boxes beside the Server-Facing and Client-Facing modules in the Select Common Components section).

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the task pane, select **Service Templates**.  
The Manage Service Templates page is displayed.
4. Click the **TLB** button.  
The list of ADC service templates is displayed.
5. Click the **Add** icon.  
The Create an ADC Planning Template window appears.
6. In the Template Name field, enter a name for the service template or profile (limit of 63 alphanumeric characters without spaces).

7. In the Instance Name field, enter a meaningful, easily-identifiable name for the service instance (limit of 255 characters). Each service instance you define can be applied to a single or multiple SDGs.
8. Click the green plus sign in the Client-Facing box. The Client facing dialog box appears.



**NOTE:** If a green tick mark is shown beside a field in the dialog box, it denotes that you can add attributes for that component. A red cross mark shows that you can delete that particular attribute for that component.

9. From the Service Gateway Name field, select the SDG group with which the service element must be associated.
10. From the Host Name field, select the SDG in the SDG high-availability pair of active and standby SDGs.
11. In the Routing Instances section, select the check box next to the routing instance of the SDG that must be used for packets arriving from clients or users. All the routing instances from the inventory of devices are listed.
12. In the Interfaces section, select the check box next to the interface instance of the SDG that must be used for packets arriving from clients or users. All of the interfaces from the inventory of devices are listed.
13. Click **OK** to save the settings. Else, click **Cancel** to discard the configuration.

## Creating a Server-Facing Interface and Routing Instance

You can configure this service element only if you create a TLB service template using the Junos OS 14.1 version.

Clients and servers can be connected through the same router port. Each port in use on the router can be configured to process client requests, server traffic, or both:

Client-facing interfaces—Router ports through which client requests to the virtual server are received.

Server-facing interfaces—Router ports to which servers are connected (directly or through routing). Responses to clients are received on the router through these ports.

To assign a server-facing instance and interface to an ADC template:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the task pane, select **Service Templates**.  
The Manage Service Templates page is displayed.
4. Click the **TLB** button.  
The list of TLB service templates is displayed.
5. Click the **Add** icon.  
The Create an ADC Planning Template window appears.
6. In the Template Name field, enter a name for the service template or profile (limit of 63 alphanumeric characters without spaces).
7. In the Instance Name field, enter a meaningful, easily-identifiable name for the service instance (limit of 255 characters). Each service instance you define can be applied to a single or multiple SDGs.
8. Click the green plus sign in the Server-Facing box. The Server facing dialog box appears.



**NOTE:** If a green tick mark is shown beside a field in the dialog box, it denotes that you can add attributes for that component. A red cross mark shows that you can delete that particular attribute for that component.

9. From the Service Gateway Name field, select the SDG group with which the service element must be associated.

10. From the Host Name field, select the SDG in the SDG high-availability pair of active and standby SDGs.
11. In the Device Inventory Routing Instances section, select the check box next to the routing instance of the SDG that must be used for packets traversing to the servers. All the routing instances from the inventory of devices are listed.
12. In the Device Inventory Interfaces section, select the check box next to the interface instance of the SDG that must be used for packets to be sent to the servers. All of the interfaces from the inventory of devices are listed.
13. Click **OK** to save the settings. Else, click **Cancel** to discard the configuration.

**Related Documentation**

- [Service Templates Overview on page 189](#)
- [Filtering Service Templates on page 189](#)
- [Viewing Service Templates on page 192](#)
- [Using the Actions Menu on the Service Template and Service Edit Pages on page 195](#)

---

## Modifying Individual Service Instances and Deploying to Devices

You can modify individual service instances, such as ADC, TLB, CGNAT, or SFW services, and create a deployment plan for such services using the Service Edit option in task pane in Gateway View of Deploy mode.

- [Modifying Service Instances on page 300](#)
- [Creating a Deployment Plan on page 302](#)

### Modifying Service Instances

On the Service Designer page, you can view the collection of service templates defined for several applications, such as stateful firewall or CGNAT.

To modify service template instances, such as ADC, SFW, CGNAT, or TLB templates:

1. From the View selector, select **Gateway View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the task pane, select **Deploy Service > Service Edit**.

The Service Instances page is displayed in the right pane, listing all the previously defined service templates.

4. From the View pane, perform one of the following tasks:

- Click the **ADC** button.

The list of ADC service templates is displayed. You need not click this button if you are launching the Service Designer page for the first time or are navigating to this page from another mode or a different page. You need to click this button only if you are viewing the other service templates, such as CGNAT or TLB.

- Click the **SFW** button.

The list of SFW templates is displayed.

- Click the **TLB** button.

The list of TLB templates is displayed.

- Click the **CGNAT** button.

The list of CGNAT templates is displayed.

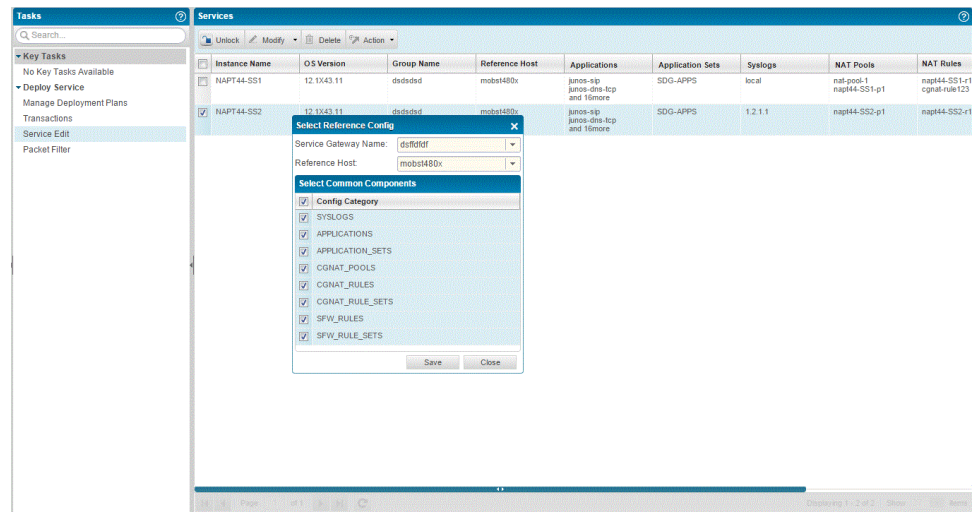
5. In the main window, click the plus sign (+) next to the SDG pairs to expand the tree and view the pair of devices in the SDG group or pair. Select the check box next to the SDG pair or individual SDG for which you want to modify settings. In an SDG pair, you can select a single SDG or both the SDGs in the in the redundancy pair of devices.



**NOTE:** Alternatively, you can also modify service templates from Service View in Build Mode by selecting the **Service Templates > Manage Service Templates** from the task pane, selecting a service instance, and clicking the **Modify** button.

6. Click the **Lock** icon above the table of listed packet filters. The Select Reference Config dialog box is displayed.

**Figure 28: Select Reference Config Dialog Box**



7. Open the **Modify** menu above the list of templates to modify an existing template, and select the component or service attribute, such as application or rule, that you want to edit.
8. Modify the service attributes, as needed, and save the changes.

## Creating a Deployment Plan

You must have previously defined service templates and policy or filter templates before you can create a deployment plan.

To create a deployment plan and assigning devices to it:

1. From the View selector, select **Gateway View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.

3. From the task pane, select **Service Edit**.

The Manage Service Templates page is displayed.

4. From the View pane, perform one of the following tasks:

- Click the **ADC** button.

The list of ADC service templates is displayed. You need not click this button if you are launching the Service Designer page for the first time or are navigating to this page from another mode or a different page. You need to click this button only if you are viewing the other service templates, such as CGNAT or TLB.

- Click the **SFW** button.

The list of SFW templates is displayed.

- Click the **TLB** button.

The list of TLB templates is displayed.

- Click the **CGNAT** button.

The list of CGNAT templates is displayed.

5. In the main window, click the plus sign (+) next to the SDG pairs to expand the tree and view the pair of devices in the SDG group or pair. Select the check box next to the SDG pair or individual SDG for which you want to modify settings. In an SDG pair, you can select a single SDG or both the SDGs in the in the redundancy pair of devices.
6. Click the down arrow in the **Actions** menu and select **Send for Deployment** to create a deployment plan for the particular service template and save the plan.

The Deployment Plan Summary dialog box appears, with the service name, type, and status listed.

Click **Send** to create a deployment plan.

A deploy plan is created for the service template with the devices that are assigned to it when you view the Deployment Plans page.

7. Alternatively, you can select **Discard changes** from the **Actions** menu to ignore the modifications done to a policy or filter template.
8. From the Deployment plans page, you can select **Reject** or **Approve** from the Actions drop-down list to reject or approve the deployment plan and make it available for commissioning to the devices.

**Related  
Documentation**

- [Service Templates Overview on page 189](#)
- [Filtering Service Templates on page 189](#)
- [Viewing Service Templates on page 192](#)
- [Using the Actions Menu on the Service Template and Service Edit Pages on page 195](#)



## CHAPTER 14

# Using the Object Builder

- [Understanding the Object Builder on page 305](#)
- [Importing All Types of Objects on page 306](#)
- [Importing SFW Rule Sets on page 308](#)
- [Importing SFW Rules on page 310](#)
- [Importing Real Server Settings on page 312](#)
- [Importing CGNAT Rule Sets on page 313](#)
- [Importing CGNAT Rules on page 315](#)
- [Importing CGNAT Pools on page 316](#)
- [Importing Applications on page 318](#)
- [Importing Application Sets on page 319](#)

### Understanding the Object Builder

---

The objects are the constituents or building blocks that are used to create service definitions and policy or filter templates. You can use the Object Builder page to retrieve and transfer the objects or components that have been previously created on the SDGs or devices. The objects might reside on the managed SDGs or SDG groups if they were defined using the appropriate configuration statements and parameters in the Junos CLI interface of the respective SDGs. Such a mechanism of importing object settings enables you to easily, quickly, and optimally use the object definitions when you create service and policy templates.

The objects that you can import from SDGs to the database of Edge Services Director comprise the following:

- Real servers
- CGNAT rules and rule sets
- CGNAT pools
- Applications and application sets
- SFW rules and rule sets
- Addresses and prefixes for CGNAT services

For example, if you have created NAT pools on an SDG device and import those objects into the Junos Space database, you can seamlessly import the pool settings during the creation of a CGNAT service or a CGNAT policy and filter template. Also, you can use the same object settings across multiple services and policies. For example, if you have imported an application into Edge Services Director, you can use the application for different services such as ADC or TLB.

You can import objects into the network management application database using two methods. One method is to import the configuration attributes and settings directly from the devices, and the other method is to import XML files that contain the configurations. The Junos XML API is an XML representation of Junos configuration statements and operational mode commands. Junos XML configuration tag elements are the content to which the Junos XML protocol operations apply. Junos XML operational tag elements are equivalent in function to operational mode commands in the CLI, which administrators use to retrieve status information for a device. With both these techniques, you can quickly obtain the objects from devices and propagate them to Edge Service Director.

**Related Documentation**

- [Importing All Types of Objects on page 306](#)

---

## Importing All Types of Objects

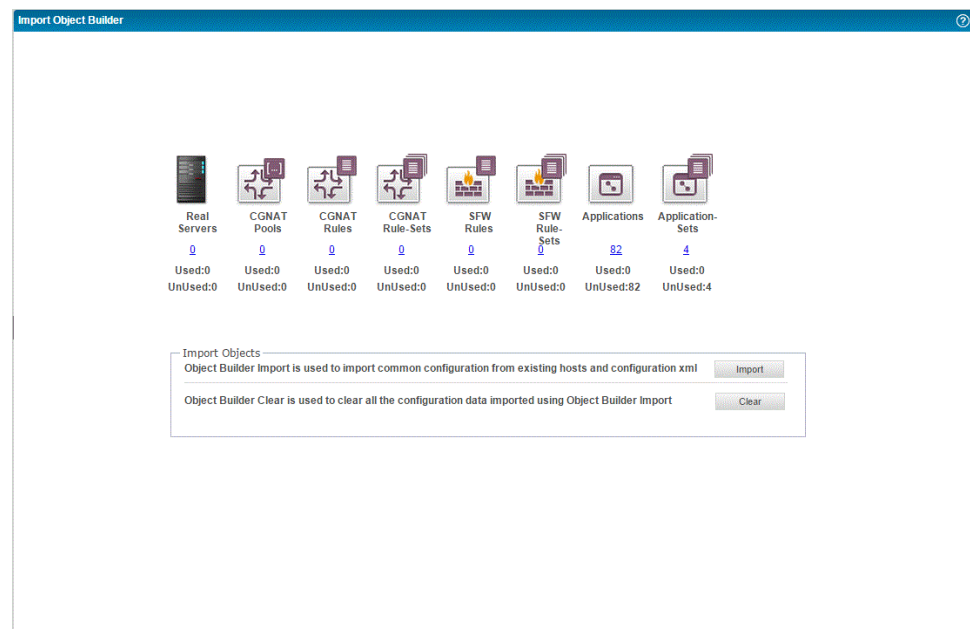
---

Although you can import objects individually based on the services or applications you are using in your deployment, you can also retrieve and add all of the object types that are supported for different services in a single, one-step operation. You can select an SDG from which you want to import all of the objects contained in it. The supported or applicable objects of CGNAT pools, CGNAT rules, CGNAT rule sets, SFW rules, SFW rule sets, applications, application sets, and real servers can be imported in a bulk manner from a device. Similarly, you can also select an XML file that contains a collection of such objects and import all object definitions to the Edge Services Director database.

To import all types of objects in a single operation:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. Select **Object Builder** from the task pane.  
The Object Builder page is displayed.

Figure 29: Object Builder Page



The Object Builder window displays the individual elements or components with a graphical icon for each of the object elements and the corresponding names in separate boxes. Beneath each of the icons that signify the object types, the number of objects of each type already imported is also displayed.

4. Click the **Import** button.

The Add to Object Builder dialog box is displayed.

5. Do one of the following for the Import section:

- Select the **From Existing Service Gateway** radio button if you want to import the real server from SDGs that are present in the Edge Services Director database.
- Select the **From XML** radio button if you want to import the real server from an XML configuration file on an external system.

6. If you selected the option to import the object from SDGs, do the following:

- Click the **Normal View** tab to view the list of SDGs. You can search for specific SDGs by entering a search item and clicking the **Search** icon.

Alternatively, click the **Group View** tab to view the list of SDG groups. You can search for specific SDG groups by entering a search item and clicking the **Search** icon.

- Click the plus sign (+) next to the All Service Gateways item to expand the tree structure that displays the list of SDGs or SDG groups. If the SDG pair is configured, you can select one of the devices, master or standby, from which you want to import the object.

Alternatively, if you selected the **Group View** tab, you can select an SDG from the groups displayed from which you want to import the object.

- Click **Import**. The object is added to the database and can be used during configuration of services or policies.
7. If you selected the option to import from an XML file, do the following:
    - Click **Browse** beside the File Name field to navigate to the path where an XML file is available to be imported.
    - Click **Import**. The object is added to the database and can be used during configuration of services or policies.
  8. Click the **Clear** button at the bottom of the Object Builder page to delete all the object definitions imported from SDGs to the database of Edge Services Director. You are prompted to confirm the deletion. Click **OK** to confirm.
  9. Click the links beneath the graphical icon of each of the configured object elements to navigate directly to the Import page of that corresponding object.

**Related  
Documentation**

- [Understanding the Object Builder on page 305](#)

---

## Importing SFW Rule Sets

---

The **rule-set** statement defines a collection of SFW rules that determine what actions the router software performs on packets in the data stream. You define each rule by specifying a rule name and configuring terms. Then, you specify the order of the rules by including the **rule-set** statement at the **[edit services srareful-firewall]** hierarchy level with a **rule** statement for each rule:

The router software processes the rules in the order in which you specify them in the configuration. If a term in a rule matches the packet, the router performs the corresponding action and the rule processing stops. If no term in a rule matches the packet, processing continues to the next rule in the rule set. If none of the rules matches the packet, the packet is dropped by default.

To import an SFW rule set:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.

The functionalities that you can configure in this mode are displayed in the task pane.
3. Select **Object Builder** from the task pane.

The Object Builder page is displayed.

4. Click the plus sign (+) next to Object Builder in the task pane to expand the tree and display the list of objects.

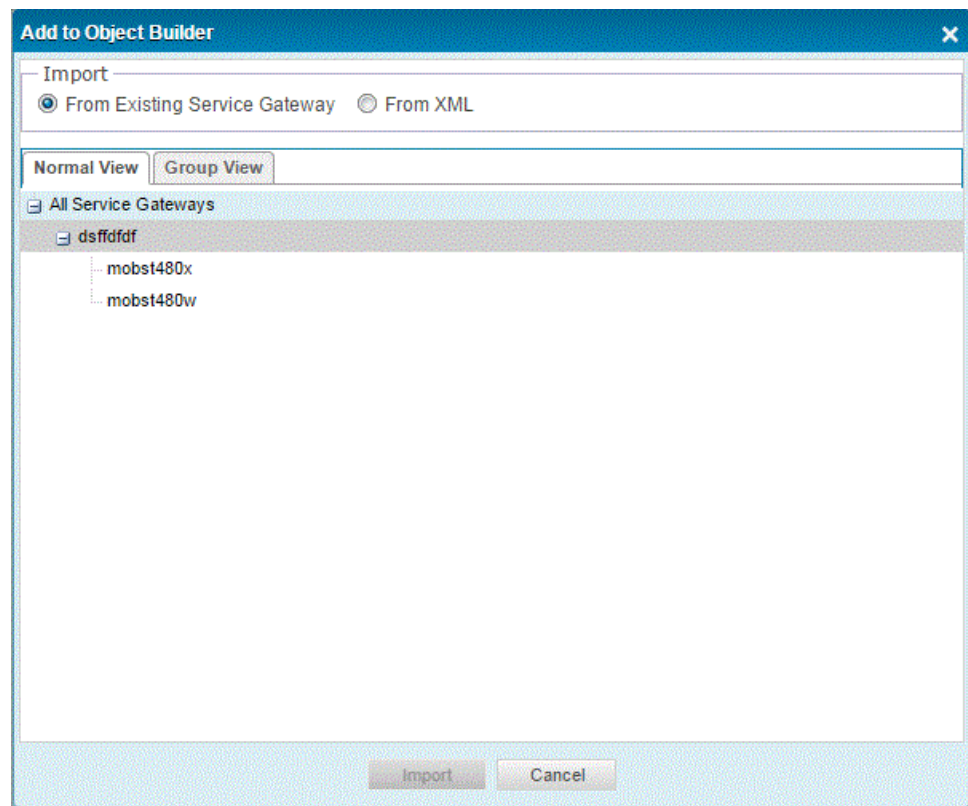
5. From the task pane, select **SFW Rule Sets** to open the SFW Rule Sets page on the right pane. The list of previously imported objects is displayed.

To filter and sort the display of objects, enter the name of the object as a match criterion in the Search box and click the Search icon. The page refreshes to display only the object names that match with the search term. You can use the paging controls to navigate across multiple pages of objects as necessary.

6. Click the **Import** icon.

The Add to Object Builder dialog box is displayed.

*Figure 30: Add to Object Builder Dialog Box*



7. Do one of the following for the Import section:
  - Select the **From Existing Service Gateway** radio button if you want to import the SFW rule set from SDGs that are present in the Edge Services Director database.
  - Select the **From XML** radio button if you want to import the SFW rule set from an XML configuration file on an external system.
8. If you selected the option to import the object from SDGs, do the following:

- Click the **Normal View** tab to view the list of SDGs. You can search for specific SDGs by entering a search item and clicking the Search icon.

Alternatively, click the **Group View** tab to view the list of SDG groups. You can search for specific SDG groups by entering a search item and clicking the Search icon.

- Click the plus sign (+) next to the All Service Gateways item to expand the tree structure that displays the list of SDGs or SDG groups. If the SDG pair is configured, you can select one of the devices, master or standby, from which you want to import the object.

Alternatively, if you selected the Group View tab, you can select an SDG from the groups displayed from which you want to import the object.

- Click **Import**. The object is added to the database and can be used during configuration of services or policies.

9. If you selected the option to import from an XML file, do the following:

- Click **Browse** beside the File Name field to navigate to the path where an XML file is available to be imported.
- Click **Import**. The object is added to the database and can be used during configuration of services or policies.

**Related  
Documentation**

- [Understanding the Object Builder on page 305](#)
- [Importing All Types of Objects on page 306](#)

---

## Importing SFW Rules

---

Each stateful firewall rule consists of a set of terms, similar to a filter configured at the [edit firewall] hierarchy level. A term consists of the following:

- from statement—Specifies the match conditions and applications that are included and excluded. The from statement is optional in stateful firewall rules.
- then statement—Specifies the actions and action modifiers to be performed by the router software. The then statement is mandatory in stateful firewall rules.

To import an SFW rule:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. Select **Object Builder** from the task pane.

The Object Builder page is displayed.

4. Click the plus sign (+) next to Object Builder in the task pane to expand the tree and display the list of objects.
5. From the task pane, select **SFW Rules** to open the SFW Rules page on the right pane. The list of previously imported objects is displayed.

To filter and sort the display of objects, enter the name of the object as a match criterion in the Search box and click the Search icon. The page refreshes to display only the object names that match with the search term. You can use the paging controls to navigate across multiple pages of objects as necessary.

6. Click the **Import** icon.

The Add to Object Builder dialog box is displayed.

7. Do one of the following for the Import section:

- Select the **From Existing Service Gateway** radio button if you want to import the SFW rule from SDGs that are present in the Edge Services Director database.
- Select the **From XML** radio button if you want to import the SFW rule from an XML configuration file on an external system.

8. If you selected the option to import the object from SDGs, do the following:

- Click the **Normal View** tab to view the list of SDGs. You can search for specific SDGs by entering a search item and clicking the Search icon.

Alternatively, click the **Group View** tab to view the list of SDG groups. You can search for specific SDG groups by entering a search item and clicking the Search icon.

- Click the plus sign (+) next to the All Service Gateways item to expand the tree structure that displays the list of SDGs or SDG groups. If the SDG pair is configured, you can select one of the devices, master or standby, from which you want to import the object.

Alternatively, if you selected the Group View tab, you can select an SDG from the groups displayed from which you want to import the object.

- Click **Import**. The object is added to the database and can be used during configuration of services or policies.

9. If you selected the option to import from an XML file, do the following:

- Click **Browse** beside the File Name field to navigate to the path where an XML file is available to be imported.
- Click **Import**. The object is added to the database and can be used during configuration of services or policies.

#### Related Documentation

- [Understanding the Object Builder on page 305](#)
- [Importing All Types of Objects on page 306](#)

## Importing Real Server Settings

---

Real servers are application servers used for traffic or server load balancing. The ADC software monitors the servers in the real-server group and the load-balanced applications running on them. If a router detects that a server or application has failed, it will not direct any new connection requests to that server. An adc-instance includes a complete set of ADC definitions: real-servers, groups of servers, virtual servers using virtual IP addresses, and virtual services accessed by clients.

Real servers are bound to real server groups. The criteria that you specify for real servers, such as weight and maximum and minimum connection thresholds, apply to the server load balancing algorithms that you specify for the real server groups. Server load balancing (SLB) uses the algorithms as it determines which real servers are to be assigned client requests. You also specify a ramp-up time, which is the period of time that it takes to reach the maximum connection threshold for the real server.

To import a real server:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. Select **Object Builder** from the task pane.  
The Object Builder page is displayed.
4. Click the plus sign (+) next to Object Builder in the task pane to expand the tree and display the list of objects.
5. From the task pane, select **Real Servers** to open the Real Servers page on the right pane. The list of previously imported objects is displayed.

To filter and sort the display of objects, enter the name of the object as a match criterion in the Search box and click the Search icon. The page refreshes to display only the object names that match with the search term. You can use the paging controls to navigate across multiple pages of objects as necessary.

6. Click the **Import** icon.  
The Add to Object Builder dialog box is displayed.
7. Do one of the following for the Import section:
  - Select the **From Existing Service Gateway** radio button if you want to import the real server from SDGs that are present in the Edge Services Director database.

- Select the **From XML** radio button if you want to import the real server from an XML configuration file on an external system.
8. If you selected the option to import the object from SDGs, do the following:
- Click the **Normal View** tab to view the list of SDGs. You can search for specific SDGs by entering a search item and clicking the Search icon.  
  
Alternatively, click the **Group View** tab to view the list of SDG groups. You can search for specific SDG groups by entering a search item and clicking the Search icon.
  - Click the plus sign (+) next to the All Service Gateways item to expand the tree structure that displays the list of SDGs or SDG groups. If the SDG pair is configured, you can select one of the devices, master or standby, from which you want to import the object.  
  
Alternatively, if you selected the Group View tab, you can select an SDG from the groups displayed from which you want to import the object.
  - Click **Import**. The object is added to the database and can be used during configuration of services or policies.
9. If you selected the option to import from an XML file, do the following:
- Click **Browse** beside the File Name field to navigate to the path where an XML file is available to be imported.
  - Click **Import**. The object is added to the database and can be used during configuration of services or policies.

**Related  
Documentation**

- [Understanding the Object Builder on page 305](#)
- [Importing All Types of Objects on page 306](#)

## Importing CGNAT Rule Sets

The **rule-set** statement defines a collection of NAT rules that determine what actions the router software performs on packets in the data stream. You define each rule by specifying a rule name and configuring terms. Then, you specify the order of the rules by including the **rule-set** statement at the **[edit services nat]** hierarchy level with a **rule** statement for each rule:

The router software processes the rules in the order in which you specify them in the configuration. If a term in a rule matches the packet, the router performs the corresponding action and the rule processing stops. If no term in a rule matches the packet, processing continues to the next rule in the rule set. If none of the rules matches the packet, no NAT action is performed on the packet. If a packet is destined to a NAT pool address, it is dropped.

To import a CGNAT rule set:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. Select **Object Builder** from the task pane. The Object Builder page is displayed.
4. Click the plus sign (+) next to Object Builder in the task pane to expand the tree and display the list of objects.

5. From the task pane, select **CGNAT Rule Sets** to open the CGNAT Rule Sets page on the right pane. The list of previously imported objects is displayed.

To filter and sort the display of objects, enter the name of the object as a match criterion in the Search box and click the Search icon. The page refreshes to display only the object names that match with the search term. You can use the paging controls to navigate across multiple pages of objects as necessary.

6. Click the **Import** icon.

The Add to Object Builder dialog box is displayed.

7. Do one of the following for the Import section:
  - Select the **From Existing Service Gateway** radio button if you want to import the CGNAT rule set from SDGs that are present in the Edge Services Director database.
  - Select the **From XML** radio button if you want to import the CGNAT rule set from an XML configuration file on an external system.
8. If you selected the option to import the object from SDGs, do the following:
  - Click the **Normal View** tab to view the list of SDGs. You can search for specific SDGs by entering a search item and clicking the Search icon.  
Alternatively, click the **Group View** tab to view the list of SDG groups. You can search for specific SDG groups by entering a search item and clicking the Search icon.
  - Click the plus sign (+) next to the All Service Gateways item to expand the tree structure that displays the list of SDGs or SDG groups. If the SDG pair is configured, you can select one of the devices, master or standby, from which you want to import the object.

Alternatively, if you selected the Group View tab, you can select an SDG from the groups displayed from which you want to import the object.

- Click **Import**. The object is added to the database and can be used during configuration of services or policies.
9. If you selected the option to import from an XML file, do the following:
    - Click **Browse** beside the File Name field to navigate to the path where an XML file is available to be imported.
    - Click **Import**. The object is added to the database and can be used during configuration of services or policies.

**Related  
Documentation**

- [Understanding the Object Builder on page 305](#)
- [Importing All Types of Objects on page 306](#)

## Importing CGNAT Rules

NAT processing centers on the evaluation of NAT rule sets and rules. A rule set determines the overall direction of the traffic to be processed. For example, a rule set can select traffic from a particular interface or to a specific zone. A rule set can contain multiple rules. Once a rule set is found that matches specific traffic, each rule in the rule set is evaluated for a match. Each rule in the rule set further specifies the traffic to be matched and the action to be taken when traffic matches the rule.

To import a CGNAT rule:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. Select **Object Builder** from the task pane. The Object Builder page is displayed.
4. Click the plus sign (+) next to Object Builder in the task pane to expand the tree and display the list of objects.
5. From the task pane, select **CGNAT Rules** to open the CGNAT Rules page on the right pane. The list of previously imported objects is displayed.

To filter and sort the display of objects, enter the name of the object as a match criterion in the Search box and click the **Search** icon. The page refreshes to display only the object names that match with the search term. You can use the paging controls to navigate across multiple pages of objects as necessary.

6. Click the **Import** icon.

The Add to Object Builder dialog box is displayed.

7. Do one of the following for the Import section:
  - Select the **From Existing Service Gateway** radio button if you want to import the CGNAT rule from SDGs that are present in the Edge Services Director database.
  - Select the **From XML** radio button if you want to import the CGNAT rule from an XML configuration file on an external system.
8. If you selected the option to import the object from SDGs, do the following:
  - Click the **Normal View** tab to view the list of SDGs. You can search for specific SDGs by entering a search item and clicking the **Search** icon.

Alternatively, click the **Group View** tab to view the list of SDG groups. You can search for specific SDG groups by entering a search item and clicking the **Search** icon.
  - Click the plus sign (+) next to the All Service Gateways item to expand the tree structure that displays the list of SDGs or SDG groups. If the SDG pair is configured, you can select one of the devices, master or standby, from which you want to import the object.

Alternatively, if you selected the **Group View** tab, you can select an SDG from the groups displayed from which you want to import the object.
  - Click **Import**. The object is added to the database and can be used during configuration of services or policies.
9. If you selected the option to import from an XML file, do the following:
  - Click **Browse** beside the File Name field to navigate to the path where an XML file is available to be imported.
  - Click **Import**. The object is added to the database and can be used during configuration of services or policies.

**Related Documentation**

- [Understanding the Object Builder on page 305](#)
- [Importing All Types of Objects on page 306](#)

---

## Importing CGNAT Pools

A Network Address Translation (NAT) pool is a continuous range of IP addresses that you can use to create a NAT policy. NAT policies perform address translation by translating internal IP addresses to the addresses in these pools.

To import a CGNAT pool:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.

The functionalities that you can configure in this mode are displayed in the task pane.

3. Select **Object Builder** from the task pane. The Object Builder page is displayed.
4. Click the plus sign (+) next to Object Builder in the task pane to expand the tree and display the list of objects.
5. From the task pane, select **CGNAT Pools** to open the Real Servers page on the right pane. The list of previously imported objects is displayed.

To filter and sort the display of objects, enter the name of the object as a match criterion in the Search box and click the **Search** icon. The page refreshes to display only the object names that match with the search term. You can use the paging controls to navigate across multiple pages of objects as necessary.

6. Click the **Import** icon.

The Add to Object Builder dialog box is displayed.

7. Do one of the following for the Import section:

- Select the **From Existing Service Gateway** radio button if you want to import the CGNAT pool from SDGs that are present in the Edge Services Director database.
- Select the **From XML** radio button if you want to import the CGNAT pool from an XML configuration file on an external system.

8. If you selected the option to import the object from SDGs, do the following:

- Click the **Normal View** tab to view the list of SDGs. You can search for specific SDGs by entering a search item and clicking the **Search** icon.

Alternatively, click the **Group View** tab to view the list of SDG groups. You can search for specific SDG groups by entering a search item and clicking the **Search** icon.

- Click the plus sign (+) next to the All Service Gateways item to expand the tree structure that displays the list of SDGs or SDG groups. If the SDG pair is configured, you can select one of the devices, master or standby, from which you want to import the object.

Alternatively, if you selected the **Group View** tab, you can select an SDG from the groups displayed from which you want to import the object.

- Click **Import**. The object is added to the database and can be used during configuration of services or policies.

9. If you selected the option to import from an XML file, do the following:

- Click **Browse** beside the File Name field to navigate to the path where an XML file is available to be imported.
- Click **Import**. The object is added to the database and can be used during configuration of services or policies.

- Related Documentation**
- [Understanding the Object Builder on page 305](#)
  - [Importing All Types of Objects on page 306](#)

---

## Importing Applications

You can define application protocols for the stateful firewall and Network Address Translation (NAT) services to use in match condition rules. An application protocol, or application layer gateway (ALG), defines application parameters using information from network Layer 3 and above. Examples of such applications are FTP and H.323. The application-protocol allows you to specify which of the supported application protocols (ALGs) to configure and include in an application set for service processing.

To import an application:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. Select **Object Builder** from the task pane.  
The Object Builder page is displayed.
4. Click the plus sign (+) next to Object Builder in the task pane to expand the tree and display the list of objects.
5. From the task pane, select **Applications** to open the Applications page on the right pane. The list of previously imported objects is displayed.

To filter and sort the display of objects, enter the name of the object as a match criterion in the Search box and click the **Search** icon. The page refreshes to display only the object names that match with the search term. You can use the paging controls to navigate across multiple pages of objects as necessary.

6. Click the **Import** icon.  
The Add to Object Builder dialog box is displayed.
7. Do one of the following for the Import section:
  - Select the **From Existing Service Gateway** radio button if you want to import the application from SDGs that are present in the Edge Services Director database.
  - Select the **From XML** radio button if you want to import the application from an XML configuration file on an external system.
8. If you selected the option to import the object from SDGs, do the following:

- Click the **Normal View** tab to view the list of SDGs. You can search for specific SDGs by entering a search item and clicking the **Search** icon.

Alternatively, click the **Group View** tab to view the list of SDG groups. You can search for specific SDG groups by entering a search item and clicking the **Search** icon.

- Click the plus sign (+) next to the All Service Gateways item to expand the tree structure that displays the list of SDGs or SDG groups. If the SDG pair is configured, you can select one of the devices, master or standby, from which you want to import the object.

Alternatively, if you selected the **Group View** tab, you can select an SDG from the groups displayed from which you want to import the object.

- Click **Import**. The object is added to the database and can be used during configuration of services or policies.

9. If you selected the option to import from an XML file, do the following:

- Click **Browse** beside the File Name field to navigate to the path where an XML file is available to be imported.
- Click **Import**. The object is added to the database and can be used during configuration of services or policies.

#### Related Documentation

- [Understanding the Object Builder on page 305](#)
- [Importing All Types of Objects on page 306](#)

## Importing Application Sets

You can define application protocols for the stateful firewall and Network Address Translation (NAT) services to use in match condition rules. You can group applications into a bundle called an application set.

To import an application set:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. Select **Object Builder** from the task pane.  
The Object Builder page is displayed.
4. Click the plus sign (+) next to Object Builder in the task pane to expand the tree and display the list of objects.

5. From the task pane, select **Application Sets** to open the Application Sets page on the right pane.

The list of previously imported objects is displayed.

To filter and sort the display of objects, enter the name of the object as a match criterion in the Search box and click the **Search** icon. The page refreshes to display only the object names that match with the search term. You can use the paging controls to navigate across multiple pages of objects as necessary.

6. Click the **Import** icon. The Add to Object Builder dialog box is displayed.
7. Do one of the following for the Import section:
  - Select the **From Existing Service Gateway** radio button if you want to import the application set from SDGs that are present in the Edge Services Director database.
  - Select the **From XML** radio button if you want to import the application set from an XML configuration file on an external system.
8. If you selected the option to import the object from SDGs, do the following:

- Click the **Normal View** tab to view the list of SDGs. You can search for specific SDGs by entering a search item and clicking the **Search** icon.

Alternatively, click the **Group View** tab to view the list of SDG groups. You can search for specific SDG groups by entering a search item and clicking the **Search** icon.

- Click the plus sign (+) next to the All Service Gateways item to expand the tree structure that displays the list of SDGs or SDG groups. If the SDG pair is configured, you can select one of the devices, master or standby, from which you want to import the object.

Alternatively, if you selected the **Group View** tab, you can select an SDG from the groups displayed from which you want to import the object.

- Click **Import**. The object is added to the database and can be used during configuration of services or policies.
9. If you selected the option to import from an XML file, do the following:
    - Click **Browse** beside the File Name field to navigate to the path where an XML file is available to be imported.
    - Click **Import**. The object is added to the database and can be used during configuration of services or policies.

- Related Documentation**
- [Understanding the Object Builder on page 305](#)
  - [Importing All Types of Objects on page 306](#)

# Managing Packet Analyzers

- [Packet Analyzer Overview on page 321](#)
- [Creating and Viewing Service Analyzers on page 323](#)

## Packet Analyzer Overview

---

Packet capture is a tool that helps you to analyze network traffic and troubleshoot network problems. The packet capture tool captures real-time data packets traveling over the network for monitoring and logging. This tool is a debugging and analysis utility that you can use to identify the problematic area in a session path. A set of counters are displayed for both forward and reverse flow for all the supported services on SDG devices. Using these statistical details and values, you can obtain adequate and useful estimates regarding the total bytes count for each service in every hop and quickly, easily locate the hop where there can be a possible packet drop.

The packet analyzer is the endpoint to which the flow collector interface sends traffic for analysis. You can process and export multiple cflowd records with a flow collector interface. You create a flow collector interface on a Monitoring Services II or Multiservices 400 PIC. The flow collector interface combines multiple cflowd records into a compressed ASCII data file and exports the file to an FTP server.

You can configure the packet analyzer filters to capture packet data flows based on a match or classification criteria to collect statistics and information only about packets that satisfy the criteria. You can define the data and control plane packet flow direction and interface settings in the filter, and the interval at which devices must be polled. You can also specify a timeout to apply a threshold on the amount of data to be collected. You can then schedule the filter to be run for different services and view the statistics as numerical values or as a graph.

Packets are captured as binary data, without modification. You can read the packet information offline with a packet analyzer such as Ethereal or tcpdump. If you need to quickly capture packets destined for, or originating from, the Routing Engine and analyze them online, you can use the packet capture diagnostic tool.

Network administrators and security engineers use packet capture to perform the following tasks:

- Monitor network traffic and analyze traffic patterns.

- Identify and troubleshoot network problems. Detect security breaches in the network, such as unauthorized intrusions, spyware activity, or ping scans.
- Packet capture operates like traffic sampling on the device, except that it captures entire packets.

Data packets are chunks of data that transit the router as they are being forwarded from a source to a destination. When a router receives a data packet on an interface, it determines where to forward the packet by looking in the forwarding table for the best route to a destination. The router then forwards the data packet toward its destination through the appropriate interface. The Packet Forwarding Engine, which is the central processing element of the router's forwarding plane, handles the flow of data packets in and out of the router's physical interfaces. Although the Packet Forwarding Engine contains Layer 3 and Layer 4 header information, it does not contain the packet data itself (the packet's payload).

You can also use the packet capture feature when you need to quickly capture and analyze control traffic on a router. Control packets refer to health check packets that are sent to examine the health and efficiency of specific URLs or paths. Health checking allows you to verify content accessibility in large websites. As content grows and information is distributed across different server farms, flexible, customizable content health checks are critical to ensure end-to-end availability.

## Pre-Service Filtering of Traffic for Service Processing

To filter IPv4 or IPv6 traffic before accepting packets for input or output service processing, include the **service-set** *service-set-name* **service-filter** *service-filter-name* at one of the following interfaces:

- **[edit interfaces *interface-name* unit *unit-number* family (inet | inet6) service input]**
- **[edit interfaces *interface-name* unit *unit-number* family (inet | inet6) service output]**

For the **service-set-name**, specify a service set configured at the **[edit services service-set]** hierarchy level.

The service set retains the input interface information even after services are applied, so that functions such as filter-class forwarding and destination class usage (DCU) that depend on input interface information continue to work.

The following requirements apply to filtering inbound or outbound traffic before accepting packets for service processing:

- You configure the same service set on the input and output sides of the interface.
- If you include the **service-set** statement without an optional **service-filter** definition, the Junos OS assumes the match condition is true and selects the service set for processing automatically.
- The service filter is applied only if a service set is configured and selected.

You can include more than one service set definition on each side of an interface. The following guidelines apply:

- If you include multiple service sets, the router (or switch) software evaluates them in the order in which they appear in the configuration. The system executes the first service set for which it finds a match in the service filter and ignores the subsequent definitions.
- A maximum of six service sets can be applied to an interface.
- When you apply multiple service sets to an interface, you must also configure and apply a service filter to the interface.

### Postservice Filtering of Returning Service Traffic

As an option to filtering of IPv4 or IPv6 input service traffic, you can apply a service filter to IPv4 or IPv6 traffic that is returning to the services interface after the service set is executed. To apply a service filter in this manner, include the **post-service-filter service-filter-name** statement at the **[edit interfaces interface-name unit unit-number family (inet | inet6) service input]** hierarchy level.

#### Related Documentation

- [Creating and Viewing Service Analyzers on page 323](#)

---

## Creating and Viewing Service Analyzers

The packet analyzer is the endpoint to which the flow collector interface sends traffic for analysis. You can process and export multiple cflowd records with a flow collector interface. You can perform the following tasks with the Service Analyzer page:

- Configure and provision filters for packet analysis.
- Configure filters for CGNAT, ADC, and TLB services.
- Start and stop the configured filters.
- View the packet analyzer details as a statistical form or a graphical form.
- [Configuring the Traffic Analyzer Filter on page 323](#)
- [Managing Service Analyzer Filter Instances on page 326](#)
- [Viewing Service Analyzer Instance Details on page 328](#)
- [Viewing the Service Analyzer Statistics in Grid Format and Graph on page 330](#)

### Configuring the Traffic Analyzer Filter

To configure the traffic analyzer filter details on packet flows for the different services and to schedule its running:

1. From the View selector, select **Gateway View**. The workspaces that are available in this view are displayed. The Gateway view displays the service delivery gateway (SDG) groups and the SDGs that are part of the high availability pair in an SDG group.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the View pane, select the All Network item in Gateway view. Click the plus sign (+) beside the All Network item in the View pane to expand the tree and select the device node you want.
4. From the task pane, do one of the following:
  - Select **Service Analyzer > ADC Filter** from the task pane. The Service Analyzer for ADC Filter page is displayed. The service instance or template that you previously configured for the ADC service type are displayed. All the instances created in the Service Templates workspace are shown.
  - Select **Service Analyzer > TLB Filter** from the task pane. The Service Analyzer for TLB Filter page is displayed. The service instance or template that you previously configured for the TLB service type are displayed. All the instances created in the Service Templates workspace are shown.
  - Select **Service Analyzer > CGNAT Filter** from the task pane. The Service Analyzer for CGNAT Filter page is displayed. The service instance or template that you previously configured for the CGNAT service type are displayed. All the instances created in the Service Templates workspace are shown.

The list of SDGs or SDG pairs in a high availability group are displayed, along with the filter instances that were configured for the different services. The number of filter instances that are currently in progress and the number of filter instances that are scheduled or planned to be run at a later time are also displayed. For information on running or clearing filter instances, see *Managing Service Analyzer Filter Instances*.

5. Select the SDGs or SDG pairs (you can select multiple rows to create and assign filters to several SDGs simultaneously) for which you want to create packet analyzer filters for services.
6. Click the plus sign (+) above the table of listed SDGs to create a new filter. The Update Service Analyzer Filter Details page is displayed.
7. In the Data Forward Flow section, do the following. A forward flow refers to packets that are sent in the forward or upward direction. A reverse flow refers to packets that are sent in the returning or backward direction.
  - From the **Egress** list, select the egress interface on which the data packets that are sent out in the forward flow must be monitored. Click **Details** beside the list to view interface details.

- From the **Ingress** list, select the input interface on which the data packets that are received in the forward flow must be monitored. Click **Details** beside the list to view interface details.
8. In the Data Reverse Flow section, do the following.
    - From the **Egress** list, select the egress interface on which the data packets that are sent out in the reverse flow must be monitored. Click **Details** beside the list to view interface details.
    - From the **Ingress** list, select the input interface on which the data packets that are received in the reverse flow must be monitored. Click **Details** beside the list to view interface details.
  9. In the Control Forward Flow section, do the following. A forward flow refers to packets that are sent in the forward or upward direction. A reverse flow refers to packets that are sent in the returning or backward direction.
    - From the **Egress** list, select the egress interface on which the control packets that are sent out in the forward flow must be monitored. Click **Details** beside the list to view interface details.
    - From the **Ingress** list, select the input interface on which the control packets that are received in the forward flow must be monitored. Click **Details** beside the list to view interface details.
  10. In the Data Reverse Flow section, do the following.
    - From the **Egress** list, select the egress interface on which the control packets that are sent out in the reverse flow must be monitored. Click **Details** beside the list to view interface details.
    - From the **Ingress** list, select the input interface on which the control packets that are received in the reverse flow must be monitored. Click **Details** beside the list to view interface details.
  11. Click **Apply** to save the filter settings. Otherwise, click **Cancel** to discard the changes. You are returned to the Service Analyzer page.
  12. If you created a new filter, the filter instance is displayed under the corresponding service type section, such as CGNAT or ADC. Such filters are provisioned filter instances. This display signifies that the filter is configured, but it needs to be scheduled to be run. Click the link that shows the number of instances under the column of the relevant service type. The Service Analyzer Instances page is shown.
  13. On this page, the names of the service instances for which filters are defined. The actions you can perform are in the form of the Clear and Run buttons, above the table of listed service instances, for each service instance with a filter.
  14. Select the check box next to a service analyzer filter and click the **Delete** button to remove a configured filter for an instance. You are prompted to confirm the deletion. If you click **OK**, a popup dialog box denotes the successful deletion.

15. Select the check box next to a service analyzer filter instance, and click the **Run** button to schedule the filter to be run. The Run Filter dialog box appears. The Run button is grayed out if the particular service filter instance is already in progress.
16. From the **Poll Interval** list, select the interval in minutes at which the data must be polled and collected. Values from 1 minute up to 59 minutes are shown in increments of 2 minutes in the list.
17. In the **Schedule Start Details** section, click **Run Now** to start the filter immediately. Alternatively, click the **Run At** radio button and select the date and time at which the filter must be run.
18. In the **Schedule End Details** section, do one of the following:
  - Click the **Stop At** radio button and select the date and time at which the filter must be stopped.
  - Click the **Stop After** radio button and specify a value for the number of polls after which the filter must be ended.
  - Click the **Run Until Stopped** radio button to continue running the test until you manually want to stop it.
19. Click **Run** to save the filter settings. Otherwise, click **Cancel** to discard the changes. You are returned to the Prepared Service Analyzer Instances dialog box. Click **Close** to return to the Service Analyzer Page.

## Managing Service Analyzer Filter Instances

To view, start, stop, or clear the configured analyzer filters:

1. From the View selector, select **Gateway View**. The workspaces that are available in this view are displayed. The Gateway view displays the service delivery gateway (SDG) groups and the SDGs that are part of the high availability pair in an SDG group.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the View pane, select the All Network item in Gateway view. Click the plus sign (+) beside the All Network item in the View pane to expand the tree and select the device node you want.
4. From the task pane, do one of the following:
  - Select **Service Analyzer > ADC Filter** from the task pane. The Service Analyzer for ADC Filter page is displayed. The service instance or template that you previously configured for the ADC service type are displayed. All the instances created in the Service Templates workspace are shown.

- Select **Service Analyzer > TLB Filter** from the task pane. The Service Analyzer for TLB Filter page is displayed. The service instance or template that you previously configured for the TLB service type are displayed. All the instances created in the Service Templates workspace are shown.
- Select **Service Analyzer > CGNAT Filter** from the task pane. The Service Analyzer for CGNAT Filter page is displayed. The service instance or template that you previously configured for the CGNAT service type are displayed. All the instances created in the Service Templates workspace are shown.

The list of SDGs or SDG pairs in a high availability group are displayed, along with the filter instances that were configured for the different services. The number of filter instances that are currently in progress and the number of filter instances that are scheduled or planned to be run at a later time are also displayed. For information on viewing filter instances, see *Viewing the Traffic Analyzer Statistics and Graph*.

5. For the SDG corresponding to a certain service, all of the previously configured service analyzer filters are displayed in the Service Analyzer Instances page with the state of the filter instance under the Status column of the relevant service type. View the Status column for the current state of the filter.

**Figure 31: Service Analyzer Instances Page**

The screenshot shows the 'Service Analyzer for TLB Filter' page. At the top, there is a toolbar with buttons: Add, View, Delete, Run, View Report, Stop, and Last Run Errors. Below the toolbar is a table with the following columns: Name, Created By, Created Time, and Status. There is one row in the table with the following data: Name: 10.5.1.7, Created By: super, Created Time: Aug 25, 2015 3:39:50 PM IST, Status: Running. At the bottom of the page, there is a pagination bar showing 'Page 1 of 1', 'Displaying 1 - 1 of 1', and 'Show 24 items'.

Name	Created By	Created Time	Status
10.5.1.7	super	Aug 25, 2015 3:39:50 PM IST	Running

You can click the links under one of the following columns:

- **View**—Click to display the traffic analyzer details on packet flows for the different services configured. For information on viewing filter instances, see *Viewing the Service Analyzer Statistics and Graph*.
- **Delete**—Click to remove the configured filter for an instance. You are prompted to confirm the deletion. If you click **OK**, a popup dialog box denotes the successful deletion.

- **Run**—Click to schedule a filter to be run. For information on scheduling a filter instance to be run, see *Configuring the Traffic Analyzer Filter*.
  - **Report**—Click to view the collection statistics and information about packets that are fetched. For information on viewing the collected details by a service analyzer, see *Viewing the Service Analyzer Collection Data*.
  - **Stop**—Click to end a running filter. You are prompted to confirm whether you want to stop the filter instance. If you click **OK**, a popup dialog box denotes the successful termination of the filter instance.
  - **Last Run Errors**—Click to display any errors that occurred during the running of the filter instance. The Last Run Status dialog box is displayed. It contains the Provisioning Errors and Decommissioning Errors tabs that describe errors that might have occurred during the initialization and start of the analyzer filters or with the decommissioning and termination. The following fields are displayed in this dialog box for both the tabs:
    - **Host Name**—Host name of the SDG device.
    - **Severity**—System logging severity level.
    - **Path**—Hierarchy level of the configuration statement corresponding to the setting in the CLI interface Info Informational message about the error that is generated.
    - **Message**—System event logging message generated that describes the error.
  - **Graph**—Click to display the packet analyzer details for monitoring as a pictorial form. The Packet Flow Graph dialog box appears.
6. In the dialog box, the Configured Instances column displays the names of the service instances for which filters are defined. The Actions column contains the Clear and Run subcolumns for each service instance with a filter.
  7. Click **Delete** to remove a configured filter for an instance. You are prompted to confirm the deletion. If you click **OK**, a popup dialog box denotes the successful deletion.
  8. Click **Run** beside the instance you want to schedule the filter to be run. The Run Filter dialog box appears to specify the schedule settings.

## Viewing Service Analyzer Instance Details

To view the service analyzer instance details:

1. From the View selector, select **Gateway View**. The workspaces that are available in this view are displayed. The Gateway view displays the service delivery gateway (SDG) groups and the SDGs that are part of the high availability pair in an SDG group.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.

3. From the View pane, select the All Network item in Gateway view. Click the plus sign (+) beside the All Network item in the View pane to expand the tree and select the device node you want.
4. From the task pane, do one of the following:
  - Select **Service Analyzer > ADC Filter** from the task pane. The Service Analyzer for ADC Filter page is displayed. The service instance or template that you previously configured for the ADC service type are displayed. All the instances created in the Service Templates workspace are shown.
  - Select **Service Analyzer > TLB Filter** from the task pane. The Service Analyzer for TLB Filter page is displayed. The service instance or template that you previously configured for the TLB service type are displayed. All the instances created in the Service Templates workspace are shown.
  - Select **Service Analyzer > CGNAT Filter** from the task pane. The Service Analyzer for CGNAT Filter page is displayed. The service instance or template that you previously configured for the CGNAT service type are displayed. All the instances created in the Service Templates workspace are shown.

The list of SDGs or SDG pairs in a high availability group are displayed, along with the filter instances that were configured for the different services. The number of filter instances that are currently in progress and the number of filter instances that are scheduled or planned to be run at a later time are also displayed. For information on running or clearing filter instances, see *Managing Service Analyzer Filter Instances*.

5. Select the SDGs or SDG pairs (you can select multiple rows to create and assign filters to several SDGs simultaneously) for which you want to create packet analyzer filters for services.
6. From the Service Analyzer page, for the SDG corresponding to a certain service, click the link under the column of the relevant service type. The Prepared Service Analyzer Instances dialog box is shown. Click **View** under the View column to view the traffic analyzer for the particular service.

The View Service Instance Analyzer Details page is displayed.

The following fields are displayed in this page:

Field	Description
Name	Name of the SDG or pair of SDGs in a high availability group.
Type	Service type for which packets collected are shown. Values are CGNAT, ADC, or TLB.
Data Packets/Control Packets	Click the <b>Data Packets</b> tab to view data packet details for the service analyzer filter. Alternatively, click the <b>Control Packets</b> tab to view control packet details for the service analyzer filter. Indicates whether data or control packet details are shown.
Forward Flow	Displays statistics for packets in forward flow direction.

Field	Description
Ingress	Number of packets that arrive in the ingress direction in forward flow.
Egress	Number of packets that are sent out in the egress direction in forward flow.
Reverse Flow	Displays statistics for packets in reverse flow direction. If a service set is a sampling service set and the reverse-flow service order is not configured, all sampled traffic is considered to be forward traffic.
Ingress	Number of packets that arrive in the ingress direction in reverse flow.
Egress	Number of packets that are sent out in the egress direction in reverse flow.

- Click **Close** after viewing the analyzer filter details. You are returned to the Prepared Service Analyzer Instances dialog box. Click **Close** to return to the Service Analyzer Page

## Viewing the Service Analyzer Statistics in Grid Format and Graph

To view the traffic analyzer details on packet flows for the different services that match the filter criteria:

- From the View selector, select **Gateway View**. The workspaces that are available in this view are displayed. The Gateway view displays the service delivery gateway (SDG) groups and the SDGs that are part of the high availability pair in an SDG group.
- From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
- From the View pane, select the All Network item in Gateway view. Click the plus sign (+) beside the All Network item in the View pane to expand the tree and select the device node you want.
- From the task pane, do one of the following:
  - Select **Service Analyzer > ADC Filter** from the task pane. The Service Analyzer for ADC Filter page is displayed. The service instance or template that you previously configured for the ADC service type are displayed. All the instances created in the Service Templates workspace are shown.
  - Select **Service Analyzer > TLB Filter** from the task pane. The Service Analyzer for TLB Filter page is displayed. The service instance or template that you previously configured for the TLB service type are displayed. All the instances created in the Service Templates workspace are shown.
  - Select **Service Analyzer > CGNAT Filter** from the task pane. The Service Analyzer for CGNAT Filter page is displayed. The service instance or template that you previously

configured for the CGNAT service type are displayed. All the instances created in the Service Templates workspace are shown.

The list of SDGs or SDG pairs in a high availability group are displayed, along with the filter instances that were configured for the different services. The number of filter instances that are currently in progress and the number of filter instances that are scheduled or planned to be run at a later time are also displayed. For information on running or clearing filter instances, see *Managing Service Analyzer Filter Instances*.

5. Select the SDGs or SDG pairs (you can select multiple rows to create and assign filters to several SDGs simultaneously) for which you want to create packet analyzer filters for services.
6. From the Service Analyzer page, for the SDG corresponding to a certain service, click the link under the column of the relevant service type. The Prepared Service Analyzer Instances dialog box is shown. Click **Report** under the View Report column to view the traffic analyzer for the particular service.

The Service Analyzer Collection Data — Grid View page is displayed.

At the top of the tabular display, select the criteria for which you want to sort and segregate the packet analyzer information to be viewed. From the Criteria section, do the following:

- a. Select **Control** or **Data** from the first drop-down list to view control or data packets.
- b. Select **Forward** or **Reverse** from the second drop-down list to view statistics for packets in forward or reverse flows.
- c. Select **IPv4** or **IPv6** from the second drop-down list to view IPv4 or IPv6 packets for the filter instance.
- d. Click the search icon to apply the filter conditions and display details matching the specified criteria.

The following fields are displayed in this page:

Field	Description
Name	Name of the SDG or pair of SDGs in a high availability group.
Type	Service type for which packets collected are shown. Values are CGNAT, ADC, or TLB.
Collection Time	Date and time at which the packet details are collected.
Ingress	Number of packets that arrive in the ingress direction in forward and reverse flow.
PreService	Number of packets in the forward flow and reverse flow before the processing of services. You can define the pre-service filter to be applied to traffic before it is accepted for service processing.

Field	Description
Post Service	Number of packets in the forward flow and reverse flow after processing of services. You can define the post-service filter to be applied to traffic after service processing. The filter is applied only if a service set is configured and selected. You can configure a postservice filter on the input side of the interface only. This setting is not supported when the service interface is on an MS-MIC or MS-MPC.
Egress	Number of packets that are sent out in the egress direction in forward flow and reverse flow.

Click **Close** after viewing the collection data in the tabular grid. You are returned to the Prepared Service Analyzer Instances dialog box. Click **Close** to return to the Service Analyzer Page.

7. Alternatively, you can view the service analyzer details in a graphical representation. Click **Graph** under the View Report column to display the packet analyzer details for monitoring as a pictorial form. The Packet Flow Graph dialog box appears.

Line graphs are displayed for data forward flow, data reverse flow, control forward flow, and control reverse flow. The number of packets is displayed on the y-axis and time is displayed along the x-axis. The legends reference the egress, pre-service, post-service, and ingress packets. Mouse over the points in the graph to highlight and view the number of packets at a particular time instance.

At the top of the graph, select the criteria for which you want to sort and segregate the packet analyzer information to be viewed. From the Criteria section, do the following:

- a. Select **Control** or **Data** from the first drop-down list to view control or data packets.
  - b. Select **IPv4** or **IPv6** from the second drop-down list to view IPv4 or IPv6 packets for the filter instance.
  - c. Select the period for which the service analyzer details must be shown from the third drop-down list. For example, you can select **Last 10 Mins** to display the service analyzer packets collected over the last 10 minutes or the **Last 1 Hr** option to display the service analyzer packets collected over the last one hour.
  - d. Click the search icon to apply the filter conditions and display details matching the specified criteria.
8. Click **Close** after viewing the graph. You are returned to the Prepared Service Analyzer Instances dialog box. Click **Close** to return to the Service Analyzer Page.

**Related Documentation**

- [Packet Analyzer Overview on page 321](#)

## PART 6

# Deploy Mode

- [About Deploy Mode on page 335](#)
- [Device Management on page 339](#)
- [Configuration File Management on page 347](#)
- [Software Image Management on page 353](#)
- [Viewing and Editing Service Instances and Packet Filters Across All Gateways on page 363](#)
- [Enhanced Editing of Services and Packet Filters on page 369](#)
- [Managing Service Instance and Policy Rule Definitions on page 375](#)
- [Managing Packet Analyzers on page 493](#)



## CHAPTER 16

# About Deploy Mode

- [Understanding Deploy Mode in Gateway and Service Views of Edge Services Director on page 335](#)
- [Understanding Deploy Mode in Location and Device Views of Edge Services Director on page 338](#)

### Understanding Deploy Mode in Gateway and Service Views of Edge Services Director

The Deploy mode in Gateway and Service views enables you to deploy configuration changes to devices. You can create a deployment plan for each of the service planning templates, such as the ones defined for ADC or SFW services, and the policy or filter templates, such as the packet filter or SFW policy, that you have created. A deploy plan contains details about the settings and configuration parameters that must be propagated and provisioned on the SDGs managed by Edge Services Director. You can also create, update, display, publish and commission of packet filters, stateful firewall and NAT policies present on discovered and managed SDGs.

This topic describes:

- [Deploying Configuration Changes on page 335](#)
- [Transactions on page 336](#)
- [Modify the Association of SDG Details and Rule Terms for a Policy Filters on page 336](#)
- [View Service Object Statistics on page 337](#)
- [Service Edit on page 337](#)
- [Policy and Filter Management on page 337](#)

### 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. 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 do the following configuration deployment tasks on devices that have pending changes:

- Run configuration deployment jobs immediately or schedule them for future times.
- Preview pending configuration changes before deploying.
- 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 Edge Services 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.

Edge Services Director does not deploy configuration to a device with a configuration that is out of sync (meaning that the device's configuration differs from Edge Services 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 Edge Services 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.

The Service Deployment page provides the following functionalities:

- Approval Management—View the details of the filters/policies and other service deployment plans which are pending for approval. Approve or reject deployment plans done to existing feature.
- Update Devices—View the details of approved filters/policies and other service deployment plans which are ready for commissioning. Commission the deployment plans or discard accordingly.

## Transactions

A transaction refers to an operation or a task that is performed on the service definitions, configuration parameters, and policy settings that are created for provisioning on the devices or Service Delivery Gateways (SDGs). When you create a deployment plan to define the services and policy filters that must be applied and propagated on the devices, the administrator can approve or reject a deploy plan. For each approved deploy plan, a transaction is automatically created by the Edge Services Director application.

## Modify the Association of SDG Details and Rule Terms for a Policy Filters

In Gateway view of Deploy mode, from the Policy & Filters page, which displays all the previously configured CGNAT and SFW service policy filters, and packet filters, you can modify the components or the parameter types that are associated with a particular

service filter. You must lock the packet filters for which you want to modify the attached rule term components or attributes before you can update the settings. You can also select a different SDG to which the packet filter must be applied.

## View Service Object Statistics

In Service view of Deploy mode, you can view a graphical representation in the form of pie charts of the configured ADC, TLB, CGNAT, SFW, and packet policies or filter.

## Service Edit

In Gateway and Service views, you can select a previously configured service template instance, such as a stateful firewall, carrier-grade NAT, traffic load balancer, or adaptive delivery controller, and lock the service instance to select the attributes or components of the service to be modified. You can publish or unpublish service template instances.

## Policy and Filter Management

The Policy and Filter Management feature in the Junos Space Edge Services Director application helps you create, update, display, publish and commission of packet filters, stateful firewall and NAT policies present on discovered and managed SDGs. The Service Management workspace displays a bar graph of draft, published and approved filters or policies for different options available under workspace:

- **Packet Filter:** This option displays packet filters present on SDGs in tabular view. It also provides the ability to create, update, and delete filters on selected SDGs.
- **Stateful Firewall:** This option displays stateful firewall policies present on SDGs in tabular view. It also provides the ability to create, update and delete stateful firewall policies on selected SDGs.
- **CGNAT:** This option displays CGNAT policies present on SDGs in tabular view. It also provides the ability to create, update and delete CGNAT policies on selected SDGs. A published filter or policy is sent for peer review and approval. After approval, the filter or policy is deployed to devices.

### Related Documentation

- [Viewing Deployment Plans on page 529](#)
- [Creating and Assigning a Deployment Plan to Devices on page 533](#)
- [Transactions Overview on page 551](#)
- [Viewing Transactions on page 552](#)

## Understanding Deploy Mode in Location and Device Views of Edge Services Director

The Deploy mode enables you to deploy configuration changes and software upgrades to devices and perform several device management and configuration file management tasks.

This topic describes:

- [Managing Software Images on page 338](#)
- [Managing Devices on page 338](#)
- [Managing Device Configuration Files on page 338](#)

### Managing Software Images

Edge Services 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 Edge Services 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 MX Series routers is supported in Edge Services Director.  
.....

### 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.

### Managing Device Configuration Files

You can back up device configuration files to the Edge Services 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.

## CHAPTER 17

# Device Management

- [Viewing the Device Inventory Page in Device View of Edge Services Director on page 340](#)
- [Resynchronizing Device Configuration on page 342](#)

## Viewing the Device Inventory Page in Device View of Edge Services Director

The Device Inventory page lists devices managed by Edge Services 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.

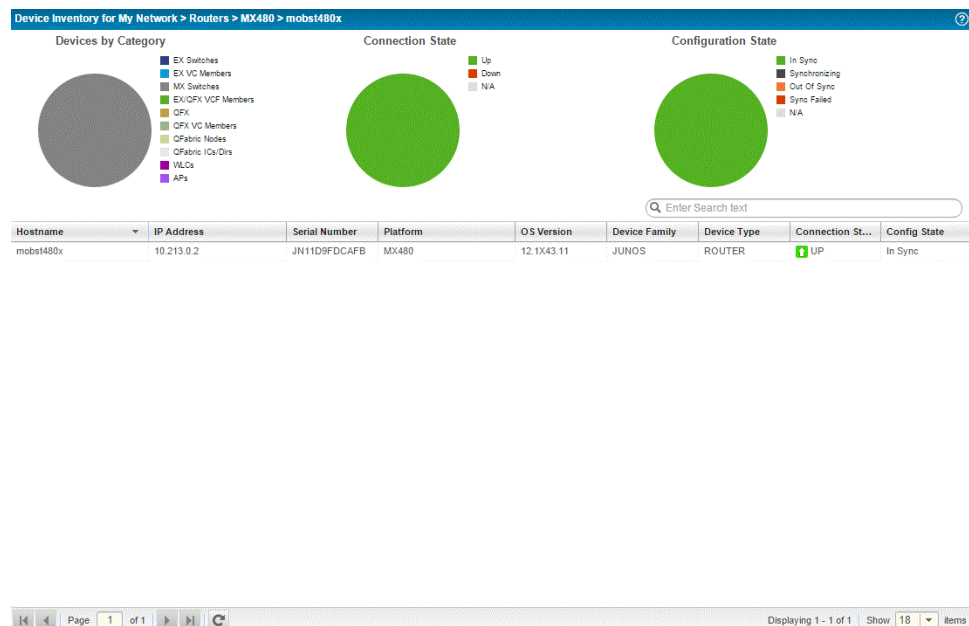
The scope you have selected in the View pane and the network view that you have selected from the View selector determines which devices are listed in the Device Inventory page. For example:

- If you are in the Device View and select My Network, all devices managed by Edge Services Director are listed.
- 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 Category—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 32 on page 134](#) for definitions of the configuration states.

Figure 32: Device Inventory Page



Mouse over a pie segment to view the actual number of devices and the percentage represented by that pie segment.

Table 32 on page 134 describes the fields in the Device Inventory table.

**Table 48: 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.
Device Family	Device family of the device, such as JUNOS for MX Series routers.
Device Type	Type of the device: <ul style="list-style-type: none"> <li>• ROUTER—MX Series routers</li> <li>• AP—Wireless LAN access point</li> <li>• Fabric Member—QFabric member switch</li> <li>• QFabric—QFabric system</li> <li>• Switch—Standalone switch</li> <li>• VC—Virtual Chassis master</li> <li>• VC Member—Virtual Chassis member switch</li> </ul>
Connection State	Connection status of the device in Edge Services Director: <ul style="list-style-type: none"> <li>• UP—Device is connected to Edge Services Director.</li> <li>• DOWN—Device is not connected to Edge Services Director.</li> <li>• N/A—Access point state is unavailable to Edge Services Director.</li> </ul>
Config State	Displays the configuration status of the device: <ul style="list-style-type: none"> <li>• In Sync—The configuration on the device is in sync with the Edge Services Director configuration for the device.</li> <li>• Out Of Sync—The configuration on the device does not match the Edge Services Director configuration for the device. This state is usually the result of the device configuration being altered outside of Edge Services Director. You cannot deploy configuration on a device from Edge Services Director when the device is Out Of Sync. To resolve this state, use the Resynchronize Device Configuration task in Deploy mode.</li> <li>• Sync failed—An attempt to resynchronize an Out Of Sync device failed.</li> <li>• Synchronizing—The device configuration is in the process of being resynchronized.</li> <li>• N/A—The device is down or is an access point.</li> </ul>

Table 48: Fields in the Device Inventory Table (continued)

Field	Description
Manageability State	Displays if the device is directly manageable or not.  This is a hidden field. To display the Manageability State field, click any column, click the down arrow to expand the list, select <b>Columns</b> from the list, and then enable <b>Manageability State</b> .

## Resynchronizing Device Configuration

A network managed by Edge Services 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 Edge Services Director.

When the configuration contained in all three repositories match, the device configuration state is shown as In Sync in Edge Services 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 Edge Services 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 Edge Services 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 Edge Services 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 Edge Services 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 Edge Services Director Build mode configuration.

For more information about out-of-band configuration changes, Junos Space SOR modes, and how Edge Services Director resynchronizes device configuration, see [“Understanding Resynchronization of Device Configuration” on page 69](#).

This topic covers:

- [The Resynchronize Device Configuration List of Devices on page 343](#)
- [Resynchronizing Devices When Junos Space Is in NSOR Mode on page 344](#)
- [Resynchronizing Devices When Junos Space Is in SSOR Mode on page 344](#)
- [Resynchronizing Devices in Manual Approval Mode on page 345](#)
- [Viewing the Network Changes on page 345](#)
- [Viewing Resynchronization Job Status on page 346](#)

## 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 49 on page 343](#) describes the fields in the list of devices.

**Table 49: 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	Connection state: <ul style="list-style-type: none"> <li>• UP—Edge Services Director is connected to the device</li> <li>• DOWN—Edge Services Director cannot connect to the device</li> </ul>
Configuration State	Shows the configuration state of the device: <ul style="list-style-type: none"> <li>• Out Of Sync—The device configuration is out of sync with either the Edge Services Director Build mode configuration or the Junos Space configuration record or both.</li> <li>• Resynchronizing—The device configuration is in the process of being resynchronized.</li> <li>• Sync Failed—The resynchronization attempt failed.</li> </ul> If the resynchronization is successful, the device is removed from the table.

Table 49: Resynchronize Device Configuration Fields (continued)

Field	Description
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"> <li>None—There are no configuration changes pending deployment.</li> <li>View—There are configuration changes that are pending deployment. Click <b>View</b> to view the changes. These changes will be lost if you resynchronize the Build mode configuration to match the device configuration.</li> </ul> <p><b>NOTE:</b> 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>
Network Changes	<p>Indicates whether you can view the out-of-band changes:</p> <ul style="list-style-type: none"> <li>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.</li> <li>View—You can view the out-of-band changes made on the device. Click <b>View</b> to view the changes presented in XML format.</li> </ul>

## 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 Edge Services 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 Edge Services 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 345](#).



**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 on page 69](#)
- [Understanding Deploy Mode in Gateway and Service Views of Edge Services Director on page 335](#)

## CHAPTER 18

# Configuration File Management

- [Managing Device Configuration Files on page 347](#)
- [Managing Jobs on page 351](#)

## Managing Device Configuration Files

---

You can back up device configuration files to the Edge Services 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 Edge Services 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 on page 347](#)
- [Backing Up Device Configuration Files on page 348](#)
- [Restoring Device Configuration Files on page 349](#)
- [Viewing Device Configuration Files on page 349](#)
- [Comparing Device Configuration Files on page 350](#)
- [Deleting Device Configuration Files on page 350](#)
- [Managing Device Configuration File Management Jobs on page 350](#)

## 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 348](#) for more information.
- Restore backup device configuration files to devices by selecting devices and clicking Restore. See [“Restoring Device Configuration Files” on page 349](#) for more information.
- View backed up configuration files by selecting a device and clicking View Configuration File. See [“Viewing Device Configuration Files” on page 349](#) for more information.
- Compare backed up device configuration files by selecting devices and clicking Compare Config Files. See [“Comparing Device Configuration Files” on page 350](#) for more information.
- Delete backup device configuration files by selecting devices and clicking Delete. See [“Deleting Device Configuration Files” on page 350](#) for more information.

Table 50 on page 348 describes the information provided in the Manage Device Configuration table.

*Table 50: 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 [“Understanding Build Mode in Location and Device Views of Edge Services Director”](#) on page 151.

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 Edge Services 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 Also**
- [Understanding Deploy Mode in Gateway and Service Views of Edge Services Director on page 335](#)

## Managing Jobs

Edge Services 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, Edge Services Director enables you to view special pre-filtered versions of this page from various other tasks, such as View Discovery Status or View Image Deployment 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 Edge Services 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 20 on page 60](#). 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 Edge Services Director will be available only from Edge Services Director. These jobs will not be listed in the Job Management pane in Junos Space platform and vice-versa.

**Table 51: Job Management Page Fields**

Field	Description
Job ID	The unique ID assigned to the job
Name	The name of the job

*Table 51: Job Management Page Fields (continued)*

Field	Description
Percent	The percentage of completion of the job
State	The status of the job: <ul style="list-style-type: none"><li>• Success—Job completed successfully</li><li>• Failure—Job failed and was terminated</li><li>• Job Scheduled—Job is scheduled but has not yet started</li><li>• In progress—Job is has started, but not completed</li><li>• Cancelled—Job is cancelled</li></ul>
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.

## CHAPTER 19

# Software Image Management

- [Managing Software Images on page 353](#)
- [Deploying Software Images on page 356](#)
- [Managing Software Image Deployment Jobs on page 359](#)

## Managing Software Images

---

This topic describes how to manage software images for managed devices.

To start managing software images:

1. Click **Deploy** in the Edge Services 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.

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 on page 353](#)
- [Adding Software Images to the Repository on page 354](#)
- [Using the Device Image Upload Window on page 354](#)
- [Viewing Software Image Details on page 355](#)
- [Using the Device Image Summary Window on page 355](#)
- [Deleting Software Images on page 355](#)

## 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 52 on page 354](#) describes the information provided in the Device Image Repository table.

*Table 52: 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 Edge Services 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 354](#) 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 355](#) 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 53 on page 355](#) describes the fields in this window.

*Table 53: 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.
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**

- [Understanding Deploy Mode in Gateway and Service Views of Edge Services Director on page 335](#)

## Deploying Software Images

---

This topic describes how to deploy software images to managed devices. You must upload software images to the Edge Services Director server before you can deploy them to devices. See *Managing Software Images* for more information.

To start deploying software images:

1. Click **Deploy** in the Edge Services Director banner.
2. Select a node in the View pane that contains the devices to 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 on page 356](#)
- [Selecting Software Images To Deploy on page 357](#)
- [Selecting Options for Software Deployment on page 358](#)
- [Summary of Software Deployment on page 359](#)

### 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 357](#).

## 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 54 on page 357](#) 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 pop-up message notifies you if you do not select a software image for all the listed devices. This is just for your information. No action will be 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 358](#).

**Table 54: Select images for devices Table**

Table Column	Description
Device Family	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.
Count	Number of devices contained within a device family.
IP Address	Device's IP address.
Device Name	Device's name.

Table 54: Select images for devices Table (continued)

Table Column	Description
State	Device's state: <ul style="list-style-type: none"> <li>• UP—Edge Services Director can communicate with the device.</li> <li>• DOWN—Edge Services Director cannot communicate with the device.</li> </ul>
Running Image Version	Software version the device is running.
Proposed Image Version/Profile	Software version that will be installed on the device when the job runs successfully.

## Selecting Options for Software Deployment

The options that you can configure in the Select Options page are described in [Table 55 on page 358](#). 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 359.

Table 55: Image Management Job Options

Option	Action
<b>Select Options</b>	
<b>All Device Types</b>	
Delete any existing image before download	Select to delete any existing software images on devices before downloading the new software image.
Reboot device after successful installation	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.  <b>NOTE:</b> This option may get disabled based on your details that you specify in the remaining fields. This indicates that for the options that you specified, the system will automatically reboot the device as per the requirement during or after the image upgrade.
<b>Wired Devices</b>	
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	Select if you want to perform a Nonstop software upgrade (NSSU) or lin-service software upgrade (ISSU).  ISSU enables you to upgrade between two different Junos OS releases with minimal disruption on the control plane and with minimal disruption of traffic.  NSSU enables you to upgrade the software running on an MX Series router with redundant Routing Engines or on most EX Series Virtual Chassis by using a single command and with minimal disruption to network traffic

Table 55: Image Management Job Options (continued)

Option	Action
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	Select to copy the new Junos OS image into the alternate root partition. This ensures that the resilient dual-root partitions feature operates correctly.  This option is available only if you select <b>Reboot device after successful installation</b> .
<b>Select Schedule</b>	
Stage now	Select <b>Stage now</b> to start staging software images to devices as soon as the job runs.
Stage later time	Select <b>Stage later time</b> 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 <b>Upgrade now</b> to start upgrading software images on devices as soon as staging finishes.
Upgrade later time	Select <b>Upgrade later time</b> to schedule the software upgrade for a later time.
Deployment Schedule	If you selected Upgrade later time, enter the date and time for upgrade to start.  If you scheduled staging, you must schedule the upgrade for at least 10 minutes after staging, to ensure that staging completes before upgrade starts.

## 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**

- [Understanding Deploy Mode in Gateway and Service Views of Edge Services Director on page 335](#)

## 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 Edge Services 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 on page 360](#)
- [Viewing Software Image Job Details on page 361](#)
- [Using the Device Image Staging Window on page 361](#)
- [Canceling Software Image Jobs on page 362](#)

## 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 361](#) for more information.
- Cancel a pending job by selecting the job and clicking Cancel Job. See [“Canceling Software Image Jobs” on page 362](#) for more information.

[Table 56 on page 360](#) describes the information provided in the of the Image Deployment Jobs table.

**Table 56: 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"> <li>• CANCELLED—The job was cancelled by a user.</li> <li>• SCHEDULED—The job is scheduled but has not run yet.</li> <li>• INPROGRESS—The job is running.</li> <li>• SUCCESS—The job completed successfully. This state is applied if all of the devices in the job completed successfully.</li> <li>• 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.</li> </ul>
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.

*Table 56: Image Deployment Jobs Table (continued)*

Table Column	Description
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 361](#) 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 57 on page 361](#) describes this window.

*Table 57: 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.
Status	Job status. The possible statuses are: <ul style="list-style-type: none"> <li>• CANCELLED—The job was cancelled by a user.</li> <li>• SCHEDULED—The job is scheduled but has not run yet.</li> <li>• INPROGRESS—The job is running.</li> <li>• SUCCESS—The job completed successfully.</li> <li>• FAILURE—The job failed.</li> </ul>
Host Name	Host name of device.
Status	Device status. The possible statuses are: <ul style="list-style-type: none"> <li>• INPROGRESS—The job is running.</li> <li>• SUCCESS—The job completed successfully.</li> <li>• FAILURE—The job failed.</li> </ul>

*Table 57: Device Image Staging Window Description (continued)*

Field	Description
% 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**

- [Understanding Deploy Mode in Gateway and Service Views of Edge Services Director on page 335](#)

## CHAPTER 20

# Viewing and Editing Service Instances and Packet Filters Across All Gateways

- [Viewing Service Object Statistics on page 363](#)
- [Modifying Service Instances on page 365](#)
- [Modifying Packet Filter Policies on page 367](#)

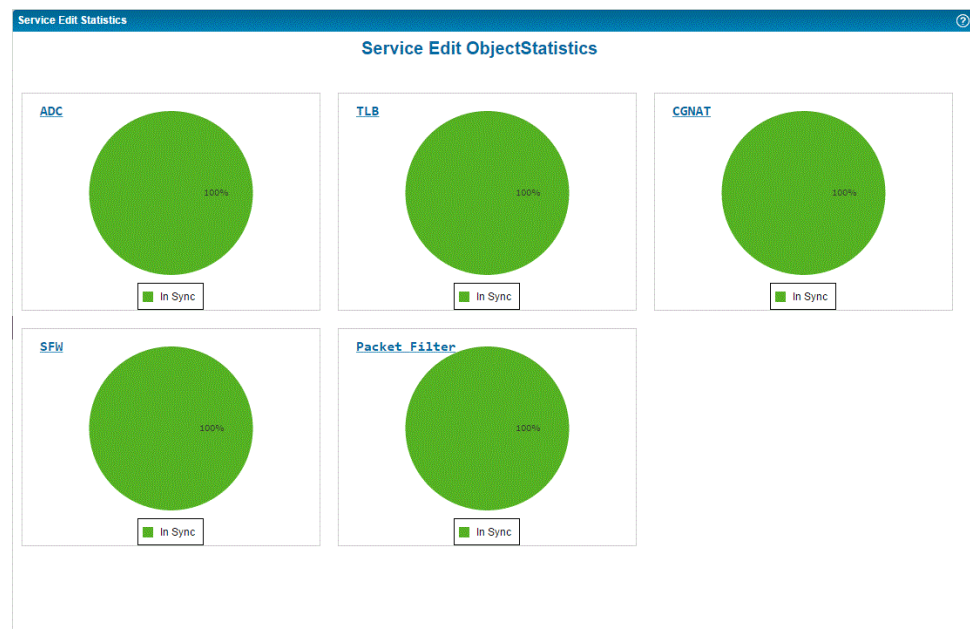
### Viewing Service Object Statistics

---

To view a graphical representation in the form of pie charts of the configured ADC, TLB, CGNAT, SFW, and packet policies or filters:

1. From the View selector, select **Gateway View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the View pane, select the All Network item in Gateway view. Click the plus sign (+) beside the All Network item in the View pane to expand the tree and select the device node you want.
4. From the task pane, select **Service Edit**. On the right pane, pie charts corresponding to the configured services and policy filters are displayed if you view the page without drilling-down the tree in the task pane to select a particular service or policy. The same Service Object Statistics page is displayed when you select **View Statistics** from the task pane.
5. In the View pane, from the tree that lists the SDGs, select **All SDG**, or the SDG or SDG pair for which you want to view the previously configured policy or filter templates.  
The page is divided into three panes. The list of SDGs are displayed on the left pane. You can drill-down to the SDG or pair of SDGs for which you want to view the service statistics.

Figure 33: Service Edit Page with Pie Charts of Configured Service Types



The Service Object Statistics page is displayed. A set of five pie charts are displayed when you select Service Edit from the task pane, without expanding the tree and selecting a policy and filter template. The pie charts are displayed for the different policy and service filters, such as ADC, TLB, CGNAT, stateful firewall, and packet filter templates. A color-code is used to denote different portions of the pie chart for the service policy filters in various states. Mouse over each portion of the pie to view the number corresponding to the percentage of each service policy filter in a particular state. The following segments are displayed in the pie chart as a percentage of the total number of service policy filters.

- In Sync—The configuration on the device is in sync with the Edge Services Director configuration for the device.
- Out Of Sync—The configuration on the device does not match the Edge Services Director configuration for the device. This state is usually the result of the device configuration being altered outside of Edge Services Director. You cannot deploy configuration on a device from Edge Services 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 or is an access point.

#### Related Documentation

- [Policy and Filter Management Overview on page 375](#)
- [Packet and Service Filters Overview on page 378](#)
- [Searching for CGNAT Policies on page 381](#)

- [Searching for Packet Filters on page 384](#)
- [Searching for SFW Policies on page 386](#)
- [Managing Service and Policy Locks on page 387](#)
- [Unlocking Locked Services and Policies on page 389](#)

## Modifying Service Instances

---

On the Service Designer page, you can view the collection of service templates defined for several applications, such as stateful firewall or CGNAT.

To modify service template instances, such as ADC, SFW, CGNAT, or TLB templates:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.

3. From the task pane, select **Service Edit**.

The Service Instances page is displayed in the right pane, listing all the previously defined service templates.

4. From the View pane, perform one of the following tasks:

- Click the **ADC** button.

The list of ADC service templates is displayed. You need not click this button if you are launching the Service Designer page for the first time or are navigating to this page from another mode or a different page. You need to click this button only if you are viewing the other service templates, such as CGNAT or TLB.

- Click the **SFW** button.

The list of SFW templates is displayed.

- Click the **TLB** button.

The list of TLB templates is displayed.

- Click the **CGNAT** button.

The list of CGNAT templates is displayed.

5. Click the **Lock** icon above the table of listed packet filters. The Select Reference Config dialog box is displayed.
6. From the Service Gateway Name drop-down list, select the SDG group to which the packet filter must be applied.

7. From the Host Name drop-down list, select the hostname of the SDG.
8. In the Select Common Components section, select the check boxes beside the service modules or components, such as packet filters, SFW rules, or CGNAT rules, that are displayed. The displayed components depend on the attributes that are previously defined for that selected packet filter. For example, if the service policy is for stateful firewall, SFW rules and SFW rule sets are shown. Select the check box beside Config Category to select all the service components.

Starting with Edge Services Director Release 1.1, you can associate multiple client-facing and server-facing VRF instances from the enhanced service edit mode (which you can access from Service View of Deploy mode, with TLB selected in View pane and Service Edit selected in the Tasks pane, and selecting the check boxes beside the Server-Facing and Client-Facing modules in the Select Common Components section).

9. Click **Save** to save the modified association.
10. Select the check box beside the template you want to modify.
11. Open the **Modify** menu above the list of templates to modify an existing template, and select the component or service attribute, such as application or rule, that you want to edit.
12. Perform one of the following from the drop-down menu displayed for each component:
  - To retrieve the service component and import into the database of Edge Services Director, select **Import Object**. The Import Services dialog box appears. You can import the service templates assigned to SDGs or choose from a list of all of the predefined templates in the database. Also, you can either import all of the components of a service or specific components.
  - To create the component afresh, select **Create New**. The Create page corresponding to the service component appears. You can define the attributes for the service component in the same manner as you define the elements during the creation of a service template.
13. Click the down arrow in the **Actions** menu and select **Send for Deployment** to create a deployment plan for the particular service template and save the plan.

If you create a deployment plan from Service view of Deploy mode, the Edit Service Instance page is displayed. You can modify the SDGs associated with the service instance and also modify the service instance attributes as necessary by either clicking the buttons corresponding to the various settings at the top of the wizard page to directly traverse to the page you want to modify or clicking the navigation buttons at the bottom of the wizard page to go to the different pages of the wizard. Click **Finish** to create a deployment plan.

A deployment plan is created for the service template with the devices that are assigned to it when you view the Deployment Plans page.

From the Deployment plans page, you can select **Reject** or **Approve** from the Actions drop-down list to reject or approve the deployment plan and make it available for commissioning to the devices.

14. Select **Discard changes** from the Actions menu to ignore the modifications done to a policy or filter template.

**Related  
Documentation**

- [Service Templates Overview on page 189](#)
- [Filtering Service Templates on page 189](#)
- [Using the Actions Menu on the Service Template and Service Edit Pages on page 195](#)

---

## Modifying Packet Filter Policies

On the Packet Filter Policies page, you can view the collection of previously configured packet filters and perform an enhanced edit to select a different SDG group and an SDG host in the group to be associated with the packet filter.

To modify packet filter services and specify the SDG group, SDG host, and service attributes to be associated with the packet filters:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the View pane, click the plus sign (+) beside All Services to expand the tree and select the type of service.
4. From the task pane, select **Deploy Service > Packet Filter**.  
The Packet Filter Policies page is displayed on the right pane, listing all the previously defined packet filters.
5. Click the **Lock** icon above the table of listed packet filters. The Select Reference Config dialog box is displayed.
6. From the Service Gateway Name drop-down list, select the SDG group to which the packet filter must be applied.
7. From the Host Name drop-down list, select the hostname of the SDG.
8. In the Select Common Components section, select the check boxes beside the service modules or components, such as packet filters, that are displayed. The displayed

components depend on the attributes that are previously defined for that selected packet filter. Select the check box beside Config Category to select all the service components.

9. Click **Save** to save the modified association.
10. Select the check box beside the packet filter you want to modify.
11. Open the **Modify** menu above the list of templates to modify an existing packet filter.

The Modify Packet Filter window is displayed. Modify the attributes that are needed and save the updated settings.
12. Click the down arrow in the **Actions** menu and select **Send for Deployment** to create a deployment plan for the particular service template and save the plan.
  - If you create a deployment plan from Gateway view of Deploy mode, the Deployment Plan Summary dialog box appears, with the service name, type, and status listed.

Click **Send** to create a deployment plan.
  - If you create a deployment plan from Service view of Deploy mode, the Edit Service Instance page is displayed. You can modify the SDGs associated with the service instance and also modify the service instance attributes as necessary by either clicking the buttons corresponding to the various settings at the top of the wizard page to directly traverse to the page you want to modify or clicking the navigation buttons at the bottom of the wizard page to go to the different pages of the wizard. Click **Finish** to create a deployment plan.

A deploy plan is created for the service template with the devices that are assigned to it when you view the Deployment Plans page.

From the Deployment plans page, you can select **Reject** or **Approve** from the Actions drop-down list to reject or approve the deployment plan and make it available for commissioning to the devices.
13. Select **Discard changes** from the Actions menu to ignore the modifications done to a packet policyfilter.

**Related  
Documentation**

- [Service Templates Overview on page 189](#)
- [Filtering Service Templates on page 189](#)
- [Using the Actions Menu on the Service Template and Service Edit Pages on page 195](#)

# Enhanced Editing of Services and Packet Filters

- [Enhanced Editing of Service Policies and Policy Filters Overview on page 369](#)
- [Modifying the Association of SDG Details and Service Components for a Packet Filter Policy on page 370](#)
- [Modifying the Association of SDG Details and Service Components for a Service Policy Filter on page 372](#)

## Enhanced Editing of Service Policies and Policy Filters Overview

---

In Gateway View of Deploy mode, with All Network selected in View pane and Policy & Filters selected in the task pane, you can select a different SDG host from the Host Name list, and a different rule term from the Term Name list from the page that lists all of the previously defined service policies. This type of inline or embedded editing enables you to quickly and optimally change the rule term in a service policy and the SDG with which the policy must be associated.

Inline modification signifies the ability to perform changes to previously defined settings in an easy and quick manner. Embedded editing is enabled, which causes the grids showing the devices and interfaces to become modifiable directly without the need to perform the process of highlighting, editing, and saving the changes every time you want to edit a particular parameter. The page that displays the configured settings presents as a form in which the fields or cells of the table are editable.

Instead of modifying an existing stateful firewall, NAT, or packet filter policy to associate a different SDG host with the policy by using the Service Edit option in the task pane in Service View of Deploy mode, you can easily and rapidly change the SDG host mapped to a policy using the enhanced editing mechanism.

### Related Documentation

- [Modifying the Association of SDG Details and Service Components for a Packet Filter Policy on page 370](#)
- [Modifying the Association of SDG Details and Service Components for a Service Policy Filter on page 372](#)

## Modifying the Association of SDG Details and Service Components for a Packet Filter Policy

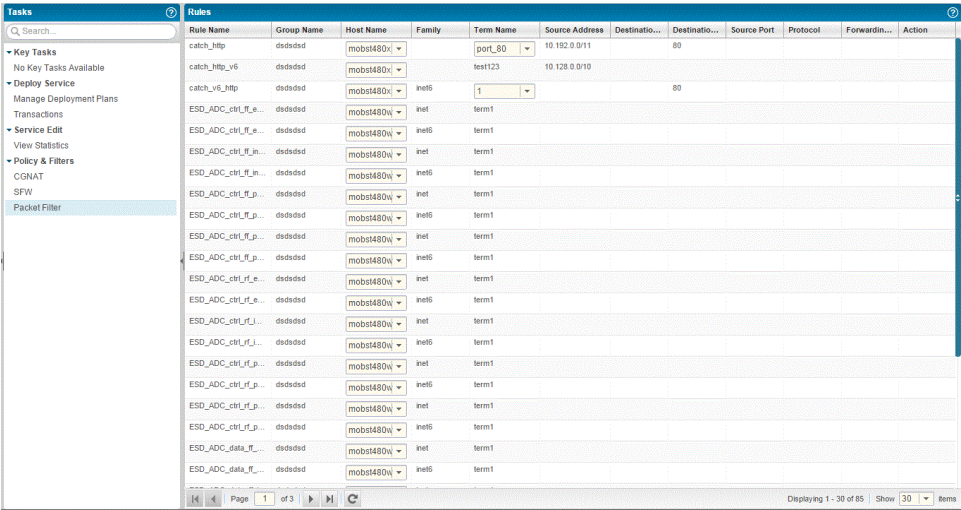
---

From the Policy & Filters page, which displays all the previously configured packet filters, you can modify the components or the parameter types that are associated with a particular service filter. You must lock the packet filters for which you want to modify the attached rule term components or attributes before you can update the settings. You can also select a different SDG to which the packet filter must be applied.

To modify the association of SDGs and the rule term component for a packet filter, such as a stateless firewall filter:

1. From the View selector, select **Service View**. The workspaces that are applicable to edge services are displayed.
2. Select All Network from the Service View pane. You can modify the association of SDGs with service policies, only if you select the All Network label in the View pane. If you expand the All Network tree and select an SDG group or an SDG in a redundancy pair, you cannot modify the association of service policies and rules with SDGs in a single-shot, one-step operation.
3. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
4. From the View pane, select the All Network item in Gateway view. Click the plus sign (+) beside the All Network item in the View pane to expand the tree and select the device node you want.
5. Select **Policy & Filters** from the task pane.  
The Services page is displayed.
6. Click the down arrow next to Policy & Filters to expand the tree in the task pane and view the list of filter templates.  
Select **Packet Filter** to open the Service Edit > Packet Filter page on the right pane.

Figure 34: Enhanced Edit Page for Packet Filters



The following fields are displayed on this page:

Table 58: Service Edit > Packet Filter Page

Field	Description
Instance Name	Name of the configured service template instance
OS Version	Junos OS release number that represents a particular revision of the software that runs on a Juniper Networks routing platform, for example, Junos OS Release 8.5, 9.1, or 9.2. Each Junos OS release has certain new features that complement the software processes that support Internet routing protocols, control the device's interfaces and the device chassis itself, and allow device system management.
Group Name	Name of the SDG group
Reference Host	Hostname of the SDG with which the service instance is associated.
Deployment Plans	Name of the deployment plan with which the service template is attached.

- From the Term Name drop-down list, select the rule term with which the packet filter must be applied.
- From the Host Name drop-down list, select the hostname of the SDG.
- In the Select Common Components section, select the check boxes beside the service modules or components, such as packet filters, SFW rules, or CGNAT rules, that are displayed. The displayed components depend on the attributes that are previously defined for that selected packet filter. For example, if the service policy is for stateful

firewall, SFW rules and SFW rule sets are shown. Select the check box beside Config Category to select all the service components.

The modified association is saved.

You can use the **Actions** menu in the Service Template pages for packet filters to publish, unpublish, export, and restore the defined policies or filters. For details, see *Using the Actions Menu in the Service Template Page*.

**Related  
Documentation**

- [Creating Service Gateway Groups on page 99](#)
- [Managing Service Gateway Groups on page 101](#)
- [Searching Unmanaged Devices on page 104](#)
- [Viewing the List of Discovered, Managed, and Unmanaged Devices on page 106](#)
- [Modifying Discovery Profiles on page 113](#)
- [Deleting Discovery Profiles on page 114](#)

---

## Modifying the Association of SDG Details and Service Components for a Service Policy Filter

---

From the Policy & Filters page, which displays all the previously configured service policy filters, you can modify the components or the parameter types that are associated with a particular service filter. You must lock the service policy filters for which you want to modify the attached service components or attributes before you can update the settings. You can also select a different SDG to which the service policy filter must be applied.

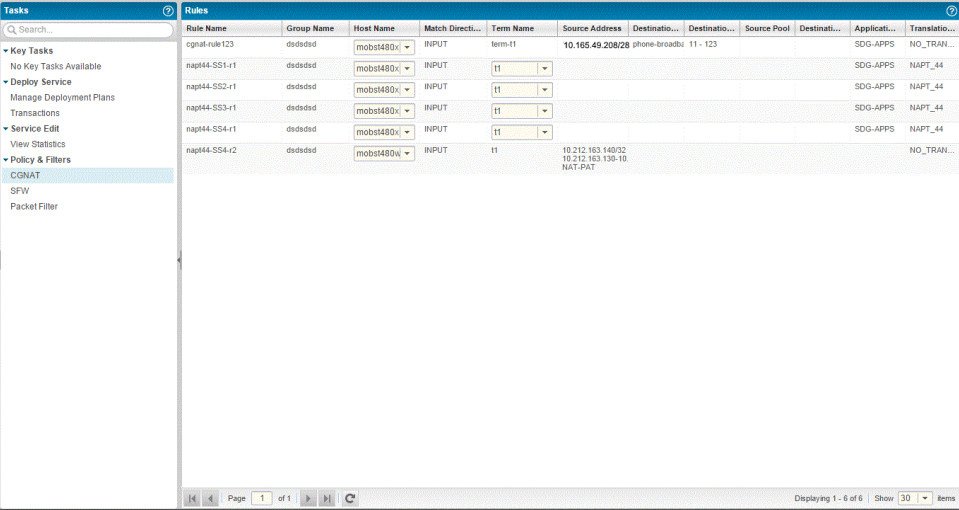
To modify the association of SDGs and service components for a service policy filter, such as a stateful firewall service, or a carrier-grade NAT service policy:

1. From the View selector, select **Service View**. The workspaces that are applicable to edge services are displayed.
2. Select All Network from the Service View pane. You can modify the association of SDGs with service policies, only if you select the All Network label in the View pane. If you expand the All Network tree and select an SDG group or an SDG in a redundancy pair, you cannot modify the association of service policies and rules with SDGs in a single-shot, one-step operation.
3. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
4. From the View pane, select the All Network item in Gateway view. Click the plus sign (+) beside the All Network item in the View pane to expand the tree and select the device node you want.

5. Select **Policy & Filters** from the task pane.

The Service Edit page is displayed.

Figure 35: Enhanced Edit Page for Service Policy Rules



6. Click the plus sign (+) next to Policy & Filters to expand the tree in the task pane and view the list of filter templates. Do one of the following:

- Select **CGNAT** to open the Service Edit > CGNAT page on the right pane.
- Select **SFW** to open the Service Edit > SFW page on the right pane.

The following fields are displayed on this page:

Table 59: Services – CGNAT and SFW Page

Field	Description
Instance Name	Name of the configured service template instance
OS Version	Junos OS release number that represents a particular revision of the software that runs on a Juniper Networks routing platform, for example, Junos OS Release 8.5, 9.1, or 9.2. Each Junos OS release has certain new features that complement the software processes that support Internet routing protocols, control the device's interfaces and the device chassis itself, and allow device system management.
Group Name	Name of the SDG group
Reference Host	Hostname of the SDG with which the service instance is associated.
Applications	Name of the applications protocols created for the service template.
Application Sets	Name of the application sets created for the service template.
SFW Rules	Name of the stateful firewall rules created for the service instance.

*Table 59: Services – CGNAT and SFW Page (continued)*

Field	Description
SFW Rule Sets	Name of the stateful firewall rule sets created for the service template.
NAT Pools	Name of the CGNAT pool created for the service template.
NAT Rules	Name of the CGNAT rules created for the service instance.
NAT Rule Sets	Name of the CGNAT rule sets created for the service template.
Syslogs	Name of the syslog created for the service template.
Deployment Plans	Name of the deployment plan with which the service template is attached.

7. From the Term Name drop-down list, select the rule term that must be assigned to the service policy filter, such as CGNAT or stateful firewall service policies.

8. From the Host Name drop-down list, select the hostname of the SDG.

The modified association is saved.

You can use the **Actions** menu in the Service Template pages for CGNAT, SFW, and packet filters to publish, unpublish, export, and restore the defined policies or filters. For details, see *Using the Actions Menu in the Service Template Page*.

#### Related Documentation

- [Creating Service Gateway Groups on page 99](#)
- [Managing Service Gateway Groups on page 101](#)
- [Searching Unmanaged Devices on page 104](#)
- [Viewing the List of Discovered, Managed, and Unmanaged Devices on page 106](#)
- [Modifying Discovery Profiles on page 113](#)
- [Deleting Discovery Profiles on page 114](#)

## CHAPTER 22

# Managing Service Instance and Policy Rule Definitions

- [Policy and Filter Management Overview on page 375](#)
- [Packet and Service Filters Overview on page 378](#)
- [Searching for CGNAT Policies on page 381](#)
- [Searching for Packet Filters on page 384](#)
- [Searching for SFW Policies on page 386](#)
- [Managing Service and Policy Locks on page 387](#)
- [Unlocking Locked Services and Policies on page 389](#)
- [Viewing Policy and Filter Instances on page 390](#)
- [Creating and Managing CGNAT Policy and Filter Instances on page 395](#)
- [Creating and Managing Packet Filter Policy Instances on page 419](#)
- [Creating and Managing SFW Policy and Filter Instances on page 434](#)
- [Viewing CGNAT Service Templates on page 456](#)
- [Viewing SFW Service Templates on page 457](#)
- [Viewing and Modifying ADC Service Instances on page 459](#)
- [Viewing and Modifying TLB Service Instances on page 471](#)
- [Using the Actions Menu on the Service Policy and Packet Filter Pages on page 483](#)
- [Tagging Junos Space Network Management Platform Objects on page 485](#)

## Policy and Filter Management Overview

---

The Policy and Filter Management feature in the Junos Space Edge Services Director application takes care of creation, update, display, publish and commission of packet filters, stateful firewall and NAT policies present on discovered and managed SDGs. The Service Management workspace displays a bar graph of draft, published, and approved filters or policies for different options available under workspace.

- **Packet Filter:** This option displays packet filters present on SDGs in a tabular layout. It also provides the ability to create, update and delete filters on selected SDGs.

- **Stateful Firewall:** This option displays stateful firewall policies present on SDGs in a tabular layout. It also provides the ability to create, update, and delete stateful firewall policies on selected SDGs.
- **CGNAT:** This option displays CGNAT policies present on SDGs in a tabular layout. It also provides the ability to create, update, and delete CGNAT policies on selected SDGs. After a filter or policy is published, it goes for peer review and approval. After approval, the filter or policy is deployed to the device.

The Service Deployment page provides the following functionalities:

- **1. Approval Management** – View the details of the filters/policies and other service deployment plans which are pending for approval. Approve or reject deployment plans done to existing feature.
- **2. Update Devices** – View the details of approved filters/policies and other service deployment plans which are ready for commissioning. Commission the deployment plans or discard accordingly.

## States and Transitions of Policies or Filters

A filter has the following states:

- New
- Updated
- Deleted

A user can carry out following operations depending on the status of a filter:

- **Add** – To create a new filter for Zone, SDG or Host.
- **Update** – Update exiting filter on SDG.
- **Delete** – Delete existing filter on SDG.
- **Send for Deployment**—Deploy the policy and filter instance on the associated standalone SDG or SDGs in a high availability pair.

You can perform the following tasks with a deployment plan created for provisioning a policy on SDGs:

- **Publish** – Publish new, updated or deleted filter for administrator or designer approval.
- **Unpublish** – Unpublish the published filter to do more changes. The filter returns to the “Draft” status.
- **Approve** – An administrator or designer approves the published filter.
- **Reject** – An administrator or designer rejects the published filter.
- **Commission** – An administrator or designer pushes updates to SDG.
- **Discard** – An administrator or designer discards an approved filter without pushing updates to SDG.

## User Roles

SDG operator is responsible for creating, modifying, and deleting a policy or filter and publishes it for approval of the designer. SDG operator can access the Service Management workspace and all options under it.

A user with the SDG designer role is responsible for review and approval of published policy or filter. Workflow for review and approval is part of another workspace called Service Deployment. As a user with the SDG designer role, you can access both Service Management' and 'Service Deployment workspaces.

SDG Administrator is responsible for commissioning of an approved policy or filter to managed SDGs. Workflow for commissioning will be part of another workspace called Service Deployment. An SDG designer can access both the Service Management and Service Deployment workspaces.

- SDG Operator – An SDG operator is responsible for creating, modifying, and deleting a policy or filter and will publish it for approval of designer. An SDG operator can access the Service Management workspace and all options under it.
- SDG Designer – An SDG designer is responsible for review and approval of a published policy or filter. The workflow for review and approval is part of another workspace called Service Deployment. An SDG designer can access both the Service Management and Service Deployment workspaces.
- SDG Administrator – An SDG administrator is responsible for commission of approved policy or filter to managed SDGs. The workflow for commissioning is part of another workspace called Service Deployment. An SDG designer can access both the Service Management and Service Deployment workspaces.

### Related Documentation

- [Policy and Filter Management Overview on page 375](#)
- [Searching for CGNAT Policies on page 381](#)
- [Searching for Packet Filters on page 384](#)
- [Searching for SFW Policies on page 386](#)
- [Managing Service and Policy Locks on page 387](#)
- [Unlocking Locked Services and Policies on page 389](#)
- [Viewing Policy and Filter Instances on page 390](#)

## Packet and Service Filters Overview

---

The Adaptive Services Physical Interface Cards (PICs), Multiservices PICs, and Multiservices Dense Port Concentrators (DPCs) provide *adaptive services interfaces*. Adaptive services interfaces enable you to coordinate a special range of services on a single PIC or DPC by configuring a set of services and applications.

A service set is an optional definition you can apply to the traffic at an adaptive services interface. A service set enables you to configure combinations of directional rules and default settings that control the behavior of each service in the service set. When you apply a service set to the traffic at an adaptive services interface, you can optionally use service filters to refine the target of the set of services and also to process traffic. Service filters enable you to manipulate traffic by performing packet filtering to a defined set of services on an adaptive services interface before the traffic is delivered to its destination. You can apply a service filter to traffic before packets are accepted for input or output service processing or after packets return from input service processing.

A service filter defines packet-filtering (a set of match conditions and a set of actions) for IPv4 or IPv6 traffic. You can apply a service filter to the inbound or outbound traffic at an adaptive services interface to perform packet filtering on traffic before it is accepted for service processing. You can also apply a service filter to the traffic that is returning to the services interface after service processing to perform postservice processing.

Service filters filter IPv4 and IPv6 traffic only and can be applied to logical interfaces on Adaptive Services PICs, MultiServices PICs, and MultiServices DPCs only.

The Junos OS standard stateless firewall filters support a rich set of packet-matching criteria that you can use to match on specific traffic and perform specific actions, such as forwarding or dropping packets that match the criteria you specify. You can configure firewall filters to protect the local router or to protect another device that is either directly or indirectly connected to the local router. For example, you can use the filters to restrict the local packets that pass from the router's physical interfaces to the Routing Engine. Such filters are useful in protecting the IP services that run on the Routing Engine, such as Telnet, SSH, and BGP, from denial-of-service attacks.



**NOTE:** If you configured targeted broadcast for virtual routing and forwarding (VRF) by including the `forward-and-send-to-re` statement, any firewall filter that is configured on the Routing Engine loopback interface (lo0) cannot be applied to the targeted broadcast packets that are forwarded to the Routing Engine. This is because broadcast packets are forwarded as flood next hop traffic and not as local next hop traffic, and you can only apply a firewall filter to local next hop routes for traffic directed toward the Routing Engine.

---

You can configure service filters to filter IPv4 traffic (**family inet**) and IPv6 traffic (**family inet6**) only. No other protocol families are supported for service filters.

Under the **family inet** or **family inet6** statement, you can include **service-filter service-filter-name** statements to create and name service filters. The filter name can

contain letters, numbers, and hyphens (-) and be up to 64 characters long. To include spaces in the name, enclose the entire name in quotation marks (" ").

Under the **service-filter** *service-filter-name* statement, you can include **term** *term-name* statements to create and name filter terms.

Service filter terms support only a subset of the IPv4 and IPv6 match conditions that are supported for standard stateless firewall filters.

If you specify an IPv6 address in a match condition (the **address**, **destination-address**, or **source-address** match conditions), use the syntax for text representations described in RFC 4291, *IP Version 6 Addressing Architecture*.

When configuring a service filter term, you must specify one of the following filter-terminating actions:

- **service**
- **skip**



**NOTE:** These actions are unique to service filters.

Service filter terms support only a subset of the IPv4 and IPv6 nonterminating actions that are supported for standard stateless firewall filters:

- **count** *counter-name*
- **log**
- **port-mirror**
- **sample**

Service filters do not support the **next** action.

## Filtering Traffic Before Accepting Packets for Service Processing

To filter IPv4 or IPv6 traffic before accepting packets for input or output service processing, include the **service-set** *service-set-name* **service-filter** *service-filter-name* at one of the following interfaces:

- [edit interfaces *interface-name* unit *unit-number* family (inet | inet6) service input]
- [edit interfaces *interface-name* unit *unit-number* family (inet | inet6) service output]

For the **service-set-name**, specify a service set configured at the [edit services **service-set**] hierarchy level.

The service set retains the input interface information even after services are applied, so that functions such as filter-class forwarding and destination class usage (DCU) that depend on input interface information continue to work.

The following requirements apply to filtering inbound or outbound traffic before accepting packets for service processing:

- You configure the same service set on the input and output sides of the interface.
- If you include the **service-set** statement without an optional **service-filter** definition, the Junos OS assumes the match condition is true and selects the service set for processing automatically.
- The service filter is applied only if a service set is configured and selected.

You can include more than one service set definition on each side of an interface. The following guidelines apply:

- If you include multiple service sets, the router (or switch) software evaluates them in the order in which they appear in the configuration. The system executes the first service set for which it finds a match in the service filter and ignores the subsequent definitions.
- A maximum of six service sets can be applied to an interface.
- When you apply multiple service sets to an interface, you must also configure and apply a service filter to the interface.

## Postservice Filtering of Returning Service Traffic

As an option to filtering of IPv4 or IPv6 input service traffic, you can apply a service filter to IPv4 or IPv6 traffic that is returning to the services interface after the service set is executed. To apply a service filter in this manner, include the **post-service-filter service-filter-name** statement at the **[edit interfaces interface-name unit unit-number family (inet | inet6) service input]** hierarchy level.

### Related Documentation

- [Policy and Filter Management Overview on page 375](#)
- [Searching for CGNAT Policies on page 381](#)
- [Searching for Packet Filters on page 384](#)
- [Searching for SFW Policies on page 386](#)
- [Managing Service and Policy Locks on page 387](#)
- [Unlocking Locked Services and Policies on page 389](#)
- [Viewing Policy and Filter Instances on page 390](#)

## Searching for CGNAT Policies

---

You can use the enhanced search utility on the Service Templates page for CGNAT policies and packet filters to effectively, quickly identify and segregate the policies and filters of relevance and interest.

The Service Templates page provides advanced search options for the CGNAT policies. Enter the term that you want to specify as the filter criterion in the search field and click the **Search** icon.

You can perform advanced searches for the following fields:

- Policy Name
- Source Address
- Destination Port
- Destination Address
- Application
- Translation Type
- NAT Pool
- Description
- Custom column

The following advanced search criteria are available:

- Wildcard search for rule names using an asterisk (\*) is allowed.
- Edge Services Director supports AND and OR operations between search items. The default behavior is OR.
- For rule name search, only the OR operation is allowed, because a policy cannot have multiple rule names.
- For zone search, only the OR operation is allowed. Wildcard search is supported.
- For service and address fields, OR and AND operations are allowed.
- Multiple groups can be grouped using parenthesis. Grouping can be used during filed or keyword searches as well.
- Negate (-) symbol can be used to exclude objects that contain a specific term name.
- The plus (+) operator can be used to specify that the term after the + symbol existing the field value to be filtered along with other searched items.
- Escaping special characters are part of the search syntax. The supported special characters are + - & | ! ( ) { } [ ] ^ " ~ \* ? : \.



**NOTE:** Use the AND operator to find rules that match all values for a given set of fields. Use the OR operator to find rules that match any of the values for a given set of fields.

---

1. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
2. Select **Service Edit > Policy and Filter** from the task pane. The Policy and Filter page is displayed.
3. Click the plus sign (+) next to the policy and filter template to expand the tree in the task pane and view the list of filter templates.
4. From the task pane, select **CGNAT Policy and Filter** to open the CGNAT and Filter page on the right pane.
5. Enter the term that you want to specify as the filter criterion in the Search field and click the **Search** icon.

**Related  
Documentation**

- [Policy and Filter Management Overview on page 375](#)
- [Packet and Service Filters Overview on page 378](#)
- [Searching for Packet Filters on page 384](#)
- [Searching for SFW Policies on page 386](#)
- [Managing Service and Policy Locks on page 387](#)
- [Unlocking Locked Services and Policies on page 389](#)
- [Viewing Policy and Filter Instances on page 390](#)

## Searching for Packet Filters

---

You can use the enhanced search utility on the Service Templates page for CGNAT policies and packet filters to effectively, quickly identify and segregate the policies and filters of relevance and interest.

The Service Templates page provides advanced search options for the packet filters. Enter the term that you want to specify as the filter criterion in the Filter field and click the **Filter** icon.

You can perform advanced searches for the following fields:

- Filter Name
- Source Port
- Source Address
- Destination Port
- Destination Address
- Action
- Description
- Custom column

The following advanced search criteria are available:

- Wildcard search for rule names using an asterisk (\*) is allowed.
- Edge Services Director supports AND and OR operations between search items. The default behavior is OR.
- For rule name search, only the OR operation is allowed, because a policy cannot have multiple rule names.
- For zone search, only the OR operation is allowed. Wildcard search is supported.
- For service and address fields, OR and AND operations are allowed.
- Multiple groups can be grouped using parenthesis. Grouping can be used during filed or keyword searches as well.
- Negate (-) symbol can be used to exclude objects that contain a specific term name.
- The plus (+) operator can be used to specify that the term after the + symbol existing the field value to be filtered along with other searched items.
- Escaping special characters are part of the search syntax. The supported special characters are + - & | ! ( ) { } [ ] ^ " ~ \* ? : \.



**NOTE:** Use the AND operator to find rules that match all values for a given set of fields. Use the OR operator to find rules that match any of the values for a given set of fields.

---

1. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
2. Select **Service Edit > Policy and Filter** from the task pane. The Policy and Filter page is displayed.
3. Click the plus sign (+) next to Service Template to expand the tree in the task pane and view the list of filter templates.
4. From the task pane, select **Packet Filter** to open the Packet Filter page on the right pane.
5. Enter the term that you want to specify as the filter criterion in the Search field and click the **Search** icon.

**Related  
Documentation**

- [Policy and Filter Management Overview on page 375](#)
- [Packet and Service Filters Overview on page 378](#)
- [Searching for CGNAT Policies on page 381](#)
- [Searching for SFW Policies on page 386](#)
- [Managing Service and Policy Locks on page 387](#)
- [Unlocking Locked Services and Policies on page 389](#)
- [Viewing Policy and Filter Instances on page 390](#)

## Searching for SFW Policies

---

You can use the enhanced search utility on the Service Templates page for SFW policies and packet filters to effectively, quickly identify and segregate the policies and filters of relevance and interest.

The Service Templates page provides advanced search options for the SFW policies. Enter the term that you want to specify as the filter criterion in the search field and click the **Search** icon.

You can perform advanced searches for the following fields:

- Policy Name
- Source Address
- Destination Port
- Destination Address
- Application
- Action
- Description
- Custom column

The following advanced search criteria are available:

- Wildcard search for rule names using an asterisk (\*) is allowed.
- Edge Services Director supports AND and OR operations between search items. The default behavior is OR.
- For rule name search, only the OR operation is allowed, because a policy cannot have multiple rule names.
- For zone search, only the OR operation is allowed. Wildcard search is supported.
- For service and address fields, OR and AND operations are allowed.
- Multiple groups can be grouped using parenthesis. Grouping can be used during filed or keyword searches as well.
- Negate (-) symbol can be used to exclude objects that contain a specific term name.
- The plus (+) operator can be used to specify that the term after the + symbol existing the field value to be filtered along with other searched items.
- Escaping special characters are part of the search syntax. The supported special characters are + - & | ! ( ) { } [ ] ^ " ~ \* ? : \.



**NOTE:** Use the AND operator to find rules that match all values for a given set of fields. Use the OR operator to find rules that match any of the values for a given set of fields.

---

1. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
2. Select **Service Edit > Policy and Filter** from the task pane. The Policy and Filter page is displayed.
3. Click the plus sign (+) next to Service Template to expand the tree in the task pane and view the list of filter templates.
4. From the task pane, select **SFW Policy and Filter** to open the SFW and Filter page on the right pane.
5. Enter the term that you want to specify as the filter criterion in the Search field and click the **Search** icon.

#### Related Documentation

- [Policy and Filter Management Overview on page 375](#)
- [Packet and Service Filters Overview on page 378](#)
- [Searching for CGNAT Policies on page 381](#)
- [Searching for Packet Filters on page 384](#)
- [Managing Service and Policy Locks on page 387](#)
- [Unlocking Locked Services and Policies on page 389](#)
- [Viewing Policy and Filter Instances on page 390](#)

## Managing Service and Policy Locks

All the locked policies can be viewed in a single page. You can display the list of SFW, CGNAT, or packet filter templates that are locked by filtering them separately. Such a page shows all the locks of users only if you have the unlock task assigned; otherwise, you see only your locks.

To view the locked services and policies:

1. From the View selector, select **Gateway View**. The devices that are organized in the entire network based on the SDG pairs and the devices in each SDG group or pair are displayed.
2. Click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. Select the All Network item in the task pane. The tree can be expanded to view all the configured SDG groups and SDGs in a high-availability or redundancy group.

4. Select **Service Edit** from the task pane. The Services page is displayed for ADC and TLB services and the Rules page is displayed for CGNAT, SFW, and packet filter policies.
5. Click the right arrow next to Service Edit to expand the tree in the task pane and view the list of filter templates.
6. From the task pane, do one of the following:
  - Select **ADC** to open the ADC Services page on the right pane.
  - Select **TLB** to open the TLB Services page on the right pane.
  - Select **CGNAT** to open the CGNAT and Filter page on the right pane.
  - Select **Packet Filter** to open the Packet Filter page on the right pane.
  - Select **SFW Policy and Filters** to open the SFW Policy and Filter page on the right pane.
7. In the Services or Rules page, from the tree that lists the SDGs, select **All Service Gateways**, or the SDG or SDG pair for which you want to lock the filter templates.
8. Select the check box next to the service or rule.
9. Click the **Lock** icon. You can select policies that are locked by you and unlock them. To unlock your policies, you do not need any administrator privileges. To unlock policies locked by other users, you must have the task LOCK assigned to you.

**Related Documentation**

- [Policy and Filter Management Overview on page 375](#)
- [Packet and Service Filters Overview on page 378](#)
- [Searching for CGNAT Policies on page 381](#)
- [Searching for Packet Filters on page 384](#)
- [Searching for SFW Policies on page 386](#)
- [Unlocking Locked Services and Policies on page 389](#)
- [Viewing Policy and Filter Instances on page 390](#)

## Unlocking Locked Services and Policies

All the locked services policies can be viewed in a single page. This page is available for a user with Manage Policy Locks tasks assigned. This page shows all the locks only if the user has the unlock task assigned; otherwise, the user sees only their locks.

To unlock the locked services and policies:

1. From the View selector, select **Gateway View**. The devices that are organized in the entire network based on the SDG pairs and the devices in each SDG group or pair are displayed.
2. Click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. Select the All Network item in the task pane. The tree can be expanded to view all the configured SDG groups and SDGs in a high-availability or redundancy group.
4. Select **Service Edit** from the task pane. The Rules page is displayed.
5. Click the right arrow next to Service Edit to expand the tree in the task pane and view the list of service and policy filter templates.
6. From the task pane, do one of the following:
  - Select **ADC** to open the ADC Services page on the right pane.
  - Select **TLB** to open the TLB Services page on the right pane.
  - Select **CGNAT** to open the CGNAT and Filter page on the right pane.
  - Select **Packet Filter** to open the Packet Filter page on the right pane.
  - Select **SFW Policy and Filters** to open the SFW Policy and Filter page on the right pane.
7. In the Services and Rules pages, respectively, from the tree that lists the SDGs, select **All Service Gateways**, or the SDG or SDG pair for which you want to view the locked filter templates.
8. Select the policy instance you want to unlock, and click the **Unlock** icon at the top of the dialog box. Click the **Close** icon to return to the services listing page. To unlock your policies, you do not need any administrator privileges. To unlock policies locked by other users, you must have the task LOCK assigned to you.

### Related Documentation

- [Policy and Filter Management Overview on page 375](#)
- [Packet and Service Filters Overview on page 378](#)
- [Searching for CGNAT Policies on page 381](#)

- [Searching for Packet Filters on page 384](#)
- [Searching for SFW Policies on page 386](#)
- [Managing Service and Policy Locks on page 387](#)
- [Viewing Policy and Filter Instances on page 390](#)

---

## Viewing Policy and Filter Instances

To view the list of CGNAT, SFW, and packet policy or filter instances:

1. From the View selector, select **Gateway View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the task pane, select **Service Edit**. On the right pane, pie charts corresponding to the configured services and policy filters are displayed if you view the page without drilling-down the tree in the task pane to select a particular service or policy.
4. Click the right arrow next to Service Edit to expand the tree in the task pane and view the list of filter instances.
5. From the task pane, do one of the following:
  - Select **ADC** to open the Service Edit > ADC page on the right pane.
  - Select **TLB** to open the Service Edit > TLB page on the right pane.
  - Select **CGNAT Policy and Filter** to open the CGNAT and Filter page on the right pane.
  - Select **Packet Filter** to open the Packet Filter page on the right pane.
  - Select **SFW Policy and Filters** to open the SFW Policy and Filter page on the right pane.
6. In the Service Edit page, from the tree that lists the SDGs, select **All Service Gateways**, or the SDG or SDG pair for which you want to view the previously configured policy or filter instances.

The page is divided into three panes. The list of SDGs are displayed on the left pane. You can drill-down to the SDG or pair of SDGs for which you want to process policies or filters. The policy and filter rules are displayed in the middle pane. The right pane lists the rule and service set details. For each rule, the terms defined are shown in a tree structure. The key value pair format can be expanded by clicking the + icon beside each term.

The following fields are displayed on the Service Edit > ADC page:

**Table 60: Service Edit > ADC Page**

Field
SDG Host
Instance Name
OS Version
Group Name
Reference Host
Real Servers
Health Check Sources
Custom Health Checks
Groups
Virtual Servers
Deployment Plans

The following fields are displayed on the Service Edit > TLB page:

**Table 61: TLB Service Edit Page**

Field
SDG Host
Instance Name
OS Version
Group Name
Reference Host
Real Servers
Network Monitoring
Groups
Virtual Servers
Deployment Plans

The following fields are displayed on the Service Gateways—CGNAT Policy and Filter page:



**TIP:** In Gateway View of Deploy mode, with All Network selected in View pane and Policy & Filters > CGNAT selected in the task pane, you can select a different SDG host from the Host Name list, and a different rule term from the Term Name list from the page that lists all of the previously defined service policies. This type of inline or embedded editing enables you to quickly and optimally change the rule term in a service policy and the SDG with which the policy must be associated.

Table 62: CGNAT Policy and Filter Page

Field
Host Name
Group Name
Rule Name
Match Direction
Term Name
Source Address
Destination Address
Destination Port
Application
Translated Packet Source
Translated Packet Destination
Translation Type

Figure 36: CGNAT Services Listing Page

Status	Service Gateway/Host	Rule Name	Match Direction	Term	Source Address	Destination Address	Source Port	Destination Port	Application	Translation
	dsf4d4d4	cgnat-rule123	INPUT							
	mobst480v	cgnat-rule123	INPUT							
	mobst480v	cgnat-rule123	INPUT							
	dsf4d4d4	nap44-SS1-r1	INPUT	term-11	10.165.49.208/28	phone-broa			SDG-APPS No-Tran...	
	dsf4d4d4	nap44-SS2-r1	INPUT							
	dsf4d4d4	nap44-SS3-r1	INPUT							
	dsf4d4d4	nap44-SS4-r1	INPUT							
	dsf4d4d4	nap44-SS4-r2	INPUT							

The following fields are displayed on the Service Gateways—Packet Filter page:



**TIP:** In Gateway view of Deploy mode, with All Network selected in the View pane and Policy & Filters > Packet Filter selected in the task pane, you can select a different SDG host from the Host Name list, and a different rule term from the Term Name list from the page that lists all of the previously defined service policies. This type of inline or embedded editing enables you to quickly and optimally change the rule term in a service policy and the SDG with which the policy must be associated.

Table 63: Packet Filter Page

Field
Host Name
Group Name
Filter Name
Term Name
Source Address
Destination Address
Destination Port

*Table 63: Packet Filter Page (continued)*

Field
Source Port
Protocol
Forwarding Class
Action
Status

The following fields are displayed on the Service Gateways—SFW Policy and Filter page:



**TIP:** In Gateway view of Deploy mode, with All Network selected in the View pane and Policy & Filters > SFW selected in the task pane, you can select a different SDG host from the Host Name list, and a different rule term from the Term Name list from the page that lists all of the previously defined service policies. This type of inline or embedded editing enables you to quickly and optimally change the rule term in a service policy and the SDG with which the policy must be associated.

*Table 64: SFW Policy and Filter Page*

Field
Host Name
Group Name
Rule Name
Term Name
Source Address
Destination Address
Destination Port
Source Port
Application Sets
Filter Outcome

Figure 37: Stateful Firewall Services Listing Page

Status	Service Gateway/Host	Rule Name	Match Direction	Term Name	Source Address	Destination Address	Destination Port	Application Sets	Filter Out
<input type="checkbox"/>	dsffdfdf	IPv6-SFW	INPUT						
<input type="checkbox"/>	mobst480x	IPv6-SFW	INPUT						
<input type="checkbox"/>	mobst480w	IPv6-SFW	INPUT						
<input type="checkbox"/>	dsffdfdf	rtsp	INPUT						
<input type="checkbox"/>	dsffdfdf	TEST_RULE	INPUT						

Select a policy or a filter and click the **Expand All** icon, and all rules corresponding to that policy or filter are expanded.

Select a policy or filter and click the **Collapse All** icon to collapse all rules.

#### Related Documentation

- [Policy and Filter Management Overview on page 375](#)
- [Packet and Service Filters Overview on page 378](#)
- [Searching for CGNAT Policies on page 381](#)
- [Searching for Packet Filters on page 384](#)
- [Searching for SFW Policies on page 386](#)
- [Managing Service and Policy Locks on page 387](#)
- [Unlocking Locked Services and Policies on page 389](#)

## Creating and Managing CGNAT Policy and Filter Instances

NAT processing centers on the evaluation of NAT rule sets and rules. A rule set determines the overall direction of the traffic to be processed. For example, a rule set can select traffic from a particular interface or to a specific zone. A rule set can contain multiple rules. Once a rule set is found that matches specific traffic, each rule in the rule set is evaluated for a match. Each rule in the rule set further specifies the traffic to be matched and the action to be taken when traffic matches the rule.



**NOTE:** Before you create a policy and filter template for packet filters, SFW, or CGNAT services, you must have previously configured the different elements or attributes of the service, such as service sets, interface sets, rule sets, and syslogs during the creation of the service template. The sections in this procedural topic that describe the creation of such service elements apply during the creation of the service template and not during the creation of the service policy filters, such as CGNAT or SFW policies.

- [Creating a NAT Policy on page 396](#)
- [Creating a Service Set on page 399](#)
- [Creating a Syslog on page 403](#)
- [Creating a Rule on page 405](#)
- [Creating a Rule Set on page 406](#)
- [Creating Addresses on page 408](#)
- [Creating Address Groups on page 409](#)
- [Address and Address Groups Overview on page 409](#)
- [Creating a NAT Rule Term on page 410](#)
- [Associating an Application and Application Set with a NAT Rule on page 414](#)
- [Creating a NAT Pool on page 414](#)
- [Associating Service Sets and Rule Sets With a NAT Rule on page 415](#)
- [Modifying NAT Policies on page 416](#)
- [Creating a Deployment Plan on page 417](#)

## Creating a NAT Policy

To configure a new CGNAT policy or filter rule:

1. From the View selector, select **Gateway View**. The View pane displays the devices in the entire network organized by the device type and device models pertaining to each device type.
2. From the View pane, select the All Network item. Expand the tree to select the device type and device node, which denotes the SDGs in a high availability pair of SDGs or an SDG group.
3. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
4. Select **Service Edit > CGNAT** from the task pane.  
The CGNAT Policies page is displayed.

5. Click the plus sign (+) next to Policy and Filter to expand the tree in the task pane and view the list of filter rules.
6. From the task pane, select **CGNAT Policy and Filter** to open the CGNAT and Filter page on the right pane.
7. Click the **Add** icon above the table of listed rules. The Create Policy and Filter window is displayed.

Figure 38: Create a CGNAT Rule Window

**Create CGNAT Rule**

Name:

Match Direction:

Service Gateway:

Associate Service Sets

Service Gateway	Rule Set	Service Set
dsffdfdf		mobst480x
		mobst480w

+    ✎    -

Host Name	Term Name
-----------	-----------

Validate    Create    Cancel

8. Enter the name of the group policy in the Name field (limit of 63 alphanumeric characters).
9. Enter a description for the group policy rules in the Description field. Edge Services Director sends the comments entered in this field to the device (limit of 255 alphanumeric characters).

10. In the Match Direction list, specify the direction in which the rule match is applied. Select one of the following options:
  - **input**—Apply the rule match on the input side of the interface.
  - **input-output**—Apply the rule match bidirectionally.
  - **output**—Apply the rule match on the output side of the interface.
11. In the SDG section, do the following:
  - From the SDG drop-down list, select the devices with which the NAT policy must be associated. Alternatively, you can select the high availability pair of SDG devices with which the NAT policy must be associated. All of the devices in the different SDG groups that were previously defined in the database are also listed in the drop-down menu.
12. Create a NAT rule term that must be added to the NAT policy. For details on configuring a NAT rule term, see *Creating a NAT Rule Term*.
13. The list of terms added, and the associated service sets and rule sets, are displayed in a tabular format in the Create Policy and Filter page. Select the check box next to the term you want to attach to the NAT policy.
14. Click **Create** to save the NAT policy.
15. Click **Validate** to perform validation checks on the configuration planned to be deployed to examine and correct any syntax errors or incompatible settings. You can also validate without deploying the configuration.



**NOTE:** In the Create Policy and Filter window, you can also do the following:

- Click the **Create** icon displayed beside the terms or attributes to add a new attribute. You can then use the newly defined attribute to add to a policy to cause the same selection for a particular term to be applied across all SDGs or groups.
  - Click the **Edit** icon displayed beside the terms or attributes to modify an attribute. You can then use the modified attribute to add to a policy to cause the same selection for a particular term to be applied across all SDGs or groups.
  - Select the check box beside the SDGs or SDG groups in the Create NAT Term page to include the devices or the SDG groups in the NAT policy for association. Deselect the check boxes beside the SDGs or groups to exclude the devices in the NAT policy..
  - Click the **Copy to All Hosts** button to apply the defined term at the system or network level and not at a particular SDG or SDG group level.
-

## Creating a Service Set

A service set is a collection of services to be performed by an Adaptive Services (AS) or Multiservices PIC. To create a service set as a component for the CGNAT rule:

1. From the View selector, select **Gateway View**. The View pane displays the devices in the entire network organized by the device type and device models pertaining to each device type.
2. From the View pane, select the All Network item. Expand the tree to select the device type and device node, which denotes the SDGs in a high availability pair of SDGs or an SDG group.
3. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
4. Select **Service Edit > CGNAT** from the task pane.  
The Service Edit > CGNAT Policies page is displayed.
5. Click the **Add** icon. The Create a CGNAT Policy and Filter Template window appears.
6. Enter the name of the rule, a description, and the direction in which the rule match must be applied in the respective fields. Also, select the SDG or SDG pair for which the syslog needs to be defined for the service set.
7. In the Create Policy and Filter page, click **Associate Service Sets/Rule Sets**. The Associate Service Sets/Rule Sets section is displayed. The SDGs and SDG groups that are part of the NAT policy filter rule term are shown in one column. Under the Association column, either the **Configure** or **Edit** icon appears. If you already created and mapped a service set with the particular SDG or group, the **Edit** icon shows.
8. Click the **Configure** or **Edit** icon. The Configure Service Sets/Rule Sets dialog box is displayed.
9. From the Type drop-down list, select **Service Set** to map a service set with the policy filter rule.
10. If you selected **Service Set** from the Type list, select a service set previously configured in the Service Designer workspace from the **Value** list.
11. Click the green plus sign next to the Value drop-down list. The Addition of Service Sets dialog box appears.



**NOTE:** If a green plus sign mark is shown beside a field in the dialog box, it denotes that you can add attributes for that component. A red minus mark shows that you can delete that particular attribute for that component.

12. In the Name field, enter the name to identify the service set. Rules are combined into rule sets, and are associated with a service set for each application such as firewall or CGNAT.

13. In the Sampling Service Choices section, do one of the following:

- Click **Interface Services** to configure an interface-style service set. An interface service set is used as an action modifier across an entire interface
- In the Service Interfaces field, specify the name for the adaptive services interface associated with an interface-wide service set.

When you have defined and grouped the service rules by configuring the service-set definition, you can apply services to one or more interfaces installed on the router. When you apply the service set to an interface, it automatically ensures that packets are directed to the PIC.

- From the **Load Balancing Options** section, configure the high availability (HA) options.

The following hash keys can be configured in the egress direction: **destination-ip** (Use the destination IP address of the flow to compute the hash used in load balancing.) and **source-ip** (Use the source IP address of the flow to compute the hash used in load balancing.)

- Click the green tick mark beside the Egress Key element to configure the hash keys to be used in the egress flow direction. The configuration is mandatory if you are using AMS for Network Address Translation (NAT). This configuration is not mandatory if you are using AMS for stateful firewall; if the hash keys are not configured, then the defaults are chosen.
- Click the green tick mark beside the Ingress Key element to configure the hash keys to be used in the ingress flow direction. The configuration is mandatory if you are using AMS for Network Address Translation (NAT). This configuration is not mandatory if you are using AMS for stateful firewall; if the hash keys are not configured, then the defaults are chosen.

Configure the hash keys used for load balancing in aggregated multiservices (AMS) for service applications (Network Address Translation [NAT], stateful firewall, application-level gateway [ALG], HTTP header enrichment, and mobility). The hash keys supported in the ingress and egress direction are the source IP address and destination IP address.

Hash keys are used to define the load-balancing behavior among the various members in the AMS group. For example, if **hash-keys** is configured as **source-ip**,

then the hashing would be performed based on the source IP address of the packet. Therefore, all packets with the same source IP address land on the same member. Hash keys must be configured with respect to the traffic direction: ingress or egress. For example, if **hash-keys** is configured as **source-ip** in the ingress direction, then it should be configured as **destination-ip** in the egress direction. This is required to ensure that the packets of the same flow reach the same member of the AMS group.

The configuration of the ingress and egress hash keys is mandatory if you are using AMS for NAT. This configuration is not mandatory if you are using AMS for stateful firewall; if the hash keys are not configured, then the defaults are chosen. Refer to [Table 46 on page 245](#) for the supported hash keys.

The resource-triggered option enables anchor session PICs to use the load or resource information from the anchor services PICs to select the AMS member will anchor the services for the subscriber for load balancing among AMS members. In addition, for mobile subscriber-aware services (such as HTTP header enrichment), you must configure the **resource-triggered** statement, which means that the load balancing is not done using the ingress and egress keys.

**Table 65: Hash Keys Supported for AMS for Service Applications**

Service Set at Ingress Interface			Service Set at Egress Interface	
Hash Keys for NAT				
NAT Type	Ingress hash key	Egress hash key	Ingress hash key	Egress hash key
source static	Destination IP address	Source IP address	Source IP address	Destination IP address
source dynamic	Source IP address	Destination IP address	Destination IP address	Source IP address
Network Address Port Translation (NAPT)	Source IP address	Destination IP address	Destination IP address	Source IP address
destination static	Source IP address	Destination IP address	Destination IP address	Source IP address
Hash Keys for Stateful Firewall				
Stateful Firewall	Destination IP address	Source IP address	Destination IP address	Source IP address
Stateful Firewall	Source IP address	Destination IP address	Source IP address	Destination IP address



**NOTE:** If NAT is used in the service set (along with stateful firewall and ALG), then the hash keys should be based on the NAT type; otherwise, the hash keys of the stateful firewall should be used.

- Click **Next Hop Services** to configure a next-hop style service set. A next-hop service set is a route-based method of applying a particular service. Only packets destined for a specific next hop are serviced by the creation of explicit static routes.

- In the **Inside Interface** list, specify the interface type of the service interface associated with the service set applied inside the network. For inline IP reassembly, set the interface type to local. Also, specify the name and logical unit number of the service interface associated with the service set applied inside the network.

When a next-hop service is configured, the AS or Multiservices PIC is considered to be a two-legged module with one leg configured to be the inside interface (inside the network) and the other configured as the outside interface (outside the network).

- In the **Outside Interface** list, specify the interface type of the service interface associated with the service set applied outside the network. For inline IP reassembly, set the interface type to local. Also, specify the name and logical unit number of the service interface associated with the service set applied outside the network.
- In the **Service Interface Pool** list, select the name of the pool of logical interfaces configured at the [edit services service-interface-pools pool pool-name] hierarchy level. You can configure a service interface pool only if the service set has a PGCP rule configured. The service set cannot contain any other type of rule.

- 

- Click **Sampling Services** to configure a sampling service set.
  - In the Service Interface field, specify the service interface, which is the interface the sampling is taken from. In the case of a sampling service set, the service interface must be a Multiservices PIC interface with a subunit number of 0 (zero). The subunit number defaults to 0. The reverse-flow statement is not mandatory. All sampled traffic is considered to be forward traffic. If you set the reverse-flow statement, it is ignored.
- Select the **Replication Service** check box to configure the services replication options for inter-chassis high availability on MS-MIC and MS-MPC.
  - In the Replication Threshold field, specify the number of seconds for the replication threshold. When a flow has been active for more than the number of seconds specified as a threshold, flow state information is replicated to the backup device. Make sure that the replication-threshold value is than the open-timeout value (the timeout period for establishing a TCP connection). The default value of the replication threshold is 180 seconds. This value is also the minimum.
  - Select the **Stateful Firewall** check box to replicate stateful firewall state information.
  - Select the **NAT** check box to replicate NAT44 information.

14. In the CGNAT Rule Sets section, select the rule set you want to associate with the service set from the Available column and click the right arrow to move to the Selected column.

15. In the CGNAT Rules section, select the rule you want to associate with the service set from the Available column and click the right arrow to move to the Selected column.

16. In the CGNAT Syslogs section, select the syslog you want to associate with the service set from the Available column and click the right arrow to move to the Selected column.
17. Click **Save** to save the service rule configuration. Else, click **Close** to discard the changes to the rule.

## Creating a Syslog

You can enable system logging. The system log information from the Adaptive Services or Multiservices PIC is passed to the kernel for logging in the /var/log directory. This setting overrides any syslog statement setting included in the service set or interface default configuration.

To create a syslog for the CGNAT rule:

1. From the View selector, select **Gateway View**. The View pane displays the devices in the entire network organized by the device type and device models pertaining to each device type.
2. From the View pane, select the All Network item. Expand the tree to select the device type and device node, which denotes the SDGs in a high availability pair of SDGs or an SDG group.
3. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
4. Select **Service Edit > CGNAT** from the task pane.  
The Service Edit > CGNAT Policies page is displayed.
5. Click the **Add** icon. The Create a CGNAT Policy and Filter Template window appears.
6. Enter the name of the rule, a description, and the direction in which the rule match must be applied in the respective fields. Also, select the SDG or SDG pair for which the syslog needs to be defined for the service set.
7. In the Create Policy and Filter page, click **Associate Service Sets/Rule Sets**. The Associate Service Sets/Rule Sets section is displayed. The SDGs and SDG groups that are part of the NAT policy filter rule term are shown in one column. Under the Association column, either the **Configure** or **Edit** icon appears. If you already created and mapped a service set with the particular SDG or group, the **Edit** icon shows.
8. Click the **Configure** or **Edit** icon. The Configure Service Sets/Rule Sets dialog box is displayed.

9. From the Type drop-down list, select **Service Set** to map a service set with the policy filter rule.
10. If you selected **Service Set** from the Type list, select a service set previously configured in the Service Designer workspace from the **Value** list.
11. Click the green plus sign next to the Value drop-down list. The Addition of Service Sets dialog box appears.



**NOTE:** If a green plus sign mark is shown beside a field in the dialog box, it denotes that you can add attributes for that component. A red minus mark shows that you can delete that particular attribute for that component.

12. Click the green plus sign next to the Syslog Settings field. The Addition of Service Sets dialog box appears.
13. In the Name field, enter the name for the syslog component. Specify the fully qualified domain name or IP address for the syslog server.
14. In the Services list, specify the system logging severity level. It assigns a severity level to the facility. Valid entries include:
  - **alert**—Conditions that should be corrected immediately.
  - **any**—Matches any level.
  - **critical**—Critical conditions.
  - **emergency**—Panic conditions.
  - **error**—Error conditions.
  - **info**—Informational messages.
  - **notice**—Conditions that require special handling.
  - **warning**—Warning messages.
15. From the Facility Override list, select the override for the default facility for system log reporting. Valid values include:

**authorization**

**daemon**

**ftp**

**kernel**

**local0 through local7**

**user**

16. In the Log Prefix field, set the system logging prefix value for all logging to the system log host.
17. In the Port field, specify the port number to be used for connection with the remote syslog server.
18. In the Class section, set the class of applications to be logged to the system log.
  - **alg-logs**—Log application-level gateway events.
  - **ids-logs**—Log intrusion detection system events.
  - **nat-logs**—Log Network Address Translation events.
  - **packet-logs**—Log general packet-related events.
  - **session-logs**—Log session open and close events.
  - **session-logs open**—Log session open events only.
  - **session-logs close**—Log session close events.
  - **stateful-firewall-logs**—Log stateful firewall events.
19. In the Source Address field, specify a source address to record in system log messages that are directed to a remote machine specified in the hostname statement. The supported interfaces are ms, rms, and mams interfaces. If you do not specify the interface parameter, the command loops on all supported interfaces. This field is available only if you selected the Junos OS 14.1 version.
20. Click **Save** to save the service rule configuration. Else, click **Close** to discard the changes to the rule.

## Creating a Rule

To create a rule for the CGNAT service:

1. From the View selector, select **Gateway View**. The View pane displays the devices in the entire network organized by the device type and device models pertaining to each device type.
2. From the View pane, select the All Network item. Expand the tree to select the device type and device node, which denotes the SDGs in a high availability pair of SDGs or an SDG group.
3. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
4. From the task pane, select **Service Edit**. The Service Edit page is displayed.

5. Click the **CGNAT** button. The list of CGNAT policies is displayed.
6. Click the **Add** icon. The Create a CGNAT Policy window appears.
7. Enter the name of the template and the service instance in the respective fields.
8. Click the green plus sign in the Rules box. The Addition of Rules dialog box appears.



**NOTE:** If a green tick mark is shown beside a field in the dialog box, it denotes that you can add attributes for that component. A red cross mark shows that you can delete that particular attribute for that component.

9. From the **Rule** list, select one of the previously configured rules. The rules that you configured in the Service Templates workspace for CGNAT, packet filter, or CGNAT are displayed.
10. Click **Save** to save the service template configuration. Else, click **Close** to discard the changes to the template.


## Creating a Rule Set

The rule-set statement defines a collection of stateful firewall rules that determine what actions the router software performs on packets in the data stream. You define each rule by specifying a rule name and configuring terms. Then, you specify the order of the rules by including the **rule-set** statement at the **[edit services stateful-firewall]** hierarchy level with a rule statement for each rule.

The router software processes the rules in the order in which you specify them in the configuration. If a term in a rule matches the packet, the router performs the corresponding action and the rule processing stops. If no term in a rule matches the packet, processing continues to the next rule in the rule set. If none of the rules matches the packet, the packet is dropped by default.

To create a rule set for the CGNAT policy:

1. From the View selector, select **Gateway View**. The View pane displays the devices in the entire network organized by the device type and device models pertaining to each device type.
2. From the View pane, select the All Network item. Expand the tree to select the device type and device node, which denotes the SDGs in a high availability pair of SDGs or an SDG group.
3. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.

4. Select **Service Edit > CGNAT** from the task pane.  
The Service Edit > CGNAT Policies page is displayed.
  5. Click the **Add** icon. The Create a CGNAT Policy and Filter Template window appears.
  6. Enter the name of the rule, a description, and the direction in which the rule match must be applied in the respective fields. Also, select the SDG or SDG pair for which the syslog needs to be defined for the service set.
  7. In the Create Policy and Filter page, click **Associate Service Sets/Rule Sets**. The Associate Service Sets/Rule Sets section is displayed. The SDGs and SDG groups that are part of the NAT policy filter rule term are shown in one column. Under the Association column, either the **Configure** or **Edit** icon appears. If you already created and mapped a service set with the particular SDG or group, the **Edit** icon shows.
  8. Click the **Configure** or **Edit** icon. The Configure Service Sets/Rule Sets dialog box is displayed.
  9. From the Type drop-down list, select **Service Set** to map a service set with the policy filter rule.
  10. If you selected **Service Set** from the Type list, select a service set previously configured in the Service Designer workspace from the **Value** list.
  11. Click the green plus sign next to the Value drop-down list. The Addition of Service Sets dialog box appears.
-  **NOTE:** If a green plus sign mark is shown beside a field in the dialog box, it denotes that you can add attributes for that component. A red minus mark shows that you can delete that particular attribute for that component.
12. In the Name field, specify a name for the rule set the router uses when applying this service.
  13. In the Rules section, select the rules that need to be added to the rule set in the from the Available column and click the right arrow to move these rules to the Selected column. All the rules that you previously configured during the creation or modification of the service rule are displayed.
  14. Click **Save** to save the rule set configuration. Else, click **Close** to discard the changes to the rule.

## Creating Addresses

To create an address:

1. In the Source and Destination Address Selector dialog box, to create a new address, click the plus sign (+).

The Create Address page appears.

2. In the Object Type section, click the **Address** radio button to create an address.
3. In the Name field, enter a name for the new address.
4. In the Description field, enter a description for the new address.
5. Direct Edge Services Director to resolve an IP address to a hostname or resolve a hostname to an IP address.
  - To specify an IP address as the address type, select **Host** from the drop-down menu and enter the **IP** address in the IP field.
  - To specify a hostname as the address type, select **Host** from the drop-down menu and enter the hostname in the Host Name field.
  - To specify an IP address range, select **Range** from the drop-down menu and enter the IP ranges in the Start IP and End IP fields.
  - To specify a network as an address type, select **Network** from the drop-down menu and enter the network address in the IP and Netmask fields.
  - To specify an IP address with a wildcard mask, select **Wildcard** from the drop-down menu and enter the IP address in the IP field and wildcard mask in the Wildcard Mask fields.
  - To specify a DNS name as an address type, select **DNS Host** from the drop-down menu and enter the DNS name in the DNS Name field.



**NOTE:** You can resolve an IP address to a hostname and a hostname to an IP address using the green arrows next to the IP and Host Name fields.

---



**NOTE:** The host and network address types support both IPv4 and IPv6 address types. These address types also supports multicast addresses. However, the range address type supports only IPv4 addresses. NAT and IPsec VPNs do not support IPv6 addressing and wildcard addresses.

---



**NOTE:** Ensure that the first 8 bits of the address are not 0 and the highest bit of the mask is 1 when you are using the wildcard address type.

6. Click **Create** to create an address.

The new address appears in the Manage Address page.

## Creating Address Groups

To create an address group:

1. In the Source and Destination Address Selector dialog box, to create a new address group, click the plus sign (+).

The Create Address Group page appears.

2. Select the Object Type as Address Group.
3. In the Name field, enter a name for the new address group.
4. In the Description field, enter a description for the new address group.
5. In the Addresses field, from the Available dialog box, select the address that you want to group, and click the right arrow to add to the Selected column.

Click **All** to move all the addresses to the Selected column. The address you have selected appears in the Selected section of the dialog box.

6. Click **Create**.

The address group appears on the Address page.

## Address and Address Groups Overview

You can use the Address Creation Wizard to create an address object that specifies an IP address or a hostname. You can specify a hostname and use the address resolution option to resolve it to an IP address. You can also resolve an IP address to the corresponding hostname.

You can group address objects to form an address group using the Address Group Creation Wizard. Junos Space creates an object in the Junos Space database to represent an address or an address group.

## Creating a NAT Rule Term

To add rules to a NAT policy:

1. In the Create Policy and Filter window, the list of rule terms already added, if any, to the NAT policy are displayed.
2. Next to the **Terms** field, click the + icon to add rules, and select the type of rule you want to add.
3. In the **Term Name** field, specify the name of the rule.

The list of SDGs with which you associated the NAT policy in the Create Policy window are displayed with the form and then sections or clauses. If you selected SDG groups to associate with the NAT policy, the SDG group names are displayed.



**NOTE:**

- Click the **Copy to All Hosts** button to apply the defined term at the system or network level and not at a particular SDG or SDG group level.
- When you create a rule or filter term, and define the name of the filter, for SDGs that are part of a high availability pair of devices, the names of the SDGs are displayed as tabs and check boxes beside the hostnames of the SDGs are displayed. If you want the policy or filter term definition to apply to both the SDGs, select the check boxes next to the SDG names.

Otherwise, when the click the SDG name tab for the SDG for which you did not select the check box, a blue highlight overlays the entire dialog box to indicate the settings are not enabled for configuration for that specific SDG.

4. In the **From** section, do the following to specify input conditions or match criteria for the NAT term :
  - In the **Source Address** field, click the down arrow in the list. The address selector dialog box appears. Select the source addresses that need to be added to the NAT policy in the from the Available column and click the right arrow to move these devices to the Selected column.

Click **OK** to confirm the selection. Click **Cancel** to discard your changes and return to the Create Policy and Filter window.

To create an address or address group from the address selector dialog box, see *Creating Addresses* and *Creating Address Groups*.
  - In the **Destination Address** field, click the down arrow in the list. The address selector dialog box appears. Select the destination addresses that need to be added to the NAT policy in the from the Available column and click the right arrow to move these devices to the Selected column.

Click **OK** to confirm the selection. Click **Cancel** to discard your changes and return to the Create Policy and Filter window.

To create an address or address group from the address selector dialog box, see *Creating Addresses* and *Creating Address Groups*.

- Specify a destination port to match the rule in the **Destination Port** field. You can configure a range of ports by specifying the upper limit and lower limit of the ports in the Start Value and End Value fields.
- Select the application protocol or name to which the NAT services apply from the **Application** drop-down menu. When you click the down arrow in the list, the application selector dialog box appears. Select the application name that needs to be added to the NAT policy.

To create a new application name or application set, see *Creating Applications and Application Sets*.

- Select the name of the target application set from the **Application Sets** drop-down menu.
5. In the **To** section, do the following to specify actions or modifiers to be performed for the NAT term :
    - In the **Translation Type** drop-down list, select the NAT translation type.

- **basic-nat44**—Translate the source address statically (IPv4 to IPv4).
  - **basic-nat66**—Translate the source address statically (IPv6 to IPv6).
  - **basic-nat-pt**—Translate the addresses of IPv6 hosts as they originate sessions to the IPv4 hosts in the external domain. The **basic-nat-pt** option is always implemented with DNS ALG.
  - **deterministic-napt44**—Translate as **napt-44**, and use deterministic port block allocation for port translation.
  - **dnat-44**—Translate the destination address statically (IPv4 to IPv4).
  - **dynamic-nat44**—Translate only the source address by dynamically choosing the NAT address from the source address pool.
  - **napt-44**—Translate the transport identifier of the IPv4 private network to a single IPv4 external address.
  - **napt-66**—Translate the transport identifier of the IPv6 private network to a single IPv6 external address.
  - **napt-pt**—Bind addresses in an IPv6 network with addresses in an IPv4 network and vice versa to provide transparent routing for the datagrams traversing between the address realms.
  - **stateful-nat64**—Implement dynamic address and port translation for source IP addresses (IPv6-to-IPv4) and prefix removal translation for the destination IP addresses (IPv6-to-IPv4).
  - **twice-basic-nat-44**—Translate the source and destination addresses statically (IPv4 to IPv4).
  - **twice-dynamic-nat-44**—Translate the source address by dynamically choosing the NAT address from the source address pool. Translate the destination address statically.
  - **twice-dynamic-napt-44**—Translate the transport identifier of the IPv4 private network to a single IPv4 external address. Translate the destination address statically.
- In the **Source Pool** field, click the down arrow in the list. The NAT pool selector dialog box appears. Select the source pools that need to be added to the NAT policy in the from the Available column and click the right arrow to move these pools to the Selected column.

Click **OK** to confirm the selection. Click **Cancel** to discard your changes and return to the Create Policy and Filter window.

To create a NAT pool from the source and destination pool selector dialog box, see *Creating a NAT Pool*.
  - In the **Destination Pool** field, click the down arrow in the list. The NAT pool selector dialog box appears. Select the destination pools that need to be added to the NAT policy in the from the Available column and click the right arrow to move these pools to the Selected column.

Click **OK** to confirm the selection. Click **Cancel** to discard your changes and return to the Create Policy and Filter window.

To create a NAT pool from the source and destination pool selector dialog box, see *Creating a NAT Pool*.

- Select the **No Translation** option to specify that traffic is not to be translated.
- Select the NAT address pooling behavior as **Paired**. Only paired address pooling is supported. Address pooling, or address pooling paired (APP) ensures assignment of the same external IP address for all sessions originating from the same internal host. You can use this feature when assigning external IP addresses from a pool. This option does not affect port utilization
- In the **Destination Prefix** field, click the down arrow in the list to specify the destination prefix for translated traffic. The address selector dialog box appears. Select the destination addresses that need to be added to the NAT policy in the from the Available column and click the right arrow to move these devices to the Selected column.

Click **OK** to confirm the selection. Click **Cancel** to discard your changes and return to the Create Policy and Filter window.

To create an address or address group from the address selector dialog box, see *Creating Addresses* and *Creating Address Groups*.

- Specify the (NAT) pool for destination translation from the **DNS ALG Pool** list.
- Set the Domain Name System (DNS) application-level gateway (ALG) 96-bit prefix for mapping IPv4 addresses to IPv6 addresses from the **DNS ALG Prefix** list.
- Select the **Endpoint Independent** check box for the Filtering Type field to specify the NAT filtering behavior for sessions initiated from outside to inside as endpoint-independent filtering (EIF).
- Select the **Endpoint Independent** check box for the Mapping Type field to specify the source NAT mapping type.
- In the **Source Prefix** field, click the down arrow in the list to specify the destination prefix for translated traffic. The address selector dialog box appears. Select the source addresses that need to be added to the NAT policy in the from the Available column and click the right arrow to move these devices to the Selected column.

Click **OK** to confirm the selection. Click **Cancel** to discard your changes and return to the Create Policy and Filter window.

To create an address or address group from the address selector dialog box, see *Creating Addresses* and *Creating Address Groups*.

- Select the **Syslog** check box to enable system logging. The system log information from the Multiservices PIC is passed to the kernel for logging in the `/var/log` directory. This field is available only if you selected the Junos OS 14.1 version to create the service template.

6. Click **Save** to create the rule. Alternatively, click **Validate** in the Create Rule page to perform validation checks on the configuration planned to be deployed to examine and correct any syntax errors or incompatible settings.
7. A new rule is added in the last row depending on the type of rule you have added. The newly added rules blink with a different color for few seconds. The behavior is same if you add a new rule before or after a rule, clone a rule, or paste a rule.

The rule is assigned a serial number based on the number of rules already added to the policy.

## Associating an Application and Application Set with a NAT Rule

To associate an application and an application set for a NAT rule term:

1. In the Add Term page, in the **Application** or **Application Set** sections, the application set selector dialog box is displayed. Select the applications or application sets that need to be added to the NAT rule term in the from the Available column and click the right arrow to move these applications or application sets to the Selected column.

## Creating a NAT Pool

A Network Address Translation (NAT) pool is a continuous range of IP addresses that you can use to create a NAT policy. NAT policies perform address translation by translating internal IP addresses to the addresses in these pools.

To create a NAT pool:

1. In the Add Term page, click the down arrow of the **Source Pool** or **Destination Pool** drop-down lists. The source and destination NAT pool selector dialog box is displayed.
2. Select a NAT pool to function as the source or destination pool from the Select NAT Pool pop-up dialog box. Click **OK** to add the selected NAT pool to the source or destination pool drop-down list in the Add Term page.
3. If a NAT address pool has not been previously created, click the plus sign (+) to create a new NAT pool. The Create NAT Pool page appears.
4. Enter the name of the NAT pool in the Name field.
5. Select the type of NAT pool as source or destination from the Pool Type menu.
6. In the Pool Address field, do one of the following
  - Select the Range radio button and enter the network address in the Prefix and Netmask fields for IPv4 or IPv6.
  - Select the Address Prefix radio button and enter the IP ranges in the Start IP and End IP fields.

7. Select the **Round Robin** check box beside the Address Allocation field if you want to use round-robin technology. When you use round-robin allocation, one port is allocated from each address in a range before repeating the process for each address in the next range. After ports have been allocated for all addresses in the last range, the allocation process wraps around and allocates the next unused port for addresses in the first range.
8. In the Auto Port Allocation field, do one of the following to specify the NAT pool port or range. You can configure an automatically assigned port or specify a range with minimum and maximum values. :
  - Select the **Automatic** radio button to use a router-assigned port.
  - Select the **Random Allocation** radio button to allocate ports within a specified range randomly. Select the **Range** check box and specify the starting and ending values for the port range in the **Low** and **High** fields.
9. Click **Create** to save the NAT address pool. The pool is now populated in the Select NAT Pool dialog box in the drop-down list. You can select the created pool as the source or destination address pool while creating a NAT rule term.

## Associating Service Sets and Rule Sets With a NAT Rule

To associate a service set and a rule set with a NAT policy filter rule term:

1. In the Create Policy and Filter page, click **Associate Service Sets/Rule Sets**. The Associate Service Sets/Rule Sets section is displayed. The SDGs and SDG groups that are part of the NAT policy filter rule term are shown in one column. Under the Association column, either the **Configure** or **Edit** icon appears. If you already created and mapped a service set with the particular SDG or group, the **Edit** icon shows.
2. Click the **Configure** or **Edit** icon. The Configure Service Sets/Rule Sets dialog box is displayed.
3. From the Type drop-down list, do either of the following:
  - Select **Service Set** to map a service set with the policy filter rule.
  - Select **Rule Set** to map a rule set with the policy filter rule.

Depending on the option selected in the Type list as service set or rule set for association with the policy filter rule, the options that are displayed in the Value list beneath the Type list varies.
4. If you selected **Service Set** from the Type list, select a service set previously configured in the Service Designer workspace from the **Value** list. If you selected **Rule Set** from the Type list, select a rule set previously configured in the Service Designer workspace from the **Value** list. Click **Add** to map the service set or rule set with the NAT policy filter rule.

5. Click **Save** to save the settings. Alternatively, click **Cancel** to abort the changes.
6. Click **Copy to All Hosts** in the Associate Service Sets dialog box to apply the defined term at the system or network level and not at a particular SDG or SDG group level. You are returned to the Add Term window.

## Modifying NAT Policies

Before you can edit the policy, you must lock it by clicking the lock icon, which is available in the policy tabular view. You can hold more than one policy lock at a given time. You can unlock the policy by clicking the unlock icon next to the lock icon in the policy tabular view. If you attempt to lock a policy that is already locked by another user, a message is displayed stating that the lock is acquired by another user.

If the Edge Services Director administrator releases the lock, you will receive the a warning message stating that the lock has been released.

The Manage Policy Locks page appears showing only those locks that can be managed by the current user. The page contains the following fields:

- Instance or Rule name
- User (IP Address)
- Lock acquired time
- Service Gateway

The policy is locked and released for the following policy operations. Also, these operations are disabled for a policy, if the policy is locked by some other user.

- Modify
- Assign devices
- Rollback
- Delete



### NOTE:

- You can unlock your policies even if they are not edited.
  - If the browser crashes when the policy is still locked, the policy is unlocked only after the timeout interval expires.
  - Policy lock is not released under the following scenario:
    - If you save or discard you changes to the locked policy.
    - if you do not make any changes to the locked policy and navigate to another policy.
-

To modify an existing CGNAT policy or filter rule:

1. From the View selector, select **Gateway View**. The View pane displays the devices in the entire network organized by the device type and device models pertaining to each device type.
2. From the View pane, select the All Network item. Expand the tree to select the device type and device node, which denotes the SDGs in a high availability pair of SDGs or an SDG group.
3. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
4. Select **Service Edit** from the task pane. The Service Templates page is displayed.
5. Click the plus sign (+) next to Policy and Filter to expand the tree in the task pane and view the list of filter rules.
6. From the task pane, select **CGNAT Policy and Filter** to open the CGNAT and Filter page on the right pane.
7. Select a policy, and click the Lock icon above the table of listed policies.
8. Click the **Modify** icon above the table of listed templates. The Modify Policy and Filter window is displayed.
9. Modify the attributes that are needed and save the updated settings.

## Creating a Deployment Plan

You must have previously defined service templates and policy or filter templates before you can create a deployment plan.

To create a deployment plan and assigning devices to it:

1. From the View selector, select **Gateway View**. The View pane displays the devices in the entire network organized by the device type and device models pertaining to each device type.
2. From the View pane, select the All Network item. Expand the tree to select the SDG in an SDG group.
3. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.

4. Select **Service Edit** from the task pane. The Service Edit page is displayed.
5. Click the right arrow next to Service Edit to expand the tree in the task pane and view the list of filter instances.
6. From the task pane, select **CGNAT Policy and Filter** to open the SFW Policy and Filter page on the right pane.
7. In the Service Edit page, from the tree that lists the SDGs, select **All Service Gateways**, or the SDG or SDG pair for which you want to view the previously configured policy or filter instances. This step is applicable only if you selected Gateway View. You can drill-down to the SDG or pair of SDGs for which you want to process policies or filters.
8. Select a rule corresponding to an SDG, and click the **Lock** icon above the table of listed policy filters.
9. Click the down arrow in the **Actions** menu and select **Send for Deployment** to create a deployment plan for the particular service template and save the plan.

The Deployment Plan Summary dialog box appears, with the service name, type, and status listed.

Click **Send** to create a deployment plan.

A deploy plan is created for the service template with the devices that are assigned to it when you view the Deployment Plans page.
10. Alternatively, you can select **Discard changes** from the **Actions** menu to ignore the modifications done to a policy or filter template.
11. From the Deployment plans page, you can select **Reject** or **Approve** from the Actions drop-down list to reject or approve the deployment plan and make it available for commissioning to the devices.

**Related Documentation**

- [Policy and Filter Management Overview on page 375](#)
- [Packet and Service Filters Overview on page 378](#)
- [Searching for CGNAT Policies on page 381](#)
- [Searching for Packet Filters on page 384](#)
- [Searching for SFW Policies on page 386](#)
- [Managing Service and Policy Locks on page 387](#)
- [Unlocking Locked Services and Policies on page 389](#)
- [Viewing Policy and Filter Instances on page 390](#)

## Creating and Managing Packet Filter Policy Instances

---

You can optionally include filters associated with each service set to refine the target and additionally process the traffic. If you include the service-set statement without a service-filter definition, the router software assumes that the match condition is true and selects the service set for processing automatically. To configure service filters, include the firewall statement at the [edit] hierarchy level. You configure service filters in a similar way to firewall filters.

If you configure **match-direction input-output**, sessions initiated from both directions might match this rule.

The match direction is used with respect to the traffic flow through the AS or Multiservices PIC. When a packet is sent to the PIC, direction information is carried along with it.

With an interface service set, packet direction is determined by whether a packet is entering or leaving the interface on which the service set is applied.

With a next-hop service set, packet direction is determined by the interface used to route the packet to the AS or Multiservices PIC. If the inside interface is used to route the packet, the packet direction is input. If the outside interface is used to direct the packet to the PIC, the packet direction is output.

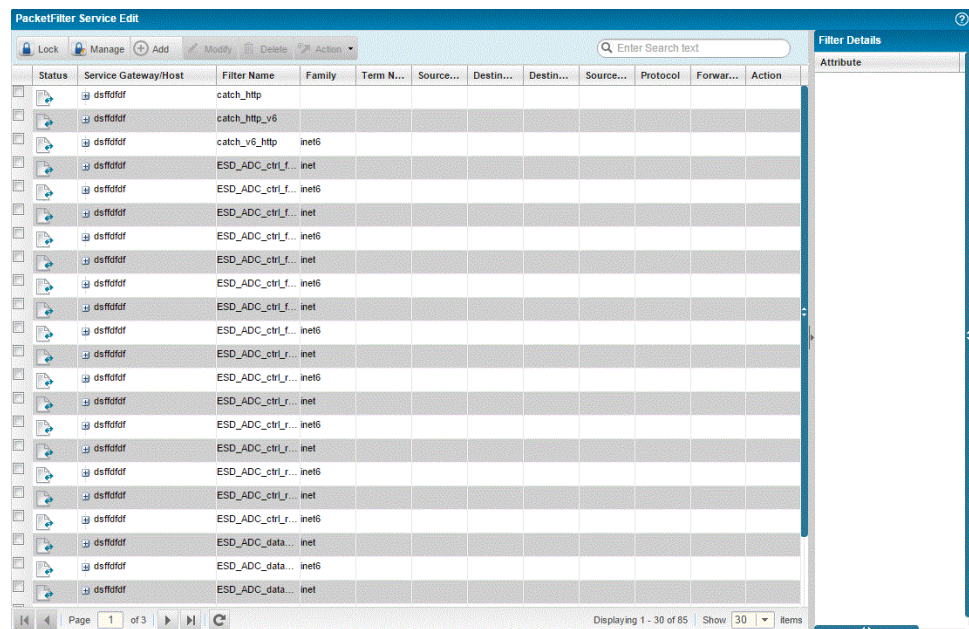
On the PIC, a flow lookup is performed. If no flow is found, rule processing is performed. Rules in this service set are considered in sequence until a match is found. During rule processing, the packet direction is compared against rule directions. Only rules with direction information that matches the packet direction are considered. Most packets result in the creation of bidirectional flows.

- [Creating a Packet Filter Policy on page 420](#)
- [Creating Addresses on page 422](#)
- [Creating Address Groups on page 423](#)
- [Address and Address Groups Overview on page 424](#)
- [Creating a Packet Filter Rule Term on page 424](#)
- [Creating an Application and Application Set on page 428](#)
- [Associating Service Sets and Rule Sets With a Packet Filter Rule on page 428](#)
- [Associating Interfaces With a Packet Filter Rule on page 429](#)
- [Modifying Packet Filter Policies on page 429](#)
- [Creating a Deployment Plan on page 432](#)

## Creating a Packet Filter Policy

To configure a new Packet Filter policy or filter instance:

1. From the View selector, select **Gateway View**. The View pane displays the devices in the entire network organized by the device type and device models pertaining to each device type.
2. From the View pane, select the All Network item. Expand the tree to select the device type and device node, which denotes the SDGs in a high availability pair of SDGs or an SDG group.
3. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
4. Select **Service Edit > Policy and Filter** from the task pane. The Service Templates page is displayed.
5. Click the plus sign (+) next to Policy and Filter to expand the tree in the task pane and view the list of filter templates.
6. From the task pane, select **Packet Filter Policy and Filter** to open the Packet Filter Policy and Filter page on the right pane.



7. Click the **Add** icon above the table of listed templates. The Create Policy and Filter window is displayed.

8. Enter the name of the group policy in the Name field.
9. Enter a description for the group policy rules in the Description field. Edge Services Director sends the comments entered in this field to the device.
10. In the Match Direction list, specify the direction in which the rule match is applied. Select one of the following options:
  - **input**—Apply the rule match on the input side of the interface.
  - **input-output**—Apply the rule match bidirectionally.
  - **output**—Apply the rule match on the output side of the interface.
11. In the SDG section, do the following:
  - From the SDG drop-down list, select the devices with which the NAT policy must be associated. Alternatively, you can select the high availability pair of SDG devices with which the NAT policy must be associated. All of the devices in the different SDG groups that were previously defined in the database are also listed in the drop-down menu.
12. Create a Packet Filter rule term that must be added to the Packet Filter policy. For details on configuring a Packet Filter rule term, see *Creating a Packet Filter Rule Term*.
13. The list of terms added, and the associated service sets and rule sets, are displayed in a tabular format in the Create Policy and Filter page. Select the check box next to the term you want to attach to the Packet Filter policy.
14. Click **Create** to save the Packet Filter policy.
15. Alternatively, click **Validate** in the Create Rule page to perform validation checks on the configuration planned to be deployed to examine and correct any syntax errors or incompatible settings.



**NOTE:** In the Create Policy and Filter window, you can also do the following:

- Click the **Create** icon displayed beside the terms or attributes to add a new attribute. You can then use the newly defined attribute to add to a policy to cause the same selection for a particular term to be applied across all SDGs or groups.
- Click the **Edit** icon displayed beside the terms or attributes to modify an attribute. You can then use the modified attribute to add to a policy to cause the same selection for a particular term to be applied across all SDGs or groups.
- Select the check box beside the SDGs or SDG groups in the Create Packet Filter Term page to include the devices or the SDG groups in the Packet Filter policy for association. Deselect the check boxes beside the SDGs or groups to exclude the devices in the Packet Filter policy..
- Click the **Copy to All Hosts** button to apply the defined term at the system or network level and not at a particular SDG or SDG group level.

---

## Creating Addresses

To create an address:

1. In the Source and Destination Address Selector dialog box, to create a new address, click the plus sign (+).

The Create Address page appears.

2. In the Object Type section, click the **Address** radio button to create an address.
3. In the Name field, enter a name for the new address.
4. In the Description field, enter a description for the new address.
5. Direct Edge Services Director to resolve an IP address to a hostname or resolve a hostname to an IP address.
  - To specify an IP address as the address type, select **Host** from the drop-down menu and enter the **IP** address in the IP field.
  - To specify a hostname as the address type, select **Host** from the drop-down menu and enter the hostname in the Host Name field.
  - To specify an IP address range, select **Range** from the drop-down menu and enter the IP ranges in the Start IP and End IP fields.
  - To specify a network as an address type, select **Network** from the drop-down menu and enter the network address in the IP and Netmask fields.

- To specify an IP address with a wildcard mask, select **Wildcard** from the drop-down menu and enter the IP address in the IP field and wildcard mask in the Wildcard Mask fields.
- To specify a DNS name as an address type, select **DNS Host** from the drop-down menu and enter the DNS name in the DNS Name field.



**NOTE:** You can resolve an IP address to a hostname and a hostname to an IP address using the green arrows next to the IP and Host Name fields.



**NOTE:** The host and network address types support both IPv4 and IPv6 address types. These address types also supports multicast addresses. However, the range address type supports only IPv4 addresses. Packet Filter and IPsec VPNs do not support IPv6 addressing and wildcard addresses.



**NOTE:** Ensure that the first 8 bits of the address are not 0 and the highest bit of the mask is 1 when you are using the wildcard address type.

6. Click **Create** to create an address.

The new address appears in the Manage Address page.

## Creating Address Groups

To create an address group:

1. In the Source and Destination Address Selector dialog box, to create a new address group, click the plus sign (+).

The Create Address Group page appears.

2. Select the Object Type as **Address Group**.
3. In the Name field, enter a name for the new address group.
4. In the Description field, enter a description for the new address group.
5. In the Addresses field, from the Available dialog box, select the address that you want to group, and click the right arrow to add to the Selected column.

Click **All** to move all the addresses to the Selected column. The address you have selected appears in the Selected section of the dialog box.

6. Click **Create**.

The address group appears on the Address page.

## Address and Address Groups Overview

You can use the Address Creation Wizard to create an address object that specifies an IP address or a hostname. You can specify a hostname and use the address resolution option to resolve it to an IP address. You can also resolve an IP address to the corresponding hostname.

You can group address objects to form an address group using the Address Group Creation Wizard. Junos Space creates an object in the Junos Space database to represent an address or an address group.

## Creating a Packet Filter Rule Term

To add rules to a Packet Filter policy:

1. In the Create Policy and Filter window, the list of rule terms already added, if any, to the Packet Filter policy are displayed.
2. Next to the **Terms** field, click the + icon to add rules, and select the type of rule you want to add.

*Figure 39: Create a Packet Filter Rule Term Window*

**Create Term**

Term Name:

mobst480x ☒ mobst480w ☐

mobst480x mobst480w

Source Address:

Available	Selected
Iwaactivate	
default-route-v6	
0/0	
phases broadband	

Source Port:

Destination Address:

Available	Selected
Iwaactivate	
default-route-v6	
0/0	
phases broadband	

Destination Port:

Count:

Forwarding Class:

Action:

Protocol:

RoutingInstance:

Syslog: ☐

Copy to All Hosts

Add Cancel

3. In the **Term Name** field, specify the name of the rule.

The list of SDGs with which you associated the Packet Filter policy in the Create Policy window are displayed with the form and then sections or clauses. If you selected SDG groups to associate with the Packet Filter policy, the SDG group names are displayed.

4. In the **From** section, do the following to specify input conditions or match criteria for the Packet Filter term :

- In the **Source Address** field, click the down arrow in the list. The address selector dialog box appears. Select the source addresses that need to be added to the Packet Filter policy from the Available column and click the right arrow to move these devices to the Selected column.

Click **OK** to confirm the selection. Click **Cancel** to discard your changes and return to the Create Policy and Filter window.

To create an address or address group from the address selector dialog box, see *Creating Addresses* and *Creating Address Groups*.

- In the **Destination Address** field, click the down arrow in the list. The address selector dialog box appears. Select the destination addresses that need to be added to the Packet Filter policy from the Available column and click the right arrow to move these devices to the Selected column.

Click **OK** to confirm the selection. Click **Cancel** to discard your changes and return to the Create Policy and Filter window.

To create an address or address group from the address selector dialog box, see *Creating Addresses* and *Creating Address Groups*.

- Specify a destination port to match the rule in the **Destination Port** field.
- Specify a source port to match the rule in the **Source Port** field.
- In the Add Term page, in the **Application** or **Application Set** sections, the application set selector dialog box is displayed. Select the applications or application sets that need to be added to the packet filter policy rule term from the Available column and click the right arrow to move these applications or application sets to the Selected column.

To create a new application name or application set, see *Creating Applications and Application Sets*.

- When you create a rule or filter term, and define the name of the filter, for SDGs that are part of a high availability pair of devices, the names of the SDGs are displayed as tabs and check boxes beside the hostnames of the SDGs are displayed. If you want the policy or filter term definition to apply to both the SDGs, select the check boxes next to the SDG names.

Otherwise, when you click the SDG name tab for the SDG for which you did not select the check box, a blue highlight overlays the entire dialog box to indicate the settings are not enabled for configuration for that specific SDG.

- Click the **Copy to All Hosts** button to apply the defined term at the system or network level and not at a particular SDG or SDG group level.

- Select the name of the target application set from the **Application Sets** selector dialog box. Select the application sets that need to be added from the Available Column and click the right arrow to move the application sets to the Selected column.
- In the **Source Prefix** field, click the down arrow in the list to specify the source prefix for rule matching traffic. The address selector dialog box appears. Select the source addresses that need to be added to the Packet Filter policy from the Available column and click the right arrow to move these devices to the Selected column.

Click **OK** to confirm the selection. Click **Cancel** to discard your changes and return to the Create Policy and Filter window.

To create an address or address group from the address selector dialog box, see *Creating Addresses* and *Creating Address Groups*.

- In the **Destination Prefix** field, click the down arrow in the list to specify the destination prefix for rule matching traffic. The address selector dialog box appears. Select the source addresses that need to be added to the packet filter policy from the Available column and click the right arrow to move these devices to the Selected column.

Click **OK** to confirm the selection. Click **Cancel** to discard your changes and return to the Create Policy and Filter window.

To create an address or address group from the address selector dialog box, see *Creating Addresses* and *Creating Address Groups*.

- Select the type of protocol from the **Protocol** drop-down menu. The Protocol selector dialog box appears. Select the protocols you want to add from the Available column, and click the right arrow to move them to the Selected column.
5. In the **To** section, do the following to specify actions or modifiers to be performed for the Packet Filter term :
- In the **Count** field, specify a name for the counter to compute the matched packet in the named counter
  - In the **Forwarding Class** list, select the name of the forwarding class that must be used to classify the packet. Select one of the following options:
    - forwarding-class-name
    - assured-forwarding
    - best-effort
    - expedited-forwarding
    - network-control
  - In the **Actions** field, click the down arrow in the list. Select one of the following options:
    - accept—Accept the traffic and send it on to its destination.
    - discard—Do not accept traffic or process it further.

reject—Do not accept the traffic and return a rejection message. Rejected traffic can be logged or sampled.

count—Add the packet to a counter total.

log—Log the packet.

port-mirror—Port-mirror the packet.

sample—Sample the packet.

service—Forward the packet for service processing.

skip—Omit the packet from service processing.

- In the **Protocol** list, select the protocol for which packets must be classified.
- In the **Routing Instance** list, select the name of the configured routing instance for the SDG or SDG group to enable the packets to be directed for processing.
- Click the **Copy to All Hosts** button to apply the defined term at the system or network level and not at a particular SDG or SDG group level.
- When you create a rule or filter term, and define the name of the filter, for SDGs that are part of a high availability pair of devices, the names of the SDGs are displayed as tabs and check boxes beside the hostnames of the SDGs are displayed. If you want the policy or filter term definition to apply to both the SDGs, select the check boxes next to the SDG names.

Otherwise, when you click the SDG name tab for the SDG for which you did not select the check box, a blue highlight overlays the entire dialog box to indicate the settings are not enabled for configuration for that specific SDG.

- Select the **Syslog** check box to enable system logging. The system log information from the Multiservices PIC is passed to the kernel for logging in the `/var/log` directory.
6. Click **Save** to create the rule. Alternatively, click **Validate** in the Create Rule page to perform validation checks on the configuration planned to be deployed to examine and correct any syntax errors or incompatible settings.
  7. A new rule is added in the last row depending on the type of rule you have added. The newly added rules blink with a different color for few seconds. The behavior is same if you add a new rule before or after a rule, clone a rule, or paste a rule.

The rule is assigned a serial number based on the number of rules already added to the policy.

## Creating an Application and Application Set

To create an application and an application set for a Packet Filter rule term:

1. In the Add Term page, in the **Application** or **Application Set** sections, the application set selector dialog box is displayed. Select the applications or application sets that need to be added to the packet filter term from the Available column and click the right arrow to move these application sets to the Selected column.

## Associating Service Sets and Rule Sets With a Packet Filter Rule

To associate a service set and a rule set with a Packet Filter rule term:

1. In the Create Policy and Filter page, click **Associate Service Sets/Rule Sets**. The Associate Service Sets/Rule Sets section is displayed. The SDGs and SDG groups that are part of the Packet Filter rule term are shown in one column. Under the Association column, either the **Configure** or **Edit** icon appears. If you already created and mapped a service set with the particular SDG or group, the **Edit** icon shows.

2. Click the **Configure** or **Edit** icon. The Configure Service Sets/Rule Sets dialog box is displayed.

3. From the Type drop-down list, do either of the following:

- Select **Service Set** to map a service set with the policy filter template.
- Select **Rule Set** to map a rule set with the policy filter template.

Depending on the option selected in the Type list as service set or rule set for association with the policy filter template, the options that are displayed in the Value list beneath the Type list varies.

4. If you selected **Service Set** from the Type list, select a service set previously configured in the Service Designer workspace from the **Value** list. If you selected **Rule Set** from the Type list, select a rule set previously configured in the Service Designer workspace from the **Value** list. Click **Add** to map the service set or rule set with the Packet Filter rule.
5. Click **Save** to save the settings. Alternatively, click **Cancel** to abort the changes.
6. Click **Copy to All Hosts** in the Associate Service Sets dialog box to apply the defined term at the system or network level and not at a particular SDG or SDG group level. You are returned to the Add Term window.

## Associating Interfaces With a Packet Filter Rule

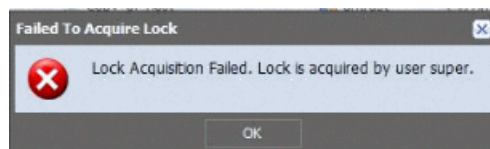
To associate a service set and a rule set with a Packet Filter rule term:

1. In the Create Policy and Filter page, click **Associate Interfaces**. The Associate Interfaces dialog box is displayed. The SDGs and SDG groups that are part of the packet filter rule term are shown in one column. Under the Association column, either the **Configure** or **Edit** link appears. If you already created and mapped a service set with the particular SDG or group, the **Edit** link shows.
2. Click the **Configure** or **Edit** link. The Associate Interfaces dialog box is displayed.
3. Select an interface previously configured in the Service Designer workspace from the **Interfaces** list. Select the logical unit number of the interface from the **Unit** list. Click **Add** to map the interface with the packet filter rule.
4. Click **Done** to save the settings. Alternatively, click **Cancel** to abort the changes.
5. Click **Done** in the Associate Interfaces dialog box. You are returned to the Add Term window.

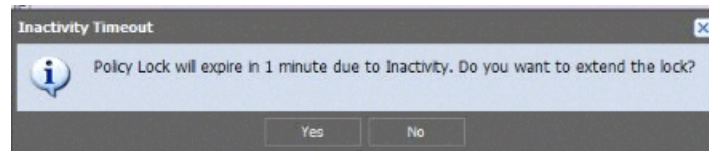
## Modifying Packet Filter Policies

Before you can edit the policy, you must lock it by clicking the lock icon, which is available in the policy tabular view. You can hold more than one policy lock at a given time. You can unlock the policy by clicking the unlock icon next to the lock icon in the policy tabular view. If you attempt to lock a policy that is already locked by another user, the following message appears, as shown in [Figure 40 on page 429](#). The tooltip shows the policy locked user information. Mouse over the policy that you want to lock to view the tooltip.

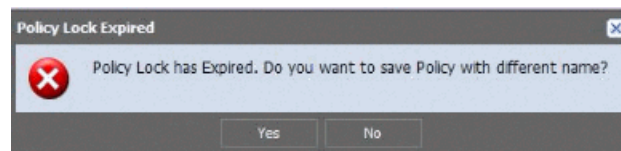
*Figure 40: Lock Failure Error Message for the Second User*



If the locked policy is inactive for the set timeout value (default 5 minutes), just 1 minute before the timeout interval expires, the following message appears, as shown in [Figure 41 on page 430](#). If the policy lock timeout interval expires for multiple locked policies, the same warning message appears for each locked policy. To understand the configuration of timeout value and session timeout value, see *Unlocking Locked Policies*

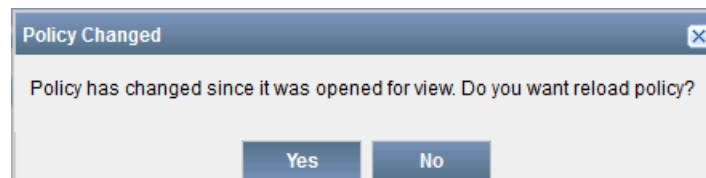
*Figure 41: Inactivity Timeout Error*

Click **Yes** to extend the locking period. If you click **No**, and if there is activity on the policy within the last minute of the lock's life, the timer will be reset and the lock will not be released. If you ignore the message, when the policy lock timeout interval expires 1 minute later, you are prompted to either save the edited policy with a different name or lose the changes, as shown in [Figure 42 on page 430](#)

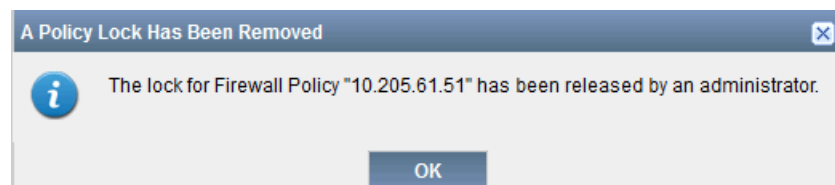
*Figure 42: Policy Lock Expired Message*

If you click **Yes** to save the edited policy with a different name, the Save As window appears. If you navigate away from the locked policy, either the policy is unlocked (when there are no changes) or you will get an option to save the edited policy with a different name.

After editing a locked policy, if you move to another policy without saving your edited policy, or if you unlock the policy without saving, the following warning message appears, as shown in [Figure 43 on page 430](#).

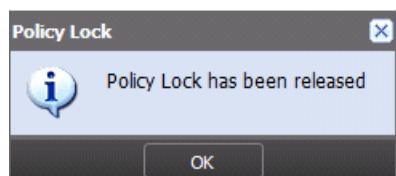
*Figure 43: Packet Filter Policy: Unsaved Changes Message*

If the Edge Services Director administrator releases the lock, you will receive the following warning message, as shown in [Figure 44 on page 430](#).

*Figure 44: Packet Filter Policy: Policy Unlock by Admin Message*

If you do not edit the locked policy and the policy lock timeout expires, the following warning message appears, as shown in [Figure 45 on page 431](#).

Figure 45: Packet Filter Policy Lock Release Message



The policy is locked and released for the following policy operations. Also, these operations are disabled for a policy, if the policy is locked by some other user.

- Modify
- Assign devices
- Rollback
- Delete



NOTE:

- You can unlock the policy by logging out of the application or when the policy lock timeout expires. You can unlock your policies even if they are not edited.
- If the browser crashes when the policy is still locked, the policy is unlocked only after the timeout interval expires.
- If there is an object conflict resolution during a migration, import, or rollback, and if you are editing any objects, you will receive a **save as** option for the edited objects. The behavior is the same when you import addresses from CSV.
- Policy lock is not released under the following scenario:
  - If you save or discard your changes to the locked policy.
  - If you do not make any changes to the locked policy and navigate to another policy.
- It is recommended to configure the session time longer than the lock timeout value.

To modify an existing Packet Filter policy or filter template:

1. From the View selector, select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.

3. Select **Service Edit > Policy and Filter** from the task pane. The Packet Filter Policies page is displayed.
4. From the task pane, select **Packet Filter Policy and Filter** to open the Packet Filter and Filter page on the right pane.
5. Select a policy, and click the Lock icon above the table of listed policies.
6. From the Service Gateway Name drop-down list, select the SDG group to which the packet filter must be applied.
7. From the Host Name drop-down list, select the hostname of the SDG.
8. In the Select Common Components section, select the check boxes beside the service modules or components, such as packet filters, SFW rules, or CGNAT rules, that are displayed. The displayed components depend on the attributes that are previously defined for that selected packet filter. For example, if the service policy is for stateful firewall, SFW rules and SFW rule sets are shown. Select the check box beside Config Category to select all the service components.
9. Click **Save** to save the modified association.
10. Select the check box beside the template you want to modify.
11. Click the **Modify** button above the table of listed templates. The Modify Policy and Filter window is displayed.
12. Modify the attributes that are needed and save the updated settings.

## Creating a Deployment Plan

You must have previously defined service instances and policy or filter instances before you can create a deployment plan.

To create a deployment plan and assigning devices to it:

1. From the View selector, select **Gateway View** or **Service View**. In Gateway view, the devices in the entire network are displayed, organized by the device types and the device models within each device type. In Service View, the different types of services are displayed in the View pane. select **Service View**. The workspaces that are applicable to this view are displayed.
2. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.

3. From the View pane, select the All Network item in Gateway view. Click the plus sign (+) beside the All Network item in the View pane to expand the tree and select the device node you want. Alternatively, from the View pane, click the plus sign (+) beside All Services to expand the tree and select the type of service.
4. From the task pane, select Service Edit. The Service Templates page is displayed.
5. If you are in Gateway view, click the plus sign (+) next to Service Edit to expand the tree in the task pane and view the list of filter templates.
6. In the Service Edit page, from the tree that lists the SDGs, select All Service Gateways, or the SDG or SDG pair for which you want to view the previously configured policy or filter templates. This step is applicable only if you selected Gateway View.

The list of SDGs are displayed on the left pane. You can drill-down to the SDG or pair of SDGs for which you want to process policies or filters. The policy and filter rules are displayed in the right pane.
7. If you are in Service View, from the View pane, select the All Services item. The Services page is displayed.
8. From the task pane, select **Deploy Service > Packet Filter**. The Packet Filter Policies page is displayed.
9. Select the check boxes next to the policy instances that you want to assign to the plan.
10. Click the down arrow in the **Actions** menu and select **Send for Deployment** to create a deployment plan for the particular service instance and save the plan.
  - If you create a deployment plan from Gateway view of Deploy mode, the Deployment Plan Summary dialog box appears, with the service name, type, and status listed.

Click **Send** to create a deployment plan.
  - If you create a deployment plan from Service view of Deploy mode, the Edit Service Instance page is displayed. You can modify the SDGs associated with the service instance and also modify the service instance attributes as necessary by either clicking the buttons corresponding to the various settings at the top of the wizard page to directly traverse to the page you want to modify or clicking the navigation buttons at the bottom of the wizard page to go to the different pages of the wizard. Click **Finish** to create a deployment plan.

A deploy plan is created for the service instance with the devices that are assigned to it when you view the Deployment Plans page.

11. From the Deployment plans page, you can select **Reject** or **Approve** from the Actions drop-down list to reject or approve the deployment plan and make it available for commissioning to the devices.

**Related  
Documentation**

- [Policy and Filter Management Overview on page 375](#)
- [Packet and Service Filters Overview on page 378](#)
- [Searching for CGNAT Policies on page 381](#)
- [Searching for Packet Filters on page 384](#)
- [Searching for SFW Policies on page 386](#)
- [Managing Service and Policy Locks on page 387](#)
- [Unlocking Locked Services and Policies on page 389](#)
- [Viewing Policy and Filter Instances on page 390](#)

---

## Creating and Managing SFW Policy and Filter Instances

---

A stateless firewall filter, often called a firewall filter or access control list (ACL), statically evaluates packet contents. In contrast, a stateful firewall filter, or stateful firewall policy, uses connection state information derived from past communications and other applications to make dynamic control decisions.

Each stateful firewall rule consists of a set of terms, similar to a filter configured at the **[edit firewall]** hierarchy level. Each rule must include a **match-direction** statement that specifies the direction in which the rule match is applied. To configure where the match is applied, include the **match-direction** statement at the **[edit services stateful-firewall rule *rule-name*]** hierarchy level:

```
[edit services stateful-firewall rule rule-name]  
match-direction (input | output | input-output);
```

If you configure **match-direction input-output**, sessions initiated from both directions might match this rule.

The match direction is used with respect to the traffic flow through the AS or Multiservices PIC. When a packet is sent to the PIC, direction information is carried along with it.

With an interface service set, packet direction is determined by whether a packet is entering or leaving the interface on which the service set is applied.

With a next-hop service set, packet direction is determined by the interface used to route the packet to the AS or Multiservices PIC. If the inside interface is used to route the packet, the packet direction is input. If the outside interface is used to direct the packet to the PIC, the packet direction is output.

On the PIC, a flow lookup is performed. If no flow is found, rule processing is performed. Rules in this service set are considered in sequence until a match is found. During rule processing, the packet direction is compared against rule directions. Only rules with direction information that matches the packet direction are considered. Most packets result in the creation of bidirectional flows.



**NOTE:** Before you create a policy and filter template for packet filters, SFW, or CGNAT services, you must have previously configured the different elements or attributes of the service, such as service sets, interface sets, rule sets, and syslogs during the creation of the service template. The sections in this procedural topic that describe the creation of such service elements apply during the creation of the service template and not during the creation of the service policy filters, such as CGNAT or SFW policies.

- [Creating an SFW Policy on page 435](#)
- [Creating a Service Set on page 438](#)
- [Creating a Syslog on page 442](#)
- [Creating a Rule on page 445](#)
- [Creating a Rule Set on page 446](#)
- [Creating Addresses on page 447](#)
- [Creating Address Groups on page 449](#)
- [Address and Address Groups Overview on page 449](#)
- [Creating an SFW Rule Term on page 449](#)
- [Creating an Application and Application Set on page 452](#)
- [Associating Service Sets and Rule Sets With an SFW Rule on page 452](#)
- [Modifying SFW Policies on page 453](#)
- [Creating a Deployment Plan on page 454](#)

## Creating an SFW Policy

To configure a new SFW policy or filter instance:

1. From the View selector, select **Gateway View**. The View pane displays the devices in the entire network organized by the device type and device models pertaining to each device type.
2. From the View pane, select the All Network item. Expand the tree to select the SDG in an SDG group.
3. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.

4. Select **Service Edit** from the task pane. The different service types are displayed in the task pane.
5. Click the right arrow next to Service Edit in the task pane to expand the tree in the task pane and view the list of filter instances.
6. From the task pane, select **SFW Policy and Filter** to open the SFW Policy and Filter page on the right pane.
7. Click the **Add** icon above the table of listed templates. The Create Policy and Filter window is displayed.

Figure 46: Create SFW Policy Window

**Create SFW Policy**

Name:

Match Direction:

Service Gateway:

Associate Service Sets

Service Gateway	Rule Set	Service Set
dsffdfdf	mobst480x	
	mobst480w	

+    ✎    -

Host Name	Term Name
-----------	-----------

Create Cancel

8. Enter the name of the group policy in the Name field.
9. Enter a description for the group policy rules in the Description field. Edge Services Director sends the comments entered in this field to the device.

10. In the Match Direction list, specify the direction in which the rule match is applied. Select one of the following options:
  - **input**—Apply the rule match on the input side of the interface.
  - **input-output**—Apply the rule match bidirectionally.
  - **output**—Apply the rule match on the output side of the interface.
11. In the SDG section, do the following:
  - From the SDG drop-down list, select the devices with which the NAT policy must be associated. Alternatively, you can select the high availability pair of SDG devices with which the NAT policy must be associated. All of the devices in the different SDG groups that were previously defined in the database are also listed in the drop-down menu.
12. Create an SFW rule term that must be added to the SFW policy. For details on configuring an SFW rule term, see *Creating an SFW Rule Term*.
13. The list of terms added, and the associated service sets and rule sets, are displayed in a tabular format in the Create Policy and Filter page. Select the check box next to the term you want to attach to the SFW policy.
14. Click **Create** to save the SFW policy.
15. Click **Validate** to perform validation checks on the configuration planned to be deployed to examine and correct any syntax errors or incompatible settings. You can also validate without deploying the configuration.



**NOTE:** In the Create Policy and Filter window, you can also do the following:

- Click the **Create** icon displayed beside the terms or attributes to add a new attribute. You can then use the newly defined attribute to add to a policy to cause the same selection for a particular term to be applied across all SDGs or groups.
- Click the **Edit** icon displayed beside the terms or attributes to modify an attribute. You can then use the modified attribute to add to a policy to cause the same selection for a particular term to be applied across all SDGs or groups.
- Select the check box beside the SDGs or SDG groups in the Create SFW Term page to include the devices or the SDG groups in the SFW policy for association. Deselect the check boxes beside the SDGs or groups to exclude the devices in the SFW policy.
- Click the **Copy to All Hosts** button to apply the defined term at the system or network level and not at a particular SDG or SDG group level.

## Creating a Service Set

A service set is a collection of services to be performed by an Adaptive Services (AS) or Multiservices PIC. To create a service set as a component for the SFW template:

1. From the View selector, select **Gateway View**. The View pane displays the devices in the entire network organized by the device type and device models pertaining to each device type.
2. From the View pane, select the All Network item. Expand the tree to select the SDG in an SDG group.
3. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
4. Select **Service Edit** from the task pane. The different types of services are displayed in the task pane.
5. Click the right arrow next to Service Edit to expand the tree in the task pane and view the list of filter instances.
6. From the task pane, select **SFW Policy and Filter** to open the SFW and Filter page on the right pane.
7. Click the **Add** icon. The Create an SFW Policy and Filter Template window appears.
8. Enter the name of the template, a description, and the direction in which the rule match must be applied in the respective fields. Also, select the SDG or SDG pair for which the syslog needs to be defined for the service set.
9. In the Create Policy and Filter page, click **Associate Service Sets/Rule Sets**. The Associate Service Sets/Rule Sets section is displayed. The SDGs and SDG groups that are part of the NAT policy filter rule term are shown in one column. Under the Association column, either the **Configure** or **Edit** icon appears. If you already created and mapped a service set with the particular SDG or group, the **Edit** icon shows.
10. Click the **Configure** or **Edit** icon. The Configure Service Sets/Rule Sets dialog box is displayed.
11. From the Type drop-down list, select **Service Set** to map a service set with the policy filter instance.
12. If you selected **Service Set** from the Type list, select a service set previously configured in the Service Designer workspace from the **Value** list.

13. Click the green plus sign next to the Value drop-down list. The Addition of Service Sets dialog box appears.



**NOTE:** If a green plus sign mark is shown beside a field in the dialog box, it denotes that you can add attributes for that component. A red minus mark shows that you can delete that particular attribute for that component.

14. In the Name field, enter the name to identify the service set. Rules are combined into rule sets, and are associated with a service set for each application such as firewall or CGNAT.

15. In the Sampling Service Choices section, do one of the following:

- Click **Interface Services** to configure an interface-style service set. An interface service set is used as an action modifier across an entire interface
  - In the Service Interfaces field, specify the name for the adaptive services interface associated with an interface-wide service set.

When you have defined and grouped the service rules by configuring the service-set definition, you can apply services to one or more interfaces installed on the router. When you apply the service set to an interface, it automatically ensures that packets are directed to the PIC.

- From the **Load Balancing Options** section, configure the high availability (HA) options.

The following hash keys can be configured in the egress direction: **destination-ip** (Use the destination IP address of the flow to compute the hash used in load balancing.) and **source-ip** (Use the source IP address of the flow to compute the hash used in load balancing.)

- Click the green tick mark beside the Egress Key element to configure the hash keys to be used in the egress flow direction. The configuration is mandatory if you are using AMS for Network Address Translation (NAT). This configuration is not mandatory if you are using AMS for stateful firewall; if the hash keys are not configured, then the defaults are chosen.
- Click the green tick mark beside the Ingress Key element to configure the hash keys to be used in the ingress flow direction. The configuration is mandatory if you are using AMS for Network Address Translation (NAT). This configuration is not mandatory if you are using AMS for stateful firewall; if the hash keys are not configured, then the defaults are chosen.

Configure the hash keys used for load balancing in aggregated multiservices (AMS) for service applications (Network Address Translation [NAT], stateful firewall, application-level gateway [ALG], HTTP header enrichment, and mobility). The hash keys supported in the ingress and egress direction are the source IP address and destination IP address.

Hash keys are used to define the load-balancing behavior among the various members in the AMS group. For example, if **hash-keys** is configured as **source-ip**, then the hashing would be performed based on the source IP address of the packet. Therefore, all packets with the same source IP address land on the same member. Hash keys must be configured with respect to the traffic direction: ingress or egress. For example, if **hash-keys** is configured as **source-ip** in the ingress direction, then it should be configured as **destination-ip** in the egress direction. This is required to ensure that the packets of the same flow reach the same member of the AMS group.

The configuration of the ingress and egress hash keys is mandatory if you are using AMS for NAT. This configuration is not mandatory if you are using AMS for stateful firewall; if the hash keys are not configured, then the defaults are chosen. Refer to [Table 46 on page 245](#) for the supported hash keys.

The resource-triggered option enables anchor session PICs to use the load or resource information from the anchor services PICs to select the AMS member will anchor the services for the subscriber for load balancing among AMS members. In addition, for mobile subscriber-aware services (such as HTTP header enrichment), you must configure the **resource-triggered** statement, which means that the load balancing is not done using the ingress and egress keys.

**Table 66: Hash Keys Supported for AMS for Service Applications**

Service Set at Ingress Interface			Service Set at Egress Interface	
Hash Keys for NAT				
NAT Type	Ingress hash key	Egress hash key	Ingress hash key	Egress hash key
source static	Destination IP address	Source IP address	Source IP address	Destination IP address
source dynamic	Source IP address	Destination IP address	Destination IP address	Source IP address
Network Address Port Translation (NAPT)	Source IP address	Destination IP address	Destination IP address	Source IP address
destination static	Source IP address	Destination IP address	Destination IP address	Source IP address
Hash Keys for Stateful Firewall				
Stateful Firewall	Destination IP address	Source IP address	Destination IP address	Source IP address
Stateful Firewall	Source IP address	Destination IP address	Source IP address	Destination IP address



**NOTE:** If NAT is used in the service set (along with stateful firewall and ALG), then the hash keys should be based on the NAT type; otherwise, the hash keys of the stateful firewall should be used.

- Click **Next Hop Services** to configure a next-hop style service set. A next-hop service set is a route-based method of applying a particular service. Only packets destined for a specific next hop are serviced by the creation of explicit static routes.

- In the **Inside Interface** list, specify the interface type of the service interface associated with the service set applied inside the network. For inline IP reassembly, set the interface type to local. Also, specify the name and logical unit number of the service interface associated with the service set applied inside the network.

When a next-hop service is configured, the AS or Multiservices PIC is considered to be a two-legged module with one leg configured to be the inside interface (inside the network) and the other configured as the outside interface (outside the network).

- In the **Outside Interface** list, specify the interface type of the service interface associated with the service set applied outside the network. For inline IP reassembly, set the interface type to local. Also, specify the name and logical unit number of the service interface associated with the service set applied outside the network.
- In the **Service Interface Pool** list, select the name of the pool of logical interfaces configured at the [edit services service-interface-pools pool pool-name] hierarchy level. You can configure a service interface pool only if the service set has a PGCP rule configured. The service set cannot contain any other type of rule.

- 

- Click **Sampling Services** to configure a sampling service set.
  - In the Service Interface field, specify the service interface, which is the interface the sampling is taken from. In the case of a sampling service set, the service interface must be a Multiservices PIC interface with a subunit number of 0 (zero). The subunit number defaults to 0. The reverse-flow statement is not mandatory. All sampled traffic is considered to be forward traffic. If you set the reverse-flow statement, it is ignored.
- Select the **Replication Service** check box to configure the services replication options for inter-chassis high availability on MS-MIC and MS-MPC. This field is available only if you selected the Junos OS 12.1 version.
  - In the Replication Threshold field, specify the number of seconds for the replication threshold. When a flow has been active for more than the number of seconds specified as a threshold, flow state information is replicated to the backup device. Make sure that the replication-threshold value is than the open-timeout value (the timeout period for establishing a TCP connection). The default value of the replication threshold is 180 seconds. This value is also the minimum.
  - Select the **Stateful Firewall** check box to replicate stateful firewall state information.
  - Select the **NAT** check box to replicate NAT44 information.

16. Select the **Service Set Options** check box to specify the service set options to apply to a service set. This field is available only if you selected the Junos OS 14.1 version.

17. In the Redundancy Set ID field, specify a unique identifier in the range of 1 through 100 for the redundancy set. The redundancy group IDs that the service redundancy daemon (srd) uses are associated with those configured for the ICCP daemon (iccpd) through

the existing ICCP configuration hierarchy by using the same redundancy group ID in the configuration of the services redundancy group. This field is available only if you selected the Junos OS 14.1 version.

The actions to be performed when configured redundancy events occur are defined in redundancy policies. Redundancy policies are associated with redundancy sets; they are analogous to rules associated with service sets. Redundancy sets are associated to redundancy groups by redundancy group IDs. Redundancy group details are defined by the underlying ICCPd configuration. Finally, service sets and redundancy sets are associated through the **redundancy-sets** statement in service sets configuration.

18. In the SFW Rule Sets section, select the rule set you want to associate with the service set from the Available column and click the right arrow to move to the Selected column.
19. In the SFW Rules section, select the rule you want to associate with the service set from the Available column and click the right arrow to move to the Selected column.
20. In the SFW Syslogs section, select the syslog you want to associate with the service set from the Available column and click the right arrow to move to the Selected column.
21. Click **Save** to save the service instance configuration. Else, click **Close** to discard the changes to the template.

## Creating a Syslog

You can enable system logging. The system log information from the Adaptive Services or Multiservices PIC is passed to the kernel for logging in the /var/log directory. This setting overrides any syslog statement setting included in the service set or interface default configuration.

To create a syslog for the SFW template:

1. From the View selector, select **Gateway View**. The View pane displays the devices in the entire network organized by the device type and device models pertaining to each device type.
2. From the View pane, select the All Network item. Expand the tree to select the SDG in an SDG group.
3. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.

4. Select **Service Edit > Policy and Filter** from the task pane. The Service Edit > Policy and Filter page is displayed.
5. Click the plus sign (+) next to Policy and Filter to expand the tree in the task pane and view the list of filter instances.
6. From the task pane, select **SFW Policy and Filter** to open the SFW and Filter page on the right pane.
7. Click the **Add** icon. The Create an SFW Policy and Filter Template window appears.
8. Enter the name of the template, a description, and the direction in which the rule match must be applied in the respective fields. Also, select the SDG or SDG pair for which the syslog needs to be defined for the service set.
9. In the Create Policy and Filter page, click **Associate Service Sets/Rule Sets**. The Associate Service Sets/Rule Sets section is displayed. The SDGs and SDG groups that are part of the NAT policy filter rule term are shown in one column. Under the Association column, either the **Configure** or **Edit** icon appears. If you already created and mapped a service set with the particular SDG or group, the **Edit** icon shows.
10. Click the **Configure** or **Edit** icon. The Configure Service Sets/Rule Sets dialog box is displayed.
11. From the Type drop-down list, select **Service Set** to map a service set with the policy filter instance.
12. If you selected **Service Set** from the Type list, select a service set previously configured in the Service Designer workspace from the **Value** list.
13. Click the green plus sign next to the Value drop-down list. The Addition of Service Sets dialog box appears.



**NOTE:** If a green plus sign mark is shown beside a field in the dialog box, it denotes that you can add attributes for that component. A red minus mark shows that you can delete that particular attribute for that component.

14. Click the green plus sign next to the Syslog Settings field. The Addition of Service Sets dialog box appears.
15. In the Host field, enter the hostname for the syslog component. Specify the fully qualified domain name or IP address for the syslog server.

16. In the Services list, specify the system logging severity level. It assigns a severity level to the facility. Valid entries include:
  - **alert**—Conditions that should be corrected immediately.
  - **any**—Matches any level.
  - **critical**—Critical conditions.
  - **emergency**—Panic conditions.
  - **error**—Error conditions.
  - **info**—Informational messages.
  - **notice**—Conditions that require special handling.
  - **warning**—Warning messages.
17. From the Facility Override list, select the override for the default facility for system log reporting. Valid values include:
  - authorization**
  - daemon**
  - ftp**
  - kernel**
  - local0 through local7**
  - user**
18. In the Log Prefix field, set the system logging prefix value for all logging to the system log host.
19. In the Port field, specify the port number to be used for connection with the remote syslog server.
20. In the Source Address field, specify a source address to record in system log messages that are directed to a remote machine specified in the hostname statement. The supported interfaces are ms, rms, and mams interfaces. If you do not specify the interface parameter, the command loops on all supported interfaces. This field is available only if you selected the Junos OS 14.1 version.
21. In the Class section, set the class of applications to be logged to the system log.
  - **alg-logs**—Log application-level gateway events.
  - **ids-logs**—Log intrusion detection system events.
  - **nat-logs**—Log Network Address Translation events.
  - **packet-logs**—Log general packet-related events.
  - **session-logs**—Log session open and close events.
  - **session-logs open**—Log session open events only.

- **session-logs close**—Log session close events.
- **stateful-firewall-logs**—Log stateful firewall events.

22. Click **Save** to save the service instance configuration. Else, click **Close** to discard the changes to the template.

## Creating a Rule

To create a rule for the SFW template:

1. From the View selector, select **Gateway View**. The View pane displays the devices in the entire network organized by the device type and device models pertaining to each device type.
2. From the View pane, select the All Network item. Expand the tree to select the SDG in an SDG group.
3. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.

4. Select **Service Edit** from the task pane.

The Service Edit page is displayed.

5. Click the right arrow next to Service Edit to expand the tree in the task pane and view the list of filter instances.
6. From the task pane, select **SFW Policy and Filter** to open the SFW and Filter page on the right pane.

7. Click the **Add** icon.

The Create an SFW Policy and Filter Template window appears.

8. Enter the name of the template and the service instance in the respective fields.

9. Click the green plus sign in the Rules box. The Addition of Rules dialog box appears.



**NOTE:** If a green tick mark is shown beside a field in the dialog box, it denotes that you can add attributes for that component. A red cross mark shows that you can delete that particular attribute for that component.

10. From the **Rule** list, select one of the previously configured rules. The rules that you configured in the Service Templates workspace for SFW, packet filter, or CGNAT are displayed.
11. Click **Save** to save the service instance configuration. Else, click **Close** to discard the changes to the template.


## Creating a Rule Set

The rule-set statement defines a collection of stateful firewall rules that determine what actions the router software performs on packets in the data stream. You define each rule by specifying a rule name and configuring terms. Then, you specify the order of the rules by including the **rule-set** statement at the **[edit services stateful-firewall]** hierarchy level with a rule statement for each rule.

The router software processes the rules in the order in which you specify them in the configuration. If a term in a rule matches the packet, the router performs the corresponding action and the rule processing stops. If no term in a rule matches the packet, processing continues to the next rule in the rule set. If none of the rules matches the packet, the packet is dropped by default.

To create a rule set for the SFW template:

1. From the View selector, select **Gateway View**. The View pane displays the devices in the entire network organized by the device type and device models pertaining to each device type.
2. From the View pane, select the All Network item. Expand the tree to select the SDG in an SDG group.
3. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
4. Select **Service Edit** from the task pane. The Service Edit page is displayed.
5. Click the right arrow next to Service Edit to expand the tree in the task pane and view the list of filter instances.
6. From the task pane, select **SFW Policy and Filter** to open the SFW and Filter page on the right pane.
7. Click the **Add** icon. The Create an SFW Policy and Filter Template window appears.
8. Enter the name of the template, a description, and the direction in which the rule match must be applied in the respective fields. Also, select the SDG or SDG pair for which the syslog needs to be defined for the service set.

9. In the Create Policy and Filter page, click **Associate Service Sets/Rule Sets**. The Associate Service Sets/Rule Sets section is displayed. The SDGs and SDG groups that are part of the NAT policy filter rule term are shown in one column. Under the Association column, either the **Configure** or **Edit** icon appears. If you already created and mapped a service set with the particular SDG or group, the **Edit** icon shows.
  10. Click the **Configure** or **Edit** icon. The Configure Service Sets/Rule Sets dialog box is displayed.
  11. From the Type drop-down list, select **Service Set** to map a service set with the policy filter instance.
  12. If you selected **Service Set** from the Type list, select a service set previously configured in the Service Designer workspace from the **Value** list.
  13. Click the green plus sign next to the Value drop-down list. The Addition of Service Sets dialog box appears.
-  **NOTE:** If a green plus sign mark is shown beside a field in the dialog box, it denotes that you can add attributes for that component. A red minus mark shows that you can delete that particular attribute for that component.
14. In the Name field, specify a name for the rule set the router uses when applying this service.
  15. In the Rules section, select the rules that need to be added to the rule set from the Available column and click the right arrow to move these rules to the Selected column. All the rules that you previously configured during the creation or modification of the service instance are displayed.
  16. Click **Save** to save the rule set configuration. Else, click **Close** to discard the changes to the template.

## Creating Addresses

To create an address:

1. In the Source and Destination Address Selector dialog box, to create a new address, click the plus sign (+).  
The Create Address page appears.
2. In the Object Type section, click the **Address** radio button to create an address.

3. In the Name field, enter a name for the new address.
4. In the Description field, enter a description for the new address.
5. Direct Edge Services Director to resolve an IP address to a hostname or resolve a hostname to an IP address.
  - To specify an IP address as the address type, select **Host** from the drop-down menu and enter the **IP** address in the IP field.
  - To specify a hostname as the address type, select **Host** from the drop-down menu and enter the hostname in the Host Name field.
  - To specify an IP address range, select **Range** from the drop-down menu and enter the IP ranges in the Start IP and End IP fields.
  - To specify a network as an address type, select **Network** from the drop-down menu and enter the network address in the IP and Netmask fields.
  - To specify an IP address with a wildcard mask, select **Wildcard** from the drop-down menu and enter the IP address in the IP field and wildcard mask in the Wildcard Mask fields.
  - To specify a DNS name as an address type, select **DNS Host** from the drop-down menu and enter the DNS name in the DNS Name field.



**NOTE:** You can resolve an IP address to a hostname and a hostname to an IP address using the green arrows next to the IP and Host Name fields.

---



**NOTE:** The host and network address types support both IPv4 and IPv6 address types. These address types also supports multicast addresses. However, the range address type supports only IPv4 addresses. NAT and IPsec VPNs do not support IPv6 addressing and wildcard addresses.

---



**NOTE:** Ensure that the first 8 bits of the address are not 0 and the highest bit of the mask is 1 when you are using the wildcard address type.

---

6. Click **Create** to create an address.

The new address appears in the Manage Address page.

## Creating Address Groups

To create an address group:

1. In the Source and Destination Address Selector dialog box, to create a new address group, click the plus sign (+).

The Create Address Group page appears.

2. Select the Object Type as Address Group.
3. In the Name field, enter a name for the new address group.
4. In the Description field, enter a description for the new address group.
5. In the Addresses field, from the Available dialog box, select the address that you want to group, and click the right arrow to add to the Selected column.

Click **All** to move all the addresses to the Selected column. The address you have selected appears in the Selected section of the dialog box.

6. Click **Create**.

The address group appears on the Address page.

## Address and Address Groups Overview

You can use the Address Creation Wizard to create an address object that specifies an IP address or a hostname. You can specify a hostname and use the address resolution option to resolve it to an IP address. You can also resolve an IP address to the corresponding hostname.

You can group address objects to form an address group using the Address Group Creation Wizard. Junos Space creates an object in the Junos Space database to represent an address or an address group.

## Creating an SFW Rule Term

To add rules to an SFW policy:

1. In the Create Policy and Filter window, the list of rule terms already added, if any, to the SFW policy are displayed.
2. Next to the **Terms** field, click the + icon to add rules, and select the type of rule you want to add.
3. In the **Term Name** field, specify the name of the rule.

The list of SDGs with which you associated the SFW policy in the Create Policy window are displayed with the form and then sections or clauses. If you selected SDG groups to associate with the SFW policy, the SDG group names are displayed.

4. In the **From** section, do the following to specify input conditions or match criteria for the SFW term :

- In the **Source Address** field, click the down arrow in the list. The address selector dialog box appears. Select the source addresses that need to be added to the SFW policy from the Available column and click the right arrow to move these devices to the Selected column.

Click **OK** to confirm the selection. Click **Cancel** to discard your changes and return to the Create Policy and Filter window.

To create an address or address group from the address selector dialog box, see *Creating Addresses* and *Creating Address Groups*.

- In the **Destination Address** field, click the down arrow in the list. The address selector dialog box appears. Select the destination addresses that need to be added to the SFW policy from the Available column and click the right arrow to move these devices to the Selected column.

Click **OK** to confirm the selection. Click **Cancel** to discard your changes and return to the Create Policy and Filter window.

To create an address or address group from the address selector dialog box, see *Creating Addresses* and *Creating Address Groups*.

- Specify a destination port to match the rule in the **Destination Port** field. You can specify a range of ports by defining the upper limit and lower limit of the range in the Start Value and End Value fields.
- In the Add Term page, in the **Application** or **Application Set** sections, the application set selector dialog box is displayed. Select the applications or application sets that need to be added to the SFW rule term from the Available column and click the right arrow to move these applications or application sets to the Selected column.

To create a new application name or application set, see *Creating Applications and Application Sets*.

- Click the **Copy to All Hosts** button to apply the defined term at the system or network level and not at a particular SDG or SDG group level.
- When you create a rule or filter term, and define the name of the filter, for SDGs that are part of a high availability pair of devices, the names of the SDGs are displayed as tabs and check boxes beside the hostnames of the SDGs are displayed. If you want the policy or filter term definition to apply to both the SDGs, select the check boxes next to the SDG names.

Otherwise, when you click the SDG name tab for the SDG for which you did not select the check box, a blue highlight overlays the entire dialog box to indicate the settings are not enabled for configuration for that specific SDG.

- Select the name of the target application set from the **Application Sets** selector dialog box. Select the application sets that need to be added from the Available Column and click the right arrow to move the application sets to the Selected column.
- In the **Source Prefix** field, click the down arrow in the list to specify the source prefix for rule matching traffic. The address selector dialog box appears. Select the source addresses that need to be added to the NAT policy from the Available column and click the right arrow to move these devices to the Selected column.

Click **OK** to confirm the selection. Click **Cancel** to discard your changes and return to the Create Policy and Filter window.

To create an address or address group from the address selector dialog box, see *Creating Addresses* and *Creating Address Groups*.

- In the **Destination Prefix** field, click the down arrow in the list to specify the destination prefix for rule matching traffic. The address selector dialog box appears. Select the source addresses that need to be added to the NAT policy from the Available column and click the right arrow to move these devices to the Selected column.

Click **OK** to confirm the selection. Click **Cancel** to discard your changes and return to the Create Policy and Filter window.

To create an address or address group from the address selector dialog box, see *Creating Addresses* and *Creating Address Groups*.

5. In the **To** section, do the following to specify actions or modifiers to be performed for the SFW term :

- In the **Actions** field, click the down arrow in the list. Select one of the following options:

accept—Accept the traffic and send it on to its destination.

discard—Do not accept traffic or process it further.

reject—Do not accept the traffic and return a rejection message. Rejected traffic can be logged or sampled.

- Click the **Copy to All Hosts** button to apply the defined term at the system or network level and not at a particular SDG or SDG group level.
- When you create a rule or filter term, and define the name of the filter, for SDGs that are part of a high availability pair of devices, the names of the SDGs are displayed as tabs and check boxes beside the hostnames of the SDGs are displayed. If you want the policy or filter term definition to apply to both the SDGs, select the check boxes next to the SDG names.

Otherwise, when the click the SDG name tab for the SDG for which you did not select the check box, a blue highlight overlays the entire dialog box to indicate the settings are not enabled for configuration for that specific SDG.

- Select the **Syslog** check box to enable system logging. The system log information from the Multiservices PIC is passed to the kernel for logging in the `/var/log` directory. This field is available only if you selected the Junos OS 14.1 version to create the service instance.

6. A new rule is added in the last row depending on the type of rule you have added. The newly added rules blink with a different color for few seconds. The behavior is same if you add a new rule before or after a rule, clone a rule, or paste a rule.

The rule is assigned a serial number based on the number of rules already added to the policy.

7. Click **Save** to create the rule. Alternatively, click **Validate** in the Create Rule page to perform validation checks on the configuration planned to be deployed to examine and correct any syntax errors or incompatible settings.

## Creating an Application and Application Set

To create an application and an application set for an SFW rule term:

1. In the Add Term page, in the **Application** or **Application Set** sections, the application set selector dialog box is displayed. Select the applications or application sets that need to be added to the SFW rule term from the Available column and click the right arrow to move these applications or application sets to the Selected column.

## Associating Service Sets and Rule Sets With an SFW Rule

To associate a service set and a rule set with an SFW policy filter rule term:

1. In the Create Policy and Filter page, click **Associate Service Sets/Rule Sets**. The Associate Service Sets/Rule Sets section is displayed. The SDGs and SDG groups that are part of the SFW policy filter rule term are shown in one column. Under the Association column, either the **Configure** or **Edit** icon appears. If you already created and mapped a service set with the particular SDG or group, the **Edit** icon shows.
2. Click the **Configure** or **Edit** icon. The Configure Service Sets/Rule Sets dialog box is displayed.
3. From the Type drop-down list, do either of the following:
  - Select **Service Set** to map a service set with the policy filter instance.
  - Select **Rule Set** to map a rule set with the policy filter instance.

Depending on the option selected in the Type list as service set or rule set for association with the policy filter instance, the options that are displayed in the Value list beneath the Type list varies.

4. If you selected **Service Set** from the Type list, select a service set previously configured in the Service Designer workspace from the **Value** list. If you selected **Rule Set** from the Type list, select a rule set previously configured in the Service Designer workspace from the **Value** list. Click **Add** to map the service set or rule set with the SFW policy filter rule.

5. Click **Save** to save the settings. Alternatively, click **Cancel** to abort the changes.
6. Click **Copy to All Hosts** in the Associate Service Sets dialog box to apply the defined term at the system or network level and not at a particular SDG or SDG group level. You are returned to the Add Term window.

## Modifying SFW Policies

Before you can edit the policy, you must lock it by clicking the lock icon, which is available in the policy tabular view. You can hold more than one policy lock at a given time. You can unlock the policy by clicking the unlock icon next to the lock icon in the policy tabular view. If you attempt to lock a policy that is already locked by another user, a message is displayed stating that the lock is acquired by another user.

If the Edge Services Director administrator releases the lock, you will receive the a warning message stating that the lock has been released.

The Manage Policy Locks page appears showing only those locks that can be managed by the current user. The page contains the following fields:

- Instance or Rule name
- User (IP Address)
- Lock acquired time
- Service Gateway

The policy is locked and released for the following policy operations. Also, these operations are disabled for a policy, if the policy is locked by some other user.

- Modify
- Assign devices
- Rollback
- Delete



### NOTE:

- You can unlock your policies even if they are not edited.
- If the browser crashes when the policy is still locked, the policy is unlocked only after the timeout interval expires.
- Policy lock is not released under the following scenario:
  - If you save or discard you changes to the locked policy.
  - if you do not make any changes to the locked policy and navigate to another policy.

To modify an existing SFW policy or filter instance:

1. From the View selector, select **Gateway View**. The workspaces that are applicable to this view are displayed. In Gateway view, the devices in the entire network are displayed, organized by the device types and the device models within each device type. In Service View, the different types of services are displayed in the View pane.
2. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the View pane, select the All Network item in Gateway view. Click the plus sign (+) beside the All Network item in the View pane to expand the tree and select the device node you want.  
  
Alternatively, from the View pane, click the plus sign (+) beside All Services to expand the tree and select the type of service.
4. From the task pane, select **Service Edit**. The Service Templates page is displayed.
5. If you are in Gateway view, click the plus sign (+) next to Service Edit to expand the tree in the task pane and view the list of filter instances.
6. From the task pane, select **SFW Policy and Filter** to open the SFW and Filter page on the right pane.
7. In the Service Edit page, from the tree that lists the SDGs, select **All Service Gateways**, or the SDG or SDG pair for which you want to view the previously configured policy or filter instances. This step is applicable only if you selected Gateway View.  
  
The page is divided into two panes. The list of SDGs are displayed on the left pane. You can drill-down to the SDG or pair of SDGs for which you want to process policies or filters. The policy and filter rules are displayed in the right pane.
8. Select a policy, and click the Lock icon above the table of listed policies.
9. Click the **Modify** icon above the table of listed templates. The Modify Policy and Filter window is displayed.
10. Modify the attributes that are needed and save the updated settings.

## Creating a Deployment Plan

You must have previously defined service templates and policy or filter templates before you can create a deployment plan.

To create a deployment plan and assigning devices to it:

1. From the View selector, select **Gateway View**. The View pane displays the devices in the entire network organized by the device type and device models pertaining to each device type.
2. From the View pane, select the All Network item. Expand the tree to select the SDG in an SDG group.
3. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
4. Select **Service Edit** from the task pane. The Service Edit page is displayed.
5. Click the right arrow next to Service Edit to expand the tree in the task pane and view the list of filter instances.
6. From the task pane, select **SFW Policy and Filter** to open the SFW Policy and Filter page on the right pane.
7. In the Service Edit page, from the tree that lists the SDGs, select **All Service Gateways**, or the SDG or SDG pair for which you want to view the previously configured policy or filter instances. This step is applicable only if you selected Gateway View. You can drill-down to the SDG or pair of SDGs for which you want to process policies or filters.
8. Select a rule corresponding to an SDG, and click the **Lock** icon above the table of listed policy filters.
9. Click the down arrow in the **Actions** menu and select **Send for Deployment** to create a deployment plan for the particular service template and save the plan.  
  
The Deployment Plan Summary dialog box appears, with the service name, type, and status listed.  
  
Click **Send** to create a deployment plan.  
  
A deploy plan is created for the service template with the devices that are assigned to it when you view the Deployment Plans page.
10. Alternatively, you can select **Discard changes** from the **Actions** menu to ignore the modifications done to a policy or filter template.
11. From the Deployment plans page, you can select **Reject** or **Approve** from the Actions drop-down list to reject or approve the deployment plan and make it available for commissioning to the devices.

**Related Documentation**

- [Policy and Filter Management Overview on page 375](#)
- [Packet and Service Filters Overview on page 378](#)
- [Searching for CGNAT Policies on page 381](#)
- [Searching for Packet Filters on page 384](#)
- [Searching for SFW Policies on page 386](#)
- [Managing Service and Policy Locks on page 387](#)
- [Unlocking Locked Services and Policies on page 389](#)
- [Viewing Policy and Filter Instances on page 390](#)

---

## Viewing CGNAT Service Templates

To view the list of CGNAT service templates:

1. From the View selector, select **Gateway View** or **Service View**. The workspaces that are applicable to this view are displayed. In Gateway view, the devices in the entire network are displayed, organized by the device types and the device models within each device type. In Service View, the different types of services are displayed in the View pane.
2. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the View pane, select the All Network item in Gateway view. Click the plus sign (+) beside the All Network item in the View pane to expand the tree and select the device node you want.  
  
Alternatively, from the View pane, click the plus sign (+) beside All Services to expand the tree and select the type of service.
4. From the task pane, select **Service Edit**. The Service Templates page is displayed.
5. If you are in Gateway view, Click the plus sign (+) next to Service Edit to expand the tree in the task pane and view the list of filter templates.
6. If you are in Gateway View, from the task pane, select **CGNAT** to open the Service Edit > CGNAT page on the right pane.
7. In the Service Edit page, from the tree that lists the SDGs, select **All Service Gateways**, or the SDG or SDG pair for which you want to view the previously configured policy or filter templates. This step is applicable only if you selected Gateway View.

The page is divided into two panes. The list of SDGs are displayed on the left pane. You can drill-down to the SDG or pair of SDGs for which you want to process policies or filters. The policy and filter rules are displayed in the right pane.

The following fields are displayed on the Service Edit > CGNAT page:

*Table 67: CGNAT Service Edit Page*

Field	Description
Instance Name	Name of the configured service template instance
OS Version	Junos OS release number that represents a particular revision of the software that runs on a Juniper Networks routing platform, for example, Junos OS Release 8.5, 9.1, or 9.2. Each Junos OS release has certain new features that complement the software processes that support Internet routing protocols, control the device's interfaces and the device chassis itself, and allow device system management.
Group Name	Name of the SDG group
Reference Host	Hostname of the SDG with which the service instance is associated.
Applications	Name of the applications protocols created for the service template.
Application Sets	Name of the application sets created for the service template.
NAT Pools	Name of the CGNAT pool created for the service template.
NAT Rules	Name of the CGNAT rules created for the service instance.
NAT Rule Sets	Name of the CGNAT rule sets created for the service template.
Syslogs	Name of the syslog created for the service template.
Deployment Plans	Name of the deployment plan with which the service template is attached.

## Viewing SFW Service Templates

To view the list of SFW service templates:

1. From the View selector, select **Gateway View** or **Service View**. The workspaces that are applicable to this view are displayed. In Gateway view, the devices in the entire network are displayed, organized by the device types and the device models within each device type. In Service View, the different types of services are displayed in the View pane.
2. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.

The functionalities that you can configure in this mode are displayed in the task pane.

- From the View pane, select the All Network item in Gateway view. Click the plus sign (+) beside the All Network item in the View pane to expand the tree and select the device node you want.

Alternatively, from the View pane, click the plus sign (+) beside All Services to expand the tree and select the type of service.

- From the task pane, select **Service Edit**. The Service Templates page is displayed.
- If you are in Gateway view, Click the plus sign (+) next to Service Edit to expand the tree in the task pane and view the list of filter templates.
- If you are in Gateway View, from the task pane, select **SFW** to open the Service Edit > SFW page on the right pane.
- In the Service Edit page, from the tree that lists the SDGs, select **All Service Gateways**, or the SDG or SDG pair for which you want to view the previously configured policy or filter templates. This step is applicable only if you selected Gateway View.

The page is divided into two panes. The list of SDGs are displayed on the left pane. You can drill-down to the SDG or pair of SDGs for which you want to process policies or filters. The policy and filter rules are displayed in the right pane.

The following fields are displayed on the Service Edit > SFW page:

**Table 68: SFW Service Edit Page**

Field	Description
Instance Name	Name of the configured service template instance
OS Version	Junos OS release number that represents a particular revision of the software that runs on a Juniper Networks routing platform, for example, Junos OS Release 8.5, 9.1, or 9.2. Each Junos OS release has certain new features that complement the software processes that support Internet routing protocols, control the device's interfaces and the device chassis itself, and allow device system management.
Group Name	Name of the SDG group
Reference Host	Hostname of the SDG with which the service instance is associated.
Applications	Name of the applications protocols created for the service template.
Application Sets	Name of the application sets created for the service template.
SFW Rules	Name of the stateful firewall rules created for the service instance.

*Table 68: SFW Service Edit Page (continued)*

Field	Description
SFW Rule Sets	Name of the stateful firewall rule sets created for the service template.
Syslogs	Name of the syslog created for the service template.
Deployment Plans	Name of the deployment plan with which the service template is attached.

## Viewing and Modifying ADC Service Instances

After you create the adaptive delivery controller (ADC) software service instance to balance user session traffic among a group of available servers that provide shared services using the Service Designer workspace, you can view and modify the components or elements of the service instance by using the Service Edit workspace.

You can perform the following tasks with the Service Edit page for ADC:

- View the list of configured ADC service instances.
- Modify an existing ADC service instance to meet the network needs and deployment scenarios.
- Delete an existing template.
- Transfer the service instance for deployment on a device.
- Discard the changes made to a service instance.
- [Viewing ADC Service Instances on page 459](#)
- [Modifying ADC Service Instances on page 461](#)
- [Creating a Deploy Plan and Provisioning Services Immediately on page 463](#)
- [Filtering ADC Service Instances on page 465](#)
- [Managing ADC Service Instance Locks on page 467](#)
- [Unlocking Locked ADC Service Instances on page 469](#)

## Viewing ADC Service Instances

To view the list of ADC service instances:

1. From the View selector, select **Gateway View** or **Service View**. The workspaces that are applicable to this view are displayed. In Gateway view, the devices in the entire network are displayed, organized by the device types and the device models within each device type. In Service View, the different types of services are displayed in the View pane.
2. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.

- From the View pane, select the All Network item in Gateway view. Click the plus sign (+) beside the All Network item in the View pane to expand the tree and select the device node you want.

Alternatively, from the View pane, click the plus sign (+) beside All Services to expand the tree and select the type of service.

- From the task pane, select **Service Edit**. The Service Instances page is displayed.
- If you are in Gateway view, Click the plus sign (+) next to Service Edit to expand the tree in the task pane and view the list of filter templates.
- If you are in Gateway View, from the task pane, select **ADC** to open the Service Edit > ADC page on the right pane.
- In the Service Edit page, from the tree that lists the SDGs, select **All Service Gateways**, or the SDG or SDG pair for which you want to view the previously configured policy or filter templates. This step is applicable only if you selected Gateway View.

The page is divided into two panes. The list of SDGs are displayed on the left pane. You can drill-down to the SDG or pair of SDGs for which you want to process policies or filters. The policy and filter rules are displayed in the right pane.

The following fields are displayed on the Service Edit > ADC page:

**Table 69: ADC Service Edit Page**

Field
SDG Host
Instance Name
OS Version
Group Name
Reference Host
Real Servers
Health Check Sources
Custom Health Checks
Groups
Virtual Servers
Deployment Plans

Select a policy or a filter and click the **Expand All** icon, and all rules corresponding to that policy or filter are expanded.

Select a policy or filter and click the **Collapse All** icon to collapse all rules.

Enter the term that you want to specify as the filter criterion in the Filter field and click the **Filter** icon to sort and display only the services that are of interest.

## Modifying ADC Service Instances

On the Service Designer page, you can view the collection of service instances defined for several applications, such as stateful firewall or CGNAT.

To modify service instance instances, such as ADC, SFW, CGNAT, or TLB templates:

1. From the View selector, select **Gateway View** or **Service View**. The workspaces that are applicable to this view are displayed. In Gateway view, the devices in the entire network are displayed, organized by the device types and the device models within each device type. In Service View, the different types of services are displayed in the View pane.

2. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.

The functionalities that you can configure in this mode are displayed in the task pane.

3. From the View pane, select the All Network item in Gateway view. Click the plus sign (+) beside the All Network item in the View pane to expand the tree and select the device node you want.

Alternatively, from the View pane, click the plus sign (+) beside All Services to expand the tree and select the type of service.

4. From the task pane, select **Service Edit**. The Service Instances page is displayed.

5. If you are in Gateway view, click the plus sign (+) next to Service Edit to expand the tree in the task pane and view the list of filter templates.

6. In the Service Edit page, from the tree that lists the SDGs, select **All Service Gateways**, or the SDG or SDG pair for which you want to view the previously configured policy or filter templates. This step is applicable only if you selected Gateway View.

The page is divided into two panes. The list of SDGs are displayed on the left pane. You can drill-down to the SDG or pair of SDGs for which you want to process policies or filters. The policy and filter rules are displayed in the right pane.

7. Alternatively, from the View selector, select **Service View**. The workspaces that are applicable to this view are displayed. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.

The functionalities that you can configure in this mode are displayed in the task pane. Click the plus sign in the View pane to expand the All Services tree and select the type of service. From the task pane, select **Manage Service Instances**.

The Service Instances page is displayed in the right pane, listing all the previously defined service instances.

8. From the View pane, perform one of the following tasks:

- Click the **ADC** button.

The list of ADC service instances is displayed. You need not click this button if you are launching the Service Designer page for the first time or are navigating to this page from another mode or a different page. You need to click this button only if you are viewing the other service instances, such as CGNAT or TLB.

- Click the **SFW** button.

The list of SFW templates is displayed.

- Click the **TLB** button.

The list of TLB templates is displayed.

- Click the **CGNAT** button.

The list of CGNAT templates is displayed.

9. Click the **Lock** icon above the table of listed packet filters. The Select Reference Config dialog box is displayed.

10. From the Service Gateway Name drop-down list, select the SDG group to which the packet filter must be applied.

11. From the Host Name drop-down list, select the hostname of the SDG.

12. In the Select Common Components section, select the check boxes beside the service modules or components, such as packet filters, SFW rules, or CGNAT rules, that are displayed. The displayed components depend on the attributes that are previously defined for that selected packet filter. For example, if the service policy is for stateful firewall, SFW rules and SFW rule sets are shown. Select the check box beside Config Category to select all the service components.

13. Click **Save** to save the modified association.

14. Select the check box beside the template you want to modify.

15. Open the **Modify** menu above the list of templates to modify an existing template, and select the component or service attribute, such as application or rule, that you want to edit.
16. Perform one of the following from the drop-down menu displayed for each component:
  - To retrieve the service component and import into the database of Edge Services Director, select **Import Object**. The Import Services dialog box appears. You can import the service instances assigned to SDGs or choose from a list of all of the predefined templates in the database. Also, you can either import all of the components of a service or specific components.
  - To create the component afresh, select **Create New**. The Create page corresponding to the service component appears. You can define the attributes for the service component in the same manner as you define the elements during the creation of a service instance.

## Creating a Deploy Plan and Provisioning Services Immediately

To deploy a deployment plan and policies immediately:

1. From the View selector, select **Gateway View** or **Service View**. The workspaces that are applicable to this view are displayed. In Gateway view, the devices in the entire network are displayed, organized by the device types and the device models within each device type. In Service View, the different types of services are displayed in the View pane.
2. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the View pane, select the All Network item in Gateway view. Click the plus sign (+) beside the All Network item in the View pane to expand the tree and select the device node you want.  
  
Alternatively, from the View pane, click the plus sign (+) beside All Services to expand the tree and select the type of service.
4. From the task pane, select **Service Edit**. The Service Instances page is displayed.
5. If you are in Gateway view, click the plus sign (+) next to Service Edit to expand the tree in the task pane and view the list of filter templates.
6. In the Service Edit page, from the tree that lists the SDGs, select **All Service Gateways**, or the SDG or SDG pair for which you want to view the previously configured policy or filter templates. This step is applicable only if you selected Gateway View.

The page is divided into two panes. The list of SDGs are displayed on the left pane. You can drill-down to the SDG or pair of SDGs for which you want to process policies or filters. The policy and filter rules are displayed in the right pane.

7. In the Service Instances page, select a service instance and click the **Lock** icon.

The corresponding service instance is locked and is available for modifications.

8. Alternatively, in Service View of Deploy mode, from the task pane, select **Service Edit**. The Service Instances page is displayed.

9. Click the plus sign (+) next to Service Instance to expand the tree in the task pane and view the list of filter templates.

10. Select **ADC** to open the Service Edit > ADC page on the right pane.

11. In the Service Instances page, select a service instance and click the **Lock** icon.

The corresponding service instance is locked and is available for modifications.

12. Click the **Send for Deployment** button.

- If you create a deployment plan from Gateway view of Deploy mode, the Deployment Plan Summary dialog box appears, with the service name, type, and status listed.

Click **Send** to create a deployment plan.

- If you create a deployment plan from Service view of Deploy mode, the Edit Service Instance page is displayed. You can modify the SDGs associated with the service instance and also modify the service instance attributes as necessary by either clicking the buttons corresponding to the various settings at the top of the wizard page to directly traverse to the page you want to modify or clicking the navigation buttons at the bottom of the wizard page to go to the different pages of the wizard. Click **Finish** to create a deployment plan.

The configuration deployment job runs. To view the status or results of the deployment job, you can view the Deployment Plans page. In the Deployment Plans page, the Provision Status and Message columns are updated indicating the progress of commission. If the deploy is successful, the status denotes Commissioned. If the deploy fails, the status changes to Commission Failed.

Alternatively, you can select **Discard changes** from the **Actions** menu to ignore the modifications done to a policy or filter template.

## Filtering ADC Service Instances

You can use the enhanced search utility on the Service Edit page for ADC service instances to effectively, quickly identify and segregate the policies and filters of relevance and interest.

The Service Edit page provides advanced search options for the ADC service instances. Enter the term that you want to specify as the filter criterion in the Filter field and click the **Filter** icon.

You can perform advanced searches for the following fields:

- SDG hostname
- Instance name of the service

The following advanced search criteria are available:

- Wildcard search for rule names using an asterisk (\*) is allowed.
- Edge Services Director supports AND and OR operations between search items. The default behavior is OR.
- For rule name search, only the OR operation is allowed, because a policy cannot have multiple rule names.
- For zone search, only the OR operation is allowed. Wildcard search is supported.
- For service and address fields, OR and AND operations are allowed.
- Multiple groups can be grouped using parenthesis. Grouping can be used during filed or keyword searches as well.
- Negate (-) symbol can be used to exclude objects that contain a specific term name.
- The plus (+) operator can be used to specify that the term after the + symbol existing the field value to be filtered along with other searched items.
- Escaping special characters are part of the search syntax. The supported special characters are + - & || ! ( ) { } [ ] ^ " ~ \* ? : \.



**NOTE:** Use the AND operator to find rules that match all values for a given set of fields. Use the OR operator to find rules that match any of the values for a given set of fields.

1. From the View selector, select **Gateway View** or **Service View**. The workspaces that are applicable to this view are displayed. In Gateway view, the devices in the entire network are displayed, organized by the device types and the device models within each device type. In Service View, the different types of services are displayed in the View pane.
2. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.

The functionalities that you can configure in this mode are displayed in the task pane.

3. From the View pane, select the All Network item in Gateway view. Click the plus sign (+) beside the All Network item in the View pane to expand the tree and select the device node you want.

Alternatively, from the View pane, click the plus sign (+) beside All Services to expand the tree and select the type of service.

4. From the task pane, select **Service Edit**. The Service Instances page is displayed.

5. If you are in Gateway view, click the plus sign (+) next to Service Edit to expand the tree in the task pane and view the list of filter templates.

6. In the Service Edit page, from the tree that lists the SDGs, select **All Service Gateways**, or the SDG or SDG pair for which you want to view the previously configured policy or filter templates. This step is applicable only if you selected Gateway View.

The page is divided into two panes. The list of SDGs are displayed on the left pane. You can drill-down to the SDG or pair of SDGs for which you want to process policies or filters. The policy and filter rules are displayed in the right pane.

7. Alternatively, from the View selector, select **Service View**. The workspaces that are applicable to this view are displayed. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.

The functionalities that you can configure in this mode are displayed in the task pane. Click the plus sign in the View pane to expand the All Services tree and select the type of service. From the task pane, select **Manage Service Templates**.

The Service Instances page is displayed in the right pane, listing all the previously defined service instances.

8. From the task pane, select **ADC** to open the ADC page on the right pane.

9. Enter the term that you want to specify as the filter criterion in the Filter field and click the **Filter** icon.

## Managing ADC Service Instance Locks

All the locked policies can be viewed in a single page. You can display the list of SFW, CGNAT, or packet filter templates that are locked by filtering them separately. Such a page shows all the locks only if the user has the unlock task assigned; otherwise, a user sees only the locks that pertain to them.

To view the locked policies:

1. From the View selector, select **Gateway View** or **Service View**. The workspaces that are applicable to this view are displayed. In Gateway view, the devices in the entire network are displayed, organized by the device types and the device models within each device type. In Service View, the different types of services are displayed in the View pane.
2. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.

3. From the View pane, select the All Network item in Gateway view. Click the plus sign (+) beside the All Network item in the View pane to expand the tree and select the device node you want.

Alternatively, from the View pane, click the plus sign (+) beside All Services to expand the tree and select the type of service.

4. From the task pane, select **Service Edit**. The Service Instances page is displayed.
5. If you are in Gateway view, click the plus sign (+) next to Service Edit to expand the tree in the task pane and view the list of filter templates.
6. In the Service Edit page, from the tree that lists the SDGs, select **All Service Gateways**, or the SDG or SDG pair for which you want to view the previously configured policy or filter templates. This step is applicable only if you selected Gateway View.

The page is divided into two panes. The list of SDGs are displayed on the left pane. You can drill-down to the SDG or pair of SDGs for which you want to process policies or filters. The policy and filter rules are displayed in the right pane.

7. Select **ADC** to open the Service Edit > ADC page on the right pane.
8. In the Service Instances page, from the tree that lists the SDGs, select **All Service Gateways**, or the SDG or SDG pair for which you want to lock the filter templates.

9. Select the check box next to the template
10. Click the **Lock** icon, or right-click the policy that you want to lock, and press **Lock**. You can select policies that are locked by you and unlock them. To unlock your policies, you do not need any administrator privileges. To unlock policies locked by other users, you must have the task LOCK assigned to you.

User with administrator privileges can configure the lock settings. To configure the lock settings:

1. Click **Application Switcher**, and go to **Network Application Platform > Administration > Manage Applications**.
2. Right-click the **Edge Services Director** application, and select **Modify Application Settings**. The Modify Edge Services Director Settings page is displayed.
3. Under the Locking option, you can configure the locking timeout value in minutes. The minimum value that you can configure is 2 minutes and the maximum 120 minutes. By default, the timeout value is configured for 5 minutes.
4. By default, the Explicit Policy Lock for edit option is enabled. You can disable this option, if you do not want to lock the policies before editing. When this option is disabled, policies can be edited by any user. The first user gets the preference of saving the changes for a policy. The next save on the same version of a policy results in the user being asked to save policy with new name.



**NOTE:** Acquiring a policy lock or releasing a lock is audit logged. Release locking will show the reason for the release, for example, an explicit release, on save, discard, timeout, or administrator release. Administrator changes of the lock configuration are also audit logged. To see the audit logs, from the Edge Services Director task bar, select Audit Logs.

---

## Unlocking Locked ADC Service Instances

All the locked policies can be viewed in a single page. This page is available for a user with Manage Policy Locks tasks assigned. Such a page shows all the locks only if the user has the unlock task assigned; otherwise, a user sees only the locks that pertain to them.

To view the unlocked policies:

1. From the View selector, select **Gateway View** or **Service View**. The workspaces that are applicable to this view are displayed. In Gateway view, the devices in the entire network are displayed, organized by the device types and the device models within each device type. In Service View, the different types of services are displayed in the View pane.
2. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.

3. From the View pane, select the All Network item in Gateway view. Click the plus sign (+) beside the All Network item in the View pane to expand the tree and select the device node you want.

Alternatively, from the View pane, click the plus sign (+) beside All Services to expand the tree and select the type of service.

4. From the task pane, select **Service Edit**. The Service Instances page is displayed.
5. If you are in Gateway view, click the plus sign (+) next to Service Edit to expand the tree in the task pane and view the list of filter templates.
6. In the Service Edit page, from the tree that lists the SDGs, select **All Service Gateways**, or the SDG or SDG pair for which you want to view the previously configured policy or filter templates. This step is applicable only if you selected Gateway View.

The page is divided into two panes. The list of SDGs are displayed on the left pane. You can drill-down to the SDG or pair of SDGs for which you want to process policies or filters. The policy and filter rules are displayed in the right pane.

7. Select **ADC** to open the Service Edit > ADC page on the right pane.
8. In the Service Instances page, from the tree that lists the SDGs, select **All Service Gateways**, or the SDG or SDG pair for which you want to view the locked filter

templates. Alternatively, click the **Filter** icon to open the list of filter options. From the list, select **All Locked**.

9. Right-click the policy that you want to unlock, and press **Unlock**. You can select policies that are locked by you and unlock them. To unlock your policies, you do not need any administrator privileges. To unlock policies locked by other users, you must have the task LOCK assigned to you.

Alternatively, select the policy you want to unlock, and click the **Manage** button. The Manage Instance Locks dialog box is displayed. The following fields are displayed in the dialog box:

*Table 70: Fields in the Manage Instance Locks Dialog Box*

Field	Description
Instance	Name of the service instance instance.
User	Name of the user that has acquired the lock.
Service Gateway Host	Name of the service gateway with which the instance is attached.
Last Acquired Time	Date and time at which the lock on the template was acquired.

Select the policy instance you want to unlock, and click the **Unlock** icon at the top of the dialog box. Click the **Close** icon to return to the services listing page.

User with administrator privileges can configure the lock settings. To configure the lock settings:

1. Click on **Application Switcher** option, and go to **Network Application Platform > Administration > Manage Applications**.
2. Right click the Edge Services Director application, and select **Modify Application Settings**. The Modify Edge Services Director Settings page is displayed.
3. Under the Locking option, you can configure the locking timeout value in minutes. The minimum value that you can configure is 2 minutes and the maximum is 120 minutes. By default, the timeout value is configured for 5 minutes.
4. By default, the Explicit Policy Lock for edit option is enabled. You can disable this option, if you do not want to lock the policies before editing. When this option is disabled, policies can be edited by any user. The behavior is the same as for concurrent editing. The first user gets the preference of saving the changes for a policy. The next save on the same version of a policy results in the user being asked to save the policy with a new name.



**NOTE:** Acquiring a policy lock or releasing lock is audit logged. Release locking will show the reason for the release, for example, an explicit release, on save, discard, timeout, or administrator release. Administrator changes of the lock configuration are also audit logged. To see the audit logs, from the Edge Services Director task bar, select Audit Logs.

#### Related Documentation

- [Policy and Filter Management Overview on page 375](#)
- [Packet and Service Filters Overview on page 378](#)
- [Searching for CGNAT Policies on page 381](#)
- [Searching for Packet Filters on page 384](#)
- [Searching for SFW Policies on page 386](#)
- [Managing Service and Policy Locks on page 387](#)
- [Unlocking Locked Services and Policies on page 389](#)
- [Viewing Policy and Filter Instances on page 390](#)

## Viewing and Modifying TLB Service Instances

After you create the traffic load balancer (TLB) software service instance to balance user session traffic among a group of available servers that provide shared services using the Service Designer workspace, you can view and modify the components or elements of the service instance by using the Service Edit workspace.

You can perform the following tasks with the Service Edit page for TLB:

- View the list of configured TLB templates.
- Modify an existing TLB template to meet the network needs and deployment scenarios.
- Delete an existing template.
- Transfer the service instance for deployment on a device.
- [Viewing TLB Service Instances on page 472](#)
- [Modifying TLB Service Instances on page 473](#)
- [Creating a Deploy Plan and Provisioning Services Immediately on page 476](#)
- [Filtering TLB Service Instances on page 478](#)
- [Managing TLB Service Instance Locks on page 480](#)
- [Unlocking Locked TLB Service Instances on page 481](#)

## Viewing TLB Service Instances

To view the list of TLB service instances:

1. From the View selector, select **Gateway View** or **Service View**. The workspaces that are applicable to this view are displayed. In Gateway view, the devices in the entire network are displayed, organized by the device types and the device models within each device type. In Service View, the different types of services are displayed in the View pane.
2. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the View pane, select the All Network item in Gateway view. Click the plus sign (+) beside the All Network item in the View pane to expand the tree and select the device node you want.

Alternatively, from the View pane, click the plus sign (+) beside All Services to expand the tree and select the type of service.

4. From the task pane, select **Service Edit**. The Service Instances page is displayed.
5. If you are in Gateway view, click the plus sign (+) next to Service Edit to expand the tree in the task pane and view the list of filter templates.
6. In the Service Edit page, from the tree that lists the SDGs, select **All Service Gateways**, or the SDG or SDG pair for which you want to view the previously configured policy or filter templates. This step is applicable only if you selected Gateway View.

The page is divided into two panes. The list of SDGs are displayed on the left pane. You can drill-down to the SDG or pair of SDGs for which you want to process policies or filters. The policy and filter rules are displayed in the right pane.

7. Alternatively, from the View selector, select **Service View**. The workspaces that are applicable to this view are displayed. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane. Click the plus sign in the View pane to expand the All Services tree and select the type of service. From the task pane, select **Manage Service Templates**.

The Service Instances page is displayed in the right pane, listing all the previously defined service instances.

8. Select **TLB** to open the Service Edit > TLB page on the right pane.
9. In the Service Instances page, from the tree that lists the SDGs, select **All Service Gateways**, or the SDG or SDG pair for which you want to view the previously configured policy or filter templates.

The page is divided into two panes. The list of SDGs are displayed on the left pane. You can drill-down to the SDG or pair of SDGs for which you want to process policies or filters. The policy and filter rules are displayed in the right pane.

The following fields are displayed on the Service Edit > TLB page:

*Table 71: TLB Service Edit Page*

Field
SDG Host
Instance Name
OS Version
Group Name
Reference Host
Real Servers
Network Monitoring
Groups
Virtual Servers
Deployment Plans

Select a policy or a filter and click the **Expand All** icon, and all rules corresponding to that policy or filter are expanded.

Select a policy or filter and click the **Collapse All** icon to collapse all rules.

Enter the term that you want to specify as the filter criterion in the Filter field and click the **Filter** icon to sort and display only the services that are of interest.

## Modifying TLB Service Instances

On the Service Designer page, you can view the collection of service instances defined for several applications, such as stateful firewall or CGNAT.

To modify service instance instances, such as ADC, SFW, CGNAT, or TLB templates:

1. From the View selector, select **Gateway View** or **Service View**. The workspaces that are applicable to this view are displayed. In Gateway view, the devices in the entire network are displayed, organized by the device types and the device models within each device type. In Service View, the different types of services are displayed in the View pane.
2. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the View pane, select the All Network item in Gateway view. Click the plus sign (+) beside the All Network item in the View pane to expand the tree and select the device node you want.  
  
Alternatively, from the View pane, click the plus sign (+) beside All Services to expand the tree and select the type of service.
4. From the task pane, select **Service Edit**. The Service Instances page is displayed.
5. If you are in Gateway view, click the plus sign (+) next to Service Edit to expand the tree in the task pane and view the list of filter templates.
6. In the Service Edit page, from the tree that lists the SDGs, select **All Service Gateways**, or the SDG or SDG pair for which you want to view the previously configured policy or filter templates. This step is applicable only if you selected Gateway View.  
  
The page is divided into two panes. The list of SDGs are displayed on the left pane. You can drill-down to the SDG or pair of SDGs for which you want to process policies or filters. The policy and filter rules are displayed in the right pane.
7. Alternatively, from the View selector, select **Service View**. The workspaces that are applicable to this view are displayed. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane. Click the plus sign in the View pane to expand the All Services tree and select the type of service. From the task pane, select **Manage Service Templates**.  
  
The Service Instances page is displayed in the right pane, listing all the previously defined service instances.
8. In Service View of Deploy mode, from the task pane, select **Deploy Service > Service Edit**.  
  
The Service Instances page is displayed in the right pane, listing all the previously defined service instances.

9. From the View pane, perform one of the following tasks:

- Click the **ADC** button.

The list of ADC service instances is displayed. You need not click this button if you are launching the Service Designer page for the first time or are navigating to this page from another mode or a different page. You need to click this button only if you are viewing the other service instances, such as CGNAT or TLB.

- Click the **SFW** button.

The list of SFW templates is displayed.

- Click the **TLB** button.

The list of TLB templates is displayed.

- Click the **CGNAT** button.

The list of CGNAT templates is displayed.

10. Click the **Lock** icon above the table of listed packet filters. The Select Reference Config dialog box is displayed.

11. From the Service Gateway Name drop-down list, select the SDG group to which the packet filter must be applied.

12. From the Host Name drop-down list, select the hostname of the SDG.

13. In the Select Common Components section, select the check boxes beside the service modules or components, such as packet filters, SFW rules, or CGNAT rules, that are displayed. The displayed components depend on the attributes that are previously defined for that selected packet filter. For example, if the service policy is for stateful firewall, SFW rules and SFW rule sets are shown. Select the check box beside Config Category to select all the service components.

14. Click **Save** to save the modified association.

15. Select the check box beside the template you want to modify.

16. Open the **Modify** menu above the list of templates to modify an existing template, and select the component or service attribute, such as application or rule, that you want to edit.

17. Perform one of the following from the drop-down menu displayed for each component:

- To retrieve the service component and import into the database of Edge Services Director, select **Import Object**. The Import Services dialog box appears. You can import the service instances assigned to SDGs or choose from a list of all of the predefined templates in the database. Also, you can either import all of the components of a service or specific components.

- To create the component afresh, select **Create New**. The Create page corresponding to the service component appears. You can define the attributes for the service component in the same manner as you define the elements during the creation of a service instance.

## Creating a Deploy Plan and Provisioning Services Immediately

To deploy a deployment plan and policies immediately:

1. From the View selector, select **Gateway View** or **Service View**. The workspaces that are applicable to this view are displayed. In Gateway view, the devices in the entire network are displayed, organized by the device types and the device models within each device type. In Service View, the different types of services are displayed in the View pane.

2. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.

The functionalities that you can configure in this mode are displayed in the task pane.

3. From the View pane, select the All Network item in Gateway view. Click the plus sign (+) beside the All Network item in the View pane to expand the tree and select the device node you want.

Alternatively, from the View pane, click the plus sign (+) beside All Services to expand the tree and select the type of service.

4. From the task pane, select **Service Edit**. The Service Instances page is displayed.

5. If you are in Gateway view, click the plus sign (+) next to Service Edit to expand the tree in the task pane and view the list of filter templates.

6. In the Service Edit page, from the tree that lists the SDGs, select **All Service Gateways**, or the SDG or SDG pair for which you want to view the previously configured policy or filter templates. This step is applicable only if you selected Gateway View.

The page is divided into two panes. The list of SDGs are displayed on the left pane. You can drill-down to the SDG or pair of SDGs for which you want to process policies or filters. The policy and filter rules are displayed in the right pane.

7. Select a service template and click the **Lock** icon above the table of listed templates.

8. Alternatively, in Service View of Deploy mode, from the task pane, select **Service Edit**. The Service Instances page is displayed.

9. Click the plus sign (+) next to Service Instance to expand the tree in the task pane and view the list of filter templates.

10. Select **TLB** to open the Service Edit > TLB page on the right pane.
11. In the Service Instances page, select a service instance and click the **Lock** icon.  
The corresponding service instance is locked and is available for modifications.
12. Click the **Send for Deployment** button.
  - If you create a deployment plan from Gateway view of Deploy mode, the Deployment Plan Summary dialog box appears, with the service name, type, and status listed.  
Click **Send** to create a deployment plan.
  - If you create a deployment plan from Service view of Deploy mode, the Edit Service Instance page is displayed. You can modify the SDGs associated with the service instance and also modify the service instance attributes as necessary by either clicking the buttons corresponding to the various settings at the top of the wizard page to directly traverse to the page you want to modify or clicking the navigation buttons at the bottom of the wizard page to go to the different pages of the wizard. Click **Finish** to create a deployment plan.

The configuration deployment job runs. To view the status or results of the deployment job, you can view the Deployment Plans page. In the Deployment Plans page, the Provision Status and Message columns are updated indicating the progress of commission. If the deploy is successful, the status denotes Commissioned. If the deploy fails, the status changes to Commission Failed.

Alternatively, you can select **Discard changes** from the **Actions** menu to ignore the modifications done to a policy or filter template.

## Filtering TLB Service Instances

You can use the enhanced search utility on the Service Edit page for TLB service instances to effectively, quickly identify and segregate the policies and filters of relevance and interest.

The Service Edit page provides advanced search options for the TLB service instances. Enter the term that you want to specify as the filter criterion in the Filter field and click the **Filter** icon.

You can perform advanced searches for the following fields:

- SDG Hostname
- Instance name of the service

The following advanced search criteria are available:

- Wildcard search for rule names using an asterisk (\*) is allowed.
- Edge Services Director supports AND and OR operations between search items. The default behavior is OR.
- For rule name search, only the OR operation is allowed, because a policy cannot have multiple rule names.
- For zone search, only the OR operation is allowed. Wildcard search is supported.
- For service and address fields, OR and AND operations are allowed.
- Multiple groups can be grouped using parenthesis. Grouping can be used during filed or keyword searches as well.
- Negate (-) symbol can be used to exclude objects that contain a specific term name.
- The plus (+) operator can be used to specify that the term after the + symbol existing the field value to be filtered along with other searched items.
- Escaping special characters are part of the search syntax. The supported special characters are + - & || ! ( ) { } [ ] ^ " ~ \* ? : \.



**NOTE:** Use the AND operator to find rules that match all values for a given set of fields. Use the OR operator to find rules that match any of the values for a given set of fields.

---

1. From the View selector, select **Gateway View** or **Service View**. The workspaces that are applicable to this view are displayed. In Gateway view, the devices in the entire network are displayed, organized by the device types and the device models within each device type. In Service View, the different types of services are displayed in the View pane.
2. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.

The functionalities that you can configure in this mode are displayed in the task pane.

3. From the View pane, select the All Network item in Gateway view. Click the plus sign (+) beside the All Network item in the View pane to expand the tree and select the device node you want.

Alternatively, from the View pane, click the plus sign (+) beside All Services to expand the tree and select the type of service.

4. From the task pane, select **Service Edit**. The Service Instances page is displayed.
5. If you are in Gateway view, click the plus sign (+) next to Service Edit to expand the tree in the task pane and view the list of filter templates.
6. In the Service Edit page, from the tree that lists the SDGs, select **All Service Gateways**, or the SDG or SDG pair for which you want to view the previously configured policy or filter templates. This step is applicable only if you selected Gateway View.

The page is divided into two panes. The list of SDGs are displayed on the left pane. You can drill-down to the SDG or pair of SDGs for which you want to process policies or filters. The policy and filter rules are displayed in the right pane.

7. Also, from the View selector, you can select select **Service View**. The workspaces that are applicable to this view are displayed. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane. Click the plus sign in the View pane to expand the All Services tree and select the type of service. From the task pane, select **Manage Service Templates**.

The Service Instances page is displayed in the right pane, listing all the previously defined service instances.

8. Alternatively, from the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.

9. Select **Service Edit** from the task pane. The Service Edit page is displayed.
10. Click the plus sign (+) next to the policy and filter template to expand the tree in the task pane and view the list of filter templates.
11. From the task pane, select **TLB** to open the TLB page on the right pane.
12. Enter the term that you want to specify as the filter criterion in the Filter field and click the **Filter** icon.

## Managing TLB Service Instance Locks

All the locked policies can be viewed in a single page. You can display the list of SFW, CGNAT, or packet filter templates that are locked by filtering them separately. Such a page shows all the locks only if the user has the unlock task assigned; otherwise, a user sees only the locks that pertain to them.

To view the locked policies:

1. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
2. From the task pane, select **Service Edit**. The Service Instances page is displayed.
3. Click the plus sign (+) next to Service Instance to expand the tree in the View pane and view the list of filter templates.
4. Select **TLB** to open the Service Edit > TLB page on the right pane.
5. In the Service Instances page, from the tree that lists the SDGs, select **All Service Gateways**, or the SDG or SDG pair for which you want to lock the filter templates.
6. Select the check box next to the template
7. Click the **Lock** icon, or right-click the policy that you want to lock, and press **Lock**. You can select policies that are locked by you and unlock them. To unlock your policies, you do not need any administrator privileges. To unlock policies locked by other users, you must have the task LOCK assigned to you.

User with administrator privileges can configure the lock settings. To configure the lock settings:

1. Click **Application Switcher**, and go to **Network Application Platform > Administration > Manage Applications**.
2. Right-click the **Edge Services Director** application, and select **Modify Application Settings**. The Modify Edge Services Director Settings page is displayed.
3. Under the Locking option, you can configure the locking timeout value in minutes. The minimum value that you can configure is 2 minutes and the maximum 120 minutes. By default, the timeout value is configured for 5 minutes.
4. By default, the Explicit Policy Lock for edit option is enabled. You can disable this option, if you do not want to lock the policies before editing. When this option is disabled, policies can be edited by any user. The first user gets the preference of saving

the changes for a policy. The next save on the same version of a policy results in the user being asked to save policy with new name.



**NOTE:** Acquiring a policy lock or releasing a lock is audit logged. Release locking will show the reason for the release, for example, an explicit release, on save, discard, timeout, or administrator release. Administrator changes of the lock configuration are also audit logged. To see the audit logs, from the Edge Services Director task bar, select Audit Logs.

## Unlocking Locked TLB Service Instances

All the locked policies can be viewed in a single page. This page is available for a user with Manage Policy Locks tasks assigned. Such a page shows all the locks only if the user has the unlock task assigned; otherwise, a user sees only the locks that pertain to them.

To view the unlocked policies:

1. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
2. From the task pane, select **Service Edit**. The Service Instances page is displayed.
3. Click the plus sign (+) next to Service Instance to expand the tree in the View pane and view the list of filter templates.
4. Select **TLB** to open the Service Edit > TLB page on the right pane.
5. In the Service Instances page, from the tree that lists the SDGs, select **All Service Gateways**, or the SDG or SDG pair for which you want to view the locked filter templates. Alternatively, click the **Filter** icon to open the list of filter options. From the list, select **All Locked**.
6. Right-click the policy that you want to unlock, and press **Unlock**. You can select policies that are locked by you and unlock them. To unlock your policies, you do not need any administrator privileges. To unlock policies locked by other users, you must have the task LOCK assigned to you.

Alternatively, select the policy you want to unlock, and click the **Manage** button. The Manage Instance Locks dialog box is displayed. The following fields are displayed in the dialog box:

*Table 72: Fields in the Manage Instance Locks Dialog Box*

Field	Description
Instance	Name of the service instance instance.
User	Name of the user that has acquired the lock.
Service Gateway Host	Name of the service gateway with which the instance is attached.
Last Acquired Time	Date and time at which the lock on the template was acquired.

Select the policy instance you want to unlock, and click the **Unlock** icon at the top of the dialog box. Click the **Close** icon to return to the services listing page.

User with administrator privileges can configure the lock settings. To configure the lock settings:

1. Click on **Application Switcher** option, and go to **Network Application Platform > Administration > Manage Applications**.
2. Right click the Edge Services Director application, and select **Modify Application Settings**. The Modify Edge Services Director Settings page is displayed.
3. Under the Locking option, you can configure the locking timeout value in minutes. The minimum value that you can configure is 2 minutes and the maximum is 120 minutes. By default, the timeout value is configured for 5 minutes.
4. By default, the Explicit Policy Lock for edit option is enabled. You can disable this option, if you do not want to lock the policies before editing. When this option is disabled, policies can be edited by any user. The behavior is the same as for concurrent editing. The first user gets the preference of saving the changes for a policy. The next save on the same version of a policy results in the user being asked to save the policy with a new name.



**NOTE:** Acquiring a policy lock or releasing lock is audit logged. Release locking will show the reason for the release, for example, an explicit release, on save, discard, timeout, or administrator release. Administrator changes of the lock configuration are also audit logged. To see the audit logs, from the Edge Services Director task bar, select Audit Logs.

#### Related Documentation

- [Policy and Filter Management Overview on page 375](#)
- [Packet and Service Filters Overview on page 378](#)
- [Searching for CGNAT Policies on page 381](#)

- [Searching for Packet Filters on page 384](#)
- [Searching for SFW Policies on page 386](#)
- [Managing Service and Policy Locks on page 387](#)
- [Unlocking Locked Services and Policies on page 389](#)
- [Viewing Policy and Filter Instances on page 390](#)

## Using the Actions Menu on the Service Policy and Packet Filter Pages

You can use the Actions menu on the Service Policy and Packet Filter Policy page for CGNAT, SFW, and packet filters to create a deployment plan for the packet filter policies or discard the modifications made to a packet filter policy.

- [Creating a Deployment Plan on page 483](#)
- [Discarding Changes Made to a Service Policy or Packet Filter Policy on page 484](#)

### Creating a Deployment Plan

You must have previously defined service templates and policy or filter templates before you can create a deployment plan.

To create a deployment plan and assigning devices to it:

1. From the View selector, select **Gateway View**. The View pane displays the devices in the entire network organized by the device type and device models pertaining to each device type.
2. From the View pane, select the All Network item. Expand the tree to select the SDG in an SDG group.
3. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
4. Select **Service Edit** from the task pane. The Service Templates page is displayed.
5. Click the right arrow next to Service Edit to expand the tree in the task pane and view the list of filter instances.
6. From the task pane, perform one of the following:
  - Select **CGNAT** to open the CGNAT and Filter page on the right pane.
  - Select **Packet Filter** to open the Packet Filter page on the right pane.
  - Select **SFW** to open the SFW Policy and Filter page on the right pane.

7. In the Service Edit page, from the tree that lists the SDGs, select **All Service Gateways**, or the SDG or SDG pair for which you want to view the previously configured policy or filter instances. This step is applicable only if you selected Gateway View. You can drill-down to the SDG or pair of SDGs for which you want to process policies or filters.
8. Select a rule corresponding to an SDG, and click the **Lock** icon above the table of listed policy filters.
9. Click the down arrow in the **Actions** menu and select **Send for Deployment** to create a deployment plan for the particular service template and save the plan.  
  
The Deployment Plan Summary dialog box appears, with the service name, type, and status listed.  
  
Click **Send** to create a deployment plan.  
  
A deploy plan is created for the service template with the devices that are assigned to it when you view the Deployment Plans page.
10. From the Deployment plans page, you can select **Reject** or **Approve** from the Actions drop-down list to reject or approve the deployment plan and make it available for commissioning to the devices.

## Discarding Changes Made to a Service Policy or Packet Filter Policy

You can discard the changes made to previously defined service templates and policy or filter templates before you can create a deployment plan.

To ignore the modifications made to a service or packet filter policy:

1. From the View selector, select **Gateway View**. The View pane displays the devices in the entire network organized by the device type and device models pertaining to each device type.
2. From the View pane, select the All Network item. Expand the tree to select the SDG in an SDG group.
3. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
4. Select **Service Edit** from the task pane. The Service Templates page is displayed.
5. Click the right arrow next to Service Edit to expand the tree in the task pane and view the list of filter instances.
6. From the task pane, perform one of the following:
  - Select **CGNAT** to open the CGNAT and Filter page on the right pane.

- Select **Packet Filter** to open the Packet Filter page on the right pane.
  - Select **SFW** to open the SFW Policy and Filter page on the right pane.
7. In the Service Edit page, from the tree that lists the SDGs, select **All Service Gateways**, or the SDG or SDG pair for which you want to view the previously configured policy or filter instances. This step is applicable only if you selected Gateway View. You can drill-down to the SDG or pair of SDGs for which you want to process policies or filters.
  8. Select a rule corresponding to an SDG, and click the **Lock** icon above the table of listed policy filters.
  9. Select **Discard changes** from the **Actions** menu to ignore the modifications done to a policy or filter template.

**Related  
Documentation**

- [Service Templates Overview on page 189](#)
- [Filtering Service Templates on page 189](#)
- [Viewing Service Templates on page 192](#)

## Tagging Junos Space Network Management Platform Objects

You can create user-defined tags on an application workspace inventory page to easily categorize and organize managed objects. Subsequently, you can view and use these tags to easily search for multiple objects to view the status or perform a bulk action on them without having to select each object individually.

The tags are classified into two categories: private tags and public tags. Private tags are those that are created by you and can be used only by you because they are not visible to others. Public tags are those that are available to all users for tagging objects that are accessible to them. You need the Tag Administrator role privileges to create, modify, or delete a public tag, manage hierarchical tags, as well as convert a private tag to a public tag. However, any Junos Space user can:

- Create, modify, and delete private tags
- View public and private tags
- Tag and untag objects by using public and private tags



**NOTE:** You cannot view or access private tags created by other users. However, if you are a user with the Tag Administrator role, you can view and access private tags of other users.

Tag names should not start with a space, cannot contain a comma, double quotation marks, and parentheses, and cannot exceed 255 characters. Also, you cannot name a tag “Untagged” because it is a reserved term.

- [Creating a Tag on page 486](#)
- [Tagging an Object on page 489](#)
- [Untagging an Object on page 490](#)

## Creating a Tag

You create tags when you want to label and categorize Junos Space Network Management Platform objects so that you can filter, monitor, or perform batch actions on them without having to select each object individually. All users can create their own private tags from the Administration > Tags inventory landing page. However, users assigned the Tag Administrator role can create public tags.

You can create tags from the Administration workspace as well as from the Device Management or Job Management inventory landing page. By default, the tags that any user creates are private tags, which means that these tags are visible only to the user who creates them. No other user can access the private tags created by other users. However, if you are a user with the Tag Administrator role, you can make these tags public, thereby allowing all users to associate objects with these tags.

To create a tag:

1. On the Junos Space Network Management Platform user interface, select **Administration > Tags**.

The Tags page appears.

2. On the toolbar, click the **Create Tag** icon.

The **Create Tag** dialog box appears.

3. If necessary, select the **Share this Tag** check box.

When you share a tag, all users can use that tag. Only users with the Tag Administrator role can publish tags to the public domain. For users without this role, the **Share this Tag** check box is disabled (grayed out).

4. In the **Tag Name** field, type a tag name.

A tag name should not:

- Exceed 255 characters
- Start with a space
- Contain special characters, such as commas, double quotation marks, and parentheses.



**NOTE:** “Untagged” is a reserved term and hence you cannot create a tag with this name.

5. Click **Create**.

The Create Tag dialog box appears, displaying that the tag is successfully created.

6. Click **OK** on the Create Tag dialog box.

The newly added tag appears on the Tags page. If the tag is shared, it is public; if not, it is private. The **Access Type** column displays whether the tag is public or private.

In addition to creating tags from the Administration workspace, you can create tags from the following inventory landing pages as well:

- Device Management
- Job Management

For example, to create a tag from the Device Management inventory landing page:

1. On the Junos Space Network Management Platform user interface, click **Devices** > **Device Management**.

The Device Management page appears.

2. If the tags are not displayed, click the **Display Tag View** icon on the toolbar located at the top of this page.

On the left side of the page, tags that are relevant to the page and the domain to which you are logged in are displayed.



**NOTE:** Tags from domains other than the domain to which the user is logged in are not displayed.

In Tags View, the tags are categorized as follows:

- **Public**—Lists public tags. Public tags are tags that are visible and available to all users and can be used by any user to tag an object in Junos Space.

You can perform the following actions on public tags:

- Mouse over a tag to view the number of objects that are associated with the specific tag.
- Click a tag to view the devices associated with the selected tag. The number displayed adjacent to the tag shows the number of devices associated with the specific tag. For example, if you have assigned this tag to two devices, then the number displayed is 2. However, this rule has the following exceptions:

- For hierarchical tags, the count on the parent tag does not include the number of objects associated with its child tags. For example, if a child tag is associated with 10 objects and its parent tag is associated with five objects, then the count displayed for the parent tag is 5 and not 15.
- You used the same tag on objects other than devices. For example, if you assigned TagC to UserA and DeviceB, then on the Device Management page, the count shown for TagC is 1. However, when you mouse over TagC, the tooltip displays a count of 2 (which includes the object type as well—in this example, the object types that are displayed are **User** and **Device**).

- **Private**—Lists private tags. Private tags are tags that you created and hence are visible only to you. No other user has access to these tags.

Click a tag to view the devices associated with the selected tag. The number displayed adjacent to the tag shows the number of devices that are associated with the specific tag. For example, if you assigned this tag to two devices, then the number displayed is 2.

- **Untagged**—Displays the number of devices that are not tagged

3. (Optional) To view all tags (that is, tags that are relevant and irrelevant to the inventory landing page to which you are currently logged in), select **Show All Tags** on the **Tags** list at the top of the Device Management inventory landing page.

By default, **Show Relevant Tags** is selected and only the tags that are relevant to the current inventory landing page are displayed.

4. To add a tag:

- a. Click the **Add Tag** icon.



**NOTE:** If you use the shortcut menu instead of the Add Tag icon, the new tag that is added is of the same type as that of the parent. For example, right-click **Private** and select **Add Tag** to create a private tag.

- b. In the **Tag Name** field, type a tag name.

A tag name should not:

- Exceed 255 characters
- Start with a space
- Contain special characters such as commas, double quotation marks, and parentheses



**NOTE:** “Untagged” is a reserved term and hence you cannot create a tag with this name.

- c. If necessary, select the **Make Public** check box to create a public tag. If left unselected, a private tag is created.

When you make a tag public, all users can use that tag. Only the Tag Administrator can publish tags to the public domain.



**NOTE:** This check box is disabled if you chose to create a tag by using the shortcut menu. The new tag that is added is of the same type as that of the parent.

- d. (Optional) In the **Description** field, add a description of the tag.
- e. Click **Add Tag**.

The tag is added to the relevant tag category and assigned to the domain to which you are currently logged in. For example, if you created a public tag, the newly added tag is placed in the **Public** category. The count is set to zero (0) because you have not assigned this tag to any object.



**NOTE:** You cannot add any tags to the **Untagged** category.

When you add a tag, an audit log entry is automatically generated.

## Tagging an Object

You can create user-defined tags on an application workspace inventory page to easily categorize and organize managed objects. Subsequently, you can view and use these tags to easily search for multiple objects to view the status or perform a bulk action on them without having to select each object individually.

By default, the tags that you create from any workspace are private tags and these private tags are visible only to you. If you want any other user to use the tag that you created, then you have to create a public tag instead of a private tag or convert the private tag to a public tag.

To tag an object:

1. Select the inventory objects that you want to tag.
2. Select **Tag It** from the Actions menu.  
The **Apply Tag** dialog box appears.
3. Select or type the tag name in the field.

If you have existing tags, start to type a tag name in the name field. Existing tags appear in the selection box.

You can also type a new tag name in the field. The new tag is automatically created and applied to the selected objects.

4. (Optional) Select the **Make Public** check box to mark the new tag created in the previous step as a public tag. If you do not select this check box, the new tag added is classified as a private tag.



**NOTE:** If you do not have permissions to create a public tag, then the **Make Public** check box is disabled.

5. (Optional) Add a comment in the **Add Description here** field.
6. Click **Apply Tag**. This action tags the object and stores the tag in the database.

## Untagging an Object

You can untag or remove a tag from an object on a workspace inventory page. You can select only one object at a time to untag.

To untag an object:

1. Navigate to the Service Templates page for CGNAT, packet filter, or SFW policies.
2. Select one object on the workspace inventory page at a time.
3. Select **UnTag It** from the Actions menu or right-click an object and select **UnTag It** from the shortcut menu.

The **UnTag The Object** dialog box appears.

4. Select the tags that you want to remove.
5. Click **Untag**.

The Untag dialog box appears, displaying that the object has been successfully untagged.

6. Click **OK**.

In this example, you are returned to the Device Management workspace.

- Related Documentation**
- [Policy and Filter Management Overview on page 375](#)
  - [Packet and Service Filters Overview on page 378](#)

- [Searching for CGNAT Policies on page 381](#)
- [Searching for Packet Filters on page 384](#)
- [Searching for SFW Policies on page 386](#)
- [Managing Service and Policy Locks on page 387](#)
- [Unlocking Locked Services and Policies on page 389](#)
- [Viewing Policy and Filter Instances on page 390](#)



# Managing Packet Analyzers

- [Packet Analyzer Overview on page 493](#)
- [Creating and Viewing Service Analyzers on page 495](#)

## Packet Analyzer Overview

---

Packet capture is a tool that helps you to analyze network traffic and troubleshoot network problems. The packet capture tool captures real-time data packets traveling over the network for monitoring and logging. This tool is a debugging and analysis utility that you can use to identify the problematic area in a session path. A set of counters are displayed for both forward and reverse flow for all the supported services on SDG devices. Using these statistical details and values, you can obtain adequate and useful estimates regarding the total bytes count for each service in every hop and quickly, easily locate the hop where there can be a possible packet drop.

The packet analyzer is the endpoint to which the flow collector interface sends traffic for analysis. You can process and export multiple cflowd records with a flow collector interface. You create a flow collector interface on a Monitoring Services II or Multiservices 400 PIC. The flow collector interface combines multiple cflowd records into a compressed ASCII data file and exports the file to an FTP server.

You can configure the packet analyzer filters to capture packet data flows based on a match or classification criteria to collect statistics and information only about packets that satisfy the criteria. You can define the data and control plane packet flow direction and interface settings in the filter, and the interval at which devices must be polled. You can also specify a timeout to apply a threshold on the amount of data to be collected. You can then schedule the filter to be run for different services and view the statistics as numerical values or as a graph.

Packets are captured as binary data, without modification. You can read the packet information offline with a packet analyzer such as Ethereal or tcpdump. If you need to quickly capture packets destined for, or originating from, the Routing Engine and analyze them online, you can use the packet capture diagnostic tool.

Network administrators and security engineers use packet capture to perform the following tasks:

- Monitor network traffic and analyze traffic patterns.

- Identify and troubleshoot network problems. Detect security breaches in the network, such as unauthorized intrusions, spyware activity, or ping scans.
- Packet capture operates like traffic sampling on the device, except that it captures entire packets.

Data packets are chunks of data that transit the router as they are being forwarded from a source to a destination. When a router receives a data packet on an interface, it determines where to forward the packet by looking in the forwarding table for the best route to a destination. The router then forwards the data packet toward its destination through the appropriate interface. The Packet Forwarding Engine, which is the central processing element of the router's forwarding plane, handles the flow of data packets in and out of the router's physical interfaces. Although the Packet Forwarding Engine contains Layer 3 and Layer 4 header information, it does not contain the packet data itself (the packet's payload).

You can also use the packet capture feature when you need to quickly capture and analyze control traffic on a router. Control packets refer to health check packets that are sent to examine the health and efficiency of specific URLs or paths. Health checking allows you to verify content accessibility in large websites. As content grows and information is distributed across different server farms, flexible, customizable content health checks are critical to ensure end-to-end availability.

## Pre-Service Filtering of Traffic for Service Processing

To filter IPv4 or IPv6 traffic before accepting packets for input or output service processing, include the **service-set** *service-set-name* **service-filter** *service-filter-name* at one of the following interfaces:

- **[edit interfaces *interface-name* unit *unit-number* family (inet | inet6) service input]**
- **[edit interfaces *interface-name* unit *unit-number* family (inet | inet6) service output]**

For the **service-set-name**, specify a service set configured at the **[edit services service-set]** hierarchy level.

The service set retains the input interface information even after services are applied, so that functions such as filter-class forwarding and destination class usage (DCU) that depend on input interface information continue to work.

The following requirements apply to filtering inbound or outbound traffic before accepting packets for service processing:

- You configure the same service set on the input and output sides of the interface.
- If you include the **service-set** statement without an optional **service-filter** definition, the Junos OS assumes the match condition is true and selects the service set for processing automatically.
- The service filter is applied only if a service set is configured and selected.

You can include more than one service set definition on each side of an interface. The following guidelines apply:

- If you include multiple service sets, the router (or switch) software evaluates them in the order in which they appear in the configuration. The system executes the first service set for which it finds a match in the service filter and ignores the subsequent definitions.
- A maximum of six service sets can be applied to an interface.
- When you apply multiple service sets to an interface, you must also configure and apply a service filter to the interface.

## Postservice Filtering of Returning Service Traffic

As an option to filtering of IPv4 or IPv6 input service traffic, you can apply a service filter to IPv4 or IPv6 traffic that is returning to the services interface after the service set is executed. To apply a service filter in this manner, include the **post-service-filter service-filter-name** statement at the **[edit interfaces interface-name unit unit-number family (inet | inet6) service input]** hierarchy level.

### Related Documentation

- [Creating and Viewing Service Analyzers on page 323](#)

---

## Creating and Viewing Service Analyzers

The packet analyzer is the endpoint to which the flow collector interface sends traffic for analysis. You can process and export multiple cflowd records with a flow collector interface. You can perform the following tasks with the Service Analyzer page:

- Configure and provision filters for packet analysis.
- Configure filters for CGNAT, ADC, and TLB services.
- Start and stop the configured filters.
- View the packet analyzer details as a statistical form or a graphical form.
- [Configuring the Traffic Analyzer Filter on page 495](#)
- [Managing Service Analyzer Filter Instances on page 498](#)
- [Viewing Service Analyzer Instance Details on page 500](#)
- [Viewing the Service Analyzer Statistics in Grid Format and Graph on page 502](#)

## Configuring the Traffic Analyzer Filter

To configure the traffic analyzer filter details on packet flows for the different services and to schedule its running:

1. From the View selector, select **Gateway View**. The workspaces that are available in this view are displayed. The Gateway view displays the service delivery gateway (SDG) groups and the SDGs that are part of the high availability pair in an SDG group.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the View pane, select the All Network item in Gateway view. Click the plus sign (+) beside the All Network item in the View pane to expand the tree and select the device node you want.
4. From the task pane, do one of the following:
  - Select **Service Analyzer > ADC Filter** from the task pane. The Service Analyzer for ADC Filter page is displayed. The service instance or template that you previously configured for the ADC service type are displayed. All the instances created in the Service Templates workspace are shown.
  - Select **Service Analyzer > TLB Filter** from the task pane. The Service Analyzer for TLB Filter page is displayed. The service instance or template that you previously configured for the TLB service type are displayed. All the instances created in the Service Templates workspace are shown.
  - Select **Service Analyzer > CGNAT Filter** from the task pane. The Service Analyzer for CGNAT Filter page is displayed. The service instance or template that you previously configured for the CGNAT service type are displayed. All the instances created in the Service Templates workspace are shown.

The list of SDGs or SDG pairs in a high availability group are displayed, along with the filter instances that were configured for the different services. The number of filter instances that are currently in progress and the number of filter instances that are scheduled or planned to be run at a later time are also displayed. For information on running or clearing filter instances, see *Managing Service Analyzer Filter Instances*.

5. Select the SDGs or SDG pairs (you can select multiple rows to create and assign filters to several SDGs simultaneously) for which you want to create packet analyzer filters for services.
6. Click the plus sign (+) above the table of listed SDGs to create a new filter. The Update Service Analyzer Filter Details page is displayed.
7. In the Data Forward Flow section, do the following. A forward flow refers to packets that are sent in the forward or upward direction. A reverse flow refers to packets that are sent in the returning or backward direction.
  - From the **Egress** list, select the egress interface on which the data packets that are sent out in the forward flow must be monitored. Click **Details** beside the list to view interface details.

- From the **Ingress** list, select the input interface on which the data packets that are received in the forward flow must be monitored. Click **Details** beside the list to view interface details.
8. In the Data Reverse Flow section, do the following.
    - From the **Egress** list, select the egress interface on which the data packets that are sent out in the reverse flow must be monitored. Click **Details** beside the list to view interface details.
    - From the **Ingress** list, select the input interface on which the data packets that are received in the reverse flow must be monitored. Click **Details** beside the list to view interface details.
  9. In the Control Forward Flow section, do the following. A forward flow refers to packets that are sent in the forward or upward direction. A reverse flow refers to packets that are sent in the returning or backward direction.
    - From the **Egress** list, select the egress interface on which the control packets that are sent out in the forward flow must be monitored. Click **Details** beside the list to view interface details.
    - From the **Ingress** list, select the input interface on which the control packets that are received in the forward flow must be monitored. Click **Details** beside the list to view interface details.
  10. In the Data Reverse Flow section, do the following.
    - From the **Egress** list, select the egress interface on which the control packets that are sent out in the reverse flow must be monitored. Click **Details** beside the list to view interface details.
    - From the **Ingress** list, select the input interface on which the control packets that are received in the reverse flow must be monitored. Click **Details** beside the list to view interface details.
  11. Click **Apply** to save the filter settings. Otherwise, click **Cancel** to discard the changes. You are returned to the Service Analyzer page.
  12. If you created a new filter, the filter instance is displayed under the corresponding service type section, such as CGNAT or ADC. Such filters are provisioned filter instances. This display signifies that the filter is configured, but it needs to be scheduled to be run. Click the link that shows the number of instances under the column of the relevant service type. The Service Analyzer Instances page is shown.
  13. On this page, the names of the service instances for which filters are defined. The actions you can perform are in the form of the Clear and Run buttons, above the table of listed service instances, for each service instance with a filter.
  14. Select the check box next to a service analyzer filter and click the **Delete** button to remove a configured filter for an instance. You are prompted to confirm the deletion. If you click **OK**, a popup dialog box denotes the successful deletion.

15. Select the check box next to a service analyzer filter instance, and click the **Run** button to schedule the filter to be run. The Run Filter dialog box appears. The Run button is grayed out if the particular service filter instance is already in progress.
16. From the **Poll Interval** list, select the interval in minutes at which the data must be polled and collected. Values from 1 minute up to 59 minutes are shown in increments of 2 minutes in the list.
17. In the **Schedule Start Details** section, click **Run Now** to start the filter immediately. Alternatively, click the **Run At** radio button and select the date and time at which the filter must be run.
18. In the **Schedule End Details** section, do one of the following:
  - Click the **Stop At** radio button and select the date and time at which the filter must be stopped.
  - Click the **Stop After** radio button and specify a value for the number of polls after which the filter must be ended.
  - Click the **Run Until Stopped** radio button to continue running the test until you manually want to stop it.
19. Click **Run** to save the filter settings. Otherwise, click **Cancel** to discard the changes. You are returned to the Prepared Service Analyzer Instances dialog box. Click **Close** to return to the Service Analyzer Page.

## Managing Service Analyzer Filter Instances

To view, start, stop, or clear the configured analyzer filters:

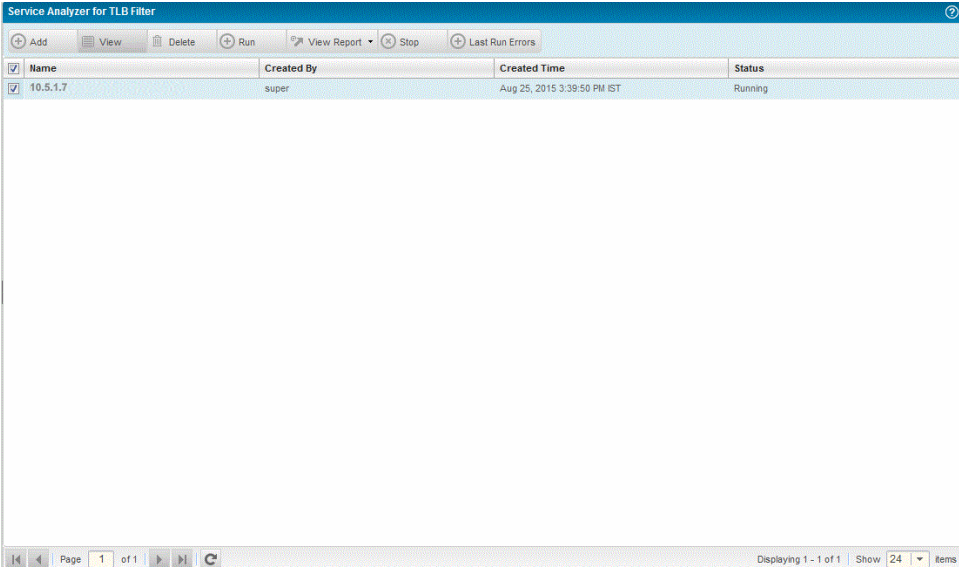
1. From the View selector, select **Gateway View**. The workspaces that are available in this view are displayed. The Gateway view displays the service delivery gateway (SDG) groups and the SDGs that are part of the high availability pair in an SDG group.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the View pane, select the All Network item in Gateway view. Click the plus sign (+) beside the All Network item in the View pane to expand the tree and select the device node you want.
4. From the task pane, do one of the following:
  - Select **Service Analyzer > ADC Filter** from the task pane. The Service Analyzer for ADC Filter page is displayed. The service instance or template that you previously configured for the ADC service type are displayed. All the instances created in the Service Templates workspace are shown.

- Select **Service Analyzer > TLB Filter** from the task pane. The Service Analyzer for TLB Filter page is displayed. The service instance or template that you previously configured for the TLB service type are displayed. All the instances created in the Service Templates workspace are shown.
- Select **Service Analyzer > CGNAT Filter** from the task pane. The Service Analyzer for CGNAT Filter page is displayed. The service instance or template that you previously configured for the CGNAT service type are displayed. All the instances created in the Service Templates workspace are shown.

The list of SDGs or SDG pairs in a high availability group are displayed, along with the filter instances that were configured for the different services. The number of filter instances that are currently in progress and the number of filter instances that are scheduled or planned to be run at a later time are also displayed. For information on viewing filter instances, see *Viewing the Traffic Analyzer Statistics and Graph*.

5. For the SDG corresponding to a certain service, all of the previously configured service analyzer filters are displayed in the Service Analyzer Instances page with the state of the filter instance under the Status column of the relevant service type. View the Status column for the current state of the filter.

**Figure 47: Service Analyzer Instances Page**



The screenshot shows the 'Service Analyzer for TLB Filter' page. At the top, there is a toolbar with buttons: Add, View, Delete, Run, View Report, Stop, and Last Run Errors. Below the toolbar is a table with the following columns: Name, Created By, Created Time, and Status. There is one row in the table with the following data: Name: 10.5.1.7, Created By: super, Created Time: Aug 25, 2015 3:39:50 PM IST, Status: Running. At the bottom of the page, there is a pagination bar showing 'Page 1 of 1', 'Displaying 1 - 1 of 1', and 'Show 24 items'.

Name	Created By	Created Time	Status
10.5.1.7	super	Aug 25, 2015 3:39:50 PM IST	Running

You can click the links under one of the following columns:

- **View**—Click to display the traffic analyzer details on packet flows for the different services configured. For information on viewing filter instances, see *Viewing the Service Analyzer Statistics and Graph*.
- **Delete**—Click to remove the configured filter for an instance. You are prompted to confirm the deletion. If you click **OK**, a popup dialog box denotes the successful deletion.

- **Run**—Click to schedule a filter to be run. For information on scheduling a filter instance to be run, see *Configuring the Traffic Analyzer Filter*.
  - **Report**—Click to view the collection statistics and information about packets that are fetched. For information on viewing the collected details by a service analyzer, see *Viewing the Service Analyzer Collection Data*.
  - **Stop**—Click to end a running filter. You are prompted to confirm whether you want to stop the filter instance. If you click **OK**, a popup dialog box denotes the successful termination of the filter instance.
  - **Last Run Errors**—Click to display any errors that occurred during the running of the filter instance. The Last Run Status dialog box is displayed. It contains the Provisioning Errors and Decommissioning Errors tabs that describe errors that might have occurred during the initialization and start of the analyzer filters or with the decommissioning and termination. The following fields are displayed in this dialog box for both the tabs:
    - **Host Name**—Host name of the SDG device.
    - **Severity**—System logging severity level.
    - **Path**—Hierarchy level of the configuration statement corresponding to the setting in the CLI interface Info Informational message about the error that is generated.
    - **Message**—System event logging message generated that describes the error.
  - **Graph**—Click to display the packet analyzer details for monitoring as a pictorial form. The Packet Flow Graph dialog box appears.
6. In the dialog box, the Configured Instances column displays the names of the service instances for which filters are defined. The Actions column contains the Clear and Run subcolumns for each service instance with a filter.
  7. Click **Delete** to remove a configured filter for an instance. You are prompted to confirm the deletion. If you click **OK**, a popup dialog box denotes the successful deletion.
  8. Click **Run** beside the instance you want to schedule the filter to be run. The Run Filter dialog box appears to specify the schedule settings.

## Viewing Service Analyzer Instance Details

To view the service analyzer instance details:

1. From the View selector, select **Gateway View**. The workspaces that are available in this view are displayed. The Gateway view displays the service delivery gateway (SDG) groups and the SDGs that are part of the high availability pair in an SDG group.
2. From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.

3. From the View pane, select the All Network item in Gateway view. Click the plus sign (+) beside the All Network item in the View pane to expand the tree and select the device node you want.
4. From the task pane, do one of the following:
  - Select **Service Analyzer > ADC Filter** from the task pane. The Service Analyzer for ADC Filter page is displayed. The service instance or template that you previously configured for the ADC service type are displayed. All the instances created in the Service Templates workspace are shown.
  - Select **Service Analyzer > TLB Filter** from the task pane. The Service Analyzer for TLB Filter page is displayed. The service instance or template that you previously configured for the TLB service type are displayed. All the instances created in the Service Templates workspace are shown.
  - Select **Service Analyzer > CGNAT Filter** from the task pane. The Service Analyzer for CGNAT Filter page is displayed. The service instance or template that you previously configured for the CGNAT service type are displayed. All the instances created in the Service Templates workspace are shown.

The list of SDGs or SDG pairs in a high availability group are displayed, along with the filter instances that were configured for the different services. The number of filter instances that are currently in progress and the number of filter instances that are scheduled or planned to be run at a later time are also displayed. For information on running or clearing filter instances, see *Managing Service Analyzer Filter Instances*.

5. Select the SDGs or SDG pairs (you can select multiple rows to create and assign filters to several SDGs simultaneously) for which you want to create packet analyzer filters for services.
6. From the Service Analyzer page, for the SDG corresponding to a certain service, click the link under the column of the relevant service type. The Prepared Service Analyzer Instances dialog box is shown. Click **View** under the View column to view the traffic analyzer for the particular service.

The View Service Instance Analyzer Details page is displayed.

The following fields are displayed in this page:

Field	Description
Name	Name of the SDG or pair of SDGs in a high availability group.
Type	Service type for which packets collected are shown. Values are CGNAT, ADC, or TLB.
Data Packets/Control Packets	Click the <b>Data Packets</b> tab to view data packet details for the service analyzer filter. Alternatively, click the <b>Control Packets</b> tab to view control packet details for the service analyzer filter. Indicates whether data or control packet details are shown.
Forward Flow	Displays statistics for packets in forward flow direction.

Field	Description
Ingress	Number of packets that arrive in the ingress direction in forward flow.
Egress	Number of packets that are sent out in the egress direction in forward flow.
Reverse Flow	Displays statistics for packets in reverse flow direction. If a service set is a sampling service set and the reverse-flow service order is not configured, all sampled traffic is considered to be forward traffic.
Ingress	Number of packets that arrive in the ingress direction in reverse flow.
Egress	Number of packets that are sent out in the egress direction in reverse flow.

- Click **Close** after viewing the analyzer filter details. You are returned to the Prepared Service Analyzer Instances dialog box. Click **Close** to return to the Service Analyzer Page

## Viewing the Service Analyzer Statistics in Grid Format and Graph

To view the traffic analyzer details on packet flows for the different services that match the filter criteria:

- From the View selector, select **Gateway View**. The workspaces that are available in this view are displayed. The Gateway view displays the service delivery gateway (SDG) groups and the SDGs that are part of the high availability pair in an SDG group.
- From the Junos Space user interface, click the **Build** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
- From the View pane, select the All Network item in Gateway view. Click the plus sign (+) beside the All Network item in the View pane to expand the tree and select the device node you want.
- From the task pane, do one of the following:
  - Select **Service Analyzer > ADC Filter** from the task pane. The Service Analyzer for ADC Filter page is displayed. The service instance or template that you previously configured for the ADC service type are displayed. All the instances created in the Service Templates workspace are shown.
  - Select **Service Analyzer > TLB Filter** from the task pane. The Service Analyzer for TLB Filter page is displayed. The service instance or template that you previously configured for the TLB service type are displayed. All the instances created in the Service Templates workspace are shown.
  - Select **Service Analyzer > CGNAT Filter** from the task pane. The Service Analyzer for CGNAT Filter page is displayed. The service instance or template that you previously

configured for the CGNAT service type are displayed. All the instances created in the Service Templates workspace are shown.

The list of SDGs or SDG pairs in a high availability group are displayed, along with the filter instances that were configured for the different services. The number of filter instances that are currently in progress and the number of filter instances that are scheduled or planned to be run at a later time are also displayed. For information on running or clearing filter instances, see *Managing Service Analyzer Filter Instances*.

5. Select the SDGs or SDG pairs (you can select multiple rows to create and assign filters to several SDGs simultaneously) for which you want to create packet analyzer filters for services.
6. From the Service Analyzer page, for the SDG corresponding to a certain service, click the link under the column of the relevant service type. The Prepared Service Analyzer Instances dialog box is shown. Click **Report** under the View Report column to view the traffic analyzer for the particular service.

The Service Analyzer Collection Data — Grid View page is displayed.

At the top of the tabular display, select the criteria for which you want to sort and segregate the packet analyzer information to be viewed. From the Criteria section, do the following:

- a. Select **Control** or **Data** from the first drop-down list to view control or data packets.
- b. Select **Forward** or **Reverse** from the second drop-down list to view statistics for packets in forward or reverse flows.
- c. Select **IPv4** or **IPv6** from the second drop-down list to view IPv4 or IPv6 packets for the filter instance.
- d. Click the search icon to apply the filter conditions and display details matching the specified criteria.

The following fields are displayed in this page:

Field	Description
Name	Name of the SDG or pair of SDGs in a high availability group.
Type	Service type for which packets collected are shown. Values are CGNAT, ADC, or TLB.
Collection Time	Date and time at which the packet details are collected.
Ingress	Number of packets that arrive in the ingress direction in forward and reverse flow.
PreService	Number of packets in the forward flow and reverse flow before the processing of services. You can define the pre-service filter to be applied to traffic before it is accepted for service processing.

Field	Description
Post Service	Number of packets in the forward flow and reverse flow after processing of services. You can define the post-service filter to be applied to traffic after service processing. The filter is applied only if a service set is configured and selected. You can configure a postservice filter on the input side of the interface only. This setting is not supported when the service interface is on an MS-MIC or MS-MPC.
Egress	Number of packets that are sent out in the egress direction in forward flow and reverse flow.

Click **Close** after viewing the collection data in the tabular grid. You are returned to the Prepared Service Analyzer Instances dialog box. Click **Close** to return to the Service Analyzer Page.

7. Alternatively, you can view the service analyzer details in a graphical representation. Click **Graph** under the View Report column to display the packet analyzer details for monitoring as a pictorial form. The Packet Flow Graph dialog box appears.  
  
Line graphs are displayed for data forward flow, data reverse flow, control forward flow, and control reverse flow. The number of packets is displayed on the y-axis and time is displayed along the x-axis. The legends reference the egress, pre-service, post-service, and ingress packets. Mouse over the points in the graph to highlight and view the number of packets at a particular time instance.  
  
At the top of the graph, select the criteria for which you want to sort and segregate the packet analyzer information to be viewed. From the Criteria section, do the following:
  - a. Select **Control** or **Data** from the first drop-down list to view control or data packets.
  - b. Select **IPv4** or **IPv6** from the second drop-down list to view IPv4 or IPv6 packets for the filter instance.
  - c. Select the period for which the service analyzer details must be shown from the third drop-down list. For example, you can select **Last 10 Mins** to display the service analyzer packets collected over the last 10 minutes or the **Last 1 Hr** option to display the service analyzer packets collected over the last one hour.
  - d. Click the search icon to apply the filter conditions and display details matching the specified criteria.
8. Click **Close** after viewing the graph. You are returned to the Prepared Service Analyzer Instances dialog box. Click **Close** to return to the Service Analyzer Page.

**Related Documentation**

- [Packet Analyzer Overview on page 321](#)

## PART 7

# Deploy Mode

- [About Deploy Mode on page 507](#)
- [Configuration File Management on page 511](#)
- [Software Image Management on page 517](#)
- [Deploying Configurations to Devices on page 527](#)
- [Viewing Transactions Associated with Deployment Jobs on page 551](#)



## CHAPTER 24

# About Deploy Mode

- [Understanding Deploy Mode in Gateway and Service Views of Edge Services Director on page 507](#)

## Understanding Deploy Mode in Gateway and Service Views of Edge Services Director

The Deploy mode in Gateway and Service views enables you to deploy configuration changes to devices. You can create a deployment plan for each of the service planning templates, such as the ones defined for ADC or SFW services, and the policy or filter templates, such as the packet filter or SFW policy, that you have created. A deploy plan contains details about the settings and configuration parameters that must be propagated and provisioned on the SDGs managed by Edge Services Director. You can also create, update, display, publish and commission of packet filters, stateful firewall and NAT policies present on discovered and managed SDGs.

This topic describes:

- [Deploying Configuration Changes on page 507](#)
- [Transactions on page 508](#)
- [Modify the Association of SDG Details and Rule Terms for a Policy Filters on page 508](#)
- [View Service Object Statistics on page 509](#)
- [Service Edit on page 509](#)
- [Policy and Filter Management on page 509](#)

## 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. 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 do the following configuration deployment tasks on devices that have pending changes:

- Run configuration deployment jobs immediately or schedule them for future times.
- Preview pending configuration changes before deploying.
- 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 Edge Services 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.

Edge Services Director does not deploy configuration to a device with a configuration that is out of sync (meaning that the device's configuration differs from Edge Services 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 Edge Services 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.

The Service Deployment page provides the following functionalities:

- Approval Management—View the details of the filters/policies and other service deployment plans which are pending for approval. Approve or reject deployment plans done to existing feature.
- Update Devices—View the details of approved filters/policies and other service deployment plans which are ready for commissioning. Commission the deployment plans or discard accordingly.

## Transactions

A transaction refers to an operation or a task that is performed on the service definitions, configuration parameters, and policy settings that are created for provisioning on the devices or Service Delivery Gateways (SDGs). When you create a deployment plan to define the services and policy filters that must be applied and propagated on the devices, the administrator can approve or reject a deploy plan. For each approved deploy plan, a transaction is automatically created by the Edge Services Director application.

## Modify the Association of SDG Details and Rule Terms for a Policy Filters

In Gateway view of Deploy mode, from the Policy & Filters page, which displays all the previously configured CGNAT and SFW service policy filters, and packet filters, you can modify the components or the parameter types that are associated with a particular

service filter. You must lock the packet filters for which you want to modify the attached rule term components or attributes before you can update the settings. You can also select a different SDG to which the packet filter must be applied.

## View Service Object Statistics

In Service view of Deploy mode, you can view a graphical representation in the form of pie charts of the configured ADC, TLB, CGNAT, SFW, and packet policies or filter.

## Service Edit

In Gateway and Service views, you can select a previously configured service template instance, such as a stateful firewall, carrier-grade NAT, traffic load balancer, or adaptive delivery controller, and lock the service instance to select the attributes or components of the service to be modified. You can publish or unpublish service template instances.

## Policy and Filter Management

The Policy and Filter Management feature in the Junos Space Edge Services Director application helps you create, update, display, publish and commission of packet filters, stateful firewall and NAT policies present on discovered and managed SDGs. The Service Management workspace displays a bar graph of draft, published and approved filters or policies for different options available under workspace:

- **Packet Filter:** This option displays packet filters present on SDGs in tabular view. It also provides the ability to create, update, and delete filters on selected SDGs.
- **Stateful Firewall:** This option displays stateful firewall policies present on SDGs in tabular view. It also provides the ability to create, update and delete stateful firewall policies on selected SDGs.
- **CGNAT:** This option displays CGNAT policies present on SDGs in tabular view. It also provides the ability to create, update and delete CGNAT policies on selected SDGs. A published filter or policy is sent for peer review and approval. After approval, the filter or policy is deployed to devices.

### Related Documentation

- [Viewing Deployment Plans on page 529](#)
- [Creating and Assigning a Deployment Plan to Devices on page 533](#)
- [Transactions Overview on page 551](#)
- [Viewing Transactions on page 552](#)



## CHAPTER 25

# Configuration File Management

- [Managing Device Configuration Files on page 511](#)
- [Managing Jobs on page 515](#)

## Managing Device Configuration Files

---

You can back up device configuration files to the Edge Services 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 Edge Services 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 on page 511](#)
- [Backing Up Device Configuration Files on page 512](#)
- [Restoring Device Configuration Files on page 513](#)
- [Viewing Device Configuration Files on page 513](#)
- [Comparing Device Configuration Files on page 514](#)
- [Deleting Device Configuration Files on page 514](#)
- [Managing Device Configuration File Management Jobs on page 514](#)

## 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 348](#) for more information.
- Restore backup device configuration files to devices by selecting devices and clicking Restore. See [“Restoring Device Configuration Files” on page 349](#) for more information.
- View backed up configuration files by selecting a device and clicking View Configuration File. See [“Viewing Device Configuration Files” on page 349](#) for more information.
- Compare backed up device configuration files by selecting devices and clicking Compare Config Files. See [“Comparing Device Configuration Files” on page 350](#) for more information.
- Delete backup device configuration files by selecting devices and clicking Delete. See [“Deleting Device Configuration Files” on page 350](#) for more information.

[Table 50 on page 348](#) describes the information provided in the Manage Device Configuration table.

**Table 73: 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 [“Understanding Build Mode in Location and Device Views of Edge Services Director”](#) on page 151.

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 Edge Services 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 Also**
- [Understanding Deploy Mode in Gateway and Service Views of Edge Services Director on page 335](#)

## Managing Jobs

Edge Services 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, Edge Services Director enables you to view special pre-filtered versions of this page from various other tasks, such as View Discovery Status or View Image Deployment 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 Edge Services 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 20 on page 60](#). 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 Edge Services Director will be available only from Edge Services Director. These jobs will not be listed in the Job Management pane in Junos Space platform and vice-versa.

**Table 74: Job Management Page Fields**

Field	Description
Job ID	The unique ID assigned to the job
Name	The name of the job

*Table 74: Job Management Page Fields (continued)*

Field	Description
Percent	The percentage of completion of the job
State	The status of the job: <ul style="list-style-type: none"><li>• Success—Job completed successfully</li><li>• Failure—Job failed and was terminated</li><li>• Job Scheduled—Job is scheduled but has not yet started</li><li>• In progress—Job is has started, but not completed</li><li>• Cancelled—Job is cancelled</li></ul>
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.

## CHAPTER 26

# Software Image Management

- [Managing Software Images on page 517](#)
- [Deploying Software Images on page 520](#)
- [Managing Software Image Deployment Jobs on page 523](#)

## Managing Software Images

---

This topic describes how to manage software images for managed devices.

To start managing software images:

1. Click **Deploy** in the Edge Services 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.

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 on page 517](#)
- [Adding Software Images to the Repository on page 518](#)
- [Using the Device Image Upload Window on page 518](#)
- [Viewing Software Image Details on page 519](#)
- [Using the Device Image Summary Window on page 519](#)
- [Deleting Software Images on page 519](#)

## 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 52 on page 354](#) describes the information provided in the Device Image Repository table.

*Table 75: 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 Edge Services 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 354](#) 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 355](#) 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 53 on page 355](#) describes the fields in this window.

*Table 76: 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.
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**

- [Understanding Deploy Mode in Gateway and Service Views of Edge Services Director on page 335](#)

## Deploying Software Images

---

This topic describes how to deploy software images to managed devices. You must upload software images to the Edge Services Director server before you can deploy them to devices. See *Managing Software Images* for more information.

To start deploying software images:

1. Click **Deploy** in the Edge Services Director banner.
2. Select a node in the View pane that contains the devices to 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 on page 520](#)
- [Selecting Software Images To Deploy on page 521](#)
- [Selecting Options for Software Deployment on page 522](#)
- [Summary of Software Deployment on page 523](#)

### 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 357](#).

## 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 54 on page 357](#) 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 pop-up message notifies you if you do not select a software image for all the listed devices. This is just for your information. No action will be 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 358](#).

*Table 77: Select images for devices Table*

Table Column	Description
Device Family	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.
Count	Number of devices contained within a device family.
IP Address	Device's IP address.
Device Name	Device's name.

Table 77: Select images for devices Table (continued)

Table Column	Description
State	Device's state: <ul style="list-style-type: none"> <li>• UP—Edge Services Director can communicate with the device.</li> <li>• DOWN—Edge Services Director cannot communicate with the device.</li> </ul>
Running Image Version	Software version the device is running.
Proposed Image Version/Profile	Software version that will be installed on the device when the job runs successfully.

## Selecting Options for Software Deployment

The options that you can configure in the Select Options page are described in [Table 55 on page 358](#). 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 359.

Table 78: Image Management Job Options

Option	Action
<b>Select Options</b>	
<b>All Device Types</b>	
Delete any existing image before download	Select to delete any existing software images on devices before downloading the new software image.
Reboot device after successful installation	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.  <b>NOTE:</b> This option may get disabled based on your details that you specify in the remaining fields. This indicates that for the options that you specified, the system will automatically reboot the device as per the requirement during or after the image upgrade.
<b>Wired Devices</b>	
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	Select if you want to perform a Nonstop software upgrade (NSSU) or lin-service software upgrade (ISSU).  ISSU enables you to upgrade between two different Junos OS releases with minimal disruption on the control plane and with minimal disruption of traffic.  NSSU enables you to upgrade the software running on an MX Series router with redundant Routing Engines or on most EX Series Virtual Chassis by using a single command and with minimal disruption to network traffic

Table 78: Image Management Job Options (continued)

Option	Action
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	Select to copy the new Junos OS image into the alternate root partition. This ensures that the resilient dual-root partitions feature operates correctly.  This option is available only if you select <b>Reboot device after successful installation</b> .
<b>Select Schedule</b>	
Stage now	Select <b>Stage now</b> to start staging software images to devices as soon as the job runs.
Stage later time	Select <b>Stage later time</b> 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 <b>Upgrade now</b> to start upgrading software images on devices as soon as staging finishes.
Upgrade later time	Select <b>Upgrade later time</b> to schedule the software upgrade for a later time.
Deployment Schedule	If you selected Upgrade later time, enter the date and time for upgrade to start.  If you scheduled staging, you must schedule the upgrade for at least 10 minutes after staging, to ensure that staging completes before upgrade starts.

## 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**
- [Understanding Deploy Mode in Gateway and Service Views of Edge Services Director on page 335](#)

## 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 Edge Services 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 on page 524](#)
- [Viewing Software Image Job Details on page 525](#)
- [Using the Device Image Staging Window on page 525](#)
- [Canceling Software Image Jobs on page 526](#)

## 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 361](#) for more information.
- Cancel a pending job by selecting the job and clicking Cancel Job. See [“Canceling Software Image Jobs” on page 362](#) for more information.

[Table 56 on page 360](#) describes the information provided in the of the Image Deployment Jobs table.

**Table 79: 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"><li>• CANCELLED—The job was cancelled by a user.</li><li>• SCHEDULED—The job is scheduled but has not run yet.</li><li>• INPROGRESS—The job is running.</li><li>• SUCCESS—The job completed successfully. This state is applied if all of the devices in the job completed successfully.</li><li>• 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.</li></ul>
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.

*Table 79: Image Deployment Jobs Table (continued)*

Table Column	Description
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 361](#) 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 57 on page 361](#) describes this window.

*Table 80: 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.
Status	Job status. The possible statuses are: <ul style="list-style-type: none"> <li>• CANCELLED—The job was cancelled by a user.</li> <li>• SCHEDULED—The job is scheduled but has not run yet.</li> <li>• INPROGRESS—The job is running.</li> <li>• SUCCESS—The job completed successfully.</li> <li>• FAILURE—The job failed.</li> </ul>
Host Name	Host name of device.
Status	Device status. The possible statuses are: <ul style="list-style-type: none"> <li>• INPROGRESS—The job is running.</li> <li>• SUCCESS—The job completed successfully.</li> <li>• FAILURE—The job failed.</li> </ul>

*Table 80: Device Image Staging Window Description (continued)*

Field	Description
% 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**

- [Understanding Deploy Mode in Gateway and Service Views of Edge Services Director on page 335](#)

## CHAPTER 27

# Deploying Configurations to Devices

- [Planning and Deployment of Service Templates Overview on page 527](#)
- [Viewing Deployment Plans on page 529](#)
- [Creating and Assigning a Deployment Plan to Devices on page 533](#)
- [Modifying the Association of SDG Details and Service Components for a Packet Filter Policy on page 544](#)
- [Modifying the Association of SDG Details and Service Components for a Service Policy Filter on page 547](#)

## Planning and Deployment of Service Templates Overview

---

The service planning functionality of Edge Services Director enables you to create service templates and deploy the same service template configuration to multiple devices. As a designer, when you create a service template, you can configure generic properties and modify it to suit your network deployment needs, thereby enabling streamlined and simplified administration of services (such as stateful firewall [SFW], carrier-grade NAT [CGNAT], application delivery controller [ADC], and traffic load balancing [TLB]) on service delivery gateways (SDGs) in your topology.

This topic contains the following sections that describe the sequence of operations performed for planning and deploying service templates:

- [Planning Workflow for Service Templates on page 527](#)
- [Deployment Workflow for Service Templates on page 528](#)

## Planning Workflow for Service Templates

The fundamental workflow of planning templates is derived from the existing devices inventory and framework:

- The designer creates the service template by using the available inventory service components and structure model.
- The designer can import the discovered service data while creating the service template for the existing service data values of the device.
- While creating the service template, the designer can add or modify service parameter values and restrict the access level for each service parameter for the operator. The

designer can set the following access levels for each service parameters to operator in the planning template:

- Read-only (The configuration parameter is read-only for operator as part of provisioning.)
- Editable (The configuration parameter is editable as part of provisioning.)
- Mandatory (The configuration parameter is part of provisioning but operator must provide the values.)
- Device-Specific (The configuration parameter value needs to be entered by the operator for each device during deployment.)

The designer must publish the service templates to the operator to use in the creation of deployment plans.

An operator can create the service deployment plan using the planning template so that one deployment plan can be applied on multiple devices. This method of deploy reduces the scope for human errors that can occur with the CLI interface.

## Deployment Workflow for Service Templates

The following workflow is used the deployment process:

1. An operator uses only published planning templates to create deployment plans for a single SDG service or multiple SDG devices.
2. The operator modifies or adds data in the allowed service specific parameters according to the access permissions specified by the designer and associate the deployment plan with a single SDG device or multiple SDG devices.
3. An operator publishes the deployment plans with the device association for the designer to review and approve.
4. The administrator must approve the configuration changes for each device for each service deployment plan.
5. The operator has a copy of the service planning template while creating the deployment plan. After creating a deployment plan, there is no association between the deployment plan and planning template. Changes made by the operator to the deployment template are maintained in their own copy and are not reflected in the original planning template and vice-versa.
6. The deployment plan is assigned to multiple devices and sent for approval. After a deployment plan is associated with a device, the device contains its own copy of the deployment plan . For example, if one deployment plan was created and associated with four devices, you see four deployment plans separately on each device in the service deployment plan. The operator can edit the deployment plan for each device if needed.

The status of a deployment plan determines the kinds of tasks that a user can perform:

- Add – Create a deployment plan; the status that immediately follows this status is the Unpublish state.
- Update – Update a deployment plan.
- Delete – Delete a deployment plan. Only plans that are in the Unpublish state can be deleted.
- Publish – Publish the deployment plan. In this state, the operator waits for an approval from the designer before the plan can be deployed to a device.
- Unpublish – Unpublish the published deployment plan to make more changes.
- Approve – The administrator or designer approves the published deployment plan.
- Reject – The administrator or designer rejects the published deployment plan.

A deployment plan can obtain any one of the following status:

- Discovered – This is the default state for filter discovered and stored in the inventory.
- Unpublished – New, updated, and deleted filters are saved in draft or Unpublished state initially.
- Published – After all the changes are done, the filter in draft status is ready for admin or designer approval and is published.
- Rejected – An administrator or designer can reject the published filter to disapprove updates.
- Approved – An administrator or designer can approve the published filter to concur changes.
- Commissioned – An administrator or designer can commission filters to push to devices.
- Commission Failed – This state is assigned to a filter if commissioning of the filter fails.

**Related  
Documentation**

- [Service Templates Overview on page 189](#)

---

## Viewing Deployment Plans

---

A deploy plan contains details about the settings and configuration parameters that must be propagated and provisioned on the SDGs managed by Edge Services Director. A deploy plan is associated with a set of SDGs.

The Deployment Plans page displays all of the created deploy plans. You can perform the following tasks on this page:

- Create a deploy plan.
- Modify or delete a previously specified deploy plan.
- Search deployment plans.
- Schedule a deployment plan to provision configuration settings on devices immediately or for a future specified time. You can also select the option to cause the configuration

set that is being propagated to devices to be rolled back if a failure occurs during the deployment operation. You can return to the most recently configured successful configuration on the device.

- Select the deploy plan for which you want to view the configuration settings contained in the plan and view the CLI format of the configuration with the statement options and hierarchy levels.
- Validate that the pending changes are compatible with the device's configuration.
- Approve or reject a published deployment plan.



**NOTE:** For details on the different states through which a deployment plan traverses, from the time of creation of a plan, see *Planning and Deployment of Service Templates Overview*.

The Deployment Plans page is divided into two panes. The top half of the page displays a bar chart. The service type is displayed on the y-axis and the number of plans corresponding to each service type is displayed on the x-axis. Mouse over the different segments of the bar chart to highlight and view the total number of deploy plans in each state. Click any of the states in the color-coding legend box to remove that particular deploy plan state from being displayed in the bar chart. The following color-coding legend denotes the deployment plans in the different states:

- Light Orange—Denotes deploy plans that are newly created
- Teal—Denotes deploy plans for which validation of configuration is successful
- Red—Denotes deploy plans for which provisioning on devices failed
- Dark orange—Denotes deploy plans currently being provisioned on devices
- Dark blue—Denotes deploy plans scheduled for deployment at a future time
- Green—Denotes deploy plans that have been successfully propagated and applied on devices
- Purple—Denotes deploy plans for which validation of the deployment plan failed

The lower half of the page displays information about the deployment plans that have been previously created, the services to which they pertain, and the status of propagation of configuration settings to the devices contained in the deploy plan.

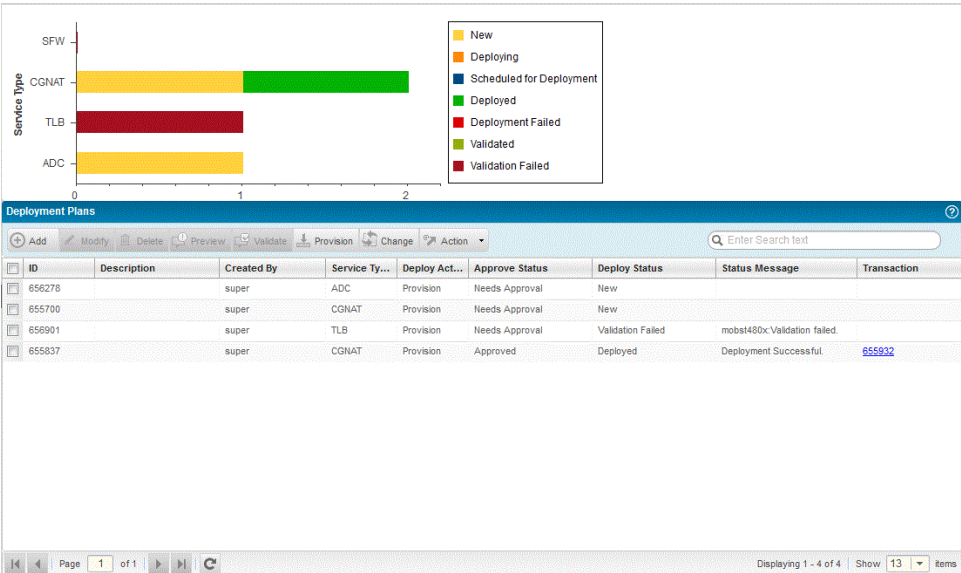
To view the configured deployment plans:

1. From the View selector, select **Gateway View** or **Service View**. The workspaces that are applicable to this view are displayed. In Gateway view, the devices in the entire network are displayed, organized by the device types and the device models within each device type. In Service View, the different types of services are displayed in the View pane.
2. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.

The functionalities that you can configure in this mode are displayed in the task pane.

- 3. From the View pane, select the All Network item in Gateway view. Click the plus sign (+) beside the All Network item in the View pane to expand the tree and select the device node you want.
- 4. Select **Deploy Service > Manage Deployment Plans** from the task pane. The Deployment Plans page is displayed.

Figure 48: Deployment Plans Page



The following fields are displayed on the page:

Field	Description
ID	Unique identifier assigned by the system for a deployment plan.
Created By	Name of the user that created the deploy plan.
Created Time	Date and time at which the deploy plan was created.
Service Type	Type of service for which the deployment plan is created.
Deploy Status	Status of the deployment plan, such as approved, needing approval, published, rejected, or unpublished.
Status Message	Information about the status of the deploy plan that you can use to modify or take appropriate steps for ensuring successful deploy of configurations to devices.

Field	Description
Transaction	<p>Unique identifier of the transaction attached to the deploy plan. Click the link in the transaction ID to open the Transactions page. A transaction is an operation that is currently running as part of the transaction to propagate configuration settings to devices. You can also view a list of all the transactions associated with this deploy plan by clicking the <b>All</b> link.</p> <p>The Transactions dialog box is displayed. A list of all of the transaction IDs, their statuses, and information about each of the statuses is displayed. These statuses denote whether a configuration change occurred on a device, whether rollback of a configuration set has been performed, and whether the provisioning of settings to a device succeeded or failed.</p>

5. Select the plan you want to modify and click the pencil icon above the table of discovery profiles to modify the deploy plan. The Modify Deployment Plan window appears. You can perform the same steps as the sequence of events that you perform to create a deploy plan.
6. Select the plan you want to delete and click the **Delete** icon at the top of the table of deploy plans that are listed. The selected deploy plan is removed from the list.
7. Select the plan for which you want to provision the configuration and click the **Provision Deployment Plan** button. See the *Scheduling Deployment of Services and Policies* section for details.
8. Select the plan for which you want to validate the configuration and click the **Validate Deployment Plan** button. You can perform validation checks on the configuration planned to be deployed to examine and correct any syntax errors or incompatible settings. You can also validate without deploying the configuration.

Configuration changes are validated for each device both in Edge Services 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. Edge Services Director will not deploy configuration to a device with a configuration that is out of sync (meaning that the device's configuration differs from Edge Services 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.

The Configuration Validation window displays the results of the verification. The 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.

The errors or warnings, if any, for the objects or components are displayed.

9. Select the deploy plan for which you want to view the configuration settings contained in the plan and click **Preview Deployment Plan** to display the CLI format of the configuration with the statement options and hierarchy levels. Close the dialog box after viewing the settings.
10. Select the deploy plan for which you want to modify the configuration settings contained in the plan and click **Modify**. Deploy plans that are not based on templates are not editable and a popup dialog box displays a message stating that only template-based deploy plans can be modified.
11. Select one or more published deploy plans or policies from the page. All the policies and deploy plans that you previously created are displayed. The devices associated with the policies and plans are also listed, categorized by the SDG pairs to which they belong.

Click the down arrow in the **Actions** menu and select **Approve**.

The state for selected plans is changed to Approved. The approved plan can be used for commissioning it to devices.

12. Select one or more published deploy plans or policies from the page. All the policies and deploy plans that you previously created are displayed. The devices associated with the policies and plans are also listed, categorized by the SDG pairs to which they belong.

Click the down arrow in the **Actions** menu and select **Reject**. The Reject Comments dialog box appears.

Enter comments for rejecting the plan to enable the operator correct and modify the plan. Click **Reject** in the dialog box.

The state for selected plans is changed to Rejected. If a plan is rejected, the operator corrects the plan and sends the updated plan to the administrator for approval. After approval, a plan cannot be modified.

#### Related Documentation

- [Creating and Assigning a Deployment Plan to Devices on page 533](#)
- [Transactions Overview on page 551](#)
- [Viewing Transactions on page 552](#)

## Creating and Assigning a Deployment Plan to Devices

You need to create a deployment plan for each of the service planning templates, such as the ones defined for ADC or SFW services, and the policy or filter templates, such as the packet filter or SFW policy, that you have created. A deploy plan contains details about the settings and configuration parameters that must be propagated and provisioned on the SDGs managed by Edge Services Director. A deploy plan is associated with a set of SDGs. A deploy plan passes through the following steps in a workflow:

1. The operator creates a deploy plan, assigns devices, and publishes it to make it available for transmission and application on the devices. A plan is initially unpublished. If you want to modify the deploy plan, it needs to be in unpublished state. The administrator can view only published plans. The operator needs to publish any unpublished plan.
  2. The administrator can approve or reject a deploy plan. If a plan is rejected, the operator corrects the plan and sends the updated plan to the administrator for approval. After approval, a plan cannot be modified. If the plan is approved, it moved to the approved state.
  3. The administrator can reject to unpublish the deploy plan. Also, the administrator can discard a plan, which causes the plan to be deleted from the database.
  4. The operator or designer can transfer the configuration to the devices in a plan. If the commissioning, which causes the configuration to be sent to the device, is successful, the relevant settings are applied. Otherwise, the configuration push fails, and the plan needs to be edited, published, approved, and commissioned again.
- [Creating a Deployment Plan on page 534](#)
  - [Publishing a Deploy Plan on page 537](#)
  - [Viewing Deploy Plans and Policies on page 538](#)
  - [Approving a Deploy Plan and Policies on page 539](#)
  - [Unpublishing a Deploy Plan and Policies on page 540](#)
  - [Deploying a Deploy Plan and Policies Immediately on page 540](#)
  - [Scheduling Deployment of Services and Policies on page 541](#)
  - [Rejecting a Deploy Plan and Policies on page 542](#)
  - [Changing a Deploy Plan Action or Decommissioning a Deploy Plan on page 543](#)
  - [Discarding a Deploy Plan and Policies on page 544](#)

## Creating a Deployment Plan

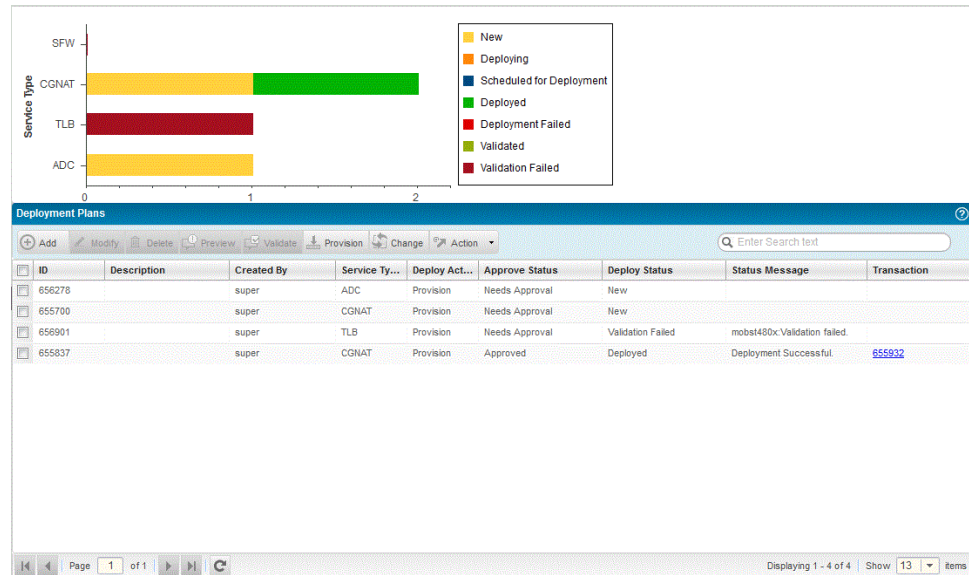
You must have previously defined service templates and policy or filter templates before you can create a deployment plan.

To create a deployment plan and assigning devices to it:

1. From the View selector, select **Gateway View** or **Service View**. The workspaces that are applicable to this view are displayed. In Gateway view, the devices in the entire network are displayed, organized by the device types and the device models within each device type. In Service View, the different types of services are displayed in the View pane.
2. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.

- From the View pane, select the All Network item in Gateway view. Click the plus sign (+) beside the All Network item in the View pane to expand the tree and select the device node you want.
- Select **Deploy Service > Manage Deployment Plans** in the task pane. The Create Deployment Plan page appears.

Figure 49: Create Deployment Plan Page



- Click the **Add** icon to begin the creation of a deployment plan
- Instead of a user-specified name for the deploy plan, a unique deployment job ID is assigned automatically by the Edge Services Director application. The deployment creation process occurs through a wizard. The first step of the wizard is to select the service type for which the deploy plan needs to be created.

Click the **Select Service Template** button. You need not click this button if you are navigating to the Deployment Plans page for the first time or are traversing from another page of the wizard to this page. You must click this button only if you are viewing other pages of the wizard, such as modifying service template settings or are assigning devices to the plan.

- Do one of the following:

From the **Service Type** list, select the service template type for which you want to assign devices and deploy. Type of the service, such as ADC, SFW, CGNAT, or TLB can be selected.

After you select the service type, the lower pane of the page displays all of the previously configured service templates for the particular service type.

8. Select the check box next to the service template that you want to provision and apply to an SDG or pair of SDGs.
9. Click **Next** to proceed to the second step of the wizard, which is to modify the service template settings. Instead, you can also click the **Service Basic Details** button.

Alternatively, click **Previous** to return to the earlier step or page of the wizard. Click **Cancel** to discard the deploy plan creation.

The service template components are displayed. You can add or update attributes or elements of the service definition. For details about managing service templates, see *Creating and Managing an ADC Service Template*, *Creating and Managing a TLB Service Template*, *Creating and Managing an SFW Service Template*, and *Creating and Managing a NAT Service Template*.

10. From the boxes that show the components of a service template, you can edit, delete, or add elements to it. If you do not have permissions to update a template, the corresponding icons are not shown.
11. Click **Next** to proceed to the third step of the wizard, which is to assign SDGs or SDG pairs to the deployment plan. Instead, you can also click the **Select Service Gateways** button. Alternatively, click **Previous** to return to the earlier step or page of the wizard. Click **Cancel** to discard the deploy plan creation.
12. Select the check boxes next to the SDGs or SDG groups that you want to assign to the plan. Based on your selection of a service or a policy template, the components or attributes are shown for the corresponding device.
13. Click **Next** to proceed to the final step of the wizard, which is to modify the service definition settings that are assigned to the devices you have selected for the deployment plan to be provisioned.

The configuration details are displayed in property view and configuration view. The property view is useful if you want a GUI, tree-based structure of display. In this view, you can drill-down the tree and view data about each of the service attributes. Property view is simple view of configuration as key value pair. The dynamic fields in form view are defined using parameters. The configuration view is beneficial if you are familiar with the CLI interface structure and want to view service attributes in the form of configuration fstanzas and hierarchy levels.

14. From the boxes that show the components of a service template, you can edit, delete, or add elements to it. If you do not have permissions to update a template, the corresponding icons are not shown.
15. Click **Finish** in the Assign Deployment Plan page to save the plan. Otherwise, click **Cancel** to discard the changes. Alternatively, click **Previous** to return to the earlier step or page of the wizard.

You are returned to the Deployment Plans page.

16. Click **Validate Configuration** in the Deployment Plans page to perform validation checks on the configuration planned to be deployed to examine and correct any syntax errors or incompatible settings. You can also validate without deploying the configuration.

Configuration changes are validated for each device both in Edge Services 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. Edge Services Director will not deploy configuration to a device with a configuration that is out of sync (meaning that the device's configuration differs from Edge Services 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.

The Configuration Validation window displays the results of the verification. The 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.

The errors or warnings, if any, for the objects or components are displayed.

17. Select **Reject** or **Approve** from the Actions menu to reject the plan or to move the plan to approved state and make it available for commissioning to the devices.

## Publishing a Deploy Plan

To publish a deployment plan:

1. From the View selector, select **Gateway View** or **Service View**. The workspaces that are applicable to this view are displayed. In Gateway view, the devices in the entire network are displayed, organized by the device types and the device models within each device type. In Service View, the different types of services are displayed in the View pane.
2. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the View pane, select the All Network item in Gateway view. Click the plus sign (+) beside the All Network item in the View pane to expand the tree and select the device node you want.
4. Select **Deploy Service > Manage Deployment Plans** in the task pane. The Deployment Plan page appears

You can search for a plan or policy by entering the search criteria in the search field at the top of the page.

5. Click **Edit** to modify a deployment plan. Modify and save the plan in the Assign Deployment Plan to Devices page.
6. Click **Publish** at the bottom of the page to move the plan to published state and make it available for commissioning to the devices.

## Viewing Deploy Plans and Policies

To view the deployment plan and policies:

1. From the View selector, select **Gateway View** or **Service View**. The workspaces that are applicable to this view are displayed. In Gateway view, the devices in the entire network are displayed, organized by the device types and the device models within each device type. In Service View, the different types of services are displayed in the View pane.
2. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the View pane, select the All Network item in Gateway view. Click the plus sign (+) beside the All Network item in the View pane to expand the tree and select the device node you want.
4. Select **Deploy Service > Manage Deployment Plans** in the task pane. The Deployment Plans page appears.

You can search for a plan or policy by entering the search criteria in the search field at the top of the page.

5. Select one or more published deploy plans or policies from the page. All the policies and deploy plans that you previously created are displayed. The devices associated with the policies and plans are also listed, categorized by the SDG pairs to which they belong.

The SDG/Feature column displays the deployment plans and policies, grouped by the pair of SDGs or SDG groups, and devices to which the plans and policies are assigned. The status of the plan, and progress of commissioning, if initiated, are also shown.

6. Click the **View** link under the Configuration column to view the CLI format of the settings.

7. Click **Actions** to perform an appropriate action on the selected plans or policies.
8. You can search and filter the displayed items on the page

## Approving a Deploy Plan and Policies

To approve a deployment plan and policies:

1. From the View selector, select **Gateway View** or **Service View**. The workspaces that are applicable to this view are displayed. In Gateway view, the devices in the entire network are displayed, organized by the device types and the device models within each device type. In Service View, the different types of services are displayed in the View pane.
2. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the View pane, select the All Network item in Gateway view. Click the plus sign (+) beside the All Network item in the View pane to expand the tree and select the device node you want.
4. Select **Deploy Service > Manage Deployment Plans** in the task pane. The Deployment Plans page appears.  
You can search for a plan or policy by entering the search criteria in the search field at the top of the page.
5. Select one or more published deploy plans or policies from the page. All the policies and deploy plans that you previously created are displayed. The devices associated with the policies and plans are also listed, categorized by the SDG pairs to which they belong.
6. Click the down arrow in the **Actions** menu and select **Approve**. The state for selected plans is changed to Approved. The approved plan can be used for commissioning it to devices.

## Unpublishing a Deploy Plan and Policies

To unpublish a deployment plan and policies:

1. From the View selector, select **Gateway View** or **Service View**. The workspaces that are applicable to this view are displayed. In Gateway view, the devices in the entire network are displayed, organized by the device types and the device models within each device type. In Service View, the different types of services are displayed in the View pane.
2. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the View pane, select the All Network item in Gateway view. Click the plus sign (+) beside the All Network item in the View pane to expand the tree and select the device node you want.
4. Select **Deploy Service > Deployment Plans** in the task pane. The Deployment Plans page appears.  
You can search for a plan or policy by entering the search criteria in the search field at the top of the page.
5. Select one or more published deploy plans or policies from the page. All the policies and deploy plans that you previously created are displayed. The devices associated with the policies and plans are also listed, categorized by the SDG pairs to which they belong.
6. Click the down arrow in the **Actions** menu and select **Unpublish**. The state for selected plans is changed to Unpublished. The operator can modify the plan and publish it again to send it for approval.

## Deploying a Deploy Plan and Policies Immediately

To deploy a deployment plan and policies immediately:

1. From the View selector, select **Gateway View** or **Service View**. The workspaces that are applicable to this view are displayed. In Gateway view, the devices in the entire network are displayed, organized by the device types and the device models within each device type. In Service View, the different types of services are displayed in the View pane.
2. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.

3. From the View pane, select the All Network item in Gateway view. Click the plus sign (+) beside the All Network item in the View pane to expand the tree and select the device node you want.
4. Select **Deploy Service > Manage Deployment Plans** in the task pane. The Deployment Plans page appears.

You can search for a plan or policy by entering the search criteria in the search field at the top of the page.

5. Select one or more published deploy plans or policies from the page. All the policies and deploy plans that you previously created are displayed. The devices associated with the policies and plans are also listed, categorized by the SDG pairs to which they belong.
6. Click the **Provision** button. The Provision Deployment Plan dialog box appears.
7. Click **Provision** to begin the commissioning of settings defined in the deploy plan to the corresponding devices immediately.

The configuration deployment job runs. To view the status or results of the deployment job, you can view the Approve and Provision columns of the Deployment Plans page.

8. Select the **Rollback Deployment Configuration in case of failure** check box to cause the configuration set that is being propagated to devices to be rolled back if a failure occurs during the deployment operation. You can return to the most recently configured successful configuration on the device.
9. In the Deployment Plans page, the Deploy Status and Message columns are updated indicating the progress of commission. The Approve Status and Deploy Action fields denote the approval status and the action being performed on the deployment plan. If the deploy is successful, the status denotes Commissioned. If the deploy fails, the status changes to Commission Failed.

## Scheduling Deployment of Services and Policies

To schedule configuration deployment to devices:

1. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
2. Select **Deploy Service > Manage Deployment Plans** in the task pane. The Deployment Plans page appears.

You can search for a plan or policy by entering the search criteria in the search field at the top of the page.

3. Select one or more published deploy plans or policies from the page. All the policies and deploy plans that you previously created are displayed. The devices associated with the policies and plans are also listed, categorized by the SDG pairs to which they belong.
4. Click the **Provision** button. The Provision Deployment Plan dialog box appears.
5. Select the **Rollback Deployment Configuration in case of failure** check box to cause the configuration set that is being propagated to devices to be rolled back if a failure occurs during the deployment operation. You can return to the most recently configured successful configuration on the device.
6. Select **Schedule at a later time**, and enter the start date and time for deployment. Click **Provision** to accept changes and exit the window. Click **Cancel** to cancel changes and exit the window.

The status changes to Commission Scheduled.

7. In the Deployment Plans page, the Deploy Status and Message columns are updated indicating the progress of commission. The Approve Status and Deploy Action fields denote the approval status and the action being performed on the deployment plan. If the deploy is successful, the status denotes Commissioned. If the deploy fails, the status changes to Commission Failed.

## Rejecting a Deploy Plan and Policies

To reject a deployment plan and policies:

1. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
2. Select **Deploy Service > Manage Deployment Plans** in the task pane. The Deployment Plans page appears.  
You can search for a plan or policy by entering the search criteria in the search field at the top of the page.
3. Select one or more published deploy plans or policies from the page. All the policies and deploy plans that you previously created are displayed. The devices associated with the policies and plans are also listed, categorized by the SDG pairs to which they belong.

4. Click the down arrow in the **Actions** menu and select **Reject**. The Reject Comments dialog box appears.
5. Enter comments for rejecting the plan to enable the operator correct and modify the plan. Click **Reject** in the dialog box.

The state for selected plans is changed to Rejected. If a plan is rejected, the operator corrects the plan and sends the updated plan to the administrator for approval. After approval, a plan cannot be modified.

## Changing a Deploy Plan Action or Decommissioning a Deploy Plan

To change a deployment plan action and policies for decommissioning or unprovisioning it when a plan is no longer needed:

1. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.

2. Select **Deploy Service > Manage Deployment Plans** in the task pane. The Deployment Plans page appears.

You can search for a plan or policy by entering the search criteria in the search field at the top of the page.

3. Select one or more published deploy plans or policies from the page. All the policies and deploy plans that you previously created are displayed. The devices associated with the policies and plans are also listed, categorized by the SDG pairs to which they belong.
4. Click the **Change** button. You can modify the action to be performed on the deployment plan, such as to not commission the settings on devices. However, you cannot change the deploy plan action for plans created from the Service Edit workspace.
5. Enter comments for decommissioning the plan. Click **OK** in the dialog box.

In the Deployment Plans page, the Provision Status and Message columns are updated indicating the progress of commission. If the deploy is successful, the status denotes Commissioned. If the deploy fails, the status changes to Commission Failed.

## Discarding a Deploy Plan and Policies

To discard a deployment plan and policies:

1. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
2. Select **Deploy Service > Manage Deployment Plans** in the task pane. The Deployment Plans page appears.  
You can search for a plan or policy by entering the search criteria in the search field at the top of the page.
3. Select one or more published deploy plans or policies from the page. All the policies and deploy plans that you previously created are displayed. The devices associated with the policies and plans are also listed, categorized by the SDG pairs to which they belong.

4. Click the down arrow in the **Actions** menu and select **Discard**.

The Discard Comments dialog box appears.

5. Enter comments for discarding the plan to enable the operator correct and modify the plan. Click **Discard** in the dialog box.

The state for selected plans is changed to Discarded. If a service or policy template is new, it is removed from the system. If a policy or service template has been updated or deleted, you can restore it using the copy from the device.

- Related Documentation**
- [Viewing Deployment Plans on page 529](#)
  - [Transactions Overview on page 551](#)
  - [Viewing Transactions on page 552](#)

---

## Modifying the Association of SDG Details and Service Components for a Packet Filter Policy

---

From the Policy & Filters page, which displays all the previously configured packet filters, you can modify the components or the parameter types that are associated with a particular service filter. You must lock the packet filters for which you want to modify the attached rule term components or attributes before you can update the settings. You can also select a different SDG to which the packet filter must be applied.

1. From the View selector, select **Service View**. The workspaces that are applicable to edge services are displayed.

3. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.

4. From the View pane, select the All Network item in Gateway view. Click the plus sign (+) beside the All Network item in the View pane to expand the tree and select the device node you want.

- The Services page is displayed.

- Select **Packet Filter** to open the Service Edit > Packet Filter page on the right pane.

**Table 81: Service Edit > Packet Filter Page**

Field	Description
Instance Name	Name of the configured service template instance
OS Version	Junos OS release number that represents a particular revision of the software that runs on a Juniper Networks routing platform, for example, Junos OS Release 8.5, 9.1, or 9.2. Each Junos OS release has certain new features that complement the software processes that support Internet routing protocols, control the device's interfaces and the device chassis itself, and allow device system management.
Group Name	Name of the SDG group
Reference Host	Hostname of the SDG with which the service instance is associated.
Deployment Plans	Name of the deployment plan with which the service template is attached.

7. From the Term Name drop-down list, select the rule term with which the packet filter must be applied.
8. From the Host Name drop-down list, select the hostname of the SDG.
9. In the Select Common Components section, select the check boxes beside the service modules or components, such as packet filters, SFW rules, or CGNAT rules, that are displayed. The displayed components depend on the attributes that are previously defined for that selected packet filter. For example, if the service policy is for stateful firewall, SFW rules and SFW rule sets are shown. Select the check box beside Config Category to select all the service components.

The modified association is saved.

You can use the **Actions** menu in the Service Template pages for packet filters to publish, unpublish, export, and restore the defined policies or filters. For details, see *Using the Actions Menu in the Service Template Page*.

#### Related Documentation

- [Creating Service Gateway Groups on page 99](#)
- [Managing Service Gateway Groups on page 101](#)
- [Searching Unmanaged Devices on page 104](#)
- [Viewing the List of Discovered, Managed, and Unmanaged Devices on page 106](#)
- [Modifying Discovery Profiles on page 113](#)
- [Deleting Discovery Profiles on page 114](#)

## Modifying the Association of SDG Details and Service Components for a Service Policy Filter

---

From the Policy & Filters page, which displays all the previously configured service policy filters, you can modify the components or the parameter types that are associated with a particular service filter. You must lock the service policy filters for which you want to modify the attached service components or attributes before you can update the settings. You can also select a different SDG to which the service policy filter must be applied.

To modify the association of SDGs and service components for a service policy filter, such as a stateful firewall service, or a carrier-grade NAT service policy:

1. From the View selector, select **Service View**. The workspaces that are applicable to edge services are displayed.
2. Select All Network from the Service View pane. You can modify the association of SDGs with service policies, only if you select the All Network label in the View pane. If you expand the All Network tree and select an SDG group or an SDG in a redundancy pair, you cannot modify the association of service policies and rules with SDGs in a single-shot, one-step operation.
3. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
4. From the View pane, select the All Network item in Gateway view. Click the plus sign (+) beside the All Network item in the View pane to expand the tree and select the device node you want.
5. Select **Policy & Filters** from the task pane.  
The Service Edit page is displayed.

Figure 51: Enhanced Edit Page for Service Policy Rules

Rule Name	Group Name	Host Name	Match Direct...	Term Name	Source Address	Destination...	Source Pool	Destination...	Application...	Translatio...
cgnat-rule123	didddd	mobst480v	INPUT	term-t1	10.165.49.208/28	phone-broadb...	11 - 123		SDG-APPS	NO_TRAN...
nap644-SS1-r1	didddd	mobst480v	INPUT	t1					SDG-APPS	NAPT_44
nap644-SS2-r1	didddd	mobst480v	INPUT	t1					SDG-APPS	NAPT_44
nap644-SS3-r1	didddd	mobst480v	INPUT	t1					SDG-APPS	NAPT_44
nap644-SS4-r1	didddd	mobst480v	INPUT	t1					SDG-APPS	NAPT_44
nap644-SS4-r2	didddd	mobst480v	INPUT	t1	10.212.163.140/32 10.212.163.130-10	NAT-PAT				NO_TRAN...

6. Click the plus sign (+) next to Policy & Filters to expand the tree in the task pane and view the list of filter templates. Do one of the following:

- Select **CGNAT** to open the Service Edit > CGNAT page on the right pane.
- Select **SFW** to open the Service Edit > SFW page on the right pane.

The following fields are displayed on this page:

Table 82: Services – CGNAT and SFW Page

Field	Description
Instance Name	Name of the configured service template instance
OS Version	Junos OS release number that represents a particular revision of the software that runs on a Juniper Networks routing platform, for example, Junos OS Release 8.5, 9.1, or 9.2. Each Junos OS release has certain new features that complement the software processes that support Internet routing protocols, control the device's interfaces and the device chassis itself, and allow device system management.
Group Name	Name of the SDG group
Reference Host	Hostname of the SDG with which the service instance is associated.
Applications	Name of the applications protocols created for the service template.
Application Sets	Name of the application sets created for the service template.
SFW Rules	Name of the stateful firewall rules created for the service instance.
SFW Rule Sets	Name of the stateful firewall rule sets created for the service template.
NAT Pools	Name of the CGNAT pool created for the service template.

*Table 82: Services – CGNAT and SFW Page (continued)*

Field	Description
NAT Rules	Name of the CGNAT rules created for the service instance.
NAT Rule Sets	Name of the CGNAT rule sets created for the service template.
Syslogs	Name of the syslog created for the service template.
Deployment Plans	Name of the deployment plan with which the service template is attached.

7. From the Term Name drop-down list, select the rule term that must be assigned to the service policy filter, such as CGNAT or stateful firewall service policies.

8. From the Host Name drop-down list, select the hostname of the SDG.

The modified association is saved.

You can use the **Actions** menu in the Service Template pages for CGNAT, SFW, and packet filters to publish, unpublish, export, and restore the defined policies or filters. For details, see *Using the Actions Menu in the Service Template Page*.

#### Related Documentation

- [Creating Service Gateway Groups on page 99](#)
- [Managing Service Gateway Groups on page 101](#)
- [Searching Unmanaged Devices on page 104](#)
- [Viewing the List of Discovered, Managed, and Unmanaged Devices on page 106](#)
- [Modifying Discovery Profiles on page 113](#)
- [Deleting Discovery Profiles on page 114](#)



## CHAPTER 28

# Viewing Transactions Associated with Deployment Jobs

- [Transactions Overview on page 551](#)
- [Viewing Transactions on page 552](#)

### Transactions Overview

---

A transaction refers to an operation or a task that is performed on the service definitions, configuration parameters, and policy settings that are created for provisioning on the devices or Service Delivery Gateways (SDGs). When you create a deployment plan to define the services and policy filters that must be applied and propagated on the devices, the administrator can approve or reject a deploy plan. For each approved deploy plan, a transaction is automatically created by the Edge Services Director application.

A transaction contains a unique identifier that denotes each deployment plan associated with it. Such an automated generation of a transaction for each deploy plan enables you to track, monitor, and maintain a comprehensive record or log of events performed on the devices. For example, if you approve a deploy plan and schedule it for transferring configuration to a set of devices, you can use the Transactions page to view the history of all of the deploy plans that were created for different devices. Also, if multiple deploy plans for the same set of devices were created, the list of transactions provides a granular, in-depth account of the operations.

This level of detail and analysis pattern is useful in diagnosis, debugging, and administration of services and device settings. You can also view the configuration that exists on the device before a deploy plan propagates and applies settings, the configuration being transmitted using the deploy plan, and a differential set of the settings that are present on the device and the settings being provisioned using the deploy plan. All of the configuration is displayed in Junos OS XML API format.

The Junos OS command-line interface (CLI) and the Junos OS infrastructure communicate using XML. When you issue an operational mode command in the CLI, the CLI converts the command into XML format for processing. After processing, Junos OS returns the output in the form of an XML document, which the CLI converts back into a readable format for display. Remote client applications also use XML-based data encoding for operational and configuration requests on devices running Junos OS. The Junos XML API is an XML representation of Junos configuration statements and operational mode

commands. It defines an XML equivalent for all statements in the Junos configuration hierarchy and many of the commands that you issue in CLI operational mode. Each operational mode command with a Junos XML counterpart maps to a request tag element and, if necessary, a response tag element.

Multiple transactions are generated for a single deployment plan with different, unique IDs for each transaction, when multiple devices are present in a single deployment plan. With transactions created for each of the devices for which configuration is propagated from Edge Services Director, you can quickly and easily view the status of the deployment plan pertaining to the transaction for diagnosis and rectification of configuration errors and discrepancies in the settings.

A configuration is stored as a hierarchy of configuration statements. In this mode, you enter statements to configure all properties of the device, including interfaces, general routing information, routing protocols, user access, and several system and hardware properties. When you specify configuration parameters on a device, you are actually viewing and changing a file called the candidate configuration. The candidate configuration file enables you to make configuration changes without causing operational changes to the current operating configuration, called the active configuration. The device does not implement the changes you added to the candidate configuration file until you commit them, which activates the configuration on the device. Candidate configurations enable you to alter your configuration without causing potential damage to your current network operations. Running configuration refers to the configuration file currently in effect on the device. The running configuration file is labeled Version 0. Candidate configuration signifies the new, not yet committed, configuration file that becomes the running configuration.

A rollback configuration refers to the previously committed configuration. The configuration set that is being propagated to devices to be rolled back if a failure occurs during the deployment operation. You can return to the most recently configured successful configuration on the device.

**Related Documentation**

- [Viewing Deployment Plans on page 529](#)
- [Creating and Assigning a Deployment Plan to Devices on page 533](#)
- [Viewing Transactions on page 552](#)

---

## Viewing Transactions

A transaction signifies the process or operation that is performed for the settings and configuration definitions that are included to be provisioned on devices in a deployment plan. A one-to-one mapping exists between the deployment plan and a transaction. A unique identifier for each transaction for easy tracking and filtering is automatically assigned by the Edge Services Director application. A transaction applies only for deploy plans that are in the approved state because only approved plans can be scheduled for provisioning configuration on devices.

The Transactions page displays all of the transactions generated by the system for approved deploy plans. You can perform the following tasks on this page:

- Delete a transaction, which causes the transaction to be removed from listing, but does not delete the deployment plan associated with it.
- View the XML API format of configurations that exist on the device, that are to be deployed on the device, or the change-set of configurations between the settings on the device and the settings to be deployed.
- View a list of all transactions

To view the transactions:

1. From the View selector, select **Gateway View** or **Service View**. The workspaces that are applicable to this view are displayed. In Gateway view, the devices in the entire network are displayed, organized by the device types and the device models within each device type. In Service View, the different types of services are displayed in the View pane.
2. From the Junos Space user interface, click the **Deploy** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the View pane, select the All Network item in Gateway view. Click the plus sign (+) beside the All Network item in the View pane to expand the tree and select the device node you want.
4. Select **Deploy Service > Transactions** from the task pane. The Transactions page is displayed.

*Figure 52: Transactions Page*

Transaction ID	Job ID	Timestamp	Deployment Plans	Service Instances	Status	Status Message
655932	1015847	Aug 20, 2015 8:51:15 AM IST	655837	NAPT44-SS1	Success	Successfully pushed configuration.

Host	Service Gateway	Existing Config	Modified Config	Rollback Config	Latest Task	Status	Status Message
mobs480w	mobs480w-mobs480x	<a href="#">View</a>	<a href="#">View</a>	<a href="#">View</a>	Pushing Configuration	ALL Success	Successfully pushed configuration.
mobs480x	mobs480w-mobs480x	<a href="#">View</a>	<a href="#">View</a>	<a href="#">View</a>	Pushing Configuration	ALL Success	Successfully pushed configuration.

The page is divided into two panes. The top pane displays the transactions that have been generated by the system for the different approved deployment plans. The deployment job ID and services to be transmitted to the devices in each transaction

or deploy plan are also displayed. The bottom pane displays extensive information about a particular transaction, such as the devices to which the configuration is to be provisioned and the exact configuration parameters in XML API form.

The following fields are displayed in the top pane of the page:

Field	Description
Transaction ID	Unique identifier assigned by the system for a transaction.
Job ID	Unique identifier assigned by the system for a deployment plan.
Timestamp	The time, day, month, and year at which the deploy plan was created. The timestamp is the UTC time in database that is mapped to the local time zone of client computer.
Deployment Plans	Names of the deploy plans, which are of the <i>DPnn</i> naming format, that are part of the job ID or transaction. The DP part of the name denotes deployment plan and <i>nn</i> signifies the number allotted to the deploy plan.
Service Instances	Names of the service instances that are to be deployed in the transaction.
Status	Current status of the deployment job that is used to provision configuration on a device.
Status Message	Information about the status of the deploy plan that you can use to modify or take appropriate steps for ensuring successful deploy of configurations to devices.

5. Select the check box next to the transaction for which you want to view detailed information. The bottom pane displays the devices and configurations of each of the devices for the selected transaction.

The following fields are displayed in the bottom pane of the page:

Field	Description
Service Gateway	Name of the service delivery gateway for which deployment of settings is being performed using the particular transaction. The SDG name follows the <i>SDG-nn</i> format, where SDG denotes that the device is a service delivery gateway and <i>nn</i> represents the unique number of the SDG that is automatically generated to differentiate each SDG in a transaction.
Host	Host names of the devices in an SDG high availability pair or the standalone SDG device.
Previous Config	Click the <b>View</b> link to open the dialog box that displays the XML form of configuration parameters and statements that are currently running on the device. Click <b>Close</b> after you complete viewing the settings.
Modified Config	Click the <b>View</b> link to open the dialog box that displays the XML form of configuration parameters and statements that are to be deployed to the device. Click <b>Close</b> after you complete viewing the settings.

Field	Description
Rollback Config	Click the <b>View</b> link to open the dialog box that displays the XML form of the previously committed (rollback) configuration parameters and statements on a device. You can revert to the last known good state before the most recent configuration change that was performed on a device. Click <b>Close</b> after you complete viewing the settings.
Latest Task	<p>The operation that is currently running as part of the transaction to propagate configuration settings to devices. You can also view a list of all the events or tasks that were performed using this transaction by clicking the <b>All</b> link.</p> <p>The Transactions Steps dialog box is displayed. A list of all of the tasks, their statuses, and information about each of the statuses is displayed. Tasks indicate whether the configuration is prepared, the currently active configuration parameters on a device have been retrieved, differential set of the configuration settings that need to be deployed, and the propagation of the configuration attributes.</p>
Status	Current status of the transaction related to a deploy job that is used to provision configuration on a device.
Status Message	Information about the status of the transaction that you can use to modify or take appropriate steps for ensuring successful deploy of configurations to devices.

**Related  
Documentation**

- [Viewing Deployment Plans on page 529](#)
- [Creating and Assigning a Deployment Plan to Devices on page 533](#)
- [Transactions Overview on page 551](#)



## PART 8

# Monitor Mode

- [About Monitor Mode on page 559](#)
- [Using Fault Management Monitors on page 561](#)
- [Using Performance Management Utilities on page 571](#)
- [General Monitoring on page 577](#)



## CHAPTER 29

# About Monitor Mode

- [Understanding Monitor Mode in Edge Services Director on page 559](#)

## Understanding Monitor Mode in Edge Services Director

---

Monitor mode in Edge Services Director provides you visibility into your network status and performance. Edge Services 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.

This topic describes:

- [General Monitoring on page 559](#)
- [Packet Analyzer on page 560](#)
- [Fault Management on page 560](#)
- [Performance Management on page 560](#)

### General Monitoring

The main purpose and benefit of monitoring functionalities is to allow the operators to quickly monitor the health (working condition), operating efficiency, traffic-handling capacity, and performance status of the managed SDGs and configured services such as ADC, TLB, CGNAT and SFW. The SDG monitoring mechanism is an extensive and ingrained tool; it allows the operator to understand the network health and status by drilling down to all the components of SDG. The SDG status is marked as Green, Red, Orange or Gray, based on the health, availability, performance and other important KPI indicators. Red denotes an emergency condition, which is a system panic or other conditions that cause the routing platform to stop functioning. It also indicates that the device is offline or turned down. Orange denotes an alert, which can be conditions that must be corrected immediately, such as a corrupted system database. Green indicates a notice, which signifies conditions that are not error conditions but are of interest or might warrant special handling. It can also include a severity level equivalent to informational or debugging messages. Gray signifies an unknown or an unconnected device that is out of synchronization.

## Packet Analyzer

Packet capture is a tool that helps you to analyze network traffic and troubleshoot network problems. The packet capture tool captures real-time data packets traveling over the network for monitoring and logging. This tool is a debugging and analysis utility that you can use to identify the problematic area in a session path. A set of counters are displayed for both forward and reverse flow for all the supported services on SDG devices. Using these statistical details and values, you can obtain adequate and useful estimates regarding the total bytes count for each service in every hop and quickly, easily locate the hop where there can be a possible packet drop.

## Fault Management

The fault management capability in Edge Services Director shows you information about the health of your network and changing conditions of your equipment. Use this diagnosis and detection mechanism to find problems with equipment, pinpoint security attacks, or to analyze trends and categories of errors. This feature includes the fault-monitoring details in the dashboard, monitoring pages, and also in a dedicated page that displays the alarms, events, and system logging messages that are generated. These charts and messages provide essential and cohesive information about system conditions, any discrepancies and malfunctioning, and protocols or components that need to be examined and debugged for better efficiency and working capacity.

## Performance Management

It is important to identify and define the performance manager KPIs that can help the operator to measure the performance and the operational status of the services running in the SDG network. Apart from services, the operator might also be interested in KPIs for the different chassis options, which are enabling the service execution in the SDG network. Different metrics such as real service instance KPIs and HA Network KPIs are collected. Such a collection causes the PM data collector and PM data aggregator to operate effectively. The following are the different sets of KPIs supported.

- Related Documentation**
- [Alarm Severities and States Overview on page 51](#)
  - *Events and Alarms Overview*

# Using Fault Management Monitors

- [Understanding Fault Management on page 561](#)
- [Viewing the Fault Management Details on page 562](#)

## Understanding Fault Management

---

The fault management capability in Edge Services Director shows you information about the health of your network and changing conditions of your equipment. Use this diagnosis and detection mechanism to find problems with equipment, pinpoint security attacks, or to analyze trends and categories of errors. This feature includes the fault-monitoring details in the dashboard, monitoring pages, and also in a dedicated page that displays the alarms, events, and system logging messages that are generated. These charts and messages provide essential and cohesive information about system conditions, any discrepancies and malfunctioning, and protocols or components that need to be examined and debugged for better efficiency and working capacity.

You can assess the performance of your network, not only at a point in time, but also over a period of time. This feature enables you to determine trending and diverse other things; for example, whether Service Level Agreements (SLAs) have been violated.

The fault management data includes SNMP traps and syslogs received from SDGs. Junos Space platform is integrated with OpenNMS, which is a network management application platform that provides solutions for enterprises and carriers, to receive SNMP Traps. Edge Services Director uses OpenNMS for SNMP trap collection and correlation.

A syslog collection mechanism, which is not available in Junos Space platform, is implemented in Edge Services Director. SNMP traps are used primarily for fault management. When the same fault information is available as both an SNMP trap and a syslog, SNMP trap takes priority.

Open NMS receives the following types of SNMP traps

- MX device-level traps
- Service level traps (ADC, TLB, SFW, CG-NAT)
- HA-related SDG traps

Edge Services Director maintains a set of defined SNMP traps that are processed and converted as alarms by using Open NMS infrastructure. The required configuration for

trap correlation (to automatically clear an existing alarms) is performed by the Open NMS infrastructure. The alarms and events details are stored in Open NMS database (postgres database) and transmitted to the Edge Services Director user interface.

- Related Documentation**
- [Alarm Severities and States Overview on page 51](#)
  - [Events and Alarms Overview](#)
  - [Understanding Monitor Mode in Edge Services Director on page 559](#)

---

## Viewing the Fault Management Details

Use the Fault Management page monitor to sort alarms, view an alarm in depth, and to assign a disposition to an alarm.

By clicking the **Fault Management** option in the task pane, you can access the alarms and syslog details as four separate graphs (alarms by Severity, Service, or State, and syslogs by severity are shown). Four quadrants are displayed on this page.

You can click on any one of the four graphs to launch the corresponding page that displays alarm, event, or syslog information. For example, while clicking on Active (green), you are navigated to alarm screen that shows only active alarms. The color-coding legend appears at the top of the window.

There are two categories of alarms: acknowledged and outstanding. Acknowledging an alarm indicates that you have taken responsibility for addressing the corresponding network or systems-related issue. Any alarm that has not been acknowledged is considered outstanding and is therefore visible to all users on the Alarms page, which displays outstanding alarms by default. If an alarm has been acknowledged in error, you can find the alarm and unacknowledge it, making it available for someone else to acknowledge.

When you acknowledge, clear, escalate, or unacknowledge an alarm, this information is displayed in the alarm's detailed view. You can click the alarm ID to view fields such as Acknowledged By, Acknowledgement Type, and Time Acknowledge. These fields display details such as who acknowledged, cleared, escalated, or unacknowledged the alarm; the acknowledgement type (acknowledge, clear, escalate, or unacknowledge); and the date and time the action was performed on the alarm.

- [Viewing Charts of Alarms and Syslogs on page 563](#)
- [Viewing Alarms, Events, and Syslogs on page 563](#)
- [Changing the Alarm State on page 566](#)
- [Searching Alarms on page 566](#)
- [Searching Events on page 567](#)
- [Searching System Log Messages on page 568](#)
- [Acknowledging Alarms on page 568](#)
- [Clearing Alarms on page 569](#)

- [Escalating Alarms on page 569](#)
- [Unacknowledging Alarms on page 570](#)

## Viewing Charts of Alarms and Syslogs

To view the charts of alarms and syslogs:

1. From the Junos Space user interface, click the **Monitor** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
2. Select **Fault Management** from the task pane. The Fault Management page is displayed with the Alarms, Events, and Syslogs tabs.
3. The page displays the alarms by severity, alarms by state, and syslogs by severity pie charts,, and the alarms by service bar chart. Click on any of the charts to navigate to the corresponding page with the appropriate filters configured. For example, if you click the green segment of the pie chart for alarms, the active alarms are sorted and displayed in the Alarms tab of the Faults page.

When you click on the **Alarms by Severity** graph, you are navigated to the Faults screen with following filter criteria:

- SDGs: All
- Severity: The clicked severity (Critical, Major, Minor, Info)

When you click on **Alarms by State** graph, you are navigated to the Faults screen with following filter criteria:

- SDGs: All
- Status: The clicked status (Active, Acknowledged, Re-assign )

When you click on the **Alarms by service** graph , you are navigated to the Faults screen with following filter criteria:

- SDGs: All
- Service: Clicked service

## Viewing Alarms, Events, and Syslogs

You can view alarms, events, and system logging messages that are triggered for various services and system conditions of SDGs. Junos OS generates system log messages (also called syslog messages) to record events that occur on the device for routine operations, failure and error conditions, and critical conditions that might require urgent resolution. Each system log message identifies the Junos OS process that generated the message and briefly describes the operation or error that occurred. Use the topics on this page to configure basic system log capabilities.

You can locate a specific alarm, research the events causing the alarm, and to assign a disposition to the alarm.

To view alarms, events, and syslogs:

1. From the Junos Space user interface, click the **Monitor** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
2. Select **Fault Management > Faults** from the task pane. The Faults page is displayed with the Alarms, Events, and Syslogs tabs.

By default, the All Service Gateways option is selected in the task pane. You can select a particular SDG or SDG pair to view alarms and syslogs for that specified device. When 'All Service Gateways' is selected the data displayed in right pane will be for all the SDGs in the system. When an individual SDG is selected, the data displayed in the right pane will be for the selected SDG. In both cases, the view of data is either at a network level (all SDGs) or for an individual SDG.

The following fields are displayed in the Alarms tab:

*Table 83: Alarms Tab Fields*

Field	Value
Acknowledged	Indicates if the alarm has been acknowledged. Select this check box and choose <b>Acknowledge</b> , <b>Clear</b> , or <b>Escalate</b> from the Actions drop-down list above the table to change the state of an alarm. Click the <b>Go</b> button to change the alarm state.
ID	A system and sequentially-generated identification number.
Severity	The severity of the alarm. Severity levels are: <ul style="list-style-type: none"> <li>• Critical—A critical condition exists; immediate action is necessary.</li> <li>• Major—A major error has occurred; escalate or notify as necessary.</li> <li>• Minor—A minor error has occurred; notify or monitor the condition.</li> <li>• Info—An informational message; no action is necessary.</li> </ul>
Node	Host name of the device or node that generated the alarm.
Count	Number of alarms of a particular severity for a specific log.
Last Event Time	Time and date at which the last event occurred.
Log Message	System logging message that indicates the event generated and its category or application.

The following fields are displayed in the Events tab:

*Table 84: Events Tab Fields*

Field	Value
ID	A system and sequentially-generated identification number.

*Table 84: Events Tab Fields (continued)*

Field	Value
Severity	<p>The severity of the alarm. Severity levels are:</p> <ul style="list-style-type: none"> <li>• Critical—A critical condition exists; immediate action is necessary.</li> <li>• Major—A major error has occurred; escalate or notify as necessary.</li> <li>• Minor—A minor error has occurred; notify or monitor the condition.</li> <li>• Info—An informational message; no action is necessary.</li> </ul>
Node	Host name of the device or node that generated the alarm.
Interface	IP address of the interface that triggered the event.
Last Event Time	Time and date at which the last alarm event occurred.
Log Message	System logging message that indicates the event generated and its category or application.

The following fields are displayed in the Syslogs tab:

*Table 85: Syslogs Tab Fields*

Field	Value
Time Stamp	Date and time at which the system event logging message was recorded.
Device IP	IP address of the node or SDG that recorded the system log.
Severity	<p>The severity of the system log: Severity levels are:</p> <p>any—Includes all severity levels</p> <p>none—Disables logging of the associated facility to a destination</p> <p>emergency—System panic or other condition that causes the router to stop functioning</p> <p>alert—Conditions that require immediate correction, such as a corrupted system database</p> <p>critical—Critical conditions, such as hard errors</p> <p>error—Error conditions that generally have less serious consequences than errors at the emergency, alert, and critical levels</p> <p>warning—Conditions that warrant monitoring</p> <p>notice—Conditions that are not errors but might warrant special handling</p> <p>info—Events or nonerror conditions of interest</p>

Table 85: Syslogs Tab Fields (continued)

Field	Value
Facility Code	<p>Name of the component or module on the device that generated the log.</p> <p>any—All (messages from all facilities)</p> <p>authorization—Authentication and authorization attempts</p> <p>change-log—Changes to the Junos OS configuration</p> <p>conflict-log—Specified configuration is invalid on the router type daemon—Actions performed or errors encountered by system processes</p> <p>dfc—Events related to dynamic flow capture</p> <p>firewall—Packet filtering actions performed by a firewall filter</p> <p>ftp—Actions performed or errors encountered by the FTP process</p> <p>interactive-commands—Commands issued at the Junos OS command-line interface (CLI) prompt or by a client application such as a Junos XML protocol or NETCONF XML client</p> <p>kernel—Actions performed or errors encountered by the Junos OS kernel</p> <p>pfe—Actions performed or errors encountered by the Packet Forwarding Engine</p> <p>user—Actions performed or errors encountered by user-space processes</p>
Service	Name of the service that triggered the log.
Log Message	System logging message that indicates the event generated and its category or application.

## 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 options from the Actions drop-down list in the Alarms tab of the Faults page. These options 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.
- **Escalate**—Use this button to increase the priority and criticality of the alarm for resolution.
- **Assign**—Use this button to assign active or acknowledged alarms to staff.

## Searching Alarms

You can filter alarms based on certain parameters, such as severity or state, to quickly identify and analyze only the alarms that are of importance or relevance. The color-coding legend appears at the top of the window. Click the color legend to display a table in a

separate window showing the full explanations and color coding for the degrees of severity.

To search and filter alarms:

1. From the Junos Space user interface, click the **Monitor** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
2. Select **Fault Management > Faults** from the task pane. The Faults page is displayed with the Alarms, Events, and Syslogs tabs.  
  
By default, the All Service Gateways option is selected in the task pane. You can select a particular SDG or SDG pair to view alarms and syslogs for that specified device. When 'All Service Gateways' is selected the data displayed in right pane will be for all the SDGs in the system. When an individual SDG is selected, the data displayed in the right pane will be for the selected SDG. In both cases, the view of data is either at a network level (all SDGs) or for an individual SDG.
3. Sort the alarms based on the following parameters from the drop-down lists:
  - Severity
  - State
  - Service
  - Time (You can choose only time spans ending now, for example, Last 12 hours.)
4. Click **Search** to filter the alarms and display the alarms based on the search criteria.

## Searching Events

You can filter events based on certain parameters, such as severity or time, to quickly identify and analyze only the events that are of importance or relevance. The color-coding legend appears at the top of the window. Click the color legend to display a table in a separate window showing the full explanations and color coding for the degrees of severity.

To search and filter alarms:

1. From the Junos Space user interface, click the **Monitor** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
2. Select **Fault Management > Faults** from the task pane. The Faults page is displayed with the Alarms, Events, and Syslogs tabs.
3. Select the **Events** tab.
4. Sort the events based on the following parameters from the drop-down lists:

- Event Text
  - Service
  - Time (You can choose only time spans ending now, for example, Last 12 hours.)
5. Click **Search** to filter the events and display the events based on the search criteria.

## Searching System Log Messages

You can filter system logs based on certain parameters, such as facility or severity, to quickly identify and analyze only the logs that are of importance or relevance. The color-coding legend appears at the top of the window. Click the color legend to display a table in a separate window showing the full explanations and color coding for the degrees of severity.

To search and filter logs:

1. From the Junos Space user interface, click the **Monitor** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
2. Select **Fault Management > Faults** from the task pane. The Faults page is displayed with the Alarms, Events, and Syslogs tabs.
3. Select the **Syslogs** tab.
4. Sort the logs based on the following parameters from the drop-down lists:
  - Severity
  - Time From (Start time in the 24-hour time format of generation of logs)
  - Time To (End time in the 24-hour time format of generation of logs)
  - Service
  - Facility
5. Click **Search** to filter the logs based on the search criteria.

## Acknowledging Alarms

To acknowledge an alarm:

1. From the Junos Space user interface, click the **Monitor** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
2. Select **Fault Management > Faults** from the task pane. The Faults page is displayed with the Alarms, Events, and Syslogs tabs.
3. Select the **Alarms** tab.

4. Select the alarm's **Ack** check box. To select all alarms, at the bottom of the page, click **Select All**.
5. At the top of the page, select **Acknowledge** from the Actions drop-down list, and click **Go**.

The alarm is removed from the default view of all users.

## Clearing Alarms

To clear an alarm:

1. From the Junos Space user interface, click the **Monitor** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
2. Select **Fault Management > Faults** from the task pane. The Faults page is displayed with the Alarms, Events, and Syslogs tabs.
3. Select the **Alarms** tab.
4. Select the alarm's **Ack** check box. To select all alarms, at the bottom of the page, click **Select All**.
5. At the top of the page, select **Clear** from the Actions drop-down list, and click **Go**.

The alarm is removed from the default view of all users.

## Escalating Alarms

To escalate an alarm:

1. From the Junos Space user interface, click the **Monitor** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
2. Select **Fault Management > Faults** from the task pane. The Faults page is displayed with the Alarms, Events, and Syslogs tabs.
3. Select the **Alarms** tab.
4. Select the alarm's **Ack** check box. To select all alarms, at the bottom of the page, click **Select All**.
5. At the top of the page, select **Escalate** from the Actions drop-down list, and click **Go**.

The alarm is escalated by one level.

6. Select the alarm's **Ack** check box. To select all alarms, at the bottom of the page, click **Select All**.
7. At the bottom of the page, select **Escalate Alarms** from the list on the left, and click **Go**. The alarm is escalated by one level.

## Unacknowledging Alarms

To unacknowledge an alarm:

1. Display the list of acknowledged alarms by toggling the Search constraint box so that it shows Alarm is acknowledged.
2. Select the **Ack** check box of the alarm you acknowledged in error. To select all alarms, at the bottom of the page, click **Select All**.
3. Select the alarm's **Ack** check box. To select all alarms, at the bottom of the page, click **Select All**.
4. At the top of the page, select **Unacknowledged** from the Actions drop-down list, and click **Go**.

The alarm appears again in the default view of All Alarms.

### Related Documentation

- [Alarm Severities and States Overview on page 51](#)
- [Events and Alarms Overview](#)
- [Understanding Monitor Mode in Edge Services Director on page 559](#)

# Using Performance Management Utilities

- [Performance Management on page 571](#)

## Performance Management

---

An important aspect of any network management system is to monitor, control and plan the network infrastructure. As the operator network increases in size, heterogeneity and complexity, effective management and planning for such network becomes more important. The main challenges in this area include:

- Identifying the data to be collected
- Measurement strategy or Interpreting the collected data
- Publishing the Threshold events in networks
- Presenting the data, which helps in analyzing the networks performance
- [The Need and Benefits of Performance Manager on page 571](#)
- [Performance Manager View After a Context-Switch from the Monitoring Page on page 576](#)

## The Need and Benefits of Performance Manager

Before understanding the Performance Manager capabilities and benefits, an operator's expectation from any performance manager tool might be the following:

- An operator looks towards the performance manager as a source that can provide a holistic insight of the network performance.
- The operator views the performance manager for data that capture the essential network data for planning, optimization and operation.
- Data that can enable to ensure a high-quality network all the time.
- A tool that can identify the performance degradations proactively.

Edge Services 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. Performance Manager is designed to address the aforementioned requirements.

The PM functionality takes care of the following aspects:

The types of features that Performance Manager presents to the operator to enable the operator to ensure service availability, verify or monitor individual services and the service network performance.

The mechanism and manner that Performance Manager uses to collect the different metrics and interpret them to achieve the functionalities.

The following table describes the operator expectations and the corresponding features or screens available in Edge Services Director:

Operator Expectation	Available Feature
Ensure high quality network	<p>Operator can view the near real time health and performance of its SDG Network</p> <p>Operator can Monitor health and performance of individual Services/Service Instances</p> <p>Operator can Monitor health and performance of individual Chassis parameters (for example, packet path through interface and port</p>
Maximize the utilization of network investments	<p>Operator can define event threshold on different Metric's, which generates alarms when a threshold limit is breached</p> <p>Analysis of trend or historical performance for a SDG or its logical and physical components when the real time data is showing some fault or performance degradation</p> <p>View the historical or trend performance for individual Service Instance</p> <p>Compare service instance performance</p>
Improve the efficiency of operations	<p>View the historical or the trend performance for individual hardware components such as traffic following through interface and port</p> <p>Analysis of Faults and Syslog for historical time period where some performance degradation is observed</p> <p>Generate consolidate PM KPI reports periodically for different services and Chassis components for easy information sharing across an operator's organization</p>

**Type of Data Collected** It is important to identify and define the performance manager KPIs that can help the operator to measure the performance and the operational status of the services running in their SDG network. Apart from services, the operator might also be interested in KPIs for the different chassis options, which are enabling the service execution in the SDG network. Different metrics such as real service instance KPIs and HA Network KPIs are collected. Such a collection causes the PM data collector and PM data aggregator to operate effectively. The following are the different sets of KPIs supported.

- SNMPv3 with MD5 and SHA authentication and DES, 3DES, AES 128, AES 192, and AES 256 privacy.
- Pluggable Message Processing Models with implementations for MPv1, MPv2c, and MPv3.
- Pluggable transport mappings. UDP, TCP, and TLS are supported out-of-the-box.

- Synchronous and asynchronous requests.
- Logging based on Log4J
- Row-based efficient asynchronous table retrieval with GETBULK.
- Multi-threading support.
- Retrieval of scalar counters with GET requests.

**Service Instance KPIs**—These KPIs are collected at a service instance level such as adc instance, tlb instances. Some KPIs from this list provide the near real time performance of the service instance. The following are the KPIs for different services supported:

```

ADC KPIs SNMP/DMI
VIP Status SNMP
Real Server Status SNMP
Connection-table count SNMP
CPU status for ADC Control for last 64 seconds SNMP and DMI
CPU status for ADC Data cores [21]for last 64 seconds/NPU SNMP and DMI
CPU status for each ADC Data cores for last 64 seconds/NPU SNMP and DMI
Allocation Failures per NPU SNMP
Allocation Failures per DP SNMP
ADC ms interface status SNMP
ADC Egress Interface SNMP
TLB KPI
TLB Routing Instance Composite next hop Index status SNMP
Real Server Status SNMP
Net-monitored & overall CPU utilization of TLB PIC SNMP and DMI
TLB ms interface status SNMP
TLB Egress Interface SNMP
CGNAT KPI
CPU status SNMP and DMI
Packet drop status [delta for every 3 polling] SNMP
Memory status SNMP
NAT pool status [ Utilization = Ports in use*100/Total configured ports] SNMP
Port Blocks in use If configured SNMP
CGNAT service pic status SNMP
CGNAT - Stateful sync CPU Utilization SNMP
Statefull Firewall KPI
CPU status SNMP and DMI
Packet drop status [delta for every 3 polling] SNMP
Memory status SNMP and DMI
SFW service pic status SNMP

```

**Chassis KPIs**—These KPIs are collected for all the interfaces, ports and other physical components which are used by the SDG service instances to perform their task such as ingress egress on a service pic, inPacket and outPacket on AE interface. Some of the KPIs are defined to provide the near real time performance for the physical component.

**HA Network KPIs**—These KPIs are applicable only when the SDG setup is a HA deployment. These KPIs are the indicators of the HA deployments performance and health. Some of the KPIs provide the near real-time performance of the HA setup (the

master and the backup SDG). The following are the KPIs identified for Master and Backup SDGs:

```

HA Master
SDG status SNMP and DMI
BGP advertising AS path information for GI-PVT  SNMP
VRRP status SNMP
CGNAT Stateful sync status SNMP
CGNAT default route Route status in GI-PVT SNMP
ADC VIP route status in radware routing instance route table SNMP
TLB routing instance default Route status SNMP
HA Backup
SDG status SNMP and DMI
BGP advertising AS path information for GI-PVT SNMP
VRRP status SNMP & DMI
CGNAT Stateful flows HA status SNMP
CGNAT default route Route status in GI-PVT SNMP
ADC VIP route status in radware routing instance route table SNMP
TLB routing instance default Route status SNMP

```

**KPI Threshold Definition**—An operator can configure on which KPIs the thresholding to be enabled. Using the Monitoring Profile, one can associate the thresholding KPIs to any in the network. Monitoring Profile is a predefined list of KPIs and threshold value for each KPI. SDG NM Fault Manager, which is an Open NMS solution integrated with Junos Space defines different UEI (Unique Event Identifiers). These UEIs are correlated by PM to the existing list of KPIs for which threshold definition exists. When there is any breach for any of these threshold KPIs, SDG NM Performance Manager will send an event trap to the Open NMS FM system.

**Throughput KPIs**—The KPIs that measure the throughput for any service instance (that is, ADC, CGNAT, SFW and TLB) running on the SDG are as follows:

```

TLB- throughput on the IRB link where firewall filter exists
ADC- associated ms interface throughput
CGNAT- Ingress Egress on the associated sp interface
SFW- Ingress Egress on the associated sp interface

```

#### Method of Collection of Data

The core capability of a PM system is to develop a robust and scalable collection mechanism. The performance manager tool supports a multiprotocol data collection using DMI and SNMP. The DMI capabilities are extended from the Junos Space Platform and the SNMP Collector is a proprietary implementation. The current collection support in Performance Manager is as follows:

- DMI based collection: SDG network deployed on Juniper's MX series router supports management interaction over a proprietary communication channel called DMI (Device Management Interface). This is an implementation of netconf protocol (RFC 6241). Edge Services Director uses the DMI capabilities from Junos Space Network Application platform and adopts an XML remote procedure call (RPC) collection mechanism. The following are the characteristics of this collector:

Data collection for KPIs marked above for real time monitoring

DMI is an xml rpc based communication which allows the Management Application to run the CLI commands and get a near real time performance data.

DMI based collection would be supported only for the features where the operator is presented with near real time data i.e. Monitor SDG and Monitor SDG Service UI. 4.

SDG NM also has a SNMP based collection available for these KPIs. Thus the data collected over the DMI channel would not be persisted as the trends for these KPIs can be read from the SNMP collected data store.

- **SNMP based Collection:** Performance manager provides a data collection for the entire SDG PM counter over SNMP. The reason for not extending DMI data collection or for implementing our proprietary SNMP collector are:

DMI channel is an SSH channel over which the two host communicate via rpc xml. This can be leveraged for collecting only a small number of counters. When the number of SDG devices in network increase, using an SSH channel is a time-consuming operation.

Build RPC XML messages and parse the reply for a large number of counters on a big network is processor-intensive that degrades the overall the data collection and representation, thereby affecting the real time aspect of monitoring the performance.

Junos Space does not support the SNMP based collection for the Utility MIB object IDs.

Performance Manager provides an SNMP based collection, offering the following advantages to SDG NM:

#### **Method of Measurement of Data**

The PM counters polled over SDG NM SNMP collector are aggregated and presented to the operator. For PM, all the counters are polled by the SDG NM SNMP collector. Data for PM is not sent to the SDG NM by the device. The following are the details on how the data collected by SDG NM SNMP collector is stored across in SDG NM.

#### **Dashboard and Monitoring View**

Few of the counters are required to display the daily performance trend of the SDG and SDG Network in the Dashboard and Monitoring View. These counters are stored in the JunosSpace MySQL database. The retrieval from MySQL is fast because of the small number of KPIs that are refreshed in frequent interval to show the current day performance. The graphical display in Dashboard and Monitoring views are more dynamic with the monitoring rules in place; therefore, the data needs to be stored for querying and application of rules.

#### **Performance Manager View**

Only the trend or past performance data are displayed, which implies a large number of KPIs and a longer duration of data.

#### **Performance Manager Functionalities**

Performance Manager provides the following functionalities:

- The trend data is available for 1 day to last 365 days with 15 minutes granularity (These numbers are proposed considering the current support for SNMP polling on the MX device and the storage capacity available on Junos Space device).
- View Service performance for each SDG.
- Compare different service instance metrics across SDGs.
- View trends for HA metrics for Master and Standby.
- View trends for the Chassis level metrics.
- Switch to FM and Syslog view for any specific time period where any peak in data is observed.
- Switch to PM view from Monitor SDG/Chassis view and Monitor SDG Services view
- View the Top 3 Talkers of the day (Based on the highest number of Threshold Alarms for the day)

There is no separate landing page for PM View. On selection, it launches the PM view as described here. Also the first SDG is selected by default in the navigation tree.

This view is split into three parts. The first pane is a navigation tree showing all the deployed SDGs. Each root node is the name of the SDG and the child nodes are the Chassis, HA (this node will appear only for SDG which are deployed as HA) and the installed services ADC, TLB, CGNAT and SFW. Each service is a root node for the service instances. The middle pane is the Graphing area which displays the trend graphs for the selected KPIs. For SDG deployed as HA there is small change in the view on selection of the Chassis and HA node in the navigation tree. In this case for Chassis and HA this section would be represented as tabbed view representing the Master and Standby's. For all other selection in the navigation tree it will be shown as single tab. Other aspects of this pane are View Service Instance Metrics, View Chassis Metrics and Compare Metrics. The last section is the KPI view which lists the KPIs for the select node in the navigation tree in the first pane. It presents different actions on the selected KPIs like Graph, Graph All, Select All, and Select None. These actions are further described below for individual views.

## Performance Manager View After a Context-Switch from the Monitoring Page

If you perform context switch to PM view from the Monitoring view or types in an SDG component by using the search utility, the Search text box displays the component selected in the Monitoring View before doing the context switch.

The navigation tree is filtered based on the search criteria. With a context switch, the navigation tree is filtered to display only the selected component for which the context switch happened and the same node is selected.

### Related Documentation

- [Alarm Severities and States Overview on page 51](#)
- [Events and Alarms Overview](#)
- [Understanding Monitor Mode in Edge Services Director on page 559](#)

# General Monitoring

- [Monitoring Capabilities Overview on page 577](#)
- [Viewing the Monitoring Page in Gateway View on page 578](#)
- [Viewing the ADC Service Details on page 583](#)
- [Viewing the TLB Service Details on page 585](#)
- [Viewing the CGNAT Service Details on page 588](#)
- [Viewing the SFW Service Details on page 591](#)

## Monitoring Capabilities Overview

---

The main purpose and benefit of monitoring functionalities is to allow the operators to quickly monitor the health (working condition), operating efficiency, traffic-handling capacity, and performance status of the managed SDGs and configured services such as ADC, TLB, CGNAT and SFW. The SDG monitoring mechanism is an extensive and ingrained tool; it allows the operator to understand the network health and status by drilling down to all the components of SDG. The SDG status is marked as Green, Red, Orange or Gray, based on the health, availability, performance and other important KPI indicators. Red denotes an emergency condition, which is a system panic or other conditions that cause the routing platform to stop functioning. It also indicates that the device is offline or turned down. Orange denotes an alert, which can be conditions that must be corrected immediately, such as a corrupted system database. Green indicates a notice, which signifies conditions that are not error conditions but are of interest or might warrant special handling. It can also include a severity level equivalent to informational or debugging messages. Gray signifies an unknown or an unconnected device that is out of synchronization.

Consider a sample scenario in which the SDG is marked as Red, which might be for various reasons. An operator that monitors using the Edge Services Director application with the Dashboard page learns that the device has been colored as red owing to the CGNAT feature (operator views the Service Status in the dashboard and it shows that CGNAT is marked red for this feature, which causes the SDG to be color-marked as red). The operator can quickly navigate to CGNAT from Dashboard using the Context Link to CGNAT. The Monitoring page provides the information that the SDG is marked as red because of the packet status drop. Therefore, the monitoring functionality helps the operator to quickly identify the issue. Edge Services Director simplifies and eases the complexity involved in monitoring the health and status of SDGs deployed across networks

through following components and visual representation. The following is the list of components available as part of the Monitoring page for which you can analyze and diagnose the device problems and associated service errors.

- Monitoring SDG
- Chassis View
- ADC
- TLB
- CGNAT
- SFW

When you log in to the Edge Services Director interface and navigate to the Monitoring page under Monitor mode, the Monitoring SDG is the landing page showing the details of a SDG. This view is split into two parts the left side is a navigation tree showing all the deployed SDGs. Each root node is the name of the SDG and the child nodes are the installed services ADC, TLB, CGNAT and SFW. Each service is a root node for the service instances.

This view is split into two parts, and the left side is a navigation tree showing all the deployed SDGs. Each root node is the name of the SDG and the child nodes are the installed services ADC, TLB, CGNAT and SFW. Each service is a root node for the service instances. You can view the health and status of the master and standby SDGs in a redundancy or high-availability pair.

The SDG and Alarms monitors are displayed at the upper half of the page.

The Statistics and Chassis View tabs are displayed at the lower half of the right pane of the Monitoring page. Under the Statistics tab, the alarms, high availability KPIs and switchover, status of services, and critical messages that are generated are displayed. The graphs are line graphs for the parameters shown in a pictorial way. Under the Chassis View tab, the hardware and line module details are displayed with a pictorial view of the slots of the SDG devices and the modules installed in these slots.

#### **Related Documentation**

- [Viewing the Monitoring Page in Gateway View on page 578](#)

---

## **Viewing the Monitoring Page in Gateway View**

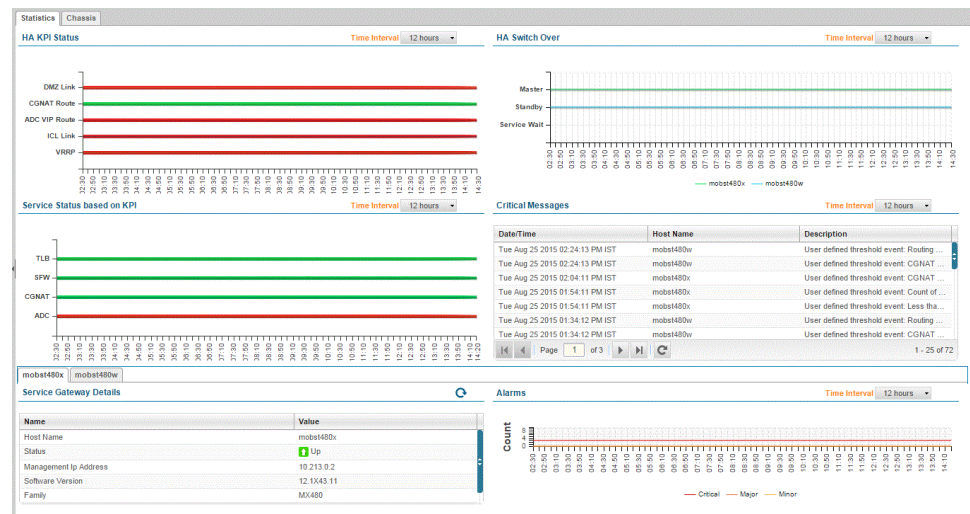
When you log in to the Edge Services Director interface and navigate to the Monitoring page under Monitor mode, the Monitoring SDG is the landing page showing the details of a SDG. This view is split into two parts the left side is a navigation tree showing all the deployed SDGs. Each root node is the name of the SDG and the child nodes are the installed services ADC, TLB, CGNAT and SFW. Each service is a root node for the service instances.

The Monitoring page is refreshed automatically every 3 minutes. Static polling occurs to obtain and display data, and asynchronous collection is not used.

To view the SDGs and associated services to examine, analyze, and troubleshoot device problems and service failures

1. From the View selector, select **Gateway View**. The workspaces that are available in this view are displayed.
2. From the Junos Space user interface, click the **Monitor** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the View pane, select the All Network item in Gateway view. Click the plus sign to expand the tree and view the SDG group or the SDGs in a high availability SDG group or SDG pair.
4. Drill down the tree and select the SDG or SDG pair for which you want to view monitoring statistics and charts of various system states. The Monitoring page is displayed.

**Figure 53: Monitoring Page**



The page is divided into two panes. The left pane displays a tree structure of the SDG pairs or SDGs configured. For SDG pairs, you can view the master and standby device information; else, the health and performance of the separate, individual SDG is shown. With an SDG pair selected, the right pane is refreshed to show the corresponding master and standby device details.

From Service View in Monitor mode, you can drill down the tree by clicking the plus sign (+) from the task pane to view the list of services, such as ADC or TLB. Drill down further to view the service template instances configured for the particular SDG or SDG pair. When you select a service instance, the right pane is refreshed to show the corresponding service details.

You can search for a particular SDG name or service instance name by using the search filter displayed at the top of the left pane.

## Viewing Device Details and Statistical Information

When you select an SDG pair on the left pane of the Monitoring page, the following monitors or quadrants are displayed on the right pane:

The Master and Standby tabs display information about the primary or master, and standby or secondary SDGs in an SDG pair. The Service Wait tab is displayed if the standby device is not fully active after a switchover.

The **Service Gateway Details** monitor provides basic information about the device, such as the running Junos OS version and release number, the management IP address, connection state of the device, hostname, the platform or model number, and the time from which the device has been up.

The **Alarms** monitor displays the SDG alarms for last 6 hours, 12 hours and 24 hours as line graphs. From the Time Interval drop-down list, select the **6 hours**, **12 hours**, or **24 hours** options to view details for the last 6 hours, last 12 hours, and last 24 hours respectively. Click the **Refresh** icon at the top of the monitor to update and display the contents of the table. The time is shown on the horizontal axis and the count of alarms is shown on the vertical axis. Red denotes a critical alarm, orange denotes a major alarm, and yellow denotes a minor alarm. Mouse over the points in the line graph to expand and show the number of alarms at a particular time.

The **SDG Availability** monitor contains details for last 6, 12 and 24 hours. This availability representation is drawn based on the Critical Status from the HA. The purpose of this monitor is to show how much time the SDG was available. It is represented with two lines for each node. This display pattern represents if any switchover happened, and it also helps to identify the reason for it. From the Time Interval drop-down list, select the **6 hours**, **12 hours**, or **24 hours** options to view details for the last 6 hours, last 12 hours, and last 24 hours respectively. Click the **Refresh** icon at the top of the monitor to update and display the contents of the table.

The lower part of the page displays the Statistics and Chassis View tabs. Under the Statistics tab, you can view the following:

The **HA KPI Status** monitor displays the KPIs configured for high availability as a line graph. The time period is shown on the x-axis and the KPI parameters are shown on the y-axis. The lines are colored red, yellow, and green based on the KPI definitions for the HA parameters. From the Time Interval drop-down list, select the **6 hours**, **12 hours**, or **24 hours** options to view details for the last 6 hours, last 12 hours, and last 24 hours respectively. Click the **Refresh** icon at the top of the monitor to update and display the contents of the table.

The **SDG Service Status** monitor displays a line graph with the SDG service status is based on the KPIs defined for a particular service. The data collected is only for the master. Assume a scenario in which a SDG switchover happened in the last 4 hours, then the present master shows status only from the fourth hour to the selected scale. From the Time Interval drop-down list, select the **6 hours**, **12 hours**, or **24 hours** options to view

details for the last 6 hours, last 12 hours, and last 24 hours respectively. Click the **Refresh** icon at the top of the monitor to update and display the contents of the table.

If the SDGs status is green, then all the services are in green. The other possibilities are red and orange. If the selected SDG is red, it might be because of service-related problems, such as 30% real servers being down in a particular ADC instance. If a certain service is not defined on an SDG, the line corresponding to it in the graph is not color-coded.

SDG availability represents the hostname of the current master. At fourth hour, if a switchover occurs, the service status shows the details from the 4th hour because this particular device was standby earlier. After the standby device transitions to be the master, the data for monitoring the service status is restarted to be collected.

The **Critical Messages** monitor displays the messages of a severity level of critical. The date and time at which the message was generated, and a description of the message to highlight the module/protocol and the problem condition are shown. From the Time Interval drop-down list, select the **6 hours**, **12 hours**, or **24 hours** options to view details for the last 6 hours, last 12 hours, and last 24 hours respectively. Click the **Refresh** icon at the top of the monitor to update and display the contents of the table. You can use the paging controls to move to a specific page, to the previous or next page, and to the first or last page.

The **HA Switchover** monitor displays line graphs for the master and standby devices. The master, standby, and service-wait states are shown on the vertical axis and the time period is shown on the horizontal axis. From the Time Interval drop-down list, select the **6 hours**, **12 hours**, or **24 hours** options to view details for the last 6 hours, last 12 hours, and last 24 hours respectively. Click the **Refresh** icon at the top of the monitor to update and display the contents of the table.

You can mouse over the line graph that shows the time intervals when a switchover has occurred to view the switchover reason in a tooltip. The reasons for a failover from the master to the standby device might be power supply failure, a switchover initiated manually, or other causes to maintain high availability in a redundancy group. Edge Services Director obtains the reason for switchover from the SNMP traps received from the router and correlates the information with the HA Status KPI in the Edge Services Director database. Knowing the switchover reason enables quick identification of the underlying problem and correcting the failure.

SDG throughput is one of the KPIs to determine exactly how much throughput or capacity used through various services. This is one of the important KPIs to mark the status of SDGs as red, green and orange. The legends are TLB, ADC, SFW and CGNAT.

SDG CPU usage across the configured services for last 6 hours, 12 hours, and 24 hours is also shown. The legends are TLB, ADC, SFW and CGNAT.

## Viewing the Chassis Image of Devices

Chassis View is enables the operator to see the details at hardware and software level in a symbolic manner. This view changes according to the type of chassis or device, and the display is modified correspondingly. In this view, the SDG is shown with the Master and Standby devices. This view helps the operator to know the health and status of a

particular SDG deployed in the network. It provides a cohesive tool for the operator to quickly see the micro-level information and take any further corrective measure required. Consider a case in which an operator has deployed n number of SDGs in the network. If the operator observes that a particular SDG status and health is not in a satisfied or fully-operational state. The operator quickly navigates to the Monitoring SDG View and selects the particular SDG to see that in one of the slots, the MS-DPC is marked as red. Clicking on the slot indicates that for the slot, configured service ADC daemon is down. In such a case, the restoration measure can be taken. In this view you can quickly view the SDG Chassis details, hardware details, interface Details, important KPIs (CPU and Memory), and the standby details.

When you select an SDG pair on the left pane of the Monitoring page, the following monitors or quadrants are displayed on the right pane:

The Master and Standby tabs display information about the primary or master, and standby or secondary SDGs in an SDG pair.

The lower part of the page displays the Statistics and Chassis View tabs. Under the Chassis View tab, you can view the following:

The **Alarms** field displays the number of alarms that are critical, major, and minor. Red denotes critical, orange denotes major, and yellow denotes minor alarms.

The master and standby chassis are represented pictorially as an entire device, with the slot numbers on the chassis and the types of modules, such as FPC, PIC, or MIC, installed in each slot. If a particular slot is having trouble then it is colored appropriately. If a card is not installed in the slot, it shows an empty slot.

The **Hardware Details** table displays the chassis properties such as the model number of the module installed in the corresponding slot, the serial number, the amount of dynamic random access memory (DRAM), the percentage of time that the CPU uses on background processes, the percentage of time for which the CPU is idle, and the temperature of the air passing by the FPC, in degrees Celsius or in both Celsius and Fahrenheit. Click the information icon to view the Hardware Details table as a pop-up dialog box. The **PIC Details** table displays any attributes that relate to the interface of that particular slot, such as the slot number, the state of the line card in the slot, the total percentage of CPU being used by the FPC's processor, the percentage of the total CPU that is used for interrupts, the percentage of heap space (dynamic memory) being used by the FPC's processor, and the percentage of buffer space being used by the FPC's processor for buffering internal messages. The **Service Details** table displays the names of different services, such as CGNAT, stateful firewall, ADC, and TLB, configured for the router chassis.

In Edge Services Director Release 1.0, the services such as CGNAT, SFW, or TLB, configured for aggregated multiservices (AMS) interfaces did not account for the services PIC (sp-interfaces) that were part of the AMS bundle. Starting with Release 1.1, you can view the service types and service names configured on AMS interfaces, which also accounts for the member interfaces of the AMS bundle.

Each FPC is shown as a bar, with the different PICs installed on the FPC slots displayed as segments within the bar. In the **PIC Details** table, the following performance monitoring

KPIs are available to measure the performance and the operational status of MS-MPCs (these KPIs are applicable only for the slots in which MS-MPCs are installed and not for other chassis slots):

- The **Heap Utilization** column displays the percentage of memory region used for microkernel or heap memory, out of the total CPU memory being used by the Routing Engine or FPC processor. The microkernel memory is generic across the different types of line cards and signifies the heap memory buffers. Because a line card or an FPC in a particular slot can contain multiple Packet Forwarding Engine complexes, the memory utilized on the application-specific integrated circuits (ASICs) are specific to a particular Packet Forwarding Engine complex.
- The **Buffer Utilization** column displays the percentage of buffer memory space being used by the Routing Engine or FPC processor for buffering internal messages.
- The **CPU Utilization** column displays the percentage of CPU memory used by the Routing Engine or FPC processor. In the chassis view, each slot also represents the number of supported PICs on the service interface. You can select the PIC or the MPC/DPC and the corresponding real-time KPIs are displayed.
- The **CPU Interrupts** column displays the percentage of CPU memory used for interrupts, out of the total CPU memory being used by the Routing Engine or FPC processor.

In Service View, click the **Monitor** icon in the Edge Services Director banner to view the workspaces available in Monitor mode. You can select the service type that you need from the View pane, and the SDG or high-availability pair of SDGs from the task pane to perform the following tasks:

- [Viewing the ADC Service Details on page 583](#)
- [Viewing the TLB Service Details on page 585](#)
- [Viewing the CGNAT Service Details on page 588](#)
- [Viewing the SFW Service Details on page 591](#)

#### Related Documentation

- [Monitoring Capabilities Overview on page 577](#)

## Viewing the ADC Service Details

Monitoring ADC is to view all the ADC-related health and status. For example, assume that operator is seeing the SDG is marked red and upon further investigation by using the SDG Monitoring page, it is detected that greater than 30% of the real servers are down, based on this criteria, the ADC is marked as red. The operator selects the particular SDG and drills down to view the service details.

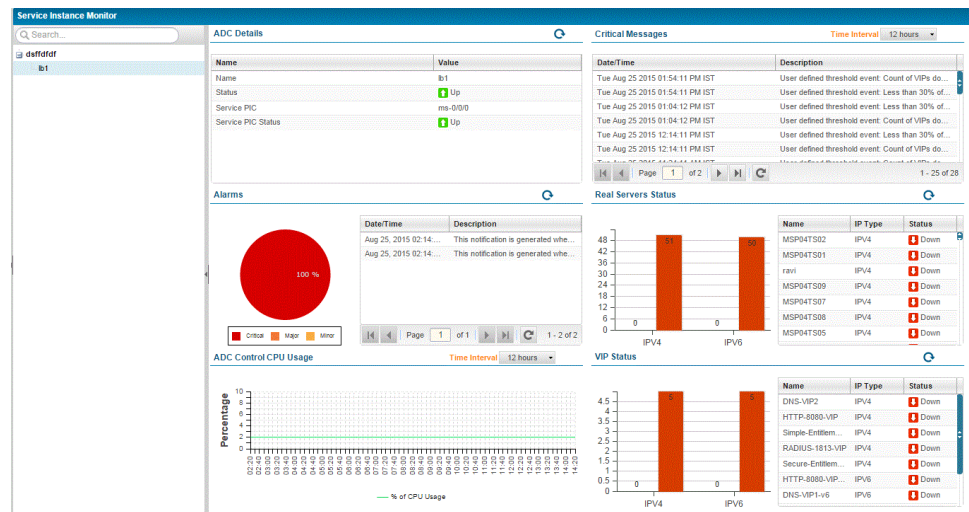
When you select an SDG pair on the left pane of the Monitoring page, you can drill down to select a specific service instance or template of a particular type. You can expand the ADC tree to view the configured ADC instances.

When you select a specific ADC instance, the right pane refreshes to show the following information:

To view the ADC services to examine, analyze, and troubleshoot device problems and service failures

1. From the View selector, select **Service View**. The workspaces that are available in this view are displayed.
2. From the Junos Space user interface, click the **Monitor** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the View pane, select the All Services item. Click the plus sign to expand the tree and select the service type, such as ADC, TLB, or SFW, for which you want to view monitoring details.
4. From the task pane, under Service Instance, click the plus sign to drill down the tree and select the SDG or SDG pair for which you want to view monitoring statistics and charts of various system states. The Monitoring page for ADC is displayed.

Figure 54: Monitoring Page for ADC Service



The **ADC Details** monitor displays the following details:

- Name of the configured service.
- Indicates whether the service is up or down.
- Name of the interface.
- Indicates whether the PIC is up or down.
- Status of the control daemon.
- Status of the data daemon.

The **Critical Messages** monitor displays the messages of a severity level of critical. The date and time at which the message was generated, and a description of the message to highlight the module/protocol and the problem condition are shown. From the Time Interval drop-down list, select the **6 hours**, **12 hours**, or **24 hours** options to view details for the last 6 hours, last 12 hours, and last 24 hours respectively. Click the **Refresh** icon at the top of the monitor to update and display the contents of the table.

The **Alarms** monitor displays a pie chart for critical, major, and minor alarms. Red denotes critical, orange denotes major, and yellow denotes minor alarms. The time at which the alarm is generated and a description of each alarm are shown. Mouse over each portion of the pie to view the number corresponding to the percentage of each alarm severity. From the Time Interval drop-down list, select the **6 hours**, **12 hours**, or **24 hours** options to view details for the last 6 hours, last 12 hours, and last 24 hours respectively. Click the **Refresh** icon at the top of the monitor to update and display the contents of the table.

The **CPU and Memory Usage** monitor displays a line chart with time along the x-axis and the percentage along the y-axis. The legends reference CPU usage and memory usage. From the Time Interval drop-down list, select the **6 hours**, **12 hours**, or **24 hours** options to view details for the last 6 hours, last 12 hours, and last 24 hours respectively. Click the **Refresh** icon at the top of the monitor to update and display the contents of the table.

The **Real Server IPv6 and IPv4 Status** monitor displays a bar graph with percentage along the y-axis and the protocol type along the x-axis. One bar is for IPv4 and the other is for IPv6. Red portion of the bar denotes the server is down. Green portion of the bar denotes the server is up. The real server names and statuses are also shown.

The **VIP Status** monitor displays a pie chart. Each segment of the pie represents the percentage of virtual services that are down or up. Red signifies that the virtual server is down and green signifies that the virtual server is up. The virtual server names and statuses are also displayed. From the Time Interval drop-down list, select the **6 hours**, **12 hours**, or **24 hours** options to view details for the last 6 hours, last 12 hours, and last 24 hours respectively. Click the **Refresh** icon at the top of the monitor to update and display the contents of the table.

#### Related Documentation

- [Monitoring Capabilities Overview on page 577](#)
- [Viewing the Monitoring Page in Gateway View on page 578](#)
- [Viewing the TLB Service Details on page 585](#)
- [Viewing the CGNAT Service Details on page 588](#)
- [Viewing the SFW Service Details on page 591](#)

## Viewing the TLB Service Details

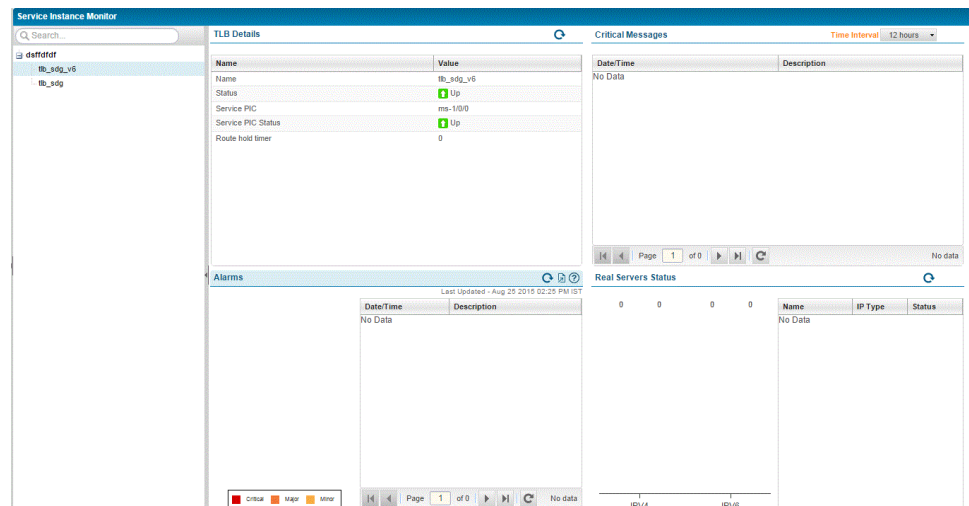
Monitoring TLB is to view all the TLB-related health and status. For example, assume that operator is seeing the SDG is marked red and upon further investigation by using the SDG Monitoring page, it is detected that greater than 30% of the real servers are down, based on this criteria, the TLB is marked as red. The operator selects the particular SDG and drills down to view the service details.

When you select an SDG pair on the left pane of the Monitoring page, you can drill down to select a specific service instance or template of a particular type. You can expand the TLB tree to view the configured TLB instances.

To view the TLB services to examine, analyze, and troubleshoot device problems and service failures

1. From the View selector, select **Service View**. The workspaces that are available in this view are displayed.
2. From the Junos Space user interface, click the **Monitor** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the View pane, select the All Services item. Click the plus sign to expand the tree and select the service type, such as ADC, TLB, or SFW, for which you want to view monitoring details.
4. From the task pane, under Service Instance, click the plus sign to drill down the tree and select the SDG or SDG pair for which you want to view monitoring statistics and charts of various system states. The Monitoring page is displayed.

*Figure 55: Monitoring Page for TLB Service*



When you select a specific TLB instance, the right pane refreshes to show the following information:

The **TLB Details** monitor displays the following details:

- Name of the configured service.
- Indicates whether the service is up or down.
- Name of a multiservices interface.

- Indicates whether the MS-PIC is up or down.
- Value of the routing hold timer.
- Name of the routing instance.
- Name of the virtual service.

The **Critical Messages** monitor displays the messages of a severity level of critical. The date and time at which the message was generated, and a description of the message to highlight the module/protocol and the problem condition are shown. From the Time Interval drop-down list, select the **6 hours**, **12 hours**, or **24 hours** options to view details for the last 6 hours, last 12 hours, and last 24 hours respectively. Click the **Refresh** icon at the top of the monitor to update and display the contents of the table.

The **Alarms** monitor displays a pie chart for critical, major, and minor alarms. Red denotes critical, orange denotes major, and yellow denotes minor alarms. The time at which the alarm is generated and a description of each alarm are shown. Mouse over each portion of the pie to view the number corresponding to the percentage of each alarm severity.

The **CPU and Memory Usage** monitor displays a line chart with time along the x-axis and the percentage along the y-axis. The legends reference CPU usage and memory usage. From the Time Interval drop-down list, select the **6 hours**, **12 hours**, or **24 hours** options to view details for the last 6 hours, last 12 hours, and last 24 hours respectively. Click the **Refresh** icon at the top of the monitor to update and display the contents of the table.

The **Real Server IPv6 and IPv4 Status** monitor displays a bar graph with percentage along the y-axis and the protocol type along the x-axis. One bar is for IPv4 and the other is for IPv6. Red portion of the bar denotes the server is down. Green portion of the bar denotes the server is up. The real server names and statuses are also shown. From the Time Interval drop-down list, select the **6 hours**, **12 hours**, or **24 hours** options to view details for the last 6 hours, last 12 hours, and last 24 hours respectively. Click the **Refresh** icon at the top of the monitor to update and display the contents of the table.

The **Real Service Packet Flow—Trend Statistics** monitor displays a line graph with time along the x-axis and count along the y-axis. The legends reference the total number of packets processed in the forward direction and in the reverse direction by the real service. From the Time Interval drop-down list, select the **6 hours**, **12 hours**, or **24 hours** options to view details for the last 6 hours, last 12 hours, and last 24 hours respectively. Click the **Refresh** icon at the top of the monitor to update and display the contents of the table.

The **Virtual Service Packet Flow—Trend Statistics** monitor displays a line graph with time along the x-axis and count along the y-axis. The legends reference the total number of packets processed in the forward direction and in the reverse direction by the virtual service. From the Time Interval drop-down list, select the **6 hours**, **12 hours**, or **24 hours** options to view details for the last 6 hours, last 12 hours, and last 24 hours respectively. Click the **Refresh** icon at the top of the monitor to update and display the contents of the table.

The **Virtual Server Status—Real Time Statistics** monitor displays a pie chart for the percentage of virtual servers that are in the up and down states. Mouse over each portion of the pie to view the number corresponding to the percentage of each virtual server

status. From the Time Interval drop-down list, select the **6 hours**, **12 hours**, or **24 hours** options to view details for the last 6 hours, last 12 hours, and last 24 hours respectively. Click the **Refresh** icon at the top of the monitor to update and display the contents of the table.

The **Server Group Status—Real Time Statistics** monitor displays a pie chart for the percentage of server groups that are in the up and down administrative statuses. Mouse over each portion of the pie to view the number corresponding to the percentage of each server group. From the Time Interval drop-down list, select the **6 hours**, **12 hours**, or **24 hours** options to view details for the last 6 hours, last 12 hours, and last 24 hours respectively. Click the **Refresh** icon at the top of the monitor to update and display the contents of the table.

**Related  
Documentation**

- [Monitoring Capabilities Overview on page 577](#)
- [Viewing the Monitoring Page in Gateway View on page 578](#)
- [Viewing the ADC Service Details on page 583](#)
- [Viewing the CGNAT Service Details on page 588](#)
- [Viewing the SFW Service Details on page 591](#)

---

## Viewing the CGNAT Service Details

---

Monitoring CGNAT is to view all the CGNAT-related health and status. For example, assume that operator is seeing the SDG is marked red and upon further investigation by using the SDG Monitoring page, it is detected that the CPU utilization is more than 80%, and based on this criteria, the CGNAT is marked as red. The operator selects the particular SDG and drills down to view the service details.

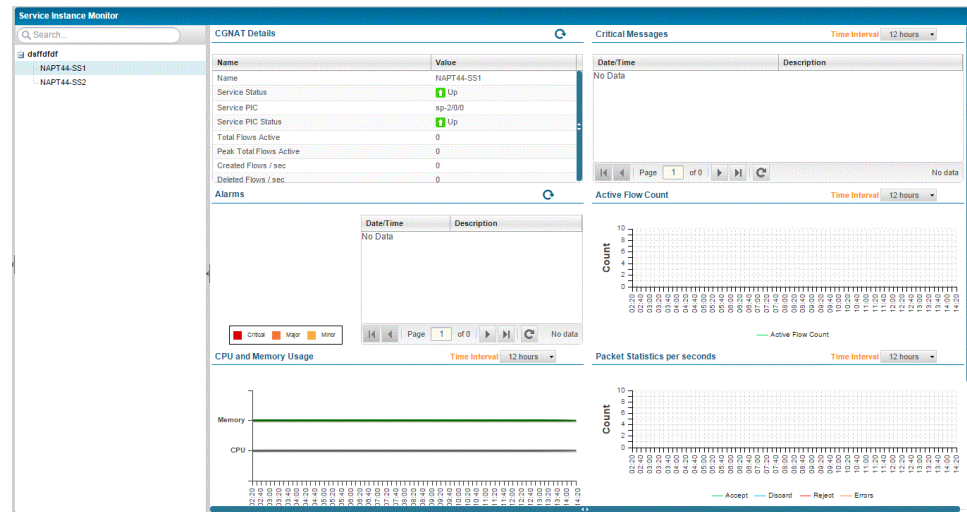
When you select an SDG pair on the left pane of the Monitoring page, you can drill down to select a specific service instance or template of a particular type. You can expand the CGNAT tree to view the configured CGNAT instances.

To view the CGNAT services to examine, analyze, and troubleshoot device problems and service failures

1. From the View selector, select **Service View**. The workspaces that are available in this view are displayed.
2. From the Junos Space user interface, click the **Monitor** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.

- From the View pane, select the All Services item. Click the plus sign to expand the tree and select the service type, such as ADC, TLB, or SFW, for which you want to view monitoring details.
- From the task pane, under Service Instance, click the plus sign to drill down the tree and select the SDG or SDG pair for which you want to view monitoring statistics and charts of various system states. The Monitoring page is displayed.

Figure 56: Monitoring Page for CGNAT Service



When you select a specific CGNAT instance, the right pane refreshes to show the following information:

The **CGNAT Details** monitor displays the following details:

- **Name**—Name of the configured service.
- **Service Status**—Indicates whether the service is up or down.
- **Services PIC Name** —Name of an adaptive services interface.
- **PIC Status**—Indicates whether the services PIC is up or down.
- **Total Flows Active**—Total number of flow sessions currently active on the service PIC.
- **Peak Total Flows Active**—Highest number of active flows since the last PIC restart or since the last time flow statistics are flushed.
- **Created Flows per Second**—Number of flows per second that were being created during the lifetime of the service PIC.
- **Deleted Flows per Second**—Number of flows per second that were being deleted during the lifetime of the service PIC.
- **Total Sessions Active**—Total number of low sessions currently active on the service PIC.

- **Peak Total Sessions Active**—Highest number of active sessions since the last PIC restart or since the last time session statistics are flushed.
- **Created Sessions per Second**—Number of sessions per second that were being created during the lifetime of the service PIC.
- **Deleted Sessions per Second**—Number of sessions per second that were being deleted during the lifetime of the service PIC.

The **CGNAT Subscriber Analysis—Trend Statistics** monitor displays the following details:

- **Total Subscribers Active**—Total number of subscribers currently active on the service PIC.
- **Peak Total Subscribers Active**—Highest number of subscribers that were active during the lifetime of the service PIC.

The **CGNAT Subscriber Analysis—Real Time Statistics** monitor displays the following details:

- **Created Subscribers per Second**—Rate at which subscribers are currently being created on the service PIC.
- **Deleted Subscribers per Second**—Rate at which subscribers are currently being deleted on the service PIC.
- **Peak Created Subscribers per Second**—Highest rate at which subscribers were being created during the lifetime of the service PIC.
- **Peak Deleted Subscribers per Second**—Highest rate at which subscribers were being deleted during the lifetime of the service PIC.

The **Critical Messages** monitor displays the messages of a severity level of critical. The date and time at which the message was generated, and a description of the message to highlight the module/protocol and the problem condition are shown. From the Time Interval drop-down list, select the **6 hours**, **12 hours**, or **24 hours** options to view details for the last 6 hours, last 12 hours, and last 24 hours respectively. Click the **Refresh** icon at the top of the monitor to update and display the contents of the table.

The **Alarms** monitor displays a pie chart for critical, major, and minor alarms. Red denotes critical, orange denotes major, and yellow denotes minor alarms. The time at which the alarm is generated and a description of each alarm are shown. Mouse over each portion of the pie to view the number corresponding to the percentage of each alarm severity.

The **CPU and Memory Usage** monitor displays a line chart with time along the x-axis and the percentage along the y-axis. The legends reference CPU usage and memory usage. From the Time Interval drop-down list, select the **6 hours**, **12 hours**, or **24 hours** options to view details for the last 6 hours, last 12 hours, and last 24 hours respectively. Click the **Refresh** icon at the top of the monitor to update and display the contents of the table.

The **Flow Count** monitor displays a line chart with time along the x-axis and the count along the y-axis. The legend references the number of flows. From the Time Interval drop-down list, select the **6 hours**, **12 hours**, or **24 hours** options to view details for the last

6 hours, last 12 hours, and last 24 hours respectively. Click the **Refresh** icon at the top of the monitor to update and display the contents of the table.

The **Packet Statistics** monitor displays a line chart with time along the x-axis and count along the y-axis. Accepted, discarded, rejected, and errored packets are shown. From the Time Interval drop-down list, select the **6 hours**, **12 hours**, or **24 hours** options to view details for the last 6 hours, last 12 hours, and last 24 hours respectively. Click the **Refresh** icon at the top of the monitor to update and display the contents of the table.

The **Session Statistics** monitor displays a line chart with time along the x-axis and the count of NAT sessions along the y-axis. From the Time Interval drop-down list, select the **6 hours**, **12 hours**, or **24 hours** options to view details for the last 6 hours, last 12 hours, and last 24 hours respectively. Click the **Refresh** icon at the top of the monitor to update and display the contents of the table.

The **Pool Utilization** monitor displays a line chart with time along the x-axis and the percentage of NAT pools and ports that are allocated or utilized for users along the y-axis. From the Time Interval drop-down list, select the **6 hours**, **12 hours**, or **24 hours** options to view details for the last 6 hours, last 12 hours, and last 24 hours respectively. Click the **Refresh** icon at the top of the monitor to update and display the contents of the table.

Until Release 1.0, this widget displayed only the percentage of pools allocated along the y axis. The following formula is used to compute the percentage of ports utilized:

$$\text{Port-Utilization \%} = (\text{Port-In-Use} / \text{Port-Available}) * 100$$

The port utilization line chart is displayed only for CGNAT service templates that are based on the Junos OS Release 14.1 and for CGNAT services that are created on SDG 2.0. Otherwise, only the CGNAT pool utilization line chart is displayed.

#### Related Documentation

- [Monitoring Capabilities Overview on page 577](#)
- [Viewing the Monitoring Page in Gateway View on page 578](#)
- [Viewing the ADC Service Details on page 583](#)
- [Viewing the TLB Service Details on page 585](#)
- [Viewing the SFW Service Details on page 591](#)

## Viewing the SFW Service Details

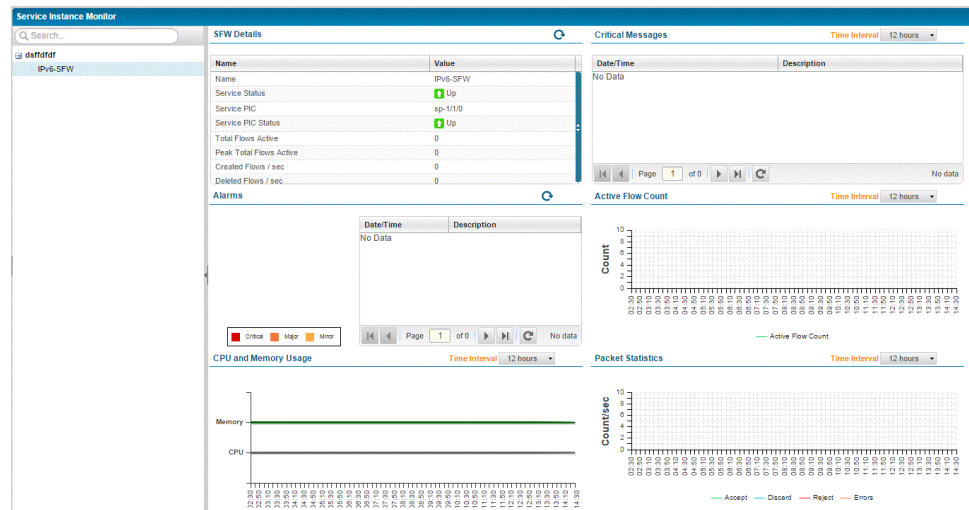
Monitoring stateful firewall is to view all the stateful firewall-related health and status. For example, assume that operator is seeing the SDG is marked red and upon further investigation by using the SDG Monitoring page, it is detected that the packet drop is more than zero and based on this criteria, the stateful firewall is marked as red. The operator selects the particular SDG and drills down to view the service details.

When you select an SDG pair on the left pane of the Monitoring page, you can drill down to select a specific service instance or template of a particular type. You can expand the stateful firewall tree to view the configured stateful firewall instances.

To view the SFW services to examine, analyze, and troubleshoot device problems and service failures

1. From the View selector, select **Service View**. The workspaces that are available in this view are displayed.
2. From the Junos Space user interface, click the **Monitor** icon on the Edge Services Director banner.  
The functionalities that you can configure in this mode are displayed in the task pane.
3. From the View pane, select the All Services item. Click the plus sign to expand the tree and select the service type, such as ADC, TLB, or SFW, for which you want to view monitoring details.
4. From the task pane, under Service Instance, click the plus sign to drill down the tree and select the SDG or SDG pair for which you want to view monitoring statistics and charts of various system states. The Monitoring page is displayed.

**Figure 57: Monitoring Page for SFW Service**



When you select a specific SFW instance, the right pane refreshes to show the following information:

The **SFW Details** monitor displays the following details:

- Name—Name of the configured service.
- Service Status—Indicates whether the service is up or down.
- Services PIC Name —Name of an adaptive services interface.
- PIC Status—Indicates whether the services PIC is up or down.
- Total Flows Active—Total number of flow sessions currently active on the service PIC.

- **Peak Total Flows Active**—Highest number of active flows since the last PIC restart or since the last time flow statistics are flushed.
- **Created Flows per Second**—Number of flows per second that were being created during the lifetime of the service PIC.
- **Deleted Flows per Second**—Number of flows per second that were being deleted during the lifetime of the service PIC.
- **Total Sessions Active**—Total number of low sessions currently active on the service PIC.
- **Peak Total Sessions Active**—Highest number of active sessions since the last PIC restart or since the last time session statistics are flushed.
- **Created Sessions per Second**—Number of sessions per second that were being created during the lifetime of the service PIC.
- **Deleted Sessions per Second**—Number of sessions per second that were being deleted during the lifetime of the service PIC.

The **SFW Subscriber Analysis—Trend Statistics** monitor displays the following details:

- **Total Subscribers Active**—Total number of subscribers currently active on the service PIC.
- **Peak Total Subscribers Active**—Highest number of subscribers that were active during the lifetime of the service PIC.

From the Time Interval drop-down list, select the **6 hours**, **12 hours**, or **24 hours** options to view details for the last 6 hours, last 12 hours, and last 24 hours respectively. Click the **Refresh** icon at the top of the monitor to update and display the contents of the table.

The **SFW Subscriber Analysis—Real Time Statistics** monitor displays the following details:

- **Created Subscribers per Second**—Rate at which subscribers are currently being created on the service PIC.
- **Deleted Subscribers per Second**—Rate at which subscribers are currently being deleted on the service PIC.
- **Peak Created Subscribers per Second**—Highest rate at which subscribers were being created during the lifetime of the service PIC.
- **Peak Deleted Subscribers per Second**—Highest rate at which subscribers were being deleted during the lifetime of the service PIC.

From the Time Interval drop-down list, select the **6 hours**, **12 hours**, or **24 hours** options to view details for the last 6 hours, last 12 hours, and last 24 hours respectively. Click the **Refresh** icon at the top of the monitor to update and display the contents of the table.

The **Critical Messages** monitor displays the messages of a severity level of critical. The date and time at which the message was generated, and a description of the message to highlight the module/protocol and the problem condition are shown. From the Time Interval drop-down list, select the **6 hours**, **12 hours**, or **24 hours** options to view details

for the last 6 hours, last 12 hours, and last 24 hours respectively. Click the **Refresh** icon at the top of the monitor to update and display the contents of the table.

The **Alarms** monitor displays a pie chart for critical, major, and minor alarms. Red denotes critical, orange denotes major, and yellow denotes minor alarms. The time at which the alarm is generated and a description of each alarm are shown. Mouse over each portion of the pie to view the number corresponding to the percentage of each alarm severity.

The **CPU and Memory Usage** monitor displays a line chart with time along the x-axis and the percentage along the y-axis. The legends reference CPU usage and memory usage. From the Time Interval drop-down list, select the **6 hours**, **12 hours**, or **24 hours** options to view details for the last 6 hours, last 12 hours, and last 24 hours respectively. Click the **Refresh** icon at the top of the monitor to update and display the contents of the table.

The **Flow and Subscriber Count** monitor displays a line chart with time along the x-axis and the count along the y-axis. The legends reference the number of flows and number of subscribers. From the Time Interval drop-down list, select the **6 hours**, **12 hours**, or **24 hours** options to view details for the last 6 hours, last 12 hours, and last 24 hours respectively. Click the **Refresh** icon at the top of the monitor to update and display the contents of the table.

The **Packet Statistics** monitor displays a line chart with time along the x-axis and count along the y-axis. Accepted, discarded, rejected, and errored packets are shown. From the Time Interval drop-down list, select the **6 hours**, **12 hours**, or **24 hours** options to view details for the last 6 hours, last 12 hours, and last 24 hours respectively. Click the **Refresh** icon at the top of the monitor to update and display the contents of the table.

#### Related Documentation

- [Monitoring Capabilities Overview on page 577](#)
- [Viewing the Monitoring Page in Gateway View on page 578](#)
- [Viewing the ADC Service Details on page 583](#)
- [Viewing the TLB Service Details on page 585](#)
- [Viewing the CGNAT Service Details on page 588](#)

## PART 9

# Fault Mode

- [About Fault Mode on page 597](#)
- [Viewing and Managing Alarms on page 601](#)
- [Alarm Monitor Reference on page 605](#)



## CHAPTER 33

# About Fault Mode

- [Understanding Fault Mode in Edge Services Director on page 597](#)
- [Understanding the Fault Mode Tasks Pane on page 598](#)

## Understanding Fault Mode in Edge Services Director

---

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? on page 597](#)
- [Alarm Severity on page 598](#)
- [Alarm State on page 598](#)
- [Threshold Alarms on page 598](#)

### 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 Edge Services Director. Edge Services 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 Edge Services 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 Edge Services 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.

**Indeterminate (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.

## 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.

## 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.

---

## Understanding the Fault Mode Tasks Pane

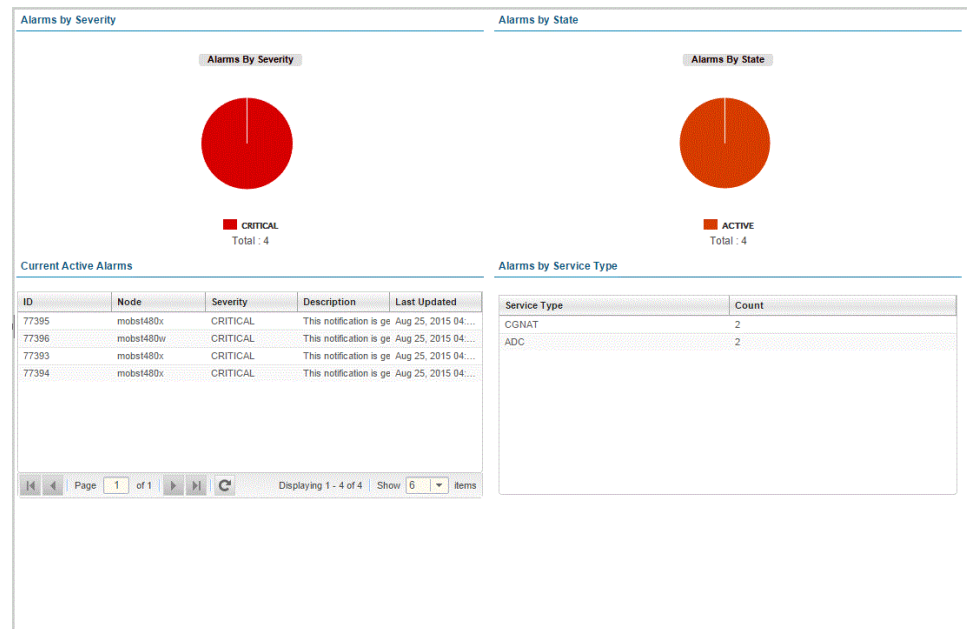
The Tasks pane in Fault mode provides you with a set of tools for effectively managing alarms on your system.

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, Edge Services 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 Edge Services 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.

**Figure 58: Alarms Page in Fault Mode**





## CHAPTER 34

# Viewing and Managing Alarms

- [Changing Alarm State on page 601](#)
- [Searching Alarms on page 601](#)

### 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 Edge Services 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.

### 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 86 on page 602](#).
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.

**Table 86: 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"> <li>• All—Alarms of all states.</li> <li>• Active—Alarms that are current and not yet acknowledged or cleared.</li> <li>• Clear—Alarms that are resolved and the device or entity has returned to normal operation.</li> </ul>
Service Type	<p>Use the list to select the service types for which you want to search the alarms:</p> <ul style="list-style-type: none"> <li>• All—Alarms of all services.</li> <li>• CGNAT—Alarms that are generated for CGNAT services.</li> <li>• SFW—Alarms that are generated for SFW services.</li> <li>• ADC—Alarms that are generated for ADC services.</li> <li>• TLB—Alarms that are generated for TLB services.</li> </ul>
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"> <li>• Critical</li> <li>• Major</li> <li>• Minor</li> <li>• Indeterminate</li> <li>• Warning</li> <li>• Normal</li> </ul>
<b>Advanced Search Criteria</b>	
(from) Date	Pull down the calendar and select the starting date of the search.
(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.

Table 86: Alarm Search Fields (continued)

Search Criteria	Description
(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.



## CHAPTER 35

# Alarm Monitor Reference

- [Alarms by State Monitor on page 605](#)
- [Alarms by Severity Monitor on page 605](#)
- [Current Active Alarms Monitor on page 606](#)
- [Alarms by Service Type Monitor on page 607](#)
- [Alarm Detail Monitor on page 607](#)

### 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.

Changing the state of an alarm using Edge Services 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.

### 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.

## 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. View [Table 87 on page 606](#) for a description of the table.

**Table 87: Current Active Alarms Monitor**

Table Column	Description	Shown in Summary by Default	Shown in Detailed View by Default
ID	A system and sequentially-generated identification number.	No	No
Node	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 device or SDG. You can correlate the alarm with the corresponding device for corrective measures..	No	Yes
Severity	<p>The severity of the alarm. Severity levels are:</p> <ul style="list-style-type: none"> <li>• Critical—A critical condition exists; immediate action is necessary.</li> <li>• Major—A major error has occurred; escalate or notify as necessary.</li> <li>• Minor—A minor error has occurred; notify or monitor the condition.</li> <li>• Indeterminate—An informational message; no action is necessary.</li> </ul>	Yes	Yes
Description	Detailed information about the alarm. This description provides more information about the probable cause or solution for the condition that caused the alarm The description also provides the date and time when the failure was detected. Note the date and time of an alarm so that you can correlate it with error messages or with the messages system log file.	Yes	Yes
Last Updated	The date and time that the information for the alarm was last modified.	Yes	Yes
Counter	Count of the alarm generated. Alarm notification dampening is performed based on the alarm counter.	No	Yes

Table 87: Current Active Alarms Monitor (continued)

Table Column	Description	Shown in Summary by Default	Shown in Detailed View by Default
Log Message	System log message generated for the alarm. Each system log message identifies the Junos OS process that generated the message and briefly describes the operation or error that occurred.	No	Yes
Acknowledged	Indicates if the alarm has been acknowledged.	No	Yes

Clicking the **Details** icon opens Alarm Details where you can sort and disposition alarms by state (Acknowledged, Clear, Active).

## Alarms by Service Type Monitor

The Alarms by Service Type monitor displays a count of alarms for each service type, such as ADC, TLB, stateful firewall, or carrier-grade NAT. The Service Type column displays the type of service for which alarms are generated, and the Count column displays the total number of alarms of different severity levels that are triggered for that particular service.

This monitor is helpful in debugging, diagnosis, and remediation of alarms triggered for the services configured on SDGs or SDG groups. format. The summarized way in which you can view statistical details enables you to examine the health and operating-efficiency of devices, and the performance of services. It provides a bird's eye, high-level view of parameters that enables effective and simplified troubleshooting and administration. You might need to examine the device and services settings to take the required corrective action for performance management to work properly.

## Alarm Detail Monitor

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 on page 608](#)
- [Sorting Alarms on page 609](#)
- [Reading Events on page 609](#)
- [Investigating Event Attributes on page 610](#)
- [Changing the Alarm State on page 610](#)

## 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 609](#) 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 88 on page 608](#).
3. Examine the events and event attributes that contributed to sending the alarm. Events and event attributes are discussed in [“Reading Events” on page 609](#) and [“Investigating Event Attributes” on page 610](#).

**Table 88: Alarm Detail Fields**

Field	Value	Shown in Detailed View by Default
ID	A system and sequentially-generated identification number.	Yes
Node	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 device or SDG. You can correlate the alarm with the corresponding device for corrective measures..	Yes
Severity	The severity of the alarm. Severity levels are: <ul style="list-style-type: none"> <li>• Critical—A critical condition exists; immediate action is necessary.</li> <li>• Major—A major error has occurred; escalate or notify as necessary.</li> <li>• Minor—A minor error has occurred; notify or monitor the condition.</li> <li>• Indeterminate—An informational message; no action is necessary.</li> </ul>	Yes
Description	Detailed information about the alarm. This description provides more information about the probable cause or solution for the condition that caused the alarm. The description also provides the date and time when the failure was detected. Note the date and time of an alarm so that you can correlate it with error messages or with the messages system log file.	Yes
Last Updated	The date and time that the information for the alarm was last modified.	Yes
Counter	Count of the alarm generated. Alarm notification dampening is performed based on the alarm counter.	No
Log Message	System log message generated for the alarm. Each system log message identifies the Junos OS process that generated the message and briefly describes the operation or error that occurred.	Yes

## 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 89 on page 609](#).

**Table 89: Sort Options for Alarms**

Alarms by Severity Sort	Alarms by State and Current Active Alarms Sort
All	Active
Indeterminate	Clear
Minor	
Major	
Critical	

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.

## 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 90 on page 609](#) lists the fields in Event Detail.

**Table 90: 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	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.
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 or controller 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.

## PART 10

# System Mode

- [About System Mode on page 613](#)



# About System Mode

- [Understanding the System Tasks Pane on page 613](#)
- [Audit Logs Overview on page 613](#)

## Understanding the System Tasks Pane

---

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

To access the System Tasks pane, click **System** in the Edge Services Director banner. The tasks are described in [Table 91 on page 613](#).

*Table 91: System Tasks*

Task	Description
View Audit Logs	View a history of user activities on Edge Services 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 Edge Services Director. You can cancel jobs that are in progress or scheduled to run in the future.
Collect Jobs for Troubleshooting	Download a zip file containing logs and troubleshooting data from both Edge Services Director and Junos Space.

## 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, Edge Services 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 Edge Services Director) or a remote network host or media.

## PART 11

# Appendix

- [Services Overview on page 617](#)



## CHAPTER 37

# Services Overview

- [Adaptive Services Overview on page 617](#)
- [Junos Address Aware Network Addressing Overview on page 619](#)
- [Packet Flow Through the Adaptive Services or Multiservices PIC on page 620](#)
- [ADC Overview on page 622](#)
- [Sample IPv6 Transition Scenarios on page 626](#)
- [Understanding Services PICs on page 628](#)
- [TLB Overview on page 631](#)
- [Installing and Configuring TLB Using the CLI Interface on page 634](#)
- [Stateful Firewall Overview for Junos OS Extension-Provider Packages on page 642](#)
- [Network Address Translation Configuration Overview on page 645](#)
- [Junos OS CGNAT Implementation Overview on page 663](#)
- [Service Redundancy Daemon Overview on page 674](#)
- [Configuring the Service Redundancy Daemon on page 676](#)
- [Application Layer Gateways Overview on page 683](#)

## Adaptive Services Overview

---

MultiServices PICs and MultiServices Dense Port Concentrators (MS-DPCs) provide *adaptive services interfaces*, which allow you to coordinate multiple services on a single PIC by configuring a set of services and applications. MultiServices PICs and MS-DPCs offer a special range of services you configure in one or more service sets.

The MultiServices PIC is available in three versions, the MultiServices 100, the MultiServices 400, and the MultiServices 500, which differ in memory size and performance. All versions offer enhanced performance in comparison with AS PICs. MultiServices PICs are supported on M Series and T Series routers except M20 routers.

The MultiServices DPC is available for MX Series routers; it includes a subset of the functionality supported on the MultiServices PIC. Currently the MultiServices DPC supports the following Layer 3 services: stateful firewall, NAT, IDS, IPsec, active flow monitoring, RPM, and generic routing encapsulation (GRE) tunnels (including GRE key and fragmentation); it also supports graceful Routing Engine switchover (GRES) and Dynamic

Application Awareness for Junos OS. For more information about supported packages, see *Enabling Service Packages*.

It is also possible to group several Multiservices PICs into an aggregated Multiservices (AMS) system. An AMS configuration eliminates the need for separate routers within a system. The primary benefit of having an AMS configuration is the ability to support load balancing of traffic across multiple services PICs. Starting with Junos OS 11.4, all MX Series routers will support high availability (HA) and Network Address Translation (NAT) on AMS infrastructure. See *Configuring Load Balancing on AMS Infrastructure* for more information.



**NOTE:** The MultiServices PICs are polling based and not interrupt based; as a result, a high value in the `show chassis pic` “Interrupt load average” field may not mean that the PIC has reached its maximum limit of processing.

The following services are configured within a service set and are available only on adaptive services interfaces:

- Stateful firewall—A type of firewall filter that considers state information derived from previous communications and other applications when evaluating traffic.
- Network Address Translation (NAT)—A security procedure for concealing host addresses on a private network behind a pool of public addresses.
- Intrusion detection service (IDS)—A set of tools for detecting, redirecting, and preventing certain kinds of network attack and intrusion.
- IP Security (IPsec)—A set of tools for configuring manual or dynamic security associations (SAs) for encryption of data traffic.
- Class of service (CoS)—A subset of CoS functionality for services interfaces, limited to DiffServ code point (DSCP) marking and forwarding-class assignment. CoS BA classification is not supported on services interfaces.

The configuration for these services comprises a series of rules that you can arrange in order of precedence as a *rule set*. Each rule follows the structure of a firewall filter, with a **from** statement containing input or match conditions and a **then** statement containing actions to be taken if the match conditions are met.

The following services are also configured on the MultiServices PICs and MS-DPCs, but do not use the rule set definition:

- Layer 2 Tunneling Protocol (L2TP)—A tool for setting up secure tunnels using Point-to-Point Protocol (PPP) encapsulation across Layer 2 networks.
- Link Services Intelligent Queuing (LSQ)—Interfaces that support Junos OS class-of-service (CoS) components, link fragmentation and interleaving (LFI) (FRF.12), Multilink Frame Relay (MLFR) user-to-network interface (UNI) network-to-network interface (NNI) (FRF.16), and Multilink PPP (MLPPP).
- Voice services—A feature that uses the Compressed Real-Time Transport Protocol (CRTP) to enable voice over IP traffic to use low-speed links more effectively.

In addition, Junos OS includes the following tools for configuring services:

- Application protocols definition—Allows you to configure properties of application protocols that are subject to processing by router services, and group the application definitions into application sets.
- Service-set definition—Allows you to configure combinations of directional rules and default settings that control the behavior of each service in the service set.



**NOTE:** Logging of adaptive services interfaces messages to an external server by means of the `fxp0` port is not supported on M Series routers. The architecture does not support system logging traffic out of a management interface. Instead, access to an external server is supported on a Packet Forwarding Engine interface.

#### Related Documentation

- *Understanding Services PICs*
- *Packet Flow Through the Adaptive Services or Multiservices PIC*
- *Enabling Service Packages*
- *Services Configuration Procedure*
- *Supported Platforms*

## Junos Address Aware Network Addressing Overview

In early 2011, the Internet Assigned Numbers Authority (IANA) allocated the last large block of IPv4 addresses. Now service providers and large enterprises, as well as cloud providers, e-tailers, and federal agencies, are evaluating technologies to help them avoid IPv4 address exhaustion and ensure uninterrupted subscriber and service growth.

*Junos Address Aware Network Addressing* is Juniper Networks' portfolio of IPv4 exhaustion avoidance, IPv4-IPv6 coexistence, and IPv6 transition technologies that include IPv6, v4/v6 dual stack, NAT44, NAT44(4), NAPT44, NAPT444, NAT-PT, NAT64, 6-to4-PMT, 6rd, and DS-Lite. These technologies help network operators improve subscriber and service scale, mitigate IPv4 address depletion, and pragmatically transition to IPv6 based on business requirements.

*Junos Address Aware Network Addressing* technologies are available on the following platforms:

- MultiServices Dense Port Concentrator (MS-DPC)
- MS-100, MS-400, and MS-500 MultiServices PICS
- MultiServices Modular Port Concentrator (MS-MPC) and MultiServices Modular Interface Card (MS-MIC)
- Modular Port Concentrator Types 1, 2, and 3 (inline NAT).

**Related Documentation** • [Adaptive Services Overview on page 617](#)

---

## Packet Flow Through the Adaptive Services or Multiservices PIC

---

You can optionally configure service sets to be applied at one of the following three points while the packets transit the router:

- An interface service set applied at the inbound interface.
- A next-hop service set applied at the forwarding table.
- An interface service set applied at the outbound interface.

The packet flow is as follows, graphically displayed in [Figure 59 on page 621](#). (You can configure a service set as either an interface service set or a next-hop service set.)

1. Packets enter the router on the inbound interface.
2. A policer, filter, service filter, service set, postservice filter, and input forwarding-table filter are applied sequentially to the traffic; these are all optional items in the configuration. If an interface service set is applied, the packets are forwarded to the AS or MultiServices PIC for services processing and then sent back to the Packet Forwarding Engine; if a service filter is also applied, only packets matching the service filter are sent to the PIC. The optional postservice filter is applied and postprocessing takes place.
3. A next-hop service set can be applied to the VPN routing and forwarding (VRF) table or to **inet.0**. If it is applied, packets are sent to the PIC for services processing and sent back to the Packet Forwarding Engine.

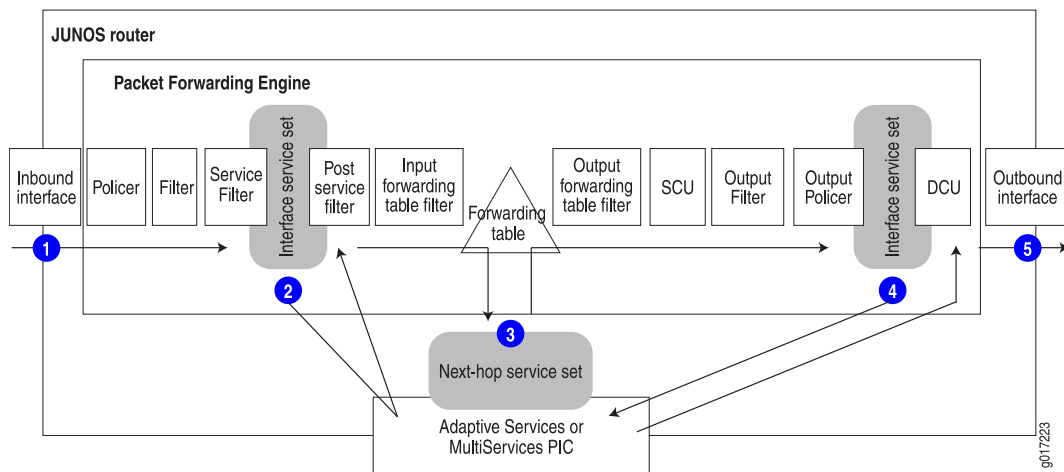


**NOTE:** For NAT, the next-hop service set can only be applied to the VRF table. For all other services, the next-hop service set can be applied to either the VRF table or to **inet.0**.

---

4. On the output interface, an output filter, output policer, and interface service set can be applied sequentially to the traffic if you have configured any of these items. If an interface service set is applied, the traffic is forwarded to the PIC for processing and sent back to the Packet Forwarding Engine, which then forwards the traffic.
5. Packets exit the router.

Figure 59: Packet Flow Through the Adaptive Services or MultiServices PIC



**NOTE:** When an AS PIC experiences persistent back pressure as a result of high traffic volume for 3 seconds, the condition triggers an automatic core dump and reboot of the PIC to help clear the blockage. A system log message at level LOG\_ERR is generated. This mechanism applies to both Layer 2 and Layer 3 service packages.

- Related Documentation**
- *Understanding Services PICs*
  - *Adaptive Services Overview*
  - *Supported Platforms*
  - *Services Configuration Procedure*

## ADC Overview

---

Juniper Networks® Application Delivery Controller (ADC) for the MX Series 3D Universal Edge Router offers advanced router-integrated ADC functions that enables service providers and enterprises to efficiently scale service capacity and increase service performance. Routers are already ubiquitously deployed throughout the network: at the network edge, in the network core, and in the data center. Integrating the advanced ADC with the carrier-grade MX 3D router promotes network consolidation and reduces the number of network elements that providers must rack, power, cool, maintain, and upgrade. Furthermore, the ADC software, which is optionally licensed, improves service resiliency by monitoring server and application health and by automatically bypassing failures.

Server load balancing (SLB) benefits your network in a number of ways:

- Increased efficiency for server use and network bandwidth With SLB, the ADC software is aware of the shared services provided by your server group and can then balance user session traffic among the available servers.
- Important session traffic gets through more easily, reducing user competition for connections on overused servers. For even greater control, traffic is distributed according to a variety of user-selectable methods.
- Increased reliability of services to users If any server in a server group fails, the remaining servers continue to provide access to vital applications and data. The failed server can be brought back up without interrupting access to services. Increased scalability of services As users are added and the server group's capabilities are saturated, new servers can be added to the group transparently.

Server load balancing (SLB) can be the right option for addressing these vital network concerns:

- A single server no longer meets the demand for its particular application.
- When servers hold critical application data and must remain available even in the event of a server failure.
- You want to use multiple servers or hot-standby servers for maximum server uptime.
- You must be able to scale your applications to meet client request capacity.
- You cannot afford to continue using an inferior load-balancing technique, such as DNS roundrobin or a software-only system.

The load-balancing module is used to efficiently deliver content and secure your servers from unauthorized intrusion, probing, and denial-of-service (DoS) attacks. The ADC software includes extensive filtering capabilities at the Layer 2 (MAC), Layer 3 (IP), and Layer 4 (TCP/UDP) levels. Traffic coming from client-facing interfaces is matched against filters. Servers must be connected to server-facing interfaces. The order of the filter term matching process is according to the order the terms appear in the configuration. You can move terms around by using Juniper Networks CLI commands. The order of matching filter terms between adc-instances is according to where the adc-instances appear in the configuration. Matches in one adc-instance are only compared with subsequent

adcinstances in the configuration. Health checking allows you to verify content accessibility in large websites. As content grows and information is distributed across different server farms, flexible, customizable content health checks are critical to ensure end-to-end availability.

## Service Instances

In a subscriber access network, each subscriber has its own set of services. You can configure a specific service instance for a particular subscriber by specifying a service name, also referred to as a service profile, and unique service parameters for that service instance. Service parameters can include a combination of policy lists, filters, rate-limit profiles, class of service (CoS) profiles, and interface profiles.

For example, filter-service (up-filter,down-filter) and filter-service (upstream-filter,downstream-filter) are considered two different instances of the same service (filter-service) because their parameters, enclosed in parentheses after the service name, are different. Each service instance is uniquely identified by the combination of its service name and service parameters. In CoA messages, the router identifies a subscriber service by its complete activation string, which consists of the service name and, if configured, one or more service parameters in the order specified.

An adc-instance is an instance of Application Delivery Software running on one or more Multiservices-DPC interfaces of a Juniper Networks device. An adc-instance includes a complete set of ADC definitions: real-servers, groups of servers, virtual servers using virtual IP addresses, and virtual services accessed by clients.

Using multiple instances on a single device allows you to create completely separate ADCs running on the same machine. Using different instances for different traffic guarantees computation power, guaranteeing no interruption between services. This can be used, for example, to load-balance traffic from different applications, where complete separation is required.

You must specify router interfaces that are bound to an adc-instance.

- **Multiservices interfaces**—The physical multiservices interfaces of a device that are used to run the load-balancing instance application. The more multiservices interfaces used for a loadbalancing instance, the more capacity and processing power the instance has. At least one MS interface must be specified for each adc-instance, up to eight interfaces can run the same instance. A multiservices interface is associated exclusively to a single load-balancing instance (it cannot be shared between instances).
- **Client-facing interfaces**—The device interfaces where client traffic is received. Traffic arriving on these interfaces is handled by the ADC software and destined to be routed to the virtual IP addresses and filter destination addresses configured in the instance. At least one client-facing interface must be specified for each adc-instance. A client-facing interface can be shared between instances.
- **Server-facing interfaces**—The device interfaces where servers are connected, usually through switches or routers. Traffic to the servers is routed to these interfaces. At least one server-facing interface must be specified for each load-balancing instance; a serverfacing interface can be shared between instances. The same device interface

can be used as a client-facing interface in one (or more) adinstances, and as a server-facing interface in other instances.

## Installing and Configuring the ADC Software

You must purchase a suitable license in order to run the ADC software. Each license is for one Multiservices-DPC (two NPUs per license). You should purchase licenses according to the number of Multiservices-DPCs that you have in your device.

As part of the ADC software integration into the Juniper Networks Junos OS system, the ADC software is using an internal configuration. The internal configuration is done in two ways: using a commit script and using the Junos OS internal API.

## Application-Based Health Checks

Application-based health checks include the following:

### SSL Server Health Checks

---

The SSL-Hello health check option on the group configuration allows the ADC to query the health of the Secure Sockets Layer (SSL) servers by sending an SSL client "Hello" packet and then verifying the contents of the server's "Hello" response. The SSL health check is performed using the server listening port configured, under the virtual service configuration, or using the virtual service port when the server listening port is not configured. The following is a summary of the SSL enhanced health check process:

- The ADC sends an SSL "Hello" packet to the SSL server.
- If it is up and running, the SSL server responds with the "Server Hello" message.
- The ADC verifies fields in the response and marks the service "Up" if the fields are OK.

During the handshake, you and the server exchange security certificates, negotiate an encryption and compression method, and establish a session ID for each session.

### DNS Health Checks

---

The ADC software supports both TCP and UDP-based DNS health checking. This health check is performed by sending a DNS query over either protocol and watching for the server reply. The domain name to be queried can be modified using the configuration.

### Ping Health Checks

---

Ping health checks verify if the real server is alive. The Layer 3 echo-reply health check is used for UDP services or when ping health checks are configured. Note: Ping health check is the default health check for a group.

### HTTP Health Checks

---

HTTP-based health checks can include the hostname for Host headers. The Host header and healthcheck URL are constructed from the virtual server hostname, domain name, and sServer group health check field components. If the Host header is required, an HTTP/1.1 GET will result. Otherwise, an HTTP/1.0 GET will result. HTTP health check is

successful if you get a return code of 200. If content is not specified, the health check is performed using the / character.

The following is an example of an HTTP-based health monitor:

```
hostname= everest
domain-name= example.com
http= index.html
Health check is performed using:
GET /index.html HTTP/1.1
Host: everest.example.com
```

### Script-Based Health Checks

Health check scripts dynamically verify application and content availability by executing a sequence of tests based on send and expect commands. Configuring Script-Based Health Checks You can configure the ADC software to send a series of health check requests to real servers or real server groups and monitor the responses. Both ASCII and binary-based scripts, for TCP and UDP protocols, can be used to verify application and content availability. The benefits of using script-based health checks are:

- Ability to send multiple commands.
- Checks for any return ASCII string or binary pattern.
- Tests availability of different applications.
- Tests availability of multiple domains or websites.
- The ADC software supports the following capacity for a single ADC: 6K bytes per script and 64 scripts per load-balancing instance

The commands are grouped together as a list so you can change their order. Each script command is made up of one or more tcp-command or udp-command containers. Commands exist to open a connection to a specific TCP or UDP port, send a request to the server, and expect an ASCII string or binary pattern. The string or pattern configured with an expect (or in the case of binary, binaryexpect) command is searched for in each response packet. If it is not seen anywhere in any response packet before the real-server health-check interval expires, the server does not pass the expect (or binary-expect) step and fails the health check. A script can contain any number of these commands, up to the allowable number of characters that a script supports.



**NOTE:** There is no need to use double slashes when configuring a script that uses special characters with single slashes. For example, the script entry GET /index.html HTTP/1.1\r\nHOST:www.hostname.com\r\n\r\n does not require the use of \\r or \\n to ensure proper functioning of the script. Only one protocol can be configured per script.

### Script Formats

Health check script formats use different commands based on whether the content to be sent is ASCII-based or binary-based. Each script should start with the command open

<protocol port number>,<protocol-name>. The next line can be either a send or expect (for ASCII-based), or bsend or bexpect (binary-based).

#### ***ASCII-Based Health Check***

The general format for TCP-based health-check scripts is as follows:

```
[edit extensions adc adc-instance demo1]
custom-health-check {
  script script1 {
    tcp-commands ASCII {
      command <name> open <port>;
      command <name> send request <text>;
      command <name> expect response <text>;
      command <name> send request <text>;
      command <name> expect response <text>;
      command <name> send request <text>;
      command <name> expect response <text>;
    }
  }
}
```

#### ***Binary-Based UDP Health Check***

The general format for UDP binary-based health check scripts is shown below. Specify the binary content in hexadecimal format. UDP-based health check scripts can use either ASCII strings or binary patterns.

```
[edit extensions adc adc-instance demo1]
custom-health-check {
  script script2 {
    udp-commands Binary-UDP {
      command <name> open <port>;
      command <name> binary-request request 1;
      command <name> binary-expect response 1 {
        offset <count>;
        depth <number of bytes from offset to count>;
      }
    }
  }
}
```

- Related Documentation**
- [Adaptive Services Overview on page 617](#)
  - [Packet Flow Through the Adaptive Services or Multiservices PIC on page 620](#)

---

## Sample IPv6 Transition Scenarios

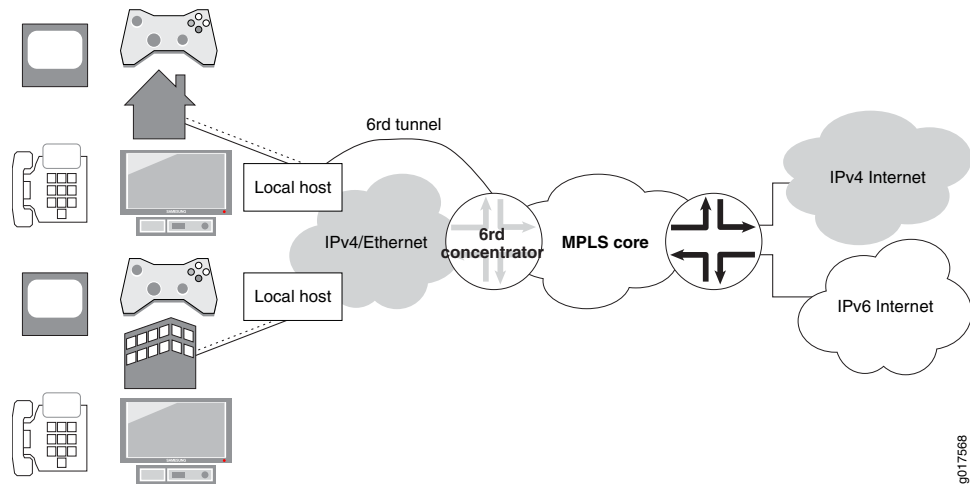
The Junos OS supports many IPv6 transition scenarios required by Junos OS customers. The following are selected examples:

- [Example 1: IPv4 Depletion with a Non-IPv6 Access Network on page 627](#)
- [Example 2: IPv4 Depletion with an IPv6 Access Network on page 627](#)
- [Example 3: IPv4 Depletion for Mobile Networks on page 628](#)

### Example 1: IPv4 Depletion with a Non-IPv6 Access Network

Figure 60 on page 627 depicts a scenario in which the Internet service provider (ISP) has not significantly changed its IPv4 network. This approach enables IPv4 hosts to access the IPv4 Internet and IPv6 hosts to access the IPv6 Internet. A dual-stack host can be treated as an IPv4 host when it uses the IPv4 access service, and as an IPv6 host when it uses the IPv6 access service.

*Figure 60: IPv4 Depletion Solution - IPv4 Access Network*

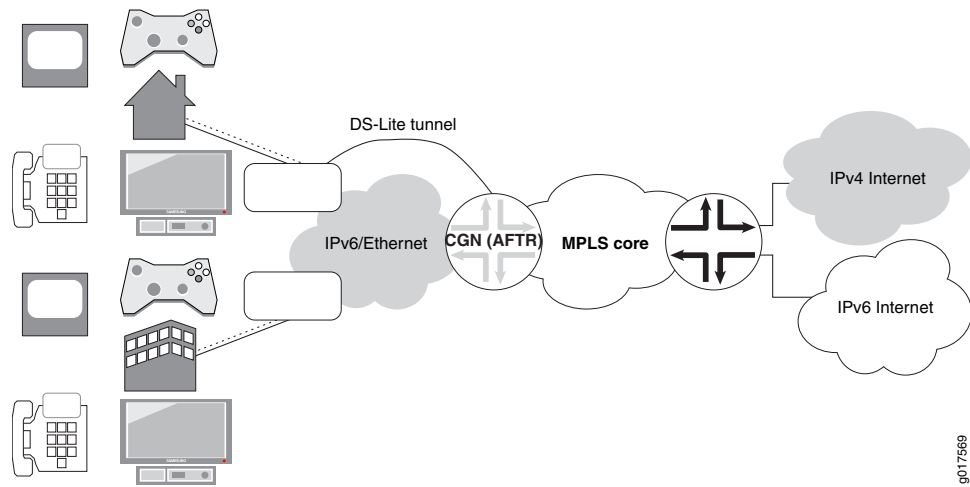


Two new types of devices must be deployed in this approach: a dual-stack home gateway and a dual-stack carrier-grade Network Address Translation (NAT). The dual-stack home gateway integrates IPv4 forwarding and v6-over-v4 tunneling functions. It can also integrate a v4-v4 NAT function. The dual-stack carrier-grade NAT (CGN) integrates v6-over-v4 tunneling and carrier-grade v4-v4 NAT functions.

### Example 2: IPv4 Depletion with an IPv6 Access Network

In the scenario shown in Figure 61 on page 628, the ISP network is IPv6-only.

Figure 61: IPv4 Depletion Solution - IPv6 Access Network



The dual-stack lite (DS-Lite) solution accommodates IPv6-only ISPs. The best business model for this approach is that the customer premises equipment (CPE) has integrated the functions for tunneling IPv6 to an IPv4 backbone, tunneling IPv4 to an IPv6 backbone, and can automatically detect which solution is required.

Not all customers of a given ISP must switch from IPv4 access to IPv6 access simultaneously; in fact, transition can be managed better by switching groups of customers (for example, all those connected to a single point of presence) on an incremental basis. Such an incremental approach should prove easier to plan, schedule, and execute than an across-the-board conversion.

### Example 3: IPv4 Depletion for Mobile Networks

The complexity of mobile networks necessitates a flexible migration approach to ensure minimal disruption and maximum backward compatibility during transition. NAT64 can be used to enable IPv6 devices to communicate to IPv4 hosts without modifying the clients.

**Related Documentation**

- [Packet Flow Through the Adaptive Services or Multiservices PIC on page 620](#)

## Understanding Services PICs

Interfaces used in router networks can be broadly classified into two:

- Networking interfaces, such as Ethernet and SONET interfaces, that primarily provide traffic connectivity. For more information on these interfaces, see the Junos® OS Network Interfaces.
- Services interfaces, such as Adaptive Services interfaces and Multiservices interfaces, that provide specific capabilities for manipulating traffic before it is delivered to its destination.

Services interfaces enable you to add services to your network incrementally. TJunos OS supports the following services interfaces:

- [Adaptive services and Multiservices PICs on page 629](#)
- [Encryption Services \(ES\) PIC on page 629](#)
- [Multilink Services and Link Services PICs on page 630](#)
- [Monitoring Services PICs on page 630](#)
- [Tunnel Services PIC on page 630](#)
- [Multiservices MIC and Multiservices MPC on page 630](#)

## Adaptive services and Multiservices PICs

Adaptive Services [AS] PICs and Multiservices PICs enable you to perform multiple services on the same PIC by configuring a set of services and applications. The AS and Multiservices PICs offer a range of services that you can configure in one or more service sets. The following are some of the services you can configure on Adaptive services or multiservices interfaces:

- Class-of-service
- Intrusion detection service (IDS)
- IP Security (IPsec)
- Layer 2 tunneling protocols
- Monitoring services
- Network Address Translation (NAT)
- Stateful firewalls
- Voice services

For more information about these services, see *Adaptive Services Overview*.



**NOTE:** On Juniper Networks MX Series 3D Universal Edge Routers, the Multiservices DPC provides essentially the same capabilities as the Multiservices PIC. The interfaces on both platforms are configured in the same way.

## Encryption Services (ES) PIC

ES PIC provides a security suite for the IP version 4 (IPv4) and IP version 6 (IPv6) network layers. The suite provides functionality such as authentication of origin, data integrity, confidentiality, replay protection, and nonrepudiation of source. It also defines mechanisms for key generation and exchange, management of security associations, and support for digital certificates. For more information about encryption interfaces, see *Configuring Encryption Interfaces*.

## Multilink Services and Link Services PICs

Multilink Services and Link Services PICs enable you to split, recombine, and sequence datagrams across multiple logical data links. The goal of multilink operation is to coordinate multiple independent links between a fixed pair of systems, providing a virtual link with greater bandwidth than any of the members. The Junos OS supports two services PICs based on the Multilink Protocol: the Multilink Services PIC and the Link Services PIC.

For more information about multilink and link services interfaces, see *Link and Multilink Services Interfaces Feature Guide for Routing Devices*.

## Monitoring Services PICs

Monitoring Services PICs enable you to monitor traffic flow and export the monitored traffic. Monitoring traffic allows you to perform the following tasks:

- Gather and export detailed information about IPv4 traffic flows between source and destination nodes in your network.
- Sample all incoming IPv4 traffic on the monitoring interface and present the data in cflowd record format.
- Perform discard accounting on an incoming traffic flow.
- Encrypt or tunnel outgoing cflowd records, intercepted IPv4 traffic, or both.
- Direct filtered traffic to different packet analyzers and present the data in its original format.

For more information about flow monitoring interfaces, see *Monitoring, Sampling, and Collection Services Interfaces Feature Guide*.

## Tunnel Services PIC

Tunnel Services PIC provides a private, secure path through an otherwise public network by encapsulating arbitrary packets inside a transport protocol. Tunnels connect discontinuous subnetworks and enable encryption interfaces, virtual private networks (VPNs), and MPLS.

For more information about tunnel interfaces, see *Tunnel Properties*.

## Multiservices MIC and Multiservices MPC

The Multiservices Modular Interfaces Card (MS-MIC) and the Multiservices Modular PIC Concentrator (MS-MPC), introduced in Junos OS Release 13.2, provide improved scaling and high performance. The MS-MIC and MS-MPC have enhanced memory (16 GB for MS-MIC, 32 GB per NPU of MS-MPC) and processing capabilities.

The services interfaces on MS-MPC and MS-MIC are identified in the configuration with an **ms-** prefix (for example, **ms-1/2/1**).

The following services packages come preinstalled and preconfigured on MS-MICs and MS-MPCs in Junos OS Release 13.2:

- Junos Traffic Vision (formerly referred to as Jflow/Flow Monitoring)
- Junos Address Aware (formerly referred to as NAT features)
- Junos VPN Site Secure (formerly referred to as IPsec features)
- Junos Network Secure (formerly referred to as the Stateful Firewall feature)

For information about MS-MIC and MS-MPC, see *Multiservices MIC and Multiservices MPC (MS-MIC and MS-MPC) Overview*.

#### Related Documentation

- *Supported Platforms*
- *Packet Flow Through the Adaptive Services or Multiservices PIC*
- *Enabling Service Packages*
- *Services Configuration Procedure*
- *Services Interface Naming Overview*

## TLB Overview

- [TLB Application Description on page 631](#)
- [TLB Topology on page 632](#)
- [TLB Key Characteristics on page 632](#)
- [TLB Application Components on page 633](#)
- [TLB Configuration Limits on page 634](#)

## TLB Application Description

Traffic load balancer (TLB) is supported on MX Series routers with Dense Port Concentrator (DPC) services PICs and Modular Port Concentrator (MPC) line cards. TLB enables you to distribute traffic among multiple next-hop servers.

The TLB solution employs a DPC-based control plane and a data plane using the Router MX 3D forwarding engine.

TLB leverages the MPC's inline functionality, based on an enhanced version of equal-cost multipath (ECMP). Enhanced ECMP facilitates the distribution of sessions across multiple next-hop servers. Enhancements to native ECMP ensure that when servers fail, only flows associated with those servers are impacted, minimizing the overall network churn on services and sessions.

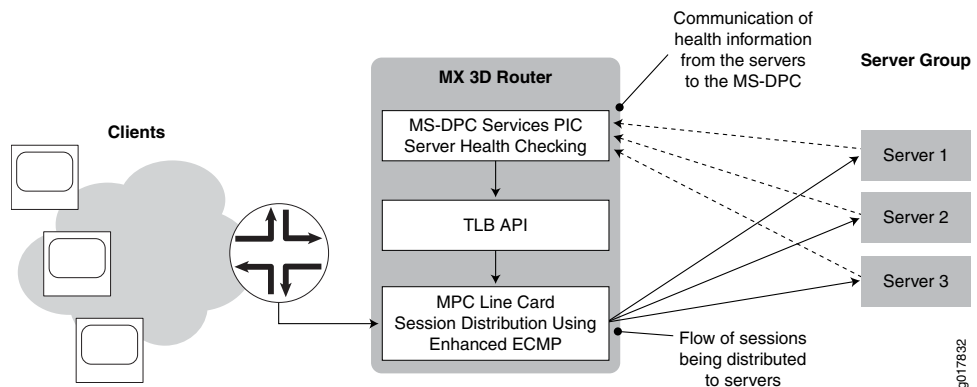
TLB uses the services PIC capabilities of the DPC to provide application-based health monitoring for up to 255 next-hop servers per group, thus providing Intelligent traffic steering based on health checking of server availability information in a next-hop server distribution table. The TLB solution uses a session distribution next-hop API to update the server distribution table and retrieve statistics.

TLB applies its session distribution processing to ingress traffic. Use firewall filters when necessary to select traffic from the ingress interface. Traffic is processed unchanged as it is moved from the ingress interface to the next-hop server. Network Address Translation (NAT) and packet modification are not applied.

## TLB Topology

TLB topology is shown in [Figure 62 on page 632](#).

*Figure 62: TLB Topology*



## TLB Key Characteristics

The following are key characteristics of TLB.

- TLB only distributes the requests for any flow; the response is expected to return directly to the client/source.
- TLB supports hash-based load balancing based on source IP, destination IP, and protocol.
- TLB enables you to configure servers offline to prevent a performance impact that might be caused by a rehashing for all existing flows. You can add a server in administrative down state and use it later for traffic distribution by disabling "admin down". This prevents traffic impact to other servers.
- When health checking determines a server to be down, only the affected flows are reshaped.
- When a previously down server is returned to service, all flows belonging to that server based on hashing return to it, impacting performance for the returned flows. For this reason, the automatic rejoining of a server to an active group can be disabled. Servers are returned to service only by issuing the *request services traffic-load-balance real-service rejoin* operational command.
- Health check monitoring application runs on an MS-DPC/NPU. TLB traffic is not forwarded to the MS-DPC/NPU.

- NAT is not be applied to the distributed sessions.
- High availability is accomplished by stateless failover between two Mx3D routers. The routers can seamlessly backup one another because they leverage the same hash algorithm which results in the same server being allocated for the same flow.

## TLB Application Components

### Servers and Server Groups

TLB enables configuration of groups of up to 255 servers (referred to in configuration as *real services*) as next-hop destinations for stateless session distribution. You can configure up to 1024 servers associated with one services PIC used for health checking. All servers used in server groups must be individually configured before assignment to groups. The session distribution hashing algorithm uses key-selectable hashing for session distribution. Distribution information is maintained in a server distribution table. Users can add and delete servers to and from the TLB server distribution table and can also change the administrative status of a server.



**NOTE:** The TLB solution uses the session distribution next-hop API to update the server distribution table and retrieve statistics. *Applications do not have direct control on the server distribution table management. They can only influence changes indirectly through the add and delete services of the TLB API.*

### Server Health Monitoring — Single Health Check and Dual Health Check

TLB supports health check protocols— ICMP, TCP, and HTTP—to monitor the health of servers in a group. You can use a single probe type for a server group, or a dual health check (TLB - DHC) configuration, which includes two probe types. The configurable health monitoring function resides on a services PIC. By default, probe requests are sent every 5 seconds. Also by default, a real server is declared down only after five consecutive probe failures and declared up only after five consecutive probe successes.

TLB provides *application stickiness*, meaning that server failures or changes do not affect traffic flows to other active servers. Changing a server's administrative status from down to up, or changing a server's administrative state to down does not impact any active flows to remaining servers in the server distribution table. Adding a server or deleting a server from a group has some traffic impact for 5 to 10 seconds.

TLB provides two levels of server health monitoring:

- Single Health Check—One probe type is attached to a server group by means of the **network-monitoring-profile** configuration statement.
- TLB Dual Health Check (TLB-DHC)—Two probe types are associated with a server group by means of the **network-monitoring-profile** configuration statement. A server's status is declared based on the result of two health check probes. This feature enhancement, allowing users to configure up to two health check profiles per server group, is in traffic-dird-12.1X43-1-A2.2 and subsequent releases. If a server group is

configured for DHC, a real-service is declared to be UP only if both health-check probes are simultaneously UP, otherwise a real-service declared to be DOWN.

### Virtual Services

The virtual service provides an address that is associated with a the group of servers to which traffic is directed as determined by hash-based session distribution and server health monitoring.

The virtual service configuration identifies:

- The group of servers to which sessions are distributed
- The session distribution hashing method



**NOTE:** TLB doesn't require a specific virtual IP. VIPs 0.0.0.0 or 0::0 are acceptable.

### TLB Configuration Limits

*Table 92: TLB Configuration Limits*

Maximum servers per group.	255
Maximum virtual services per services PIC.	32
Maximum real servers per services PIC	1024
Maximum groups per virtual service.	1
Maximum network monitoring profiles per group.	2
Maximum number of TLB instances per service interface unit.	1
Maximum number of VIPs per virtual service.	1
Supported health checking protocols.	ICMP, TCP, HTTP

**Related Documentation** • [Configuring TLB](#)

### Installing and Configuring TLB Using the CLI Interface

Before you configure TLB, install the TLB application package.

To install the TLB application:

1. Download the TLB package.
2. Install the TLB package in the router.

```
user@host> request system software add
traffic-dird-14.1I20150205_0332_rsbu-builder-A4.0-signed.tgz no-validate
```

3. At the **[edit chassis fpc slot-number pic pic-number adaptive-services service-package extension-provider]** hierarchy level, configure the desired level of syslog information from the TLB daemon.

```
[edit chassis fpc 1 pic 0 adaptive-services service-package extension-provider]
user@host# set syslog daemon info
```

4. Enable TLB network monitoring daemon on the PIC.

```
[edit chassis fpc 1 pic 0 adaptive-services service-package extension-provider]
user@host# set package traffic-dird-services
```

After you have installed the application package, you can configure or reconfigure TLB as needed. The following topics describe how to configure TLB. To create a complete application, you must also define interfaces and routing information. You can optionally define firewall filters and policy options in order to differentiate TLB traffic.

- [Configuring a TLB Instance on page 635](#)
- [Configuring Interface and Routing Information on page 635](#)
- [Configuring Servers on page 637](#)
- [Configuring Network Monitoring Profiles on page 638](#)
- [Configuring Server Groups on page 639](#)
- [Configuring Virtual Services on page 640](#)

## Configuring a TLB Instance

TLB configurations are saved in named instances.

## Configuring Interface and Routing Information

To configure interface and routing information:

1. At the **[edit services traffic-load-balance instance *instance-name*]** hierarchy level, identify the service interface associated with this instance.

```
user@host# [edit services traffic-load-balance instance instance-name]
user@host# set interface interface-name
```

For example:

```
[edit services traffic-load-balance instance tlb-instance1]
```

```
user@host# set interface ms-1/0/0
```

2. Specify the client interface for which an implicit filter is defined to direct traffic in the forward direction. This is required only for translated mode.

```
user@host# [edit services traffic-load-balance instance instance-name]  
user@host# set client-interface interface-name
```

For example:

```
[edit services traffic-load-balance instance tlb-instance1]  
user@host# set client-interface ge-5/2/0.0
```

3. Specify the virtual routing instance, **server-vrf**, used to route data traffic in the forward direction to servers. This is required for SLT and Layer 3 DSR; it is optional for Layer 2 DSR.

```
user@host# [edit services traffic-load-balance instance instance-name]  
user@host# set server-vrf server-vrf-name
```

For example:

```
[edit services traffic-load-balance instance tlb-instance1]  
user@host# set server-vrf server-vrf
```

4. Specify the server interface for which implicit filters are defined to direct return traffic to the client.



**NOTE:** Implicit filters for return traffic are not used for DSR.

```
user@host# [edit services traffic-load-balance instance instance-name]  
user@host# set server-interface server-interface
```

For example:

```
[edit services traffic-load-balance instance tlb-instance1]  
user@host# set server-interface ge-5/2/1.0
```

5. (Optional) Specify the filter used to bypass rephrase as health-check traffic from real servers.

```
user@host# [edit services traffic-load-balance instance instance-name]  
user@host# set server-inet-bypass-filter server-inet-bypass-filter
```

For example:

```
[edit services traffic-load-balance instance tlb-instance1]
```

```
user@host# set server-inet-bypass-filter tlb-ipv4-bypass
```

- Specify the virtual routing instance **client-vrf** in which you want the data in the reverse direction to be routed to the clients.



**NOTE:** Virtual routing instances for routing data in the reverse direction are not used with DSR.

```
user@host# [edit services traffic-load-balance instance instance-name]
user@host# set client-vrf client-vrf
```

For example:

```
[edit services traffic-load-balance instance tlb-instance1]
user@host# set client-vrf client-vrf
```

## Configuring Servers

To configure servers for the TLB instance:

- Go to the **[edit services traffic-load-balance instance *instance-name*]** hierarchy level.

```
user@host# [edit services traffic-load-balance instance instance-name]
```

- At the **[edit services traffic-load-balance instance *instance-name*]** hierarchy level, configure a logical name and IP address for each server to be made available for next-hop distribution.

```
[edit services traffic-load-balance instance instance-name]
user@host# set real-service name address ip-address
```

For example:

```
[edit services traffic-load-balance instance tlb-instance1]
user@host# set real-service rs138 address 172.26.99.138
user@host# set real-service rs139 address 172.26.99.139
user@host# set real-service rs140 address 172.26.99.140
```

## Configuring Network Monitoring Profiles

A network monitoring profile configures a health check probe, which you assign to a server group to which session traffic is distributed. To configure a network monitoring profile:

1. Configure the type of probe to use for health monitoring — **icmp**, **tcp**, **http**, or **custom**.

- For an ICMP probe:

```
[edit services network-monitoring profile profile-name]  
user@host.com# set icmp
```

- For a TCP probe:

```
[edit services network-monitoring profile profile-name]  
user@host.com# set tcp port port-number
```

- For an HTTP probe:

```
[edit services network-monitoring profile profile-name]  
user@host.com# set http host hostname url url-name port port-number method (get  
| option)
```

- For an SSL probe:

```
[edit services network-monitoring profile profile-name]  
user@host.com# set ssl-hello port ssl-version
```

- For a custom probe:

```
[edit services network-monitoring profile profile-name]  
user@host.com# set custom cmd priority default-rs-status (down | up) expect (ascii  
| binary) receive-string port port rs-action (down | up) send (ascii | binary)  
send-string
```

2. Configure the interval for probe attempts, in seconds (1 through 180).

```
[edit services network-monitoring profile profile-name]  
user@host.com# set probe-interval interval
```

For example:

```
[edit services network-monitoring profile profile1-icmp]  
user@host.com# set probe-interval 2
```

The default value is 5.

3. Configure the number of failure retries, after which the real server is tagged as down.

```
[edit services network-monitoring profile profile-name]  
user@host.com# set failure-retries number-of-retries
```

For example:

```
[edit services network-monitoring profile profile1-icmp]
user@host.com# set failure-retries 3
```

4. Configure the number of recovery retries, which is the number of successful probe attempts after which the server is declared up.

```
[edit services network-monitoring profile profile-name]
user@host.com# set recovery-retries number-of-retries
```

For example:

```
[edit services network-monitoring profile profile1-icmp]
user@host.com# set recovery-retries 1
```

## Configuring Server Groups

Server groups consist of servers to which traffic is distributed by means of stateless, hash-based session distribution and server health monitoring.

To configure a server group:

1. At the **[edit services traffic-load-balance instance *instance-name*]** hierarchy level, specify the names of one or more configured real servers.

```
[edit services traffic-load-balance instance instance-name groups group-name]
user@host.com# set real-services real-service-name, ...
```

For example:

```
[edit services traffic-load-balance instance tlb-instance1 groups tlb-group1]
user@host.com# set real-services [ rs138 rs139 rs140 ]
```

2. Configure the routing instance for the group when you do not want to use the default instance, **inet.0**.

```
[edit services traffic-load-balance instance instance-name groups group-name]
user@host.com# set routing-instance routing-instance-name
```

For example:

```
[edit services traffic-load-balance instance tlb-instance1 groups tlb-group1]
user@host.com# set routing-instance tlb-routing-instance1
```

3. (Optional) Disable the default option that allows a server to rejoin the group automatically when it comes up.

```
[edit services traffic-load-balance instance instance-name group group-name]
user@host.com# set real-service-rejoin-options no-auto-rejoin
```

4. Configure one or two network monitoring profiles to be used to monitor the health of servers in this group.

```
[edit services traffic-load-balance instance instance-name groups group-name]  
user@host.com# set network-monitoring-profile profile-name1 profile-name2
```

For example:

```
[edit services traffic-load-balance instance tlb-instance1 groups tlb-group1]  
user@host.com# set network-monitoring-profile profile1-icmp profile2-http
```

## Configuring Virtual Services

A virtual service provides an address that is associated with a group of servers to which traffic is directed as determined by hash-based or random session distribution and server health monitoring. You may optionally specify filters and routing instances to steer traffic for TLB.

To configure a virtual service:

1. At the **[edit services traffic-load-balance instance *instance-name*]** hierarchy level, specify a non-zero address for the virtual service.

```
[edit services traffic-load-balance instance instance-name virtual-service  
  virtual-service-name]  
user@host# set address local-address
```

For example:

```
[edit services traffic-load-balance instance tlb-instance1 virtual-service virtual-service1]  
user@host# set address 10.1.1.1
```

2. Specify the server group used for this virtual service.

```
[edit services traffic-load-balance instance instance-name virtual-service  
  virtual-service-name]  
user@host# set group group-name
```

For example:

```
[edit services traffic-load-balance instance tlb-instance1 virtual-service virtual-service1]  
user@host# set group tlb-group1
```

3. Specify a routing instance for the virtual service. If you do not specify a routing instance, the default routing instance is used.

```
[edit services traffic-load-balance instance instance-name virtual-service  
  virtual-service-name]  
user@host# set routing-instance routing-instance
```

For example:

```
[edit services traffic-load-balance instance tlb-instance1 virtual-service virtual-service1]
user@host# set routing-instance msp-tproxy-server-vrf31
```

4. (Optional) Specify a routing metric for the virtual service.

```
[edit services traffic-load-balance instance instance-name virtual-service
virtual-service-name]
user@host# set routing-metric routing-metric
```

For example:

```
[edit services traffic-load-balance instance tlb-instance1 virtual-service virtual-service1]
user@host# set routing-metric 128
```

5. Specify the method used for load balancing. You can specify a hash method, **source-ip**, **destination-ip**, or **protocol**, or specify **random**.

```
[edit services traffic-load-balance instance tlb-instance1 virtual-service virtual-service1]
user@host# set load-balancing-method hash hash-key (source-ip | destination-ip |
proto)
```

For example:

```
[edit services traffic-load-balance instance tlb-instance1 virtual-service virtual-service1]
user@host# set load-balancing-method hash hash-key source-ip
```

or

```
[edit services traffic-load-balance instance tlb-instance1 virtual-service virtual-service1]
user@host# set load-balancing-method random
```

6. For mode translated services, specify a service for translation, including a virtual-port, server-listening-port, and protocol.

```
[edit services traffic-load-balance instance instance-name virtual-service
virtual-service-name]
user@host# set service service-name virtual-port virtual-port server-listening-port
server-listening-port protocol protocol
```

For example:

```
[edit services traffic-load-balance instance tlb-instance1 virtual-service virtual-service1]
user@host# set service fast-track-service virtual-port 1111 server-listening-port 22
protocol tcp
```

7. Commit the configuration.

```
[edit services traffic-load-balance instance instance-name virtual-service
virtual-service-name]
```

```
user@host# commit
```



**NOTE:** In the absence of a client-interface configuration under the TLB instance, the implicit client filter (for VIP) is attached to the client-vrf configured under the TLB instance. In this case, the routing-instance under a translate mode virtual service cannot be the same as the client-vrf configured under the TLB instance. If it is, the commit fails.

**Related Documentation**

- [Traffic Load Balancer Overview](#)

---

## Stateful Firewall Overview for Junos OS Extension-Provider Packages

---

Routers use firewalls to track and control the flow of traffic. Adaptive Services and MultiServices PICs employ a type of firewall called a *stateful firewall*. Contrasted with a *stateless* firewall that inspects packets in isolation, a stateful firewall provides an extra layer of security by using state information derived from past communications and other applications to make dynamic control decisions for new communication attempts.

Stateful Firewall (SFW) is supported on Junos OS extension-provider packages (known as Junos Services Framework (JSF) in Junos OS Releases earlier than 12.3). Junos OS extension-provider packages are suites of applications or features that enable the integration of various services on Junos-based platforms.

Stateful firewalls group relevant *flows* into *conversations*. A flow is identified by the following five properties:

- Source address
- Source port
- Destination address
- Destination port
- Protocol



**NOTE:** The protocols that are not supported on top of TCP/UDP can have the source port and destination port mapped to other fields.

A typical Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) conversation consists of two flows: the initiation flow and the responder flow. However, some conversations, such as an FTP conversation, might consist of two control flows and many data flows.

Firewall rules govern whether the conversation is allowed to be established. If a conversation is allowed, all flows within the conversation are permitted, including flows that are created during the life cycle of the conversation.

You configure stateful firewalls using a powerful rule-driven conversation handling path. A *rule* consists of direction, source address, source port, destination address, destination port, IP protocol value, and application protocol or service. In addition to the specific values you configure, you can assign the value **any** to rule objects, addresses, or ports, which allows them to match any input value. Finally, you can optionally negate the rule objects, which negates the result of the type-specific match.

Firewall rules are directional. For each new conversation, the router software checks the initiation flow matching the direction specified by the rule.

Firewall rules are ordered. The software checks the rules in the order in which you include them in the configuration. The first time the firewall discovers a match, the router implements the action specified by that rule. Rules still unchecked are ignored.

## Stateful Firewall Support for Application Protocols

By inspecting the application protocol data, the AS or MultiServices PIC firewall can intelligently enforce security policies and allow only the minimal required packet traffic to flow through the firewall.

The firewall rules are configured in relation to an interface. By default, the stateful firewall allows all sessions initiated from the hosts behind the interface to pass through the router.

## Stateful Firewall Anomaly Checking

The stateful firewall recognizes the following events as anomalies and sends them to the IDS software for processing:

- IP anomalies:
  - IP version is not correct.
  - IP header length field is too small.
  - IP header length is set larger than the entire packet.
  - Bad header checksum.
  - IP total length field is shorter than header length.
  - Packet has incorrect IP options.
  - Internet Control Message Protocol (ICMP) packet length error.
  - Time-to-live (TTL) equals 0.
- IP address anomalies:
  - IP packet source is a broadcast or multicast.
  - Land attack (source IP equals destination IP).
- IP fragmentation anomalies:
  - IP fragment overlap.
  - IP fragment missed.

- IP fragment length error.
- IP packet length is more than 64 kilobytes (KB).
- Tiny fragment attack.
- TCP anomalies:
  - TCP port 0.
  - TCP sequence number 0 and flags 0.
  - TCP sequence number 0 and FIN/PSH/RST flags set.
  - TCP flags with wrong combination (TCP FIN/RST or SYN/(URG|FIN|RST)).
  - Bad TCP checksum.
- UDP anomalies:
  - UDP source or destination port 0.
  - UDP header length check failed.
  - Bad UDP checksum.
- Anomalies found through stateful TCP or UDP checks:
  - SYN followed by SYN-ACK packets without ACK from initiator.
  - SYN followed by RST packets.
  - SYN without SYN-ACK.
  - Non-SYN first flow packet.
  - ICMP unreachable errors for SYN packets.
  - ICMP unreachable errors for UDP packets.
- Packets dropped according to stateful firewall rules.

If you employ stateful anomaly detection in conjunction with stateless detection, IDS can provide early warning for a wide range of attacks, including these:

- TCP or UDP network probes and port scanning
- SYN flood attacks
- IP fragmentation-based attacks such as teardrop, bonk, and boink

**Related  
Documentation**

- [Creating and Managing CGNAT Service Templates on page 234](#)

## Network Address Translation Configuration Overview

- [Configuring Source and Destination Addresses Network Address Translation Overview on page 645](#)
- [Configuring Pools of Addresses and Ports for Network Address Translation Overview on page 646](#)
- [Configuring Address Pools for Network Address Port Translation \(NAPT\) Overview on page 648](#)
- [Network Address Translation Rules Overview on page 656](#)
- [Configuring Service Sets for Network Address Translation on page 661](#)

### Configuring Source and Destination Addresses Network Address Translation Overview

You must configure a specific address, a prefix, or the address-range boundaries:

- The following addresses, while valid in **inet.0**, cannot be used for NAT translation:
  - **0.0.0.0/32**
  - **127.0.0.0/8** (loopback)
  - **128.0.0.0/16** (martian)
  - **191.255.0.0/16** (martian)
  - **192.0.0.0/24** (martian)
  - **223.255.255.0/24** (martian)
  - **224.0.0.0/4** (multicast)
  - **240.0.0.0/4** (reserved)
  - **255.255.255.255** (broadcast)
- You can specify one or more IPv4 address prefixes in the **pool** statement and in the **from** clause of the NAT rule term. This enables you to configure source translation from a private subnet to a public subnet without defining a rule term for each address in the subnet. Destination translation cannot be configured by this method. For more information, see *Examples: Configuring NAT Rules*.
- When you configure static source NAT, the **address** prefix size you configure at the **[edit services nat pool pool-name]** hierarchy level must be larger than the **source-address** prefix range configured at the **[edit services nat rule rule-name term term-name from]** hierarchy level. The **source-address** prefix range must also map to a single subnet or range of IPv4 or IPv6 addresses in the **pool** statement. Any pool addresses that are not used by the **source-address** prefix range are left unused. Pools cannot be shared.



**NOTE:** When you include a NAT configuration that changes IP addresses, it might affect forwarding path features elsewhere in your router configuration, such as source class usage (SCU), destination class usage (DCU), filter-based forwarding, or other features that target specific IP addresses or prefixes.

NAT configuration might also affect routing protocol operation, because the protocol peering, neighbor, and interface addresses can be altered when routing protocols packets transit the Adaptive Services (AS) or Multiservices PIC.

**See Also** • [Network Address Translation Overview for MS-DPC, MS-MPC, and MS-MIC Line Cards](#)

## Configuring Pools of Addresses and Ports for Network Address Translation Overview

- [Configuring NAT Pools on page 646](#)
- [Preserve Range and Preserve Parity on page 647](#)
- [Specifying Destination and Source Prefixes without Configuring a Pool on page 647](#)

### Configuring NAT Pools

You can use the **pool** statement to define the addresses (or prefixes), address ranges, and ports used for Network Address Translation (NAT). To configure the information, include the **pool** statement at the **[edit services nat]** hierarchy level:

```
[edit services nat]
pool nat-pool-name {
  address ip-prefix </prefix-length>;
  address-range low minimum-value high maximum-value;
  port (automatic | range low minimum-value high maximum-value);
  preserve-parity;
  preserve-range {
  }
}
```

To configure pools for traditional NAT, specify either a destination pool or a source pool.

With static source NAT and dynamic source NAT, you can specify multiple IPv4 addresses (or prefixes) and IPv4 address ranges. Up to 32 prefixes or address ranges (or a combination) can be supported within a single pool.

With static destination NAT, you can also specify multiple address prefixes and address ranges in a single term. Multiple destination NAT terms can share a destination NAT pool. However, the netmask or range for the **from** address must be smaller than or equal to the netmask or range for the destination pool address. If you define the pool to be larger than required, some addresses will not be used. For example, if you define the pool size as 100 addresses and the rule specifies only 80 addresses, the last 20 addresses in the pool are not used.

For constraints on specific translation types, see *Network Address Translation Rules Overview*.

With source static NAT, the prefixes and address ranges cannot overlap between separate pools.

In an address range, the **low** value must be a lower number than the **high** value. When multiple address ranges and prefixes are configured, the prefixes are depleted first, followed by the address ranges.

When you specify a port for dynamic source NAT, address ranges are limited to a maximum of 65,000 addresses, for a total of (65,000 x 65,535) or 4,259,775,000 flows. A dynamic NAT pool with no address port translation supports up to 65,535 addresses. There is no limit on the pool size for static source NAT.

### Preserve Range and Preserve Parity

You can configure your carrier-grade NAT (CGN) to preserve the range or parity of the packet source port when it allocates a source port for an outbound connection. You can configure the preserve parity and preserve range options under the NAT pool definition by including the **preserve-range** and **preserve-parity** configuration statements at the **[edit services nat pool poolname port]** hierarchy level.

- Preserve range—RFC 4787, *Network Address Translation (NAT) Behavioral Requirements for Unicast UDP*, defines two ranges: 0 through 1023, and 1024 through 65,535. When the **preserve-range** knob is configured and the incoming port falls into one of these ranges, CGN allocates a port from that range only. However, if there is no available port in the range, the port allocation request fails and that session is not created. The failure is reflected on counters and system logging, but no Internet Control Message Protocol (ICMP) message is generated. If this knob is not configured, allocation is based on the configured port range without regard to the port range that contains the incoming port. The exception is some application-level gateways (ALGs), such as hello, that have special zones.
- Preserve parity—When the **preserve-parity** knob is configured, CGN allocates a port with the same even or odd parity as the incoming port. If the incoming port number is odd or even, the outgoing port number should correspondingly be odd or even. If a port number of the desired parity is not available, the port allocation request fails, the session is not created, and the packet is dropped.

### Specifying Destination and Source Prefixes without Configuring a Pool

You can directly specify the destination or source prefix used in NAT without configuring a pool.

To configure the information, include the **rule** statement at the **[edit services nat]** hierarchy level:

```
[edit services nat]
rule rule-name {
  term term-name {
    then {
```

```
        translated {  
            destination-prefix prefix;  
        }  
    }  
}
```

## Configuring Address Pools for Network Address Port Translation (NAPT) Overview

With Network Address Port Translation (NAPT), you can configure up to 32 address ranges with up to 65,536 addresses each.

The **port** statement specifies port assignment for the translated addresses. To configure automatic assignment of ports, include the **port automatic** statement at the **[edit services nat pool *nat-pool-name*]** hierarchy level. To configure a specific range of port numbers, include the **port range low *minimum-value* high *maximum-value*** statement at the **[edit services nat pool *nat-pool-name*]** hierarchy level.

The Junos OS provides several alternatives for allocating ports:

- [Round-Robin Allocation for NAPT on page 648](#)
- [Sequential Allocation for NAPT on page 649](#)
- [Preserve Parity and Preserve Range for NAPT on page 649](#)
- [Address Pooling and Endpoint Independent Mapping for NAPT on page 650](#)
- [Port Block Allocation for NAPT on page 651](#)
- [Deterministic Port Block Allocation for NAPT on page 651](#)
- [Comparison of NAPT Implementation Methods on page 656](#)

### Round-Robin Allocation for NAPT

---

To configure round-robin allocation for NAT pools, include the **address-allocation round-robin** configuration statement at the **[edit services nat pool *pool-name*]** hierarchy level. When you use round-robin allocation, one port is allocated from each address in a range before repeating the process for each address in the next range. After ports have been allocated for all addresses in the last range, the allocation process wraps around and allocates the next unused port for addresses in the first range.

- The first connection is allocated to the address:port 100.0.0.1:3333.
- The second connection is allocated to the address:port 100.0.0.2:3333.
- The third connection is allocated to the address:port 100.0.0.3:3333.
- The fourth connection is allocated to the address:port 100.0.0.4:3333.
- The fifth connection is allocated to the address:port 100.0.0.5:3333.
- The sixth connection is allocated to the address:port 100.0.0.6:3333.
- The seventh connection is allocated to the address:port 100.0.0.7:3333.
- The eighth connection is allocated to the address:port 100.0.0.8:3333.

- The ninth connection is allocated to the address:port 100.0.0.9:3333.
- The tenth connection is allocated to the address:port 100.0.0.10:3333.
- The eleventh connection is allocated to the address:port 100.0.0.11:3333.
- The twelfth connection is allocated to the address:port 100.0.0.12:3333.
- Wraparound occurs and the thirteenth connection is allocated to the address:port 100.0.0.1:3334.

### Sequential Allocation for NAT

With sequential allocation, the next available address in the NAT pool is selected only when all the ports available from an address are exhausted.



**NOTE:** This legacy implementation provides backward compatibility and is no longer a recommended approach..

The NAT pool called **napt** in the following configuration example uses the sequential implementation:

```
pool napt {
  address-range low 100.0.0.1 high 100.0.0.3;
  address-range low 100.0.0.4 high 100.0.0.6;
  address-range low 100.0.0.8 high 100.0.0.10;
  address-range low 100.0.0.12 high 100.0.0.13;
  port {
    range low 3333 high 3334;
  }
}
```

In this example, the ports are allocated starting from the first address in the first address-range, and allocation continues from this address until all available ports have been used. When all available ports have been used, the next address (in the same address-range or in the following address-range) is allocated and all its ports are selected as needed. In the case of the example **napt** pool, the tuple address, port 100.0.0.4:3333, is allocated only when all ports for all the addresses in the first range have been used.

- The first connection is allocated to the address:port 100.0.0.1:3333.
- The second connection is allocated to the address:port 100.0.0.1:3334.
- The third connection is allocated to the address:port 100.0.0.2:3333.
- The fourth connection is allocated to the address:port 100.0.0.2:3334, and so on.

### Preserve Parity and Preserve Range for NAT

The following options are available for NAT:

- Preserving parity—Use the **preserve-parity** command to allocate even ports for packets with even source ports and odd ports for packets with odd source ports.

- Preserving range—Use the **preserve-range** command to allocate ports within a range from 0 to 1023, assuming the original packet contains a source port in the reserved range. This applies to control sessions, not data sessions.

---

### Address Pooling and Endpoint Independent Mapping for NAPT

- [Address Pooling and Endpoint Independent Mapping for NAPT on page 650](#)

#### ***Address Pooling and Endpoint Independent Mapping for NAPT***

##### ***Address Pooling***

Address pooling, or address pooling paired (APP) ensures assignment of the same external IP address for all sessions originating from the same internal host. You can use this feature when assigning external IP addresses from a pool. This option does not affect port utilization

Address pooling solves the problems of an application opening multiple connections. For example, when Session Initiation Protocol (SIP) client sends Real-Time Transport Protocol (RTP) and Real-Time Control Protocol (RTCP) packets, the SIP generally server requires that they come from the same IP address, even if they have been subject to NAT. If RTP and RTCP IP addresses are different, the receiving endpoint might drop packets. Any point-to-point (P2P) protocol that negotiates ports (assuming address stability) benefits from address pooling paired.

The following are use cases for address pooling.

- A site that offers instant messaging services requires that chat and their control sessions come from the same public source address. When the user signs on to chat, a control session authenticates the user. A different session begins when the user starts a chat session. If the chat session originates from a source address that is different from the authentication session, the instant messaging server rejects the chat session, because it originates from an unauthorized address.
- Certain websites such as online banking sites require that all connections from a given host come from the same IP address.

##### ***Endpoint Independent Mapping and Endpoint Independent Filtering***

Endpoint independent mapping (EIM) ensures the assignment of the same external address *and* port for all connections from a given host if they use the same internal port. This means if they come from a different source port, you are free to assign a different external address.

EIM and APP differ as follows:

- APP ensures assigning the same external IP address.
- EIM provides a stable external IP address and port (for a period of time) to which external hosts can connect. Endpoint independent filtering (EIF) controls which external hosts can connect to an internal host.

### Port Block Allocation for NAPT

Carriers track subscribers using the IP address (RADIUS or DHCP) log. If they use NAPT, an IP address is shared by multiple subscribers, and the carrier must track the IP address and port, which are part of the NAT log. Because ports are used and reused at a very high rate, tracking subscribers using the log becomes difficult due to the large number of messages, which are difficult to archive and correlate. By enabling the allocation of ports in blocks, port block allocation can significantly reduce the number of logs, making it easier to track subscribers.

Port block allocation is supported on MX series routers with MultiServices Dense Port Concentrators (MS-DPCs).

- [Secured Port Block Allocation for NAPT on page 651](#)

#### ***Secured Port Block Allocation for NAPT***

When allocating blocks of ports, the most recently allocated block is the current active block. New requests for NAT ports are served from the active block. Ports are allocated randomly from the current active block.

When you configure secured port block allocation, you can specify the following:

- **block-size**
- **max-blocks-per-address**
- **active-block-timeout**

**See Also** • [Configuring Secured Port Block Allocation](#)

### Deterministic Port Block Allocation for NAPT

You can configure NAT algorithm-based allocation of blocks of destination ports. By specifying **deterministic-port-block-allocation blocksize *blocksize*** at the **[edit services nat pool *poolname* port]** hierarchy level, you ensure that an incoming (source) IP address and port always map to the same destination IP address and port, thus eliminating the need for the address translation logging. You can also specify **include-boundary-addresses** if you want the lowest and highest addresses in the source address range of a NAT rule to be translated when the NAT pool is used. When you use deterministic port block allocation, you must specify **deterministic-nat44** as the **translation-type** in your NAT rule.

For detailed information on how to configure deterministic port block allocation, see *Configuring Deterministic Port Block Allocation*.

- [Understanding Deterministic Port Block Allocation Algorithms on page 652](#)
- [Deterministic Port Block Allocation Algorithm Usage on page 652](#)
- [Deterministic Port Block Allocation Restrictions on page 655](#)

### ***Understanding Deterministic Port Block Allocation Algorithms***

The effectiveness of your implementation of deterministic port block allocation depends on your analysis of your subscriber requirements. The block size you provide indicates how many ports will be made available for each incoming subscriber address in the range the **from** clause specified in the applicable NAT rule. The allocation algorithm computes an offset value to determine the outgoing port. A reverse algorithm is used to derive the originating subscriber address.



**NOTE:** In order to track subscribers without using logs, an ISP must use a reverse algorithm to derive a subscriber (source) addresses from translated addresses.

### ***Deterministic Port Block Allocation Algorithm Usage***

When you have configured deterministic port block allocation, you can use the ***show services nat deterministic-nat internal-host*** and ***show services nat deterministic-nat nat-port-block*** commands to show forward and reverse mapping. However, mappings will change if you reconfigure your deterministic port block allocation block size or the **from** clause for your NAT rule. In order to provide historical information on mappings, we recommend that you write scripts that can show specific mappings for prior configurations.

The following variables are used in forward calculation (private subscriber IP address to public IP address) and reverse calculation (public IP address to private subscriber IP address):

- Pr\_Prefix—Any pre-NAT IPv4 subscriber address
- Pr\_Port—Any pre-NAT protocol port
- Block\_Size—Number of ports configured to be available for each Pr\_Prefix
- Base\_PR\_Prefix—First usable pre-NAT IPv4 subscriber address in a “from” clause match condition
- Base\_PU\_Prefix—First usable post-NAT IPv4 subscriber address configured in the NAT pool.
- Pu\_Port\_Range\_Start—1024 (ports 0 through 1023 are not used when **port automatic** is configured)
- Pr\_Offset—Pr\_Prefix – Base\_Pr\_Prefix
- PR\_Port\_Offset—Pr\_Offset \* Block\_Size
- Pu\_Prefix—Post-NAT address for a given Pr\_Prefix
- Pu\_Start\_Port—Post-NAT start port for a flow from a given Pr\_Prefix
- Pu\_Actual\_Port—Post-NAT port seen on a reverse flow
- Nr\_Addr\_PR\_Prefix — Number of usable pre-NAT IPv4 subscriber addresses in a “from” clause match condition

- $Nr\_Addr\_PU\_Prefix$  — Number of usable post-NAT IPv4 addresses configured in the NAT pool
- $Rounded\_Port\_Range\_Per\_IP = \lceil (Nr\_Addr\_PR\_Prefix / Nr\_Addr\_PU\_Prefix) \rceil * Block\_Size$
- $Pu\_Offset = Pu\_Prefix - Base\_Pu\_Prefix$
- $Pu\_Port\_Offset = (Pu\_Offset * Port\_Range\_Per\_Pu\_IP) + (Pu\_Actual\_Port - Pu\_Port\_Start\_Port)$



**NOTE:** If block-size is configured as zero, the method for computing the block size has changed and is computed as follows:

$$block-size = \text{int}(64512 / \lceil (Nr\_Addr\_PR\_Prefix / Nr\_Addr\_PU\_Prefix) \rceil)$$

where 64512 is the maximum available port range per public IP address.

**Algorithm Usage**—Assume the following configuration:

```
services {
  nat {
    pool src-pool {
      address-range low 32.32.32.1 high 32.32.32.254;
      port {
        automatic {
          random-allocation;
        }
        deterministic-block-allocation {
          block-size 249;
        }
      }
    }
  }
  rule det-nat {
    match-direction input;
    term t1 {
      from {
        source-address {
          10.1.0.0/16;
        }
      }
      then {
        translated {
          source-pool src-pool;
          translation-type {
            deterministic-napt44;
          }
        }
      }
    }
  }
}
```

**Forward Translation**

1.  $Pr\_Offset = Pr\_Prefix - Base\_Pr\_Prefix$
2.  $Pr\_Port\_Offset = Pr\_Offset * Block\_Size$
3.  $Rounded\_Port\_Range\_Per\_IP = \lceil (Nr\_Addr\_PR\_Prefix / Nr\_Addr\_PU\_Prefix) \rceil * Block\_Size$
4.  $Pu\_Prefix = Base\_Public\_Prefix + \text{floor}(Pr\_Port\_Offset / Rounded\_Port\_Range\_Per\_IP)$
5.  $Pu\_Start\_Port = Pu\_Port\_Range\_Start + (Pr\_Port\_Offset \% Rounded\_Port\_Range\_Per\_IP)$

Using the sample configuration and assuming a subscriber flow sourced from 10.1.1.250:5000:

1.  $Pr\_Offset = 10.1.1.250 - 10.1.0.1 = 505$
2.  $Pu\_Port\_Offset = 505 * 249 = 125,745$
3.  $Rounded\_Port\_Range\_Per\_IP = \lceil (65,533 / 254) \rceil * 249 = 259 * 249 = 64,491$
4.  $Pu\_Prefix = 32.32.32.1 + \text{floor}(125,745 / 64,491) = 32.32.32.1 + 1 = 32.32.32.2$
5.  $Pu\_Start\_Port = 1,024 + (125,745 \% 64,491) = 62278$ 
  - 10.1.1.250 is translated to 32.32.32.2.
  - The starting port is 62278. There are 249 ports available to the subscriber based on the configured block size. The available port range spans ports 62278 through 62526 (inclusive).
  - The specific flow 10.1.1.250:5000 randomly assigns any of the ports in its range because random allocation was specified.

**Reverse Translation**

1.  $Pu\_Offset = Pu\_Prefix - Base\_Pu\_Prefix$
2.  $Pu\_Port\_Offset = (Pu\_Offset * Rounded\_Port\_Range\_Per\_IP) + (Pu\_Actual\_Port - Pu\_Port\_Range\_Start)$
3.  $Subscriber\_IP = Base\_Pr\_Prefix + \text{floor}(Pu\_Port\_Offset / Block\_Size)$

The reverse translation is determined as follows. Assume a flow returning to 32.32.32.2:62278.

1.  $Pu\_Offset = 32.32.32.2 - 32.32.32.1 = 1$
2.  $Pu\_Port\_Offset = (1 * 64,491) + (62,280 - 1024) = 125,747$
3.  $Subscriber\_IP = 10.1.0.1 + \text{floor}(125,747 / 249) = 10.1.0.1 + 505 = 10.1.1.250$



**NOTE:** In reverse translation, only the original private IP address can be derived, and not the original port in use. This is sufficiently granular for law enforcement requirements.

### ***Deterministic Port Block Allocation Restrictions***

When you configure deterministic port block allocation, you must be aware of the following restrictions. Violation of any restriction results in a commit error. The restrictions and their error messages are shown in [Table 93 on page 655](#)

**Table 93: Deterministic Port Block Allocation Commit Constraints**

Restriction	Error Message
The total number of deterministic NAT blocks must be greater than or equal to the 'from' clause addresses configured. This means that the Rounded_Port_Range_Per_IP value must be less than or equal to 64,512.	Number of addresses and port blocks combination in the NAT pool is less than number of addresses in 'from' clause
IPv6 addresses should not be used in deterministic NAT pool/from clause.	Invalid IP address in pool p1 with translation type deterministic-napt44  OR  There is already a range configured with v4 address range
The <b>from</b> clause addresses should be same if the same deterministic NAT pool is used across multiple terms/rules. Only one <b>from</b> clause address/range should be specified if the same deterministic NAT pool is used across multiple terms/rules.	With translation-type deterministic-napt44, same 'from' address/range should be configured if pool is shared by multiple rules or terms
There shouldn't be address overlap between <b>except</b> entries in the <b>from</b> clause addresses.	overlapping address, in the 'from' clause between 'except' entries
A deterministic NAT pool cannot be used with other translation-types	Deterministic NAT pool cannot be used with other translation-types
Deterministic NAPT44 must use a source pool with deterministic-port-block-allocation configuration	Deterministic NAPT44 must use a source pool with deterministic-port-block-allocation configuration
If <b>address-allocation round-robin</b> is configured, a commit results in display of a warning indicating that this technique is not needed with translation-type deterministic-napt44 and is ignored.	Address allocation round-robin is not needed with translation-type deterministic-napt44

Table 93: Deterministic Port Block Allocation Commit Constraints (continued)

Restriction	Error Message
The total number of IP addresses assigned to a deterministic NAT pool should be less than or equal to $2^{24}$ (16777216).	Number of addresses in pool with deterministic-napt44 translation are limited to at most 16777216( $2^{24}$ )

### Comparison of NAPT Implementation Methods

Table 94 on page 656 provides a feature comparison of available NAPT implementation methods.

Table 94: Comparison of NAPT Implementation Methods

Feature/Function	Dynamic Port Allocation	Secured Port Block Allocation	Deterministic Port Block Allocation
Users per IP	High	Medium	Low
Security Risk	Low	Medium	Medium
Log Utilization	High	Low	None (no logs necessary)
Security Risk Reduction	Random allocation	<b>active-block-timeout</b> feature	n/a
Increasing Users per IP	n/a	Configure multiples of smaller port blocks to maximize users/ public IP	Algorithm-based port allocation

## Network Address Translation Rules Overview

To configure a NAT rule, include the **rule** *rule-name* statement at the **[edit services nat]** hierarchy level:

```
[edit services nat]
rule (Services NAT) rule-name {
  match-direction (Services NAT) (input | output);
  term (Services NAT) term-name {
    from (Services NAT) {
      application-sets (Services NAT) set-name;
      applications (Services NAT) [ application-names ];
      destination-address (Services NAT) (address | any-unicast) <except>;
      destination-address-range (Services NAT) low minimum-value high maximum-value
        <except>;
      destination-prefix-list (Services NAT) list-name <except>;
      source-address (address | any-unicast) <except>;
      source-address-range (Services NAT) low minimum-value high maximum-value
        <except>;
      source-prefix-list (Services Stateful Firewall) list-name <except>;
    }
    then (Services NAT) {
      no-translation;
      translated {
        address-pooling paired;
      }
    }
  }
}
```

```

destination-pool nat-pool-name;
destination-prefix (Services NAT) destination-prefix;
dns-alg-pool dns-alg-pool;
dns-alg-prefix dns-alg-prefix;
filtering-type endpoint-independent;
mapping-type endpoint-independent;
overload-pool overload-pool-name;
overload-prefix overload-prefix;
source-pool nat-pool-name;
source-prefix (Services NAT) source-prefix;
translation-type {
    (basic-nat-pt | basic-nat44 | basic-nat66 | dnat-44 | dynamic-nat44 | napt-44 |
     napt-66 | napt-pt | stateful-nat64 | twice-basic-nat-44 | twice-dynamic-nat-44
     | twice-napt-44);
}
}
syslog (Services NAT) ;
}
}
}

```

Each rule must include a **match-direction** statement that specifies the direction in which the match is applied.

In addition, each NAT rule consists of a set of terms, similar to a firewall filter. A term consists of the following:

- **from** statement—Specifies the match conditions and applications that are included and excluded.
- **then** statement—Specifies the actions and action modifiers to be performed by the router software.

The following sections explain how the components of NAT rules:

- [Configuring Match Direction for NAT Rules on page 657](#)
- [Configuring Match Conditions in NAT Rules on page 658](#)
- [Configuring Actions in NAT Rules on page 658](#)
- [Configuring Translation Types on page 659](#)

### Configuring Match Direction for NAT Rules

Each rule must include a **match-direction** statement that specifies the direction in which the match is applied. To configure where the match is applied, include the **match-direction** statement at the **[edit services nat rule *rule-name*]** hierarchy level:

```

[edit services nat rule rule-name]
match-direction (input | output);

```

The match direction is used with respect to the traffic flow through the Multiservices DPC and Multiservices PICs. When a packet is sent to the PIC, direction information is carried along with it. The packet direction is determined based on the following criteria:

- With an interface service set, packet direction is determined by whether a packet is entering or leaving the interface on which the service set is applied.
- With a next-hop service set, packet direction is determined by the interface used to route the packet to the Multiservices DPC or Multiservices PIC. If the inside interface is used to route the packet, the packet direction is input. If the outside interface is used to direct the packet to the PIC or DPC, the packet direction is output. For more information about inside and outside interfaces, see “*Configuring Service Sets to be Applied to Services Interfaces*.”
- On the Multiservices DPC and Multiservices PIC, a flow lookup is performed. If no flow is found, rule processing is performed. All rules in the service set are considered. During rule processing, the packet direction is compared against rule directions. Only rules with direction information that matches the packet direction are considered.

---

### Configuring Match Conditions in NAT Rules

To configure NAT match conditions, include the **from** statement at the **[edit services nat rule *rule-name* term *term-name*]** hierarchy level:

```
[edit services nat rule rule-name term term-name]  
from {  
  application-sets set-name;  
  applications [ application-names ];  
  destination-address (address | any-unicast) <except>;  
  destination-address-range low minimum-value high maximum-value <except>;  
  destination-prefix-list list-name <except>;  
  source-address (address | any-unicast) <except>;  
  source-address-range low minimum-value high maximum-value <except>;  
  source-prefix-list list-name <except>;  
}
```

To configure traditional NAT, you can use the destination address, a range of destination addresses, the source address, or a range of source addresses as a match condition, in the same way that you would configure a firewall filter; for more information, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide*.

Alternatively, you can specify a list of source or destination prefixes by including the **prefix-list** statement at the **[edit policy-options]** hierarchy level and then including either the **destination-prefix-list** or **source-prefix-list** statement in the NAT rule. For an example, see “*Examples: Configuring Stateful Firewall Rules*.”

---

### Configuring Actions in NAT Rules

To configure NAT actions, include the **then** statement at the **[edit services nat rule *rule-name* term *term-name*]** hierarchy level:

```
[edit services nat rule rule-name term term-name]  
then {  
  no-translation;  
  syslog;  
  translated {
```

```

destination-pool nat-pool-name;
destination-prefix destination-prefix;
source-pool nat-pool-name;
source-prefix source-prefix;
translation-type (basic-nat-pt | basic-nat44 | basic-nat66 | dnat-44 | dynamic-nat44
| napt-44 | napt-66 | napt-pt | stateful-nat64 | twice-basic-nat-44 |
twice-dynamic-nat-44 | twice-napt-44);
    }
}
}

```

The **no-translation** statement allows you to specify addresses that you want excluded from NAT.

The **system log** statement enables you to record an alert in the system logging facility.

The **destination-pool**, **destination-prefix**, **source-pool**, and **source-prefix** statements specify addressing information that you define by including the **pool** statement at the **[edit services nat]** hierarchy level.

### Configuring Translation Types

The **translation-type** statement specifies the type of NAT used for source or destination traffic. The options are **basic-nat-pt**, **basic-nat44**, **basic-nat66**, **dnat-44**, **dynamic-nat44**, **napt-44**, **napt-66**, **napt-pt**, **stateful-nat64**, **twice-basic-nat-44**, **twice-dynamic-nat-44**, and **twice-napt-44**.

The implementation details of the nine options of the **translation-type** statement are as follows:

- **basic-nat44**—This option implements the static translation of source IP addresses without port mapping. You must configure the **from source-address** statement in the match condition for the rule. The size of the address range specified in the statement must be the same as or smaller than the source pool. You must specify either a source pool or a destination prefix. The referenced pool can contain multiple addresses but you cannot specify ports for translation.



**NOTE:** In an interface service set, all packets destined for the source address specified in the match condition are automatically routed to the services PIC, even if no service set is associated with the interface.



**NOTE:** Prior to Junos OS Release 11.4R3, you could only use a source NAT pool in a single service set. As of Junos OS Release 11.4R3 and subsequent releases, you can reuse a source NAT pool in multiple service sets.

- **basic-nat66**—This option implements the static translation of source IP addresses without port mapping in IPv6 networks. The configuration is similar to the **basic-nat44** implementation, but with IPv6 addresses.

- **basic-nat-pt**—This option implements translation of addresses of IPv6 hosts, as they originate sessions to the IPv4 hosts in an external domain and vice versa. This option is always implemented with DNS ALG. You must define the source and destination pools of IPv4 addresses. You must configure one rule and define two terms. Configure the IPv6 addresses in the **from** statement in both **term** statements. In the **then** statement of the first term within the rule, reference both the source and destination pools and configure **dns-alg-prefix**. Configure the source prefix in the **then** statement of the second term within the same rule.
- **dnat-44**—This option implements static translation of destination IP addresses without port mapping. The size of the pool address space must be greater than or equal to the destination address space. You must specify a name for the **destination-pool** statement. The referenced pool can contain multiple addresses, ranges, or prefixes, as long as the number of NAT addresses in the pool is larger than the number of destination addresses in the **from** statement. You must include exactly one **destination-address** value at the **[edit services nat rule rule-name term term-name from]** hierarchy level; if it is a prefix, the size must be less than or equal to the pool prefix size. Any addresses in the pool that are not matched in the **yvalue** remain unused, because a pool cannot be shared among multiple terms or rules.
- **dynamic-nat44**—This option implements dynamic translation of source IP addresses without port mapping. You must specify a **source-pool**. The referenced pool must include an **address** configuration (for address-only translation).

The **dynamic-nat44** address-only option supports translating up to 16,777,216 addresses to a smaller size pool. The requests from the source address range are assigned to the addresses in the pool until the pool is used up, and any additional requests are rejected. A NAT address assigned to a host is used for all concurrent sessions from that host. The address is released to the pool only after all the sessions for that host expire. This feature enables the router to share a few public IP addresses between several private hosts. Because all the private hosts might not simultaneously create sessions, they can share a few public IP addresses.

- **napt-44**—This option implements dynamic translation of source IP addresses with port mapping. You must specify a name for the **source-pool** statement. The referenced pool must include a **port** configuration. If the port is configured as automatic or a port range is specified, then it implies that Network Address Port Translation (NAPT) is used.
- **napt-66**—This option implements dynamic address translation of source IP addresses with port mapping for IPv6 addresses. The configuration is similar to the **napt-44** implementation, but with IPv6 addresses.
- **napt-pt**—This option implements dynamic address and port translation for source and static translation of destination IP address. You must specify a name for the **source-pool** statement. The referenced pool must include a port configuration (for NAPT). Additionally, you must configure two rules, one for the DNS traffic and the other for the rest of the traffic. The rule meant for the DNS traffic should be DNS ALG enabled and the **dns-alg-prefix** statement should be configured. Moreover, the prefix configured in the **dns-alg-prefix** statement must be used in the second rule to translate the destination IPv6 addresses to IPv4 addresses.

- **stateful-nat64**—This option implements dynamic address and port translation for source IP addresses and prefix removal translation for destination IP addresses. You must specify the IPv4 addresses used for translation at the **[edit services nat pool]** hierarchy level. This pool must be referenced in the rule that translates the IPv6 addresses to IPv4.
- **twice-basic-nat-44**—This option implements static source and static destination translation for IPv4 addresses, thus combining **basic-nat44** for source and **dnat-44** for destination addresses.
- **twice-dynamic-nat-44**—This option implements source dynamic and destination static translation for IPv4 addresses, combining **dynamic-nat44** for source and **dnat-44** for destination addresses.
- **twice-napt-44**—This option implements source NAPT and destination static translation for IPv4 addresses, combining **napt-44** for source and **dnat-44** for destination addresses.



**NOTE:** When configuring NAT, if any traffic is destined for the following addresses and does not match a NAT flow or NAT rule, the traffic is dropped:

- Addresses specified in the **from destination-address** statement when you are using destination translation
- Addresses specified in the source NAT pool when you are using source translation

For more information on NAT methods, see RFC 2663, *IP Network Address Translator (NAT) Terminology and Considerations*.

## Configuring Service Sets for Network Address Translation

When configuring a service set for NAT processing, make sure you have defined:

- Service interface(s) for handling inbound and outbound traffic



**NOTE:** Prior to Junos OS Release 11.4R3, you could only use a source NAT pool in a single service set. As of Junos OS Release 11.4R3 and subsequent releases, you can reuse a source or destination NAT pool in multiple service sets, provided that the service interfaces associated with the service sets are in different virtual routing and forwarding (VRF) instances.

- For interface style service sets, when a NAT pool is reused in multiple service sets, the service interfaces used in the `interface-service` `service-interface` option of each service set must be in different VRFs.
- For next-hop style service sets, when a NAT pool is reused in multiple service sets, the service interfaces used in the `outside-interface` option of each service set must be in different VRFs.

*Not adhering to these service interface restrictions will cause multiple routes to be installed in the same VRF for the same NAT addresses, causing reverse traffic to be processed incorrectly.*

To enable sharing of source NAT pools, include the `allow-overlapping-nat-pools` statement at the `[edit services nat]` hierarchy level.

- A NAT rule or ruleset



**NOTE:** To configure an MX-DPC interface to be used exclusively for carrier-grade NAT (CGN) or related services (intrusion detection, stateful firewall, and software), include the `cg-pic` statement at the `[edit interfaces interface-name services-options]` hierarchy level.

To configure a NAT service set:

1. At the `[edit services]` hierarchy level, define the service set.

```
[edit services]
user@host# edit service-set service-set-name
```

2. Configure either an interface service, which requires a single service interface, or a next-hop service, which requires an inside and outside service interface.

```
[edit services service-set service-set-name]
user@host# set interface-service service-interface interface-name
```

Or

```
[edit services service-set service-set-name]
user@host# set next-hop-service inside-service-interface interface-name
outside-service-interface interface-name
```



**NOTE:** If you have a Trio-based line card (MPC/MIC), you can use an inline-services interface that was configured on that card, as shown in this example:

```
[edit]
user@host# set interfaces si-0/0/0
[edit services service-set s1]
user@host# set interface-service service-interface si-0/0/0
```

For more information on interface service and next-hop service, see “*Configuring Service Sets to be Applied to Services Interfaces*.”

3. Configure a reference to the NAT rules or ruleset to be used with the service set.

```
[edit services service-set service-set-name]
user@host# set nat-rules rule-or-ruleset-name
```

4. (Optional) For NAT64, specify that the don't fragment (DF) bit for IPv4 packet headers is cleared when packet length is less than 1280 bytes.

```
[edit services service-set service-set-name]
user@host# set nat-options stateful-nat64 clear-dont-fragment-bit
```

**See Also** • *Configuring Service Sets to be Applied to Services Interfaces*

## Junos OS CGNAT Implementation Overview

The Junos OS enables its users to implement and scale their CGNAT (Carrier-Grade Network Address Translation) solutions based on the type of services interfaces used for the implementation.

- MultiServices Denser Port Concentrator (MS-DPC)—The layer 3 services package is used to configure NAT for MS-DPC adaptive services PICs. You must configure the layer-3 services package before implementing NAT on the MS-DPC. This solution provides NAT functionality described in *Network Address Translation Overview for MS-DPC, MS-MPC, and MS-MIC Line Cards*.
- MultiServices Modular Port Concentrator (MS-MPC) and MultiServices Modular Interface Card (MS-MIC)—MS-MPCs and MS-MICs are pre-configured to enable configuration of carrier-grade NAT. This solution provides NAT functionality also described in *Network Address Translation Overview for MS-DPC, MS-MPC, and MS-MIC Line Cards*.
- Inline NAT for Type 1, 2, and 3 Modular Port Concentrator (MPC Line Cards)—Inline NAT leverages the services capabilities of TRIO-based MPC line cards, allowing

cost-effective implementation NAT functionality on the data plane, as described in *Inline Network Address Translation Overview*

- [Network Address Translation Overview for MS-DPC, MS-MPC, and MS-MIC Line Cards on page 664](#)
- [Inline Network Address Translation Overview for MPC Types 1, 2, and 3 on page 668](#)
- [CGNAT Implementations Feature Comparison for Junos Address Aware by Type of Interface Card on page 669](#)
- [ALGs Available by Default for Junos OS Address Aware NAT on page 671](#)

## Network Address Translation Overview for MS-DPC, MS-MPC, and MS-MIC Line Cards

- [Types of NAT on page 664](#)

### Types of NAT

---

The types of Network Address Translation (NAT) supported by the Junos OS are described in the following sections:

- [NAT Concept and Facilities Overview on page 664](#)
- [IPv4-to-IPv4 Basic NAT on page 665](#)
- [Static Destination NAT on page 666](#)
- [Twice NAT on page 666](#)
- [IPv6 NAT on page 666](#)
- [Application-level gateway \(ALG\) Support on page 666](#)
- [NAT-PT with DNS ALG on page 667](#)
- [Dynamic NAT on page 667](#)
- [Stateful NAT64 on page 668](#)

### **NAT Concept and Facilities Overview**

NAT is a mechanism for translating IP addresses. NAT provides the technology used to support a wide range of networking goals, including:

- Concealing a set of host addresses on a private network behind a pool of public addresses.
- Providing a security measure to protect the host addresses from direct targeting in network attacks.
- Providing a tool set for coping with IPv4 address depletion and IPv6 transition issues.

The Junos OS provides carrier-grade NAT (CGN) for IPv4 and IPv6 networks, and facilitates the transit of traffic between different types of networks.



**NOTE:** The Junos OS supports a diverse set of NAT translation options. Not all types of NAT are supported on all interface types.

- Static-source translation—Allows you to hide a private network. It features a one-to-one mapping between the original address and the translated address; the mapping is configured statically.
- Dynamic-source translation—Includes two options: dynamic address-only source translation and Network Address Port Translation (NAPT):
  - Dynamic address-only source translation—A NAT address is picked up dynamically from a source NAT pool and the mapping from the original source address to the translated address is maintained as long as there is at least one active flow that uses this mapping. For more information, see [“Dynamic NAT” on page 667](#).
  - NAPT—Both the original source address and the source port are translated. The translated address and port are picked up from the corresponding NAT pool. For more information, see [“NAPT” on page 666](#).
- Static destination translation—Allows you to make selected private servers accessible. It features a one-to-one mapping between the translated address and the destination address; the mapping is configured statically.
- Protocol translation—Allows you to assign addresses from a pool on a static or dynamic basis as sessions are initiated across IPv4 or IPv6 boundaries.
- Encapsulation of IPv4 packets into IPv6 packets using softwires—Enables packets to travel over softwires to a carrier-grade NAT endpoint where they undergo source-NAT processing to hide the original source address.

The Junos OS supports NAT functionality described in IETF RFCs and Internet drafts, as shown in “Supported NAT and SIP Standards” in [Standard supported in Junos 13.2](#).

#### ***IPv4-to-IPv4 Basic NAT***

Basic Network Address Translation or Basic NAT is a method by which IP addresses are mapped from one group to another, transparent to end users. Network Address Port Translation or NAPT is a method by which many network addresses and their TCP/UDP ports are translated into a single network address and its TCP/UDP ports. Together, these two operations, referred to as traditional NAT, provide a mechanism to connect a realm with private addresses to an external realm with globally unique registered addresses.

Traditional NAT, specified in RFC 3022, *Traditional IP Network Address Translator*, is fully supported by the Junos OS. In addition, NAPT is supported for source addresses.

#### ***Basic NAT***

With Basic NAT, a block of external addresses is set aside for translating addresses of hosts in a private domain as they originate sessions to the external domain. For packets outbound from the private network, Basic NAT translates source IP addresses and related fields such as IP, TCP, UDP, and ICMP header checksums. For inbound packets, Basic NAT translates the destination IP address and the checksums listed above.

### **NAPT**

Use NAPT to enable the components of the private network to share a single external address. NAPT translates the transport identifier (for example, TCP port number, UDP port number, or ICMP query ID) of the private network into a single external address. NAPT can be combined with Basic NAT to use a pool of external addresses in conjunction with port translation.

For packets outbound from the private network, NAPT translates the source IP address, source transport identifier (TCP/UDP port or ICMP query ID), and related fields, such as IP, TCP, UDP, and ICMP header checksums. For inbound packets, NAPT translates the destination IP address, the destination transport identifier, and the IP and transport header checksums.

### **Static Destination NAT**

Use static destination NAT to translate the destination address for external traffic to an address specified in a destination pool. The destination pool contains one address and no port configuration.

For more information about static destination NAT, see RFC 2663, *IP Network Address Translator (NAT) Terminology and Considerations*.

### **Twice NAT**

In Twice NAT, both the source and destination addresses are subject to translation as packets traverse the NAT router. The source information to be translated can be either address only or address and port. For example, you would use Twice NAT when you are connecting two networks in which all or some addresses in one network overlap with addresses in another network (whether the network is private or public). In traditional NAT, only one of the addresses is translated.

To configure Twice NAT, you must specify both a destination address and a source address for the match direction, pool or prefix, and translation type.

You can configure application-level gateways (ALGs) for ICMP and traceroute under stateful firewall, NAT, or class-of-service (CoS) rules when Twice NAT is configured in the same service set. These ALGs cannot be applied to flows created by the Packet Gateway Control Protocol (PGCP). Twice NAT does not support other ALGs. By default, the Twice NAT feature can affect IP, TCP, and UDP headers embedded in the payload of ICMP error messages.

Twice NAT, specified in RFC 2663, *IP Network Address Translator (NAT) Terminology and Considerations*, is fully supported by the Junos OS.

### **IPv6 NAT**

IPv6-to-IPv6 NAT (NAT66), defined in Internet draft draft-mrw-behave-nat66-01, *IPv6-to-IPv6 Network Address Translation (NAT66)*, is fully supported by the Junos OS.

### **Application-level gateway (ALG) Support**

The Junos OS supports a number of ALGs. You can use NAT rules to filter incoming traffic based on ALGS. For more information, see *Network Address Translation Rules Overview*

### NAT-PT with DNS ALG

NAT-PT and Domain Name System (DNS) ALG are used to facilitate communication between IPv6 hosts and IPv4 hosts. Using a pool of IPv4 addresses, NAT-PT assigns addresses from that pool to IPv6 nodes on a dynamic basis as sessions are initiated across IPv4 or IPv6 boundaries. Inbound and outbound sessions must traverse the same NAT-PT router so that it can track those sessions. RFC 2766, *Network Address Translation - Protocol Translation (NAT-PT)*, recommends the use of NAT-PT for translation between IPv6-only nodes and IPv4-only nodes, and *not* for IPv6-to-IPv6 translation between IPv6 nodes or IPv4-to-IPv4 translation between IPv4 nodes.

DNS is a distributed hierarchical naming system for computers, services, or any resource connected to the Internet or a private network. The DNS ALG is an application-specific agent that allows an IPv6 node to communicate with an IPv4 node and vice versa.

When DNS ALG is employed with NAT-PT, the DNS ALG translates IPv6 addresses in DNS queries and responses to the corresponding IPv4 addresses and vice versa. IPv4 name-to-address mappings are held in the DNS with “A” queries. IPv6 name-to-address mappings are held in the DNS with “AAAA” queries.



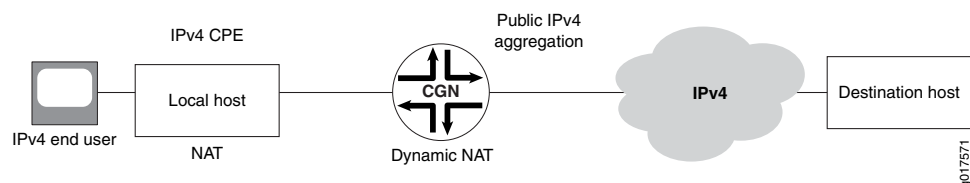
**NOTE:** For IPv6 DNS queries, use the `do-not-translate-AAAA-query-to-A-query` statement at the `[edit applications application application-name]` hierarchy level.

- See Also**
- *Configuring NAT-PT*
  - *Example: Configuring NAT-PT*

### Dynamic NAT

Dynamic NAT flow is shown in [Figure 63 on page 667](#).

**Figure 63: Dynamic NAT Flow**



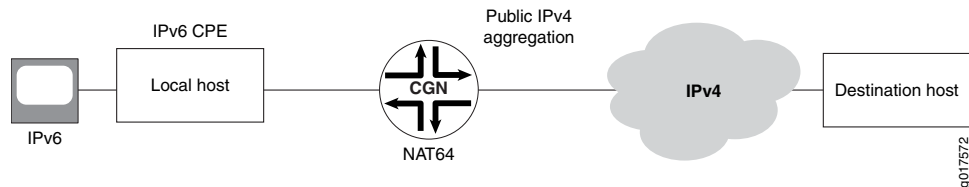
With dynamic NAT, you can map a private IP address (source) to a public IP address drawing from a pool of registered (public) IP addresses. NAT addresses from the pool are assigned dynamically. Assigning addresses dynamically also allows a few public IP addresses to be used by several private hosts, in contrast with an equal-sized pool required by source static NAT.

For more information about dynamic address translation, see RFC 2663, *IP Network Address Translator (NAT) Terminology and Considerations*.

### Stateful NAT64

Stateful NAT64 flow is shown in [Figure 64 on page 668](#).

**Figure 64: Stateful NAT64 Flow**



Stateful NAT64 is a mechanism to move to an IPv6 network and at the same time deal with IPv4 address depletion. By allowing IPv6-only clients to contact IPv4 servers using unicast UDP, TCP, or ICMP, several IPv6-only clients can share the same public IPv4 server address. To allow sharing of the IPv4 server address, NAT64 translates incoming IPv6 packets into IPv4 (and vice versa).

When stateful NAT64 is used in conjunction with DNS64, no changes are usually required in the IPv6 client or the IPv4 server. DNS64 is out of scope of this document because it is normally implemented as an enhancement to currently deployed DNS servers.

Stateful NAT64, specified in RFC 6146, *Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers*, is fully supported by the Junos OS.

### Inline Network Address Translation Overview for MPC Types 1, 2, and 3

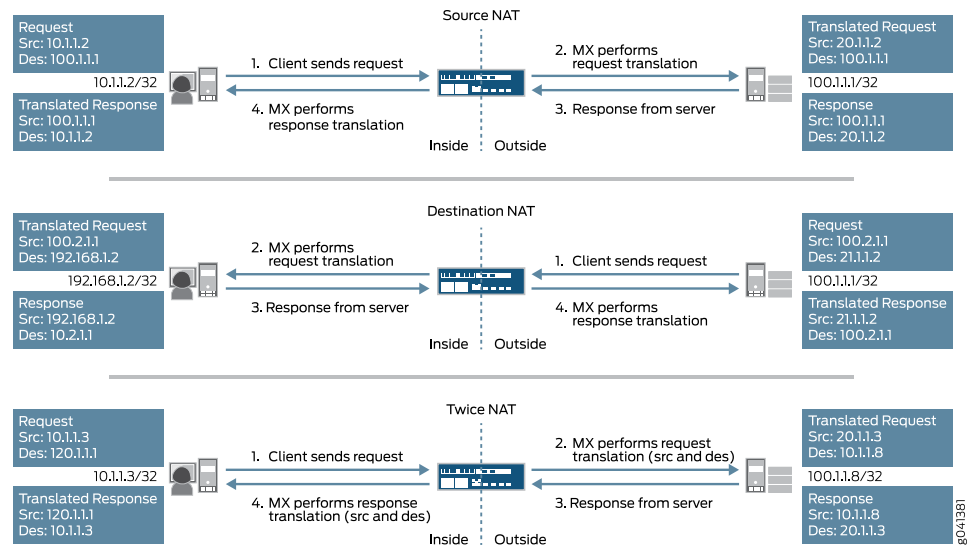
Inline network address translation (NAT) uses the capabilities of the Modular Port Concentrator (MPC) line card, eliminating the need for a MultiServices Dense Port Concentrator (MS-DPC) for NAT. Consequently, you can achieve line-rate, low-latency address translations (up to 120 Gbps per slot). The current implementation provides:

- 1:1 static address mapping
- Bidirectional mapping - source NAT for outbound traffic and destination NAT for inbound traffic
- No limit on number of flows
- Support for Source, destination, and twice NAT, as shown in [Figure 65 on page 669](#)



**NOTE:** Inline NAT generally only the `basic-nat44` translation type, and implements destination NAT and twice NAT by applying NAT at the egress interface or to back-to-back, as shown in the following figure.

Figure 65: Supported Inline NAT Types



To configure inline NAT, you define your service interface as type **si-** (service-inline) interface. You must also reserve adequate bandwidth for the inline interface. This enables you to configure both interface or next-hop service-sets used for NAT. The **si-** interface serves as a “virtual service PIC”.



**NOTE:** Only static source NAT is supported. Port translation and dynamic NAT are not supported. An MS-DPC or MS-PIC will still be needed for any stateful-firewall processing.

## CGNAT Implementations Feature Comparison for Junos Address Aware by Type of Interface Card

Table 95 on page 669 summarizes feature differences between the Junos OS carrier-grade NAT implementations

Table 95: CGNAT Implementation—Feature Comparison by Platform

Feature	MS-DPC	MS-MPC MS-MIC	MPC Types 1, 2, 3 <i>Inline NAT</i>
	MS-100		
	MS-400		
	MS-500		
Static Source NAT	yes	yes	yes
DynamicSource NAT - Address Only	yes	yes	no

Table 95: CGNAT Implementation—Feature Comparison by Platform (continued)

Feature	MS-DPC		
	MS-100		
	MS-400	MS-MPC	MPC Types 1, 2, 3
	MS-500	MS-MIC	<i>Inline NAT</i>
Dynamic Source NAT - NAPT Port Translation with Secured Port Block Allocation	yes	yes	no
Dynamic Source NAT - NAPT Port Translation with Deterministic Port Block Allocation	yes	yes	no
Static Destination NAT	yes	yes	yes  <i>NOTE: Destination NAT can be implemented indirectly. See <a href="#">Inline Network Address Translation Overview</a></i>
Twice NAT	yes	no	yes  <i>NOTE: Twice NAT can be implemented indirectly. See <a href="#">Inline Network Address Translation Overview</a></i>
NAPT - Preserve Parity and Port	yes	no	no
NAPT - EIM/EIF/APP	yes	yes	no
NAT64	yes	yes	no
NAT64 with APP/EIM/EIF	no	yes	no
DS-Lite	yes	no	no
6rd	yes	no	no
Overload Pool/Overlap Address Across NAT Pool	yes	no	no
Port Control Protocol	yes	no	no
CGN-PIC	yes	no	no
AMS Support	no	yes	no

[Table 96 on page 671](#) summarizes availability of translation types by type of interface card.

**Table 96: CGNAT Translation Types**

Translation Type	MS-DPC MS-100 MS-400 MS-500	MS-MPC MS-MIC	MPC Types 1, 2, 3 <i>Inline NAT</i>
<b>basic-nat44</b>	yes	yes	yes
<b>basic-nat66</b>	yes	no	no
<b>basic-nat-pt</b>	yes	no	no
<b>deterministic-napt44</b>	yes	no	no
<b>dnat-44</b>	yes	yes	no
<b>dynamic-nat44</b>	yes	yes	no
<b>napt-44</b>	yes	yes	no
<b>napt-66</b>	yes	no	no
<b>napt-pt</b>	yes	no	no
<b>stateful-nat64</b>	yes	yes	no
<b>twice-basic-nat-44</b>	yes	no	no
<b>twice-dynamic-nat-44</b>	yes	no	no
<b>twice-dynamic-napt-44</b>	yes	no	no

**See Also** • [Junos OS Carrier-Grade NAT Implementation Overview](#)

## ALGs Available by Default for Junos OS Address Aware NAT

The following application-level gateways (ALGs) listed in [Table 97 on page 672](#) are supported for NAT processing on the listed platforms.

To view the implementation details (port, protocol, and so on) for these Junos OS default applications, locate the Junos OS Default ALG Name in the table and then look up the listed name in the **groups**. For example, for details about TFTP, look up **junos-tftp** as shown.



**TIP:** The Junos OS provides the `junos-alg`, which enables other ALGs to function by handling ALG registrations, causing slow path packets to flow through registered ALGs, and transferring ALG events to the ALG plug-ins. The `junos-alg` ALG is automatically available on the MS-MPC and MS-MIC platforms and does not require further configuration.

```
user@host# show groups junos-defaults applications application junos-tftp
application-protocol tftp;
protocol udp;
destination-port 69;
```

**Table 97: ALGs Available by Default**

ALG	MS-DPC	MS-MPC, MS-MIC	Junos OS Default ALG Name
Basic TCP ALG	yes	yes	<b>NOTE:</b> Specific Junos ALGs are not supported. However, a feature called TCP tracker, available by default, performs segment ordering and retransmit and connection tracking, validations for TCP connections.
Basic UDP ALG	yes	yes	<b>NOTE:</b> TCP tracker performs limited integrity and validation checks for UDP.
BOOTP	yes	no	<ul style="list-style-type: none"> <li>• <code>junos-bootpc</code></li> <li>• <code>junos-bootps</code></li> </ul>
DCE RPC Services	yes	yes	<ul style="list-style-type: none"> <li>• <code>junos-dce-rpc-portmap</code></li> <li>• <code>junos-dcerpc-endpoint-mapper-service</code></li> <li>• <code>junos-dcerpc-msexchange-directory-nsp</code></li> <li>• <code>junos-dcerpc-msexchange-directory-rfr</code></li> <li>• <code>junos-dcerpc-msexchange-information-store</code></li> </ul>
DNS	yes	yes	<ul style="list-style-type: none"> <li>• <code>junos-dns-tcp</code></li> <li>• <code>junos-dns-udp</code></li> </ul>
FTP	yes	yes	<ul style="list-style-type: none"> <li>• <code>junos-ftp</code></li> </ul>
H323	yes	no	<ul style="list-style-type: none"> <li>• <code>junos-h323</code></li> </ul>
ICMP	yes	yes	<ul style="list-style-type: none"> <li>• <code>junos-icmp-all</code></li> <li>• <code>junos-icmp-ping</code></li> </ul> <p><b>NOTE:</b> ICMP messages are handled by default, but PING ALG support is not provided.</p>

Table 97: ALGs Available by Default (continued)

ALG	MS-DPC	MS-MPC, MS-MIC	Junos OS Default ALG Name
IIOp	yes	no	<ul style="list-style-type: none"> <li>• junos-iio-p-java</li> <li>• junos-iio-p-orbix</li> </ul>
IP	yes	The TCP tracker, available by default on these platforms, performs limited integrity and validation checks.	<ul style="list-style-type: none"> <li>• junos-ip</li> </ul>
NETBIOS	yes	no	<ul style="list-style-type: none"> <li>• junos-netbios-datagram</li> <li>• junos-netbios-name-tcp</li> <li>• junos-netbios-name-udp</li> <li>• junos-netbios-session</li> </ul>
NETSHOW	yes	no	<ul style="list-style-type: none"> <li>• junos-netshow</li> </ul>
PPTP	yes	yes	<ul style="list-style-type: none"> <li>• junos-pptp</li> </ul>
REALAUDIO	yes	no	<ul style="list-style-type: none"> <li>• junos-realaudio</li> </ul>
Sun RPC and RPC Port Map Services	yes	yes	<ul style="list-style-type: none"> <li>• junos-rpc-portmap-tcp</li> <li>• junos-rpc-portmap-udp</li> </ul>
RTSP	yes	yes	<ul style="list-style-type: none"> <li>• junos-rtsp</li> </ul>
SIP	yes	Yes	<ul style="list-style-type: none"> <li>• junos-sip</li> </ul>
SNMP	yes	No	<ul style="list-style-type: none"> <li>• junos-snmp-get</li> <li>• junos-snmp-get-next</li> <li>• junos-snmp-response junos-snmp-trap</li> </ul>
SQLNET	yes	yes	<ul style="list-style-type: none"> <li>• junos-sqlnet</li> </ul>
TFTP	yes	yes	<ul style="list-style-type: none"> <li>• junos-tftp</li> </ul>
Traceroute	yes	no	<ul style="list-style-type: none"> <li>• junos-traceroute</li> </ul>
Unix Remote Shell Service	yes	Yes	<ul style="list-style-type: none"> <li>• junos-rsh</li> </ul>
WINFrame	yes	No	<ul style="list-style-type: none"> <li>• junos-citrix-winframe</li> <li>• junos-citrix-winframe-udp</li> </ul>
TALK-UDP	No	Yes	<ul style="list-style-type: none"> <li>• junos-talk-udp</li> </ul>

Table 97: ALGs Available by Default (continued)

ALG	MS-DPC	MS-MPC, MS-MIC	Junos OS Default ALG Name
MS RPC	No	Yes	<ul style="list-style-type: none"> <li>• <code>junos-rpc-portmap-tcp</code></li> <li>• <code>junos-rpc-portmap-udp</code></li> <li>• <code>junos-rpc-services-tcp</code></li> <li>• <code>junos-rpc-services-udp</code></li> </ul>

**Related Documentation** • [Network Address Translation Configuration Overview on page 645](#)

## Service Redundancy Daemon Overview

- [Introduction to the Service Redundancy Daemon on page 674](#)
- [Service Redundancy Daemon Components on page 674](#)
- [Service Redundancy Daemon Constraints on page 675](#)
- [Service Redundancy Daemon Operation on page 676](#)

### Introduction to the Service Redundancy Daemon

- The service redundancy daemon (srd) provides configurable redundancy across multiple gateways on MX Series routers with MPC. You can configure redundancy based on monitored events, including:
  - Link down events.
  - FPC and PIC reboots.
  - Routing protocol daemon (rpd) aborts and restarts.
  - Peer gateway events, including requests to acquire or release mastership, or to broadcast warnings.

The srd also enables you to manage stateful sync session synchronization across gateways.

### Service Redundancy Daemon Components

The following configurable components control srd processing:

- **Redundancy Event (RE)**—A monitored critical event that triggers the srd to acquire or release mastership for redundancy peers, or to trigger warning-only events, and to add or delete signal routes. Monitored events include interface or link down events, rpd events, and acquire or release mastership events from peers.
- **Redundancy Policy (RP)**—A policy that defines the set of actions taken when a redundancy event occurs. Available actions include acquisition or release of mastership, and addition or deletion of signal routes.
- **Redundancy Set (RS)**—A collection of more or more service sets with a common redundancy policy or policies. A redundancy set applies to two or more system

gateways. Only one of the gateways is master and the peer or peers are standby at any time. Redundancy policies define the actions to be taken for an RS when the srd detects a triggering event.

- **Redundancy Group (RG)**—A collection of redundancy sets that defines common peering properties across a set of gateways. Redundancy groups allow for different peering settings across same peers.



**NOTE:** In the current implementation, a one-to-one relationship exists between redundancy set and redundancy group.

- **Signal routes**—Static routes that are added or deleted by the srd based on mastership state changes.
- **Routing Policies**—Policies that are configured to advertise routes based on the existence or non-existence of signal routes using the **if-route-exists** condition.
- **VRRP (Virtual Router Redundancy Protocol) route tracking**—A standard Junos OS VRRP feature, but optional srd component, that tracks whether a reachable route exists in the routing table of the routing instance included in the configuration and dynamically changes the priority of the VRRP group based on the reachability of the tracked route, triggering a new master router election. The route to be tracked is the a signal route.

## Service Redundancy Daemon Constraints

The following constraints apply to srd processing configurations:

- A one-to-one relationship exists between a redundancy set (RS) and a redundancy group (RG). One RS can be part of only one RG.
- One redundancy policy (RP) can be part of only one redundancy set (RS), but one redundancy set can have multiple redundancy policies. For example, redundancy set RS1 can include redundancy policies RP1 and RP2. Redundancy policies RP1 and RP2 cannot be included in redundancy sets other than RS1.
- One redundancy event (RE) can be part of only one redundancy policy (RP), but one redundancy policy can have multiple redundancy events. For example, redundancy policy RP1 can include redundancy events RE1 and RE2. Redundancy events RE1 and RE2 cannot be included in redundancy policies other than RP1.
- One monitored interface or link can be part of only one redundancy event (RE), but one redundancy event can have multiple monitored interfaces.
- One service set (SS) can be part of only one redundancy set (RS), but one redundancy set may have multiple service sets.

## Service Redundancy Daemon Operation

The srd operates as follows:

1. The srd runs on the Routing Engine. It continuously monitors configured redundancy events.
2. When a redundancy event is detected, the srd:
  - a. Adds or removes signal routes specified in the redundancy policy.
  - b. Switches services to the next preferred standby gateway.
  - c. Updates stateful sync roles as needed.
3. Resulting route changes cause:
  - a. The routing policy connected to this route to advertise routes differently.
  - b. VRRP to change advertised priorities.

To summarize the switchover process:

1. A critical event occurs.
2. srd adds or removes a signal route.
3. A routing policy advertises routes differently. VRRP changes advertised priorities.
4. Services switch over to the next preferred standby gateway.
5. Stateful sync is updated accordingly.



**NOTE:** The order of routing priorities must match the order of services mastership.

---

## Configuring the Service Redundancy Daemon

---

Before you configure srd processing, we recommend that you be familiar with [Configuring ICCP for Multichassis Link Aggregation](#), which explains peer relationships between gateways that are enabled to exchange master and standby roles.

You use the following configuration statements:

- **redundancy-policy** under the **[edit policy-options]** hierarchy level
- **redundancy-event** under the **[edit event-options]** hierarchy level
- **redundancy-set** under the **[edit services]** hierarchy level

The actions to be performed when configured redundancy events occur are defined in redundancy policies. Redundancy policies are associated with redundancy sets; they are analogous to rules associated with service sets. Redundancy sets are associated to redundancy groups by redundancy group IDs. Redundancy group details are defined by

the underlying ICCPd configuration. Finally, service sets and redundancy sets are associated through the **redundancy-sets** statement in service sets configuration.

To configure **srd**, perform the following configuration tasks in the recommended sequence. Configurations are show for two gateways for which mastership may change.

The procedures that follow, redundancy events that are configured and associated with a redundancy policy. The redundancy policy is associated with a redundancy set to take appropriate action of mastership-release or mastership-acquire. If an event is associated with a policy that takes the **release-mastership** action, **srd** checks whether the redundancy peer's state is ready or warned. If the standby is in a warned state, then the **release-mastership** action fails. You can take restore the healthcheck and manually execute the **release-mastership** action.

To release mastership in any case, you can either configure the policy action as **release-mastership-force** or use **force** option in the operational CLI. Even if your configuration specifies the **force** option, using the force option in the CLI takes precedence and mastership is released. Similarly, if a redundancy event is configured with a policy with an **acquire-mastership** action, then **srd** checks the local redundancy set state. In the case of a wait state, the action fails unless the **force** option is used. We recommend that you determine why health checks fail and take action to correct the failure. After that, when the redundancy set state returns to STANDBY, then this mastership change action succeeds.

- [Configuring Redundancy Events on page 677](#)
- [Configuring Redundancy Policies on page 678](#)
- [Configuring Redundancy Set and Group on page 680](#)
- [Configuring Routing Policies Supporting Redundancy on page 681](#)
- [Configuring Service Sets on page 682](#)

## Configuring Redundancy Events

To configure redundancy events:

1. Configure any link-down redundancy events for the master gateway.

```
user@gateway1# set event-options redundancy-event redundancy-event monitor
link-down link-down
```

For example:

```
user@gateway1# set event-options redundancy-event RELS_MSHIP_CRIT_EV monitor
link-down ms-2/3/0.0
user@gateway1# set event-options redundancy-event RELS_MSHIP_CRIT_EV monitor
link-down xe-3/0/0.0
```

2. Configure any process redundancy events for the master gateway.

```
user@gateway1# set event-options redundancy-event redundancy-event monitor
process routing restart
```

For example:

```
user@gateway1# set event-options redundancy-event RELS_MSHIP_CRIT_EV monitor
process routing restart
```

3. Configure any link-down redundancy events for the standby gateway.

```
user@gateway2# set event-options redundancy-event redundancy-event monitor
link-down link-down
```

For example:

```
user@gateway2# set event-options redundancy-event WARN_EV monitor link-down
ms-2/3/0.0
user@gateway2# set event-options redundancy-event WARN_EV monitor link-down
xe-3/0/0.0
```

4. Configure any process redundancy events for the standby gateway.

```
user@gateway2# set event-options redundancy-event redundancy-event monitor
process routing restart
```

For example:

```
user@gateway2# set event-options redundancy-event WARN_EV monitor process
routing restart
```

5. Configure any peer redundancy events for the standby gateway.

```
user@gateway2# set event-options redundancy-event redundancy-event monitor peer
(mastership-acquire | mastership-release)
```

For example:

```
user@gateway2# set event-options redundancy-event PEER_MSHIP_ACQU_EV monitor
peer mastership-acquire
user@gateway2# set event-options redundancy-event PEER_MSHIP_RELS_EV monitor
peer mastership-release
```

## Configuring Redundancy Policies

Service redundancy policies specify actions triggered by monitored redundancy events.

To configure redundancy policies:

1. Specify a redundancy policy and redundancy event for the master gateway. Follow the same steps for the standby gateway.

```
user@gateway1# edit policy-options redundancy-policy policy-name redundancy-event
event-name then
```

- Specify an action of acquiring or releasing mastership.

```
user@gateway1# set acquire-mastership
```

or

```
user@gateway1# set (release-mastership | release-mastership-force |
release-mastership-if-standby-clear
```

- (Optional) Specify an action of adding a static route.

```
user@gateway1# set add-static-route destination (receive | next-hop next-hop)
routing-instance vrf-name
```



**BEST PRACTICE:** We recommend using the receive option.

- (Optional) Specify an action of deleting a static route.

```
user@gateway1# set delete-static-route destination routing-instance vrf-name
```

The following example demonstrates configuring redundancy policies for two peer gateways:

```
user@gateway1# edit policy-options redundancy-policy ACQU_MSHIP_POL
redundancy-events ACQU_MSHIP_MANUAL_EV then
```

```
[edit policy-options redundancy-policy ACQU_MSHIP_POL redundancy-event
ACQU_MSHIP_MANUAL_EV then]
```

```
user@gateway1# set acquire-mastership add-static-route 10.45.45.0/24 receive
routing-instance SGI-PRIVATE
```

```
user@gateway1# top
```

```
user@gateway1# edit policy-options redundancy-policy RELS_MSHIP_POL
redundancy-events PEER_MSHIP_ACQU_EV then
```

```
[edit policy-options redundancy-policy RELS_MSHIP_POL redundancy-events
PEER_MSHIP_ACQU_EV then]
```

```
user@gateway1# set release-mastership-force delete-static-route 10.45.45.0/24 receive
routing-instance SGI-PRIVATE
```

```
user@gateway2# edit policy-options redundancy-policy RELS_MSHIP_POL
redundancy-events PEER_MSHIP_ACQU_EV then
```

```
[edit policy-options redundancy-policy ACQU_MSHIP_POL redundancy-event
ACQU_MSHIP_MANUAL_EV then]
```

```
user@gateway2# set release-mastership-force add-static-route 10.45.45.0/24 receive
routing-instance SGI-PRIVATE
```

```

user@gateway2# top
user@gateway2# edit policy-options redundancy-policy ACQU_MSHIP_POL
redundancy-events PEER_MSHIP_RELS_EV then

[edit policy-options redundancy-policy ACQU_MSHIP_POL redundancy-events
PEER_MSHIP_RELS_EV then]
user@gateway2# set acquire-mastership delete-static-route 10.45.45.0/24 receive
routing-instance SGI-PRIVATE
user@gateway2# top
user@gateway2# edit policy-options redundancy-policy WARN_POL redundancy-events
WARN_EV then

[edit policy-options redundancy-policy WARN_POL redundancy-events WARN_EV then]
user@gateway2# set broadcast-warning

```

## Configuring Redundancy Set and Group

The redundancy group IDs that `srd` uses are associated with those configured for the ICCP daemon (`iccpd`) through the existing ICCP configuration hierarchy by using the same redundancy group ID in the configuration of the services redundancy group.

```

iccp {
  local-ip-addr 10.1.1.1;
  peer 10.2.2.2 {
    redundancy-group-id-list 1;
    liveness-detection {
      minimum-interval 1000;
    }
  }
}

```

To configure redundancy sets:

1. Specify redundancy set and group for the master gateway.

```

user@gateway1# set redundancy-set redundancy-set redundancy-group
redundancy-group

```

For example:

```

user@gateway1# set redundancy-set 1 redundancy-group 1

```

2. Specify redundancy policies for the redundancy set.

```

user@gateway1# set redundancy-set redundancy-set redundancy-policy
[redundancy-policy-list]

```

For example:

```

user@gateway1# set redundancy-set 1 redundancy-policy ACQU_MSHIP_POL
RELS_MSHIP_POL WARN_POL

```

3. Specify redundancy set and group for the peer gateway.

```
user@gateway2# set redundancy-set redundancy-set redundancy-group
redundancy-group
```

For example:

```
user@gateway2# set redundancy-set 1 redundancy-group 1
```

4. Specify redundancy policies for the redundancy set.

```
user@gateway2# set redundancy-set redundancy-set redundancy-policy
[redundancy-policy-list]
```

For example:

```
user@gateway1# set redundancy-set 1 redundancy-policy [ACQU_MSHIP_POL
RELS_MSHIP_POL WARN_POL]
```

## Configuring Routing Policies Supporting Redundancy

To configure routing policies that support redundancy:

1. At the **[edit policy-options condition]** hierarchy level, use the **if-route-exists** configuration statement set a condition based on the existence of signal routes that requires redundancy-related routing changes. Specify the routing table that includes

```
[edit policy-options condition condition-name]
user@gateway# set if-route-exists signal-route table routing-table
```

For example:

```
[edit policy-options condition switchover-route-exists]
user@gateway# set if-route-exists 10.45.45.0/24 table bgp1_table
```

2. At the **[edit policy-options policy-statement *statement-name*]** hierarchy level, specify routing changes based on the condition indicating the existence of the signal route. For BGP, routing changes typically include change to local-preference and as-path-prepend values.
  - a. To change local-preference, specify local-preference in the **then** clause of the policy statement.

```
[edit policy-options policy-statement policy-name]
user@gateway# set term term from protocol [protocol variables] prefix-list prefix-list
condition condition-name then local-preference preference-value accept
```

For example:

```
[edit policy-options policy-statement ha-export-v6-policy]
```

```
user@gateway# set term update-local-pref from protocol static bgp prefix-list
ipv4-default-route condition switchover-route-exists then local-preference 350
accept
```

- b. To change **as-path-prepend** values, specify **as-path-prepend** in the **then** clause of the policy statement.

```
[edit policy-options policy-statement policy-name]
user@gateway# set term term from prefix-list prefix-list condition condition-name
then as-path-prepend [as-prepend-values] next-hop self accept
```

For example:

```
[edit policy-options policy-statement ha-export-v6-policy]
user@gateway# set term update-as-prepend prefix-list ipv6-default-route condition
switchover-route-exists then as-path-prepend "64674 64674 64674 64674"
next-hop self accept
```

## Configuring Service Sets

Specify stateful sync of services for a service set.

1. Specify the service set and redundancy-set.

```
[edit]
user@gateway1# set services service-set service-set redundancy-set redundancy-set
```

For example:

```
[edit]
user@gateway1# set services service-set CGN4_SP-7-0-0 redundancy-set 1
```

2. Specify the replication threshold and services to be replicated.

```
[edit]
user@gateway1# set services service-set service-set replicate-services
replication-threshold replication-threshold <stateful-firewall> <nat>
```

For example:

```
[edit]
user@gateway1# set services service-set service-set replicate-services
replication-threshold 360 stateful-firewall nat
```

**Related Documentation**

- [Service Redundancy Daemon Overview on page 674](#)

## Application Layer Gateways Overview

This topic describes the Application Layer Gateways (ALGs) supported by Junos OS. ALG support includes managing pinholes and parent-child relationships for the supported ALGs. This topic includes the following sections:

- [Supported ALGs on page 683](#)
- [ALG Support Details on page 684](#)
- [Juniper Networks Defaults on page 693](#)
- [Examples: Referencing the Preset Statement from the Junos Default Group on page 704](#)

### Supported ALGs

Table 98 on page 683 lists ALGs supported by Junos OS.

**Table 98: ALGs Supported by Junos OS**

ALGs Supported	v4 - v4	v4 - v6	v6 - v6	DS-Lite
Basic TCP ALG	Yes	Yes	Yes	Yes
Basic UPD ALG	Yes	Yes	Yes	Yes
BOOTP	Yes	No	No	No
DCE RPC Services	Yes	No	No	No
DNS	Yes	Yes	No	No
FTP	Yes	No	No	Yes
H323	Yes	No	No	No
ICMP	Yes	Yes	Yes	Yes
IIOP	Yes	No	No	No
IP	Yes	No	No	No
NETBIOS	Yes	No	No	No
NETSHOW	Yes	No	No	No
PPTP	Yes	No	No	Yes
REALAUDIO	Yes	No	No	No
Sun RPC and RPC Port Map Services	Yes	No	No	No

*Table 98: ALGs Supported by Junos OS (continued)*

ALGs Supported	v4 - v4	v4 - v6	v6 - v6	DS-Lite
RTSP	Yes	No	No	Yes
SIP	Yes	No	No	No
SNMP	Yes	No	No	No
SQLNET	Yes	No	No	No
TFTP	Yes	No	No	Yes
Traceroute	Yes	Yes	No	Yes
Unix Remote Shell Service	Yes	No	No	No
WINFrame	Yes	No	No	No

## ALG Support Details

This section includes details about the ALGs. It includes the following:

- [Basic TCP ALG on page 685](#)
- [Basic UDP ALG on page 685](#)
- [BOOTP on page 686](#)
- [DCE RPC Services on page 686](#)
- [DNS on page 686](#)
- [FTP on page 686](#)
- [H323 on page 687](#)
- [ICMP on page 687](#)
- [IIOP on page 688](#)
- [IP on page 688](#)
- [NetBIOS on page 688](#)
- [NetShow on page 688](#)
- [ONC RPC Services on page 688](#)
- [PPTP on page 689](#)
- [RealAudio on page 689](#)
- [Sun RPC and RPC Portmap Services on page 690](#)
- [RTSP on page 691](#)
- [SIP on page 691](#)
- [SNMP on page 692](#)
- [SQLNet on page 692](#)

- [TFTP on page 692](#)
- [Traceroute on page 692](#)
- [UNIX Remote-Shell Services on page 693](#)
- [Winframe on page 693](#)

---

### Basic TCP ALG

This ALG performs basic sanity checking on TCP packets. If it finds errors, it generates the following anomaly events and system log messages:

- TCP source or destination port zero
- TCP header length check failed
- TCP sequence number zero and no flags are set
- TCP sequence number zero and FIN/PSH/RST flags are set
- TCP FIN/RST or SYN(URG|FIN|RST) flags are set

The TCP ALG performs the following steps:

1. When the router receives a SYN packet, the ALG creates TCP forward and reverse flows and groups them in a *conversation*. It tracks the TCP three-way handshake.
2. The SYN-defense mechanism tracks the TCP connection establishment state. It expects the TCP session to be established within a small time interval (currently 4 seconds). If the TCP three-way handshake is not established in that period, the session is terminated.
3. A keepalive mechanism detects TCP sessions with nonresponsive endpoints.
4. ICMP errors are allowed only if there is a flow that matches the selector information specified in the ICMP data.

---

### Basic UDP ALG

This ALG performs basic sanity checking on UDP headers. If it finds errors, it generates the following anomaly events and system log messages:

- UDP source or destination port 0
- UDP header length check failed

The UDP ALG performs the following steps:

1. When it receives the first packet, the ALG creates bidirectional flows to accept forward and reverse UDP session traffic.
2. If the session is idle for more than the maximum allowed idle time (the default is 30 seconds), the flows are deleted.
3. ICMP errors are allowed only if there is a flow that matches the selector information specified in the ICMP data.

## BOOTP

---

The Bootstrap Protocol (BOOTP) client retrieves its networking information from a server across the network. It sends out a general broadcast message to request the information, which is returned by the BOOTP server. For the protocol specification, see <ftp://ftp.isi.edu/in-notes/rfc951.txt>.

Stateful firewall support requires that you configure the BOOTP ALG on UDP server port 67 and client port 68. If the client sends a broadcast message, you should configure the broadcast address in the **from** statement of the service rule. Network Address Translation (NAT) is not performed on the BOOTP traffic, even if the NAT rule matches the traffic. If the BOOTP relay feature is activated on the router, the remote BOOTP server is assumed to assign addresses for clients masked by NAT translation.

## DCE RPC Services

---

Distributed Computing Environment (DCE) Remote Procedure Call (RPC) services are mainly used by Microsoft applications. The ALG uses well-known TCP port 135 for port mapping services, and uses the universal unique identifier (UUID) instead of the program number to identify protocols. The main application-based DCE RPC is the Microsoft Exchange Protocol.

Support for stateful firewall and NAT services requires that you configure the DCE RPC portmap ALG on TCP port 135. The DCE RPC ALG uses the TCP protocol with application-specific UUIDs.

## DNS

---

The Domain Name Service (DNS) ALG handles data associated with locating and translating domain names into IP addresses. The ALG typically runs on port 53. The ALG monitors DNS query and reply packets and supports only UDP traffic. The ALG does not support payload translations. The DNS ALG will only close the session when a reply is received or an idle timeout is reached.

## FTP

---

FTP is the File Transfer Protocol, specified in RFC 959. In addition to the main control connection, data connections are also made for any data transfer between the client and the server; and the host, port, and direction are negotiated through the control channel.

For non-passive-mode FTP, Junos OS stateful firewall service scans the client-to-server application data for the PORT command, which provides the IP address and port number to which the server connects. For passive-mode FTP, Junos OS stateful firewall service scans the client-to-server application data for the PASV command and then scans the server-to-client responses for the 227 response, which contains the IP address and port number to which the client connects.

There is an additional complication: FTP represents these addresses and port numbers in ASCII. As a result, when addresses and ports are rewritten, the TCP sequence number might be changed, and thereafter the NAT service needs to maintain this delta in SEQ and ACK numbers by performing sequence NAT on all subsequent packets.

Support for stateful firewall and NAT services requires that you configure the FTP ALG on TCP port 21 to enable the FTP control protocol. The ALG performs the following tasks:

- Automatically allocates data ports and firewall permissions for dynamic data connection
- Creates flows for the dynamically negotiated data connection
- Monitors the control connection in both active and passive modes
- Rewrites the control packets with the appropriate NAT address and port information

On MS-MPCs and MS-MICs, for passive FTP to work properly without FTP application layer gateway (ALG) enabled (by not specifying the **application junos-ftp** statement at the **[edit services stateful-firewall rule rule-name term term-name from]** and the **[edit services nat rule rule-name term term-name from]** hierarchy levels), you must enable the address pooling paired (APP) functionality enabled (by including the **address-pooling** statement at the **[edit services nat rule rule-name term term-name then translated]** hierarchy level). Such a configuration causes the data and control FTP sessions to receive the same NAT address.

### H323

H323 is a suite of ITU protocols for audio and video conferencing and collaboration applications. H323 consists of H.225 call signaling protocols and H.245 control protocol for media communication. During H.225 negotiation, the endpoints create a call by exchanging call signaling messages on the control channel and negotiate a new control channel for H.245. A new control connection is created for H.245 messages. Messages are exchanged on the H.245 control channel to open media channels.

Stateful firewall monitors the H.225 control channel to open the H.245 control channel. After the H.245 channel is created, stateful firewall also monitors this channel for media channel information and allows the media traffic through the firewall.

H323 ALG supports static destination, static and dynamic source NAT by rewriting the appropriate addresses and ports in the H.225 and H.245 messages.

### ICMP

The Internet Control Message Protocol (ICMP) is defined in RFC 792. The Junos OS stateful firewall service allows ICMP messages to be filtered by specific type or specific type code value. ICMP error packets that lack a specifically configured type and code are matched against any existing flow in the opposite direction to check for the legitimacy of the error packet. ICMP error packets that pass the filter matching are subject to NAT translation.

The ICMP ALG always tracks ping traffic statefully using the ICMP sequence number. Each echo reply is forwarded only if there is an echo request with the corresponding sequence number. For any ping flow, only 20 echo requests can be forwarded without receiving an echo reply. When you configure dynamic NAT, the PING packet identifier is translated to allow additional hosts in the NAT pool to use the same identifier.

Support for stateful firewall and NAT services requires that you configure the ICMP ALG if the protocol is needed. You can configure the ICMP type and code for additional filtering.

## **IIOp**

---

The Oracle Application Server NameServer Internet Inter-ORB Protocol (IIOp). This ALG is used in Common Object Request Broker Architecture (CORBA) based on distributed computing. Even though CORBA and IIOp are Object Management Group (OMG) standards, there is no fixed port assigned for IIOp. Each vendor implementing CORBA chooses a port. Java Virtual machine uses port 1975 by default, while ORBIX uses port 3075 as a default.

Stateful firewall and NAT require ALG IIOp be configured for TCP port 1975 for Java VM IIOp, and 3075 for CORBA applications ORBIX, a CORBA framework from Iona Technologies.

## **IP**

---

The IP ALG is used to create uni-directional flows only. In case of TCP traffic, it does not check the 3-way handshake process. This ALG is useful in case of stateful firewall only service sets, where it allows traffic to flow uni-directionally only. When configuring in conjunction with **match-direction input-output** it allows the return traffic to flow through the stateful firewall as well. Typical scenarios are static NAT, destination NAT or scenarios where traffic is expected to traverse the stateful firewall in the presence of asymmetric routing. The Junos IP ALG is not intended for use with NAPT, which will cause matching traffic to be discarded through the creation of a drop flow.

## **NetBIOS**

---

A NetBIOS ALG translates NetBIOS IP addresses and port numbers when NAT is used.

NetBIOS supports the TCP and UDP transport protocols. Support for stateful firewall and NAT services requires that you configure the NetBIOS ALG on UDP port 138 and TCP port 139.

## **NetShow**

---

The Microsoft protocol ms-streaming is used by NetShow, the Microsoft media server. This protocol supports several transport protocols: TCP, UDP, and HTTP. The client starts a TCP connection on port 1755 and sends the PORT command to the server. The server then starts UDP on that port to the client. Support for stateful firewall and NAT services requires that you configure the NetShow ALG on UDP port 1755.

## **ONC RPC Services**

---

Open Networks Computing (ONC) RPC services function similarly to DCE RCP services. However, the ONC RPC ALG uses TCP/UDP port 111 for port mapping services, and uses the program number to identify protocols rather than the UUID.

Support for stateful firewall and NAT services requires that you configure the ONC RPC portmap ALG on TCP port 111. The ONC RPC ALG uses the TCP protocol with application-specific program numbers.

## PPTP

The Point-to-Point Tunneling Protocol (PPTP) ALG is a TCP-based ALG. PPTP allows the Point-to-Point Protocol (PPP) to be tunneled through an IP network. PPTP defines a client-server architecture, a PPTP Network Server, and a PPTP Access Concentrator. The PPTP ALG requires a control connection and a data tunnel. The control connection uses TCP to establish and disconnect PPP sessions, and runs on port 1723. The data tunnel carries PPP traffic in generic routing encapsulated (GRE) packets that are carried over IP.

## RealAudio

Real Networks PNA protocol RealVideo is not a separate service. It is part of the RealPlayer and most likely uses another channel for video. The RealPlayer versions G2, 7, and 8 use PNA and RTSP. For this version to work, the ALG must allow both PNA(7070) and RTSP(554). For the media, the server selects from a range of UDP ports(6970 through 7170), or TCP port 7071, or HTTP. The client can be configured to use a particular port. The RealPlayer versions 4.0 and 5.0 use control channel 7070 media UDP ports 6970 through 7170, or TCP port 7071, or HTTP. RealAudio player version 3.0 uses control channel 7070 media, UDP ports 6770-7170, or TCP port 7071.

Real products use the ports and ranges of ports shown in [Table 99 on page 689](#).

**Table 99: RealAudio Product Port Usage**

Real Product	Port Usage
4.0 and 5.0 Servers/Players	Control channel (bidirectional) on TCP port 7070. Data channel from server to player on TCP port 7070 or UDP port 6970-7170.
4.0 and 5.0 Servers/Encoders	Control channel (bidirectional) on TCP port 7070. Data channel from encoder or server on TCP port 7070.
G2 Servers/Players	Control channel (bidirectional) on TCP port 80, 554, 7070, or 8080. Data channel from server to player on TCP port 80, 554, 7070, 8080 or UDP port 6970-32,000.
G2 Server/3.1, and 5.x Encoders	Control channel (bidirectional) on TCP port 7070. Data channel from encoder to server on TCP port 7070.
G2 Server/G2 Producer	Control channel (bidirectional) on TCP port 4040. Data channel from encoder to server on TCP port 4040 and UDP port 6970-32,000.
2 Server/G2 Producer (TCP ONLY)	Control channel (bidirectional) on TCP port 4040 Data channel from encoder to server on TCP port 4040. Note: TCP-ONLY option available in version 6.1 or above.



**NOTE:** RealAudio was the original protocol by RealPlayers. Newer versions of RealPlayer use RTSP. Stateful firewall and NAT require ALG RealAudio to be programmed on TCP port 7070.

## Sun RPC and RPC Portmap Services

The Remote Procedure Call (RPC) ALG uses well-known ports TCP 111 and UDP 111 for port mapping, which dynamically assigns and opens ports for RPC services. The RPC Portmap ALG keeps track of port requests and dynamically opens the firewall for these requested ports. The RPC ALG can further restrict the RPC protocol by specifying allowed program numbers.

The ALG includes the RPC services listed in [Table 100 on page 690](#).

**Table 100: Supported RPC Services**

Name	Description	Comments
<b>rpc.mountd</b>	Network File Server (NFS) mount daemon; for details, see the UNIX man page for <b>rpc.mountd(8)</b> .	The base support is RPC v2 and the port mapper service on port 111 (see RFC 1050).
<b>rpc.nfsprog</b>	Used as part of NFS. For details, see RFC 1094. See also RFC1813 for NFS v3.	The base support is RPC v2 and the port mapper service on port 111 (see RFC 1050).
<b>rpc.nisplus</b>	Network Information Service Plus (NIS+), designed to replace NIS; it is a default naming service for Sun Solaris and is not related to the old NIS. No protocol information is available.	The base support is RPC v2 and the port mapper service on port 111 (see RFC 1050).
<b>rpc.nlockmgr</b>	Network lock manager.	The base support is RPC v2 and the port mapper service on port 111 (see RFC 1050). Once the RPC program table is built, <b>rpc.nlockmgr</b> service can be allowed or blocked based on RPC program 100021.
<b>rpc.pcnfsd</b>	Kernel statistics server. For details, see the UNIX man pages for <b>rstatd</b> and <b>rpc.rstatd</b> .	The base support is RPC v2 and the port mapper service on port 111 (see RFC 1050). Once the RPC program table is built, <b>rpc.rstatd</b> service can be allowed or blocked based on RPC program 150001.
<b>rpc.rwall</b>	Used to write a message to users; for details, see the UNIX man page for <b>rpc.rwalld</b> .	The base support is RPC v2 and the port mapper service on port 111 (see RFC 1050). Once the RPC program table is built, <b>rpc.rwall</b> service can be allowed or blocked based on RPC program 150008.
<b>rpc.yplibd</b>	NIS binding process. For details, see the UNIX man page for <b>ypbind</b> .	The base support is RPC v2 and the port mapper service on port 111 (see RFC 1050). Once the RPC program table is built, <b>rpc.yplibd</b> service can be allowed or blocked based on RPC program 100007.
<b>rpc.yppasswd</b>	NIS password server. For details, see the UNIX man page for <b>yppasswd</b> .	The base support is RPC v2 and the port mapper service on port 111 (see RFC 1050). Once the RPC program table is built, <b>rpc.yppasswd</b> service can be allowed or blocked based on RPC program 100009.
<b>rpc.ypserv</b>	NIS server. For details, see the UNIX man page for <b>ypserv</b> .	The base support is RPC v2 and the port mapper service on port 111 (see RFC 1050). Once the RPC program table is built, <b>rpc.ypserv</b> service can be allowed or blocked based on RPC program 100004.

Table 100: Supported RPC Services (continued)

Name	Description	Comments
<b>rpc-ypupdated</b>	Network updating tool.	The base support is RPC v2 and the port mapper service on port 111 (see RFC 1050). Once the RPC program table is built, <b>rpc-ypupdated</b> service can be allowed or blocked based on RPC program 100028.
<b>rpc-ypxfrd</b>	NIS map transfer server. For details, see the UNIX man page for <b>rpc.ypxfrd</b> .	The base support is RPC v2 and the port mapper service on port 111 (see RFC 1050). Once the RPC program table is built, <b>rpc-ypxfrd</b> service can be allowed or blocked based on RPC program 100069.

Support for stateful firewall and NAT services that use port mapping requires that you configure the RPC portmap ALG on TCP/UDP destination port 111 and the RPC ALG for both TCP and UDP. You can specify one or more **rpc-program-number** values to further restrict allowed RPC protocols.

### RTSP

The Real-Time Streaming Protocol (RTSP) controls the delivery of data with real-time properties such as audio and video. The streams controlled by RTSP can use RTP, but it is not required. Media can be transmitted on the same RTSP control stream. This is an HTTP-like text-based protocol, but client and server maintain session information. A session is established using the SETUP message and terminated using the TEARDOWN message. The transport (the media protocol, address, and port numbers) is negotiated in the setup and the setup-response.

Support for stateful firewall and NAT services requires that you configure the RTSP ALG for TCP port 554.

The ALG monitors the control connection, opens flows dynamically for media (RTP/RTSP) streams, and performs NAT address and port rewrites.

### SIP

The Session Initiation Protocol (SIP) is an application layer protocol that can establish, maintain, and terminate media sessions. It is a widely used voice over IP (VoIP) signaling protocol. The SIP ALG monitors SIP traffic and dynamically creates and manages pinholes on the signaling and media paths. The ALG only allows packets with the correct permissions. The SIP ALG also performs the following functions:

- Manages parent-child session relationships.
- Enforces security policies.
- Manages pinholes for VoIP traffic.

The SIP ALG supports the following features:

- Stateful firewall
- Static source NAT

- Dynamic address only source NAT
- Network Address Port Translation (NAPT)



**NOTE:** SIP sessions are limited to 12 hours (720 minutes) for NAT processing on the MS-MIC and MS-MPC interface cards. There is no time limit for SIP sessions on the MS-DPC.

---

## SNMP

SNMP is a communication protocol for managing TCP/IP networks, including both individual network devices and aggregated devices. The protocol is defined by RFC 1157. SNMP runs on top of UDP.

The Junos OS stateful firewall service implements the SNMP ALG to inspect the SNMP type. SNMP does not enforce stateful flow. Each SNMP type needs to be specifically enabled. Full SNMP support of stateful firewall services requires that you configure the SNMP ALG on UDP port 161. This enables the SNMP **get** and **get-next** commands, as well as their response traffic in the reverse direction: UDP port 161 enables the SNMP **get-response** command. If SNMP traps are permitted, you can configure them on UDP port 162, enabling the SNMP **trap** command.

---

## SQLNet

The SQLNet protocol is used by Oracle SQL servers to execute SQL commands from clients, including load balancing and application-specific services.

Support of stateful firewall and NAT services requires that you configure the SQLNet ALG for TCP port 1521.

The ALG monitors the control packets, opens flows dynamically for data traffic, and performs NAT address and port rewrites.

---

## TFTP

The Trivial File Transfer Protocol (TFTP) is specified in RFC 1350. The initial TFTP requests are sent to UDP destination port 69. Additional flows can be created to **get** or **put** individual files. Support of stateful firewall and NAT services requires that you configure the TFTP ALG for UDP destination port 69.

---

## Traceroute

Traceroute is a tool for displaying the route that packets take to a network host. It uses the IP time-to-live (TTL) field to trigger ICMP time-exceeded messages from routers or gateways. It sends UDP datagrams to destination ports that are believed to be not in use; destination ports are numbered using the formula:  $+ n\text{hops} - 1$ . The default base port is 33434. To support traceroute through the firewall, two types of traffic must be passed through:

1. UDP probe packets (UDP destination port  $> 33000$ , IP TTL  $< 30$ )
2. ICMP response packets (ICMP type time-exceeded)

When NAT is applied, the IP address and port within the ICMP error packet also must be changed.

Support of stateful firewall and NAT services requires you to configure the Traceroute ALG for UDP destination port 33434 to 33450. In addition, you can configure the TTL threshold to prevent UDP flood attacks with large TTL values.

### UNIX Remote-Shell Services

Three protocols form the basis for UNIX remote-shell services:

- **Exec**—Remote command execution; enables a user on the client system to execute a command on the remote system. The first command from client (**rcmd**) to server (**rshd**) uses well-known TCP port 512. A second TCP connection can be opened at the request of **rcmd**. The client port number for the second connection is sent to the server as an ASCII string.
- **Login**—Better known as **rlogin**; uses well-known TCP port 513. For details, see RFC 1282. No special firewall processing is required.
- **Shell**—Remote command execution; enables a user on the client system to execute a command on the remote system. The first command from client (**rcmd**) to server (**rshd**) uses well-known TCP port 514. A second TCP connection can be opened at the request of **rcmd**. The client port number for the second connection is sent to the server as an ASCII string.

Support of stateful firewall services requires that you configure the Exec ALG on TCP port 512, the Login ALG on TCP port 513, and the Shell ALG on TCP port 514. NAT remote-shell services require that any dynamic source port assigned be within the port range 512 to 1023. If you configure a NAT pool, this port range is reserved exclusively for remote shell applications.

### Winframe

WinFrame application server software provides access to virtually any Windows application, across any type of network connection to any type of client.

This protocol is mainly used by Citrix Windows applications.

Stateful firewall and NAT require the ALG Winframe to be configured on TCP destination port 1494 and UDP port 1604.

## Juniper Networks Defaults

The Junos OS provides a default, hidden configuration group called **junos-defaults** that is automatically applied to the configuration of your router. The **junos-defaults** group contains preconfigured statements that contain predefined values for common applications. Some of the statements must be referenced to take effect, such as applications like FTP or Telnet. Other statements are applied automatically, such as terminal settings. All of the preconfigured statements begin with the reserved name **junos-**.



**NOTE:** You can override the Junos default configuration values, but you cannot delete or edit them. If you delete a configuration, the defaults return when a new configuration is added.

You cannot use the `apply-groups` statement with the Junos defaults group.

To view the full set of available preset statements from the Junos default group, issue the `show groups junos-defaults` configuration mode command. The following example displays the list of Junos default groups that use application protocols (ALGs).

```
user@host# show groups junos-defaults
applications {
  #
  # File Transfer Protocol
  #
  application junos-ftp {
    application-protocol ftp;
    protocol tcp;
    destination-port 21;
  }
  #
  # Trivial File Transfer Protocol
  #
  application junos-tftp {
    application-protocol tftp;
    protocol udp;
    destination-port 69;
  }
  #
  # RPC portmapper on TCP
  #
  application junos-rpc-portmap-tcp {
    application-protocol rpc-portmap;
    protocol tcp;
    destination-port 111;
  }
  #
  # RPC portmapper on UDP
  #
  application junos-rpc-portmap-udp {
    application-protocol rpc-portmap;
    protocol udp;
    destination-port 111;
  }
  #
  # SNMP get
  #
  application junos-snmp-get {
    application-protocol snmp;
    protocol udp;
    destination-port 161;
    snmp-command get;
  }
}
```

```
}  
#  
# SNMP get next  
#  
application junos-snmp-get-next {  
    application-protocol snmp;  
    protocol udp;  
    destination-port 161;  
    snmp-command get-next;  
}  
#  
# SNMP response  
#  
application junos-snmp-response {  
    application-protocol snmp;  
    protocol udp;  
    source-port 161;  
    snmp-command get-response;  
}  
#  
# SNMP trap  
#  
application junos-snmp-trap {  
    application-protocol snmp;  
    protocol udp;  
    destination-port 162;  
    snmp-command trap;  
}  
#  
# remote exec  
#  
application junos-rexec {  
    application-protocol exec;  
    protocol tcp;  
    destination-port 512;  
}  
#  
# remote login  
#  
application junos-rlogin {  
    application-protocol shell;  
    protocol tcp;  
    destination-port 513;  
}  
#  
# remote shell  
#  
application junos-rsh {  
    application-protocol shell;  
    protocol tcp;  
    destination-port 514;  
}  
#  
# Real Time Streaming Protocol  
#
```

```
application junos-rtsp {
    application-protocol rtsp;
    protocol tcp;
    destination-port 554;
}
#
# Citrix windows application server protocol
# windows applications remotely on windows/non-windows clients
#
# citrix needs udp 1604 to be open
#
application junos-citrix-wiframe {
    application-protocol wiframe;
    protocol tcp;
    destination-port 1494;
}
application junos-citrix-wiframe-udp {
    protocol udp;
    destination-port 1604;
}
#
# Oracle SQL servers use this protocol to execute sql commands
# from clients, load balance, use application-specific servers, etc
#
application junos-sqlnet {
    application-protocol sqlnet;
    protocol tcp;
    destination-port 1521;
}
#
# H.323 Protocol for audio/video conferencing
#
application junos-h323 {
    application-protocol h323;
    protocol tcp;
    destination-port 1720;
}
#
# Internet Inter-ORB Protocol - used for CORBA applications
# The ORB protocol in Java virtual machines uses port 1975 as default
#
application junos-iiop-java {
    application-protocol iiop;
    protocol tcp;
    destination-port 1975;
}
#
# Internet Inter-ORB Protocol - used for CORBA applications
# ORBIX is a CORBA framework from Iona Technologies that uses port
# 3075 as default
#
application junos-iiop-orbix {
    application-protocol iiop;
    protocol tcp;
    destination-port 3075;
}
```

```
}  
#  
# Real players use this protocol for real time streaming  
# This was the original protocol for real players.  
# RTSP is more widely used by real players  
# but they still support realaudio.  
#  
application junos-realaudio {  
    application-protocol realaudio;  
    protocol tcp;  
    destination-port 7070;  
}  
#  
# traceroute application.  
#  
application junos-traceroute {  
    application-protocol traceroute;  
    protocol udp;  
    destination-port 33435-33450;  
    ttl-threshold 30;  
}  
#  
# The full range of known RPC programs using UDP  
# The program numbers can be more specific to certain applications.  
#  
application junos-rpc-services-udp {  
    application-protocol rpc;  
    protocol udp;  
    rpc-program-number 100000-400000;  
}  
#  
# The full range of known RPC programs using TCP  
# The program numbers can be more specific to certain applications.  
#  
application junos-rpc-services-tcp {  
    application-protocol rpc;  
    protocol tcp;  
    rpc-program-number 100000-400000;  
}  
#  
# All ICMP traffic  
# This can be made to be more restrictive by specifying ICMP type  
# and code.  
#  
application junos-icmp-all {  
    application-protocol icmp;  
}  
#  
# Protocol used by Windows media server and windows media player  
#  
application junos-netshow {  
    application-protocol netshow;  
    protocol tcp;  
    destination-port 1755;  
}
```

```
#
# NetBIOS - networking protocol used on
# Windows networks name service port, both UDP and TCP
#
application junos-netbios-name-udp {
    application-protocol netbios;
    protocol udp;
    destination-port 137;
}
application junos-netbios-name-tcp {
    protocol tcp;
    destination-port 137;
}
#
# NetBIOS - networking protocol used on
# Windows networks datagram service port
#
application junos-netbios-datagram {
    application-protocol netbios;
    protocol udp;
    destination-port 138;
}
#
# NetBIOS - networking protocol used on
# Windows networks session service port
#
application junos-netbios-session {
    protocol tcp;
    destination-port 139;
}
#
# DCE-RPC portmapper on TCP
#
application junos-dce-rpc-portmap {
    application-protocol dce-rpc-portmap;
    protocol tcp;
    destination-port 135;
}
#
# DCE-RPC application on TCP sample UUID
# This application requires user to specify the UUID value
#
# application junos-dcerpc {
#     # application-protocol dce-rpc;
#     # protocol tcp;
#     #
#     # # UUID also needs to be defined as shown below
#     # UUID 11223344 22334455 33445566 44556677;
#     #
# }
#
# ms-exchange needs these 3 UUIDs
#
application junos-dcerpc-endpoint-mapper-service {
    application-protocol dce-rpc;
```

```
    protocol tcp;
    uuid e1af8308-5d1f-11c9-91a4-08002b14a0fa;
}
application junos-dcerpc-msexchange-directory-rfr {
    application-protocol dce-rpc;
    protocol tcp;
    uuid 1544f5e0-613c-11d1-93df-00c04fd7bd09;
}
application junos-dcerpc-msexchange-information-store {
    application-protocol dce-rpc;
    protocol tcp;
    uuid a4f1db00-ca47-1067-b31f-00dd010662da;
}
application junos-ssh {
    protocol tcp;
    destination-port 22;
}
application junos-telnet {
    protocol tcp;
    destination-port 23;
}
application junos-smtp {
    protocol tcp;
    destination-port 25;
}
application junos-dns-udp {
    protocol udp;
    destination-port 53;
}
application junos-dns-tcp {
    protocol tcp;
    destination-port 53;
}
application junos-tacacs {
    protocol tcp;
    destination-port 49;
}
# TACACS Database Service
application junos-tacacs-ds {
    protocol tcp;
    destination-port 65;
}
application junos-dhcp-client {
    protocol udp;
    destination-port 68;
}
application junos-dhcp-server {
    protocol udp;
    destination-port 67;
}
application junos-bootpc {
    protocol udp;
    destination-port 68;
}
application junos-bootps {
```

```
    protocol udp;
    destination-port 67;
  }
  application junos-finger {
    protocol tcp;
    destination-port 79;
  }
  application junos-http {
    protocol tcp;
    destination-port 80;
  }
  application junos-https {
    protocol tcp;
    destination-port 443;
  }
  application junos-pop3 {
    protocol tcp;
    destination-port 110;
  }
  application junos-ident {
    protocol tcp;
    destination-port 113;
  }
  application junos-nntp {
    protocol tcp;
    destination-port 119;
  }
  application junos-ntp {
    protocol udp;
    destination-port 123;
  }
  application junos-imap {
    protocol tcp;
    destination-port 143;
  }
  application junos-imaps {
    protocol tcp;
    destination-port 993;
  }
  application junos-bgp {
    protocol tcp;
    destination-port 179;
  }
  application junos-ldap {
    protocol tcp;
    destination-port 389;
  }
  application junos-snpp {
    protocol tcp;
    destination-port 444;
  }
  application junos-biff {
    protocol udp;
    destination-port 512;
  }
```

```
# UNIX who
application junos-who {
    protocol udp;
    destination-port 513;
}
application junos-syslog {
    protocol udp;
    destination-port 514;
}
# line printer daemon, printer, spooler
application junos-printer {
    protocol tcp;
    destination-port 515;
}
# UNIX talk
application junos-talk-tcp {
    protocol tcp;
    destination-port 517;
}
application junos-talk-udp {
    protocol udp;
    destination-port 517;
}
application junos-ntalk {
    protocol udp;
    destination-port 518;
}
application junos-rip {
    protocol udp;
    destination-port 520;
}
# INA sanctioned RADIUS port numbers
application junos-radius {
    protocol udp;
    destination-port 1812;
}
application junos-radacct {
    protocol udp;
    destination-port 1813;
}
application junos-nfsd-tcp {
    protocol tcp;
    destination-port 2049;
}
application junos-nfsd-udp {
    protocol udp;
    destination-port 2049;
}
application junos-cvspserver {
    protocol tcp;
    destination-port 2401;
}
#
# Label Distribution Protocol
#
```

```
application junos-ldp-tcp {
    protocol tcp;
    destination-port 646;
}
application junos-ldp-udp {
    protocol udp;
    destination-port 646;
}
#
# JUNOScript and JUNOScope management
#
application junos-xnm-ssl {
    protocol tcp;
    destination-port 3220;
}
application junos-xnm-clear-text {
    protocol tcp;
    destination-port 3221;
}
#
# IPsec tunnel
#
application junos-ipsec-esp {
    protocol esp;
}
application junos-ike {
    protocol udp;
    destination-port 500;
}
#
# 'junos-algs-outbound' defines a set of all applications
# requiring an ALG. Useful for defining rule to the the public
# internet allowing private network users to use all JUNOS OS
# supported ALGs initiated from the private network.
#
# NOTE: the contents of this set might grow in future JUNOS OS versions.
#
application-set junos-algs-outbound {
    application junos-ftp;
    application junos-tftp;
    application junos-rpc-portmap-tcp;
    application junos-rpc-portmap-udp;
    application junos-snmp-get;
    application junos-snmp-get-next;
    application junos-snmp-response;
    application junos-snmp-trap;
    application junos-rexec;
    application junos-rlogin;
    application junos-rsh;
    application junos-rtsp;
    application junos-citrix-winframe;
    application junos-citrix-winframe-udp;
    application junos-sqlnet;
    application junos-h323;
    application junos-iiop-java;
```

```

application junos-iiop-orbix;
application junos-realaudio;
application junos-traceroute;
application junos-rpc-services-udp;
application junos-rpc-services-tcp;
application junos-icmp-all;
application junos-netshow;
application junos-netbios-name-udp;
application junos-netbios-datagram;
application junos-dcerpc-endpoint-mapper-service;
application junos-dcerpc-msexchange-directory-rfr;
application junos-dcerpc-msexchange-information-store;
}
#
# 'junos-management-inbound' represents the group of applications
# that might need access the router from public network for
# for management purposes.
#
# Set is intended for a UI to display management choices.
#
# NOTE: It is not recommended the user to use the entire set
# directly in a firewall rule and open up firewall to all
# of these applications. Also, the user should always
# specify the source and destination prefixes when using
# each application.
#
# NOTE: the contents of this set may grow in future JUNOS versions.
#
application-set junos-management-inbound {
    application junos-snmp-get;
    application junos-snmp-get-next;
    application junos-snmp-response;
    application junos-snmp-trap;
    application junos-ssh;
    application junos-telnet;
    application junos-http;
    application junos-https;
    application junos-xnm-ssl;
    application junos-xnm-clear-text;
}
#
# 'junos-routing-inbound' represents routing protocols that might
# need to access the router from public network.
#
# Set is intended for a UI to display routing involvement choices.
#
# NOTE: It is not recommended the user to use the entire set
# directly in a firewall rule and open up firewall to all
# of these applications. Also, the user should always
# specify the source and destination prefixes when using
# each application.
#
# NOTE: the contents of this set might grow in future JUNOS OS versions.
#
application-set junos-routing-inbound {

```

```

    application junos-bgp;
    application junos-rip;
    application junos-ldp-tcp;
    application junos-ldp-udp;
  }
}

```

To reference statements available from the **junos-defaults** group, include the selected **junos-default-name** statement at the applicable hierarchy level.

## Examples: Referencing the Preset Statement from the Junos Default Group

The following example is a preset statement from the Junos default groups that is available for FTP in a stateful firewall:

```

[edit]
groups {
  junos-defaults {
    applications {
      application junos-ftp { # Use FTP default configuration
        application-protocol ftp;
        protocol tcp;
        destination-port 21;
      }
    }
  }
}

```

To reference a preset Junos default statement from the Junos default groups, include the **junos-default-name** statement at the applicable hierarchy level. For example, to reference the Junos default statement for FTP in a stateful firewall, include the **junos-ftp** statement at the **[edit services stateful-firewall rule rule-name term term-name from applications]** hierarchy level.

```

[edit]
services {
  stateful-firewall {
    rule my-rule {
      term my-term {
        from {
          applications junos-ftp; #Reference predefined statement, junos-ftp,
        }
      }
    }
  }
}

```

The following example shows configuration of the default Junos IP ALG:

```

[edit]
services {
  stateful-firewall {
    rule r1 {
      match-direction input;
    }
  }
}

```

```

term t1 {
  from {
    applications junos-ip;
  }
  then {
    accept;
    syslog;
  }
}
}
}

```

If you configure the IP ALG in the stateful firewall rule, it is matched by any IP traffic, but if there is any other more specific application that matches the same traffic, the IP ALG will not be matched. For example, in the following configuration, both the ICMP ALG and the IP ALG are configured, but traffic is matched for ICMP packets, because it is the more specific match.

```

[edit]
services {
  stateful-firewall {
    rule r1 {
      match-direction input;
      term t1 {
        from {
          applications [ junos-ip junos-icmp-all ];
        }
        then {
          accept;
          syslog;
        }
      }
    }
  }
}
}

```

