




Junos[®] OS

High Availability Feature Guide



Modified: 2019-09-19



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® OS High Availability Feature Guide

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	xxv
	Documentation and Release Notes	xxv
	Using the Examples in This Manual	xxv
	Merging a Full Example	xxvi
	Merging a Snippet	xxvi
	Documentation Conventions	xxvii
	Documentation Feedback	xxix
	Requesting Technical Support	xxix
	Self-Help Online Tools and Resources	xxx
	Creating a Service Request with JTAC	xxx
Part 1	Overview	
Chapter 1	High Availability Overview	3
	Understanding High Availability Features on Juniper Networks Routers	3
	Routing Engine Redundancy	3
	Graceful Routing Engine Switchover	3
	Nonstop Bridging	4
	Nonstop Active Routing	4
	Graceful Restart	5
	Nonstop Active Routing Versus Graceful Restart	6
	Effects of a Routing Engine Switchover	7
	VRRP	7
	Unified ISSU	7
	Interchassis Redundancy for MX Series Routers Using Virtual Chassis	8
	High Availability-Related Features in Junos OS	9
	High Availability Features for EX Series Switches Overview	9
	VRRP	9
	Graceful Protocol Restart	10
	Redundant Routing Engines	10
	Virtual Chassis	11
	Graceful Routing Engine Switchover	11
	Link Aggregation	12
	Nonstop Active Routing and Nonstop Bridging	12
	Nonstop Software Upgrade	12
	Redundant Power System	13

Part 2	Configuring Switching Control Board Redundancy	
Chapter 2	Understanding How Switching Control Board Redundancy Prevents Network Failures	17
	Understanding Switching Control Board Redundancy	17
	Redundant CFEBs on the M10i Router	17
	Redundant FEBs on the M120 Router	18
	Redundant SSBs on the M20 Router	20
	Redundant SFMs on the M40e and M160 Routers	20
Chapter 3	Configuring Switching Control Board Redundancy	23
	Configuring CFEB Redundancy on the M10i Router	23
	Configuring FEB Redundancy on the M120 Router	24
	Example: Configuring FEB Redundancy on M120 Routers	25
	Configuring SFM Redundancy on M40e and M160 Routers	26
	Configuring SSB Redundancy on the M20 Router	27
	Configuring the Junos OS to Support Redundancy on Routers Having Multiple Routing Engines or Switching Boards	27
Part 3	Configuring Bidirectional Forwarding Detection (BFD)	
Chapter 4	Understanding How BFD Detects Network Failures	31
	Understanding BFD for Static Routes for Faster Network Failure Detection	31
	Understanding BFD for BGP	36
	Understanding BFD for OSPF	38
	Understanding BFD for IS-IS	41
	Understanding BFD for RIP	43
	Understanding Independent Micro BFD Sessions for LAG	44
	Understanding Distributed BFD	47
	Understanding Static Route State When BFD is in Admin Down State	52
Chapter 5	Configuring BFD	53
	Example: Configuring BFD for Static Routes for Faster Network Failure Detection	53
	Example: Configuring BFD on Internal BGP Peer Sessions	59
	Example: Configuring BFD for OSPF	70
	Example: Configuring BFD for IS-IS	74
	Example: Configuring BFD for RIP	81
	Configuring Independent Micro BFD Sessions for LAG	87
	Example: Configuring Independent Micro BFD Sessions for LAG	93
	Configuring BFD for PIM	103
	Enabling Dedicated and Real-Time BFD	106
Part 4	Configuring Routing Engine Redundancy	
Chapter 6	Understanding How Routing Engine Redundancy Prevents Network Failures	111
	Understanding Routing Engine Redundancy on Juniper Networks Routers	111
	Routing Engine Redundancy Overview	111
	Conditions That Trigger a Routing Engine Failover	112
	Default Routing Engine Redundancy Behavior	113

	Routing Engine Redundancy on a TX Matrix Router	114
	Routing Engine Redundancy on a TX Matrix Plus Router	115
	Situations That Require You to Halt Routing Engines	116
Chapter 7	Configuring Routing Engine Redundancy	119
	Configuring Routing Engine Redundancy	119
	Modifying the Default Routing Engine Mastership	119
	Configuring Automatic Failover to the Backup Routing Engine	120
	Without Interruption to Packet Forwarding	120
	On Detection of a Hard Disk Error on the Master Routing Engine	120
	On Detection of a Broken LCMD Connectivity Between the VM and RE	120
	On Detection of a Loss of Keepalive Signal from the Master Routing Engine	121
	On Detection of the em0 Interface Failure on the Master Routing Engine	122
	When a Software Process Fails	122
	Manually Switching Routing Engine Mastership	122
	Verifying Routing Engine Redundancy Status	123
	Initial Routing Engine Configuration Example	124
	Copying a Configuration File from One Routing Engine to the Other	125
	Loading a Software Package from the Other Routing Engine	127
Chapter 8	Configuring Load Balancing	129
	Understanding Aggregated Ethernet Load Balancing	129
	Configuring Adaptive Load Balancing	131
Part 5	Configuring Graceful Routing Engine Switchover (GRES)	
Chapter 9	Understanding How GRES Enables Uninterrupted Packet Forwarding During a Routing Engine Switchover	135
	Understanding Graceful Routing Engine Switchover	135
	Graceful Routing Engine Switchover Concepts	135
	Effects of a Routing Engine Switchover	140
	Graceful Routing Engine Switchover on Aggregated Services interfaces	141
Chapter 10	GRES System Requirements	143
	Graceful Routing Engine Switchover System Requirements	143
	Graceful Routing Engine Switchover Platform Support	143
	Graceful Routing Engine Switchover Feature Support	144
	Graceful Routing Engine Switchover DPC Support	145
	Graceful Routing Engine Switchover and Subscriber Access	145
	Graceful Routing Engine Switchover PIC Support	146
Chapter 11	Configuring GRES	147
	Requirements for Routers with a Backup Router Configuration	147
	Configuring Graceful Routing Engine Switchover	148
	Enabling Graceful Routing Engine Switchover	148
	Configuring Graceful Routing Engine Switchover with Graceful Restart	148
	Synchronizing the Routing Engine Configuration	148

	Verifying Graceful Routing Engine Switchover Operation	150
	Configuring Graceful Routing Engine Switchover in a Virtual Chassis	151
	Preventing Graceful Routing Engine Switchover in the Case of Slow Disks	152
	Resetting Local Statistics	152
	Example: Configuring IS-IS for GRES with Graceful Restart	153
Chapter 12	Configuring Ethernet Automatic Protection Switching for High Availability	157
	Ethernet Automatic Protection Switching Overview	157
	Unidirectional and Bidirectional Switching	158
	Selective and Merging Selectors	158
	Revertive and Nonrevertive Switching	158
	Protection Switching Between VPWS Pseudowires	158
	CLI Configuration Statements	159
	Mapping of CCM Defects to APS Events	160
	Example: Configuring Protection Switching Between Psuedowires	161
Chapter 13	Configuring Ethernet Ring Protection Switching for High Availability	165
	Ethernet Ring Protection Switching Overview	165
	Understanding Ethernet Ring Protection Switching Functionality	166
	Acronyms	167
	Ring Nodes	167
	Ring Node States	167
	Default Logging of Basic State Transitions on EX Series Switches	168
	Logical Ring	168
	FDB Flush	168
	Traffic Blocking and Forwarding	169
	RPL Neighbor Node	169
	RAPS Message Blocking and Forwarding	169
	Dedicated Signaling Control Channel	171
	RAPS Message Termination	171
	Revertive and Non-revertive Modes	171
	Multiple Rings	171
	Node ID	171
	Ring ID	172
	Bridge Domains with the Ring Port (MX Series Routers Only)	172
	Wait-to-Block Timer	172
	Adding and Removing a Node	173
	Configuring Ethernet Ring Protection Switching	173
	Example: Ethernet Ring Protection Switching Configuration on MX Routers	174
Part 6	Configuring Nonstop Bridging	
Chapter 14	Understanding How Nonstop Bridging Preserves Layer 2 Protocol Information During a Routing Engine Switchover	185
	Nonstop Bridging Concepts	185
	Understanding Nonstop Bridging on EX Series Switches	187

Chapter 15	Nonstop Bridging System Requirements	189
	Nonstop Bridging System Requirements	189
	Platform Support	189
	Protocol Support	189
Chapter 16	Configuring Nonstop Bridging	191
	Configuring Nonstop Bridging	191
	Enabling Nonstop Bridging	191
	Synchronizing the Routing Engine Configuration	191
	Verifying Nonstop Bridging Operation	192
	Configuring Nonstop Bridging on Switches (CLI Procedure)	193
	Configuring Nonstop Bridging on EX Series Switches (CLI Procedure)	195
Part 7	Configuring Nonstop Active Routing (NSR)	
Chapter 17	Understanding How Nonstop Active Routing Preserves Routing Protocol Information During a Routing Engine Switchover	199
	Nonstop Active Routing Concepts	199
	Understanding Nonstop Active Routing on EX Series Switches	203
Chapter 18	Nonstop Active Routing System Requirements	205
	Nonstop Active Routing System Requirements	205
	Nonstop Active Routing Platform and Switching Platform Support	205
	Nonstop Active Routing Protocol and Feature Support	206
	Nonstop Active Routing BFD Support	209
	Nonstop Active Routing BGP Support	210
	Nonstop Active Routing Layer 2 Circuit and VPLS Support	211
	Nonstop Active Routing PIM Support	212
	Nonstop Active Routing MSDP Support	214
	Nonstop Active Routing Support for RSVP-TE LSPs	215
Chapter 19	Configuring Nonstop Active Routing	217
	Configuring Nonstop Active Routing	217
	Enabling Nonstop Active Routing	217
	Synchronizing the Routing Engine Configuration	218
	Verifying Nonstop Active Routing Operation	219
	Configuring Nonstop Active Routing on Switches	220
	Preventing Automatic Reestablishment of BGP Peer Sessions After NSR Switchovers	221
	Example: Configuring Nonstop Active Routing	222
	Tracing Nonstop Active Routing Synchronization Events	224
	Resetting Local Statistics	226
	Example: Configuring Nonstop Active Routing on Switches	226
Part 8	Configuring Graceful Restart	
Chapter 20	Understanding How Graceful Restart Enables Uninterrupted Packet Forwarding When a Router Is Restarted	233
	Graceful Restart Concepts	233
	Graceful Restart for Aggregate and Static Routes	234

	Graceful Restart and Routing Protocols	235
	BGP	235
	IS-IS	235
	OSPF and OSPFv3	235
	PIM Sparse Mode	236
	RIP and RIPng	237
	Graceful Restart and MPLS-Related Protocols	237
	LDP	237
	RSVP	238
	CCC and TCC	238
	Understanding Restart Signaling-Based Helper Mode Support for OSPF Graceful Restart	238
	Graceful Restart and Layer 2 and Layer 3 VPNs	239
	Graceful Restart on Logical Systems	240
Chapter 21	Graceful Restart System Requirements	241
	Graceful Restart System Requirements	241
Chapter 22	Configuring Graceful Restart	243
	Enabling Graceful Restart	243
	Configuring Graceful Restart	244
	Configuring Routing Protocols Graceful Restart	270
	Enabling Graceful Restart	270
	Configuring Graceful Restart Options for BGP	271
	Using Control Plane Dependent BFD along with Graceful Restart Helper Mode	272
	Configuring Graceful Restart Options for ES-IS	273
	Configuring Graceful Restart Options for IS-IS	273
	Configuring Graceful Restart Options for OSPF and OSPFv3	274
	Configuring Graceful Restart Options for RIP and RIPng	275
	Configuring Graceful Restart Options for PIM Sparse Mode	276
	Tracking Graceful Restart Events	277
	Configuring Graceful Restart for MPLS-Related Protocols	277
	Configuring Graceful Restart Globally	278
	Configuring Graceful Restart Options for RSVP, CCC, and TCC	278
	Configuring Graceful Restart Options for LDP	279
	Configuring VPN Graceful Restart	280
	Configuring Graceful Restart Globally	280
	Configuring Graceful Restart for the Routing Instance	280
	Configuring Logical System Graceful Restart	281
	Enabling Graceful Restart Globally	281
	Configuring Graceful Restart for a Routing Instance	282
	Configuring Graceful Restart for QFabric Systems	282
	Enabling Graceful Restart	283
	Configuring Graceful Restart Options for BGP	284
	Configuring Graceful Restart Options for OSPF and OSPFv3	285
	Tracking Graceful Restart Events	286
	Example: Managing Helper Modes for OSPF Graceful Restart	286
	Tracing Restart Signaling-Based Helper Mode Events for OSPF Graceful Restart	289

	Verifying Graceful Restart Operation	290
	Graceful Restart Operational Mode Commands	290
	Verifying BGP Graceful Restart	291
	Verifying IS-IS and OSPF Graceful Restart	291
	Verifying CCC and TCC Graceful Restart	292
Part 9	Power Management Overview	
Chapter 23	Understanding Power Management	295
	Understanding Power Management on EX Series Switches	295
	Power Priority of Line Cards	296
	How a Line Card's Power Priority Is Determined	297
	Line Card Priority and Line Card Power	297
	Line Card Priority and PoE Power	297
	Line Card Priority and Changes in the Power Budget	298
	Power Supply Redundancy	299
	Configuring the Power Priority of Line Cards (CLI Procedure)	302
	Configuring Power Supply Redundancy (CLI Procedure)	303
Chapter 24	Redundant Power System Overview	305
	EX Series Redundant Power System Hardware Overview	305
	Benefits of the EX Series Redundant Power System	306
	Switch Models and Configurations Supported by the RPS	306
	When a Switch's Power Supply Fails	307
	Components of the RPS	307
	Understanding How Power Priority Is Determined and Set for Switches Connected to the EX Series Redundant Power System	308
	Default RPS Priority	308
	Changing the Priority of Switches on an EX Series RPS	309
	Determining and Setting Priority for Switches Connected to an EX Series RPS	310
	Using RPS Default Configuration	311
	Setting the EX Series RPS Priority for a Switch (CLI)	311
Part 10	Configuring Virtual Router Redundancy Protocol (VRRP)	
Chapter 25	Understanding How the VRRP Router Failover Mechanism Prevents Network Failures	315
	Understanding VRRP	315
	Understanding VRRP Between QFabric Systems	319
	VRRP Differences on QFabric Systems	319
	Configuration Details	320
	Junos OS Support for VRRPv3	322
	Junos OS VRRP Support	322
	IPv6 VRRP Checksum Behavioral Differences	323
	VRRP Interoperability	324
	Upgrading from VRRPv2 to VRRPv3	324
	Functionality of VRRPv3 Features	326
	VRRPv3 Authentication	326
	VRRPv3 Advertisement Intervals	326

	Unified ISSU for VRRPv3	326
	VRRP failover-delay Overview	328
	When failover-delay Is Not Configured	328
	When failover-delay Is Configured	329
Chapter 26	Configuring VRRP	331
	Configuring Basic VRRP Support	332
	Configuring VRRP	337
	VRRP and VRRP for IPv6 Overview	339
	Configuring VRRP and VRRP for IPv6	339
	Configuring VRRP for IPv6 (CLI Procedure)	341
	Example: Configuring VRRP for IPv6	342
	Configuring VRRP Authentication (IPv4 Only)	348
	Configuring VRRP Preemption and Hold Time	349
	Configuring VRRP Preemption	350
	Configuring the Preemption Hold Time	350
	Configuring the Advertisement Interval for the VRRP Master Router	350
	Modifying the Advertisement Interval in Seconds	351
	Modifying the Advertisement Interval in Milliseconds	352
	Configuring the Startup Period for VRRP Operations	353
	Configuring a Backup Router to Preempt the VRRP Master Router	353
	Configuring a Backup to Accept Packets Destined for the Virtual IP Address	354
	Modifying the Preemption Hold-Time Value for the VRRP Master Router	355
	Configuring the Asymmetric Hold Time for VRRP Routers	355
	Configuring Passive ARP Learning for Backup VRRP Routers	356
	Configuring VRRP Route Tracking	357
	Configuring a Logical Interface to Be Tracked for a VRRP Group	358
	Configuring a Route to Be Tracked for a VRRP Group	361
	Example: Configuring Multiple VRRP Owner Groups	363
	Configuring Inheritance for a VRRP Group	369
	Configuring an Interface to Accept All Packets Destined for the Virtual IP Address of a VRRP Group	371
	Configuring the Silent Period to Avoid Alarms Due to Delay in Receiving VRRP Advertisement Packets	372
	Enabling the Distributed Periodic Packet Management Process for VRRP	373
	Improving the Convergence Time for VRRP	374
	Configuring VRRP to Improve Convergence Time	376
	Tracing VRRP Operations	377
	Example: Configuring VRRP for Load Sharing	378
	Troubleshooting VRRP	383

Part 11	Performing Unified In-Service Software Upgrade (ISSU)	
Chapter 27	Getting Started with Unified ISSU and Understanding How Unified ISSU Works	387
	Getting Started with Unified In-Service Software Upgrade	387
	Understanding the Unified ISSU Process	388
	Understanding the Unified ISSU Process on a Router	388
	Unified ISSU Process on a Router	389
	Understanding the Unified ISSU Process on the TX Matrix Router	393
	Unified ISSU Process on the TX Matrix Router	394
	Understanding the Unified ISSU Process on the TX Matrix Plus Router and on the TX Matrix Plus Router with 3D SIBs	396
	Unified ISSU Process on the TX Matrix Plus Router and on the TX Matrix Plus Router with 3D SIBs	397
	Understanding In-Service Software Upgrade (ISSU)	399
	In-Service Software Upgrade Process	399
	Understanding In-Service Software Upgrade (ISSU) in ACX5000 Series Routers	400
	In-Service Software Upgrade Process	400
Chapter 28	Unified ISSU System Requirements	401
	Unified ISSU System Requirements	401
	General Unified ISSU Considerations for All Platforms	402
	Unified ISSU Considerations for MX Series Routers	403
	Unified ISSU Considerations for PTX Series Routers	404
	Unified ISSU Considerations for T Series Routers	404
	Unified ISSU Considerations for EX Series Switches	405
	Unified ISSU Platform Support	405
	Unified ISSU Protocol Support for M Series, MX Series, and T Series Routers and EX9200 Switches	406
	Unified ISSU Feature Support	407
	Unified ISSU PIC Support Considerations	407
	PIC Considerations	408
	SONET/SDH PICs	409
	Fast Ethernet and Gigabit Ethernet PICs	410
	Channelized PICs	413
	Tunnel Services PICs	414
	ATM PICs	414
	Serial PICs	415
	DS3, E1, E3, and T1 PICs	415
	Enhanced IQ PICs	416
	Enhanced IQ2 Ethernet Services Engine (ESE) PIC	416
	Unified ISSU FPC Support on T4000 Routers	417
	Unified ISSU Support on MX Series 3D Universal Edge Routers	417

Chapter 29	Performing a Unified ISSU	423
	Best Practices for Performing a Unified ISSU	423
	Example: Performing a Unified ISSU	424
	Performing an In-Service Software Upgrade (ISSU) with Non-Stop Routing . . .	452
	Preparing the Switch for Software Installation	453
	Upgrading the Software Using ISSU	454
	Performing an In-Service Software Upgrade (ISSU) in ACX5000 Series	
	Routers	457
	Preparing the Router for Software Installation	457
	Upgrading the Software Using ISSU	459
	Verifying a Unified ISSU	461
	Verifying a Unified ISSU	461
	Troubleshooting Unified ISSU Problems	462
	Managing and Tracing BFD Sessions During Unified ISSU Procedures	462
Chapter 30	Performing an ISSR	465
	Performing an In-Service Software Reboot	465
Part 12	Performing Nonstop Software Upgrade (NSSU)	
Chapter 31	Getting Started with NSSU and Understanding How NSSU Works	471
	Understanding Nonstop Software Upgrade on EX Series Switches	471
	Requirements for Performing an NSSU	473
	How an NSSU Works	474
	EX3300, EX3400, EX4200, EX4300, EX4500, EX4600, and Mixed	
	Virtual Chassis	474
	EX6200 and EX8200 Switches	475
	EX8200 Virtual Chassis	476
	NSSU Limitations	477
	NSSU and Junos OS Release Support	477
	Overview of NSSU Configuration and Operation	478
Chapter 32	Performing a NSSU	481
	Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade	482
	Upgrading Software on an EX6200 or EX8200 Standalone Switch Using Nonstop	
	Software Upgrade (CLI Procedure)	485
	Preparing the Switch for Software Installation	485
	Upgrading Both Routing Engines Using NSSU	487
	Upgrading One Routing Engine Using NSSU (EX8200 Switch Only)	490
	Upgrading the Original Master Routing Engine (EX8200 Switch Only) . . .	492
	Example: Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade	
	on EX Series Switches	494
Part 13	Configuration Statements and Operational Commands	
Chapter 33	Configuration Statements: Adaptive Load Balancing	499
	adaptive	500

Chapter 34	Configuration Statements: Bidirectional Forwarding Detection	503
	dedicated-ukern-cpu (BFD)	504
	realtime-ukern-thread (BFD)	505
	authentication (LAG)	506
	bfd-liveness-detection (LAG)	507
	detection-time (LAG)	509
	traceoptions (Protocols BFD)	510
	transmit-interval (LAG)	512
Chapter 35	Configuration Statements: Graceful Routing Engine Switchover	513
	graceful-switchover	513
	graceful-switchover	514
	redundancy (Graceful Switchover)	515
Chapter 36	Configuration Statements: Graceful Restart	517
	disable	518
	disable (BGP Graceful Restart)	519
	dont-help-shared-fate-bfd-down	520
	graceful-restart (Enabling Globally)	521
	graceful-restart (Multicast Snooping)	522
	graceful-restart (Protocols BGP)	523
	graceful-restart (Protocols OSPF)	525
	helper-disable (Multiple Protocols)	527
	helper-disable (OSPF)	528
	kernel-replication	529
	maximum-helper-recovery-time	530
	maximum-helper-restart-time (RSVP)	531
	maximum-neighbor-reconnect-time	532
	maximum-neighbor-recovery-time	533
	no-strict-lsa-checking	534
	notify-duration	535
	not-on-disk-underperform	536
	reconnect-time	537
	recovery-time	538
	restart-duration	539
	restart-time (BGP Graceful Restart)	541
	stale-routes-time	542
	traceoptions (Protocols)	543
	warm-standby	544
Chapter 37	Configuration Statements: Nonstop Active Routing	545
	nonstop-routing	546
	switchover-on-routing-crash	547
	synchronize	548
	traceoptions	550
Chapter 38	Configuration Statements: Nonstop Bridging	553
	nonstop-bridging	553
	nonstop-bridging (Ethernet Switching)	554

Chapter 39	Configuration Statements: NSSU	555
	fpcs (NSSU Upgrade Groups)	556
	member (NSSU Upgrade Groups)	558
	nssu	559
	upgrade-group	560
Chapter 40	Configuration Statements: Power Management	561
	power-budget-priority	562
	n-plus-n (Power Management)	563
	psu	563
	redundancy (Power Management)	564
Chapter 41	Configuration Statements: Redundant Power System	565
	member (Redundant Power System)	565
	priority (Redundant Power System)	566
	redundant-power-system	567
Chapter 42	Configuration Statements: Routing Engine and Switching Control Board Redundancy	569
	cfcb	570
	description (Chassis Redundancy)	571
	disk-failure-action	572
	failover (Chassis)	573
	failover (Chassis)	574
	failover (System Process)	575
	feb (Creating a Redundancy Group)	576
	feb (Assigning a FEB to a Redundancy Group)	577
	keepalive-time	578
	keepalive-time	579
	no-auto-failover	580
	on-disk-failure (Chassis Redundancy Failover)	580
	on-disk-failure	581
	on-loss-of-keepalives	582
	on-loss-of-keepalives	583
	redundancy	584
	redundancy-group	585
	routing-engine (Chassis Redundancy)	586
	routing-engine	587
	sfm (Chassis Redundancy)	588
	ssb	589
	vcp-no-hold-time	590
Chapter 43	Configuration Statements: Unified ISSU	593
	no-issu-timer-negotiation	593
	traceoptions (Protocols BFD)	594
Chapter 44	Configuration Statements: VRRP	597
	accept-data	599
	advertise-interval	600
	asymmetric-hold-time	601
	asymmetric-hold-time	602

authentication-key	603
authentication-type	604
bandwidth-threshold	605
delegate-processing (VRRP)	606
failover-delay	607
fast-interval	608
global-advertisements-threshold	609
hold-time (VRRP)	610
hold-time	611
inherit-advertisement-interval	611
inet6-advertise-interval	612
inet6-advertise-interval	613
interface	614
preempt (VRRP)	615
preempt	616
priority (Protocols VRRP)	617
priority	618
priority-cost (VRRP)	619
priority-hold-time	620
route (Interfaces)	621
skew-timer-disable	622
startup-silent-period	623
traceoptions (Protocols VRRP)	624
traceoptions	626
track (VRRP)	628
version-3	629
virtual-address	630
virtual-inet6-address	631
virtual-inet6-address	632
virtual-link-local-address	633
virtual-link-local-address	634
vrrp-group	635
vrrp-inet6-group	637
vrrp-inet6-group	638
vrrp-inherit-from	639
Chapter 45 Administration	641
Administration Tasks	641
Upgrading Software on an EX6200 or EX8200 Standalone Switch Using	
Nonstop Software Upgrade (CLI Procedure)	641
Preparing the Switch for Software Installation	641
Upgrading Both Routing Engines Using NSSU	643
Upgrading One Routing Engine Using NSSU (EX8200 Switch Only)	646
Upgrading the Original Master Routing Engine (EX8200 Switch	
Only)	648
Upgrading Software on an EX8200 Virtual Chassis Using Nonstop Software	
Upgrade (CLI Procedure)	650
Preparing the Switch for Software Installation	650
Upgrading the Software Using NSSU	651

	Upgrading Software Using Nonstop Software Upgrade on EX Series Virtual Chassis and Mixed Virtual Chassis (CLI Procedure)	654
	Preparing the Switch for Software Installation	655
	Upgrading the Software Using NSSU	656
	Verification Tasks	658
	Verifying Power Configuration and Use	658
Chapter 46	Operational Commands	661
	show bgp neighbor	663
	show log	683
	show (ospf ospf3) overview	687
	show chassis dedicated-ukern-cpu	693
	show chassis realtime-ukern-thread	694
	clear vrrp	695
	request chassis ssb master switch	696
	request redundant-power-system multi-backup	698
	request system software in-service-upgrade	699
	request system software in-service-upgrade (MX Series 5G Universal Routing Platforms and EX9200 Switches)	712
	request system software nonstop-upgrade	729
	request system software validate in-service-upgrade	740
	show chassis nonstop-upgrade	743
	show chassis nonstop-upgrade node-group	745
	show chassis power-budget-statistics	746
	show chassis redundant-power-system	750
	show redundant-power-system led	752
	show redundant-power-system multi-backup	754
	show redundant-power-system network	755
	show redundant-power-system power-supply	756
	show redundant-power-system status	758
	show redundant-power-system upgrade	760
	show redundant-power-system version	761
	show chassis ssb	763
	show nonstop-routing	765
	show pfe ssb	768
	show system switchover	774
	show task replication	779
	show vrrp	782
	show vrrp track	793
Chapter 47	Troubleshooting	797
	Tracing Nonstop Active Routing Synchronization Events	797
	Troubleshooting the EX Series Redundant Power System Power On and Power Backup Issues	799
	The EX Series RPS Is Not Powering On	799
	A Switch Is Not Recognized by the RPS	799
	An Error Message Indicates That an RPS Power Supply is Not Supported	800
	The EX Series Redundant Power System Is Not Providing Power Backup to a Connected Switch	800

The Wrong Switches Are Being Backed Up	801
Six Switches That Do Not Require PoE Are Not All Being Backed Up	801

List of Figures

Part 3	Configuring Bidirectional Forwarding Detection (BFD)	
Chapter 5	Configuring BFD	53
	Figure 1: Customer Routes Connected to a Service Provider	54
	Figure 2: Typical Network with IBGP Sessions	62
	Figure 3: Configuring BFD for IS-IS	75
	Figure 4: RIP BFD Network Topology	83
	Figure 5: Configuring an Independent Micro BFD Session for LAG	93
Part 5	Configuring Graceful Routing Engine Switchover (GRES)	
Chapter 9	Understanding How GRES Enables Uninterrupted Packet Forwarding During a Routing Engine Switchover	135
	Figure 6: Preparing for a Graceful Routing Engine Switchover	138
	Figure 7: Graceful Routing Engine Switchover Process	139
Chapter 12	Configuring Ethernet Automatic Protection Switching for High Availability	157
	Figure 8: Connections Terminating on Single PE	158
	Figure 9: Connections Terminating on a Different PE	159
	Figure 10: Understanding APS Events	160
	Figure 11: Topology of a Network Using VPWS Psuedowires	161
Chapter 13	Configuring Ethernet Ring Protection Switching for High Availability	165
	Figure 12: Protocol Packets from the Network to the Router	169
	Figure 13: Protocol Packets from the Router or Switch to the Network	169
	Figure 14: Example of a Three-Node Ring Topology	175
Part 6	Configuring Nonstop Bridging	
Chapter 14	Understanding How Nonstop Bridging Preserves Layer 2 Protocol Information During a Routing Engine Switchover	185
	Figure 15: Nonstop Bridging Switchover Preparation Process	186
	Figure 16: Nonstop Bridging During a Switchover	187
Part 7	Configuring Nonstop Active Routing (NSR)	
Chapter 17	Understanding How Nonstop Active Routing Preserves Routing Protocol Information During a Routing Engine Switchover	199
	Figure 17: Nonstop Active Routing Switchover Preparation Process	201
	Figure 18: Nonstop Active Routing During a Switchover	202

Part 8	Configuring Graceful Restart	
Chapter 22	Configuring Graceful Restart	243
	Figure 19: Layer 3 VPN Graceful Restart Topology	245
Part 9	Power Management Overview	
Chapter 24	Redundant Power System Overview	305
	Figure 20: Default PoE Switch Priority Is Determined by Connector Port Location	309
	Figure 21: Switch Priority After CLI Configuration	309
Part 10	Configuring Virtual Router Redundancy Protocol (VRRP)	
Chapter 25	Understanding How the VRRP Router Failover Mechanism Prevents Network Failures	315
	Figure 22: Basic VRRP	317
	Figure 23: Basic VRRP on EX Series Switches	318
	Figure 24: VRRP on Virtual Chassis Switches	318
Chapter 26	Configuring VRRP	331
	Figure 25: VRRP Load-Sharing Configuration	379
Part 11	Performing Unified In-Service Software Upgrade (ISSU)	
Chapter 27	Getting Started with Unified ISSU and Understanding How Unified ISSU Works	387
	Figure 26: Device Status Before Starting a Unified ISSU	390
	Figure 27: Device Status After the Backup Routing Engine Is Upgraded	391
	Figure 28: Device Status After One Packet Forwarding Engine Downloads the New Software	391
	Figure 29: Device Status Before the Routing Engine Switchover	392
	Figure 30: Device Status After the Routing Engine Switchover	392
	Figure 31: Device Status After the Unified ISSU Is Complete	393
Chapter 29	Performing a Unified ISSU	423
	Figure 32: Unified ISSU Example Topology	425
Part 12	Performing Nonstop Software Upgrade (NSSU)	
Chapter 32	Performing a NSSU	481
	Figure 33: Example Line-Card Upgrade Group Topology	495

List of Tables

	About the Documentation	xxv
	Table 1: Notice Icons	xxvii
	Table 2: Text and Syntax Conventions	xxviii
Part 3	Configuring Bidirectional Forwarding Detection (BFD)	
Chapter 4	Understanding How BFD Detects Network Failures	31
	Table 3: Configuring BFD for IS-IS	42
	Table 4: BFD Modes Supported on SRX Series Devices	49
Part 4	Configuring Routing Engine Redundancy	
Chapter 7	Configuring Routing Engine Redundancy	119
	Table 5: Routing Engine Mastership Log	123
Part 5	Configuring Graceful Routing Engine Switchover (GRES)	
Chapter 9	Understanding How GRES Enables Uninterrupted Packet Forwarding During a Routing Engine Switchover	135
	Table 6: Effects of a Routing Engine Switchover	140
Chapter 10	GRES System Requirements	143
	Table 7: Graceful Routing Engine Switchover Feature Support	144
Part 7	Configuring Nonstop Active Routing (NSR)	
Chapter 18	Nonstop Active Routing System Requirements	205
	Table 8: Nonstop Active Routing Platform Support	205
	Table 9: Nonstop Active Routing Protocol and Feature Support	206
Part 9	Power Management Overview	
Chapter 23	Understanding Power Mangement	295
	Table 10: Available Operating Power in N+1 and N+N Redundancy Configurations	300
Chapter 24	Redundant Power System Overview	305
	Table 11: Sample Requirements and RPS Solutions	306
	Table 12: Redundant Power System Components	307

Part 10	Configuring Virtual Router Redundancy Protocol (VRRP)	
Chapter 25	Understanding How the VRRP Router Failover Mechanism Prevents Network Failures	315
	Table 13: RVIs on QFabric systems in example VRRP configuration	320
	Table 14: Sample VRRP configuration each RVI	321
	Table 15: Interfaces on QFabric system A. All interfaces are members of VLAN 100.	321
	Table 16: Interfaces on QFabric system B. All interfaces are members of VLAN 100 (same as on QFabric system A).	322
	Table 17: VRRPv2 to VRRPv3 Transition Steps and Events	325
Chapter 26	Configuring VRRP	331
	Table 18: Interface State and Priority Cost Usage	360
	Table 19: Settings for VRRP Load-Sharing Example	379
Part 11	Performing Unified In-Service Software Upgrade (ISSU)	
Chapter 27	Getting Started with Unified ISSU and Understanding How Unified ISSU Works	387
	Table 20: Locating the Information You Need to Work With ISSU	387
Chapter 28	Unified ISSU System Requirements	401
	Table 21: Unified ISSU Support for Dual Routing Engine Platforms	405
	Table 22: Unified ISSU PIC Support: SONET/SDH	409
	Table 23: Unified ISSU PIC Support: Fast Ethernet and Gigabit Ethernet	411
	Table 24: Unified ISSU PIC Support: Channelized	413
	Table 25: Unified ISSU PIC Support: Tunnel Services	414
	Table 26: Unified ISSU PIC Support: ATM	414
	Table 27: Unified ISSU Support: Enhanced IQ2 Ethernet Services Engine (ESE) PIC	416
	Table 28: Unified ISSU Support: MX Series Router MPCs	418
	Table 29: Unified ISSU Support: MX Series Router MICs	419
Chapter 29	Performing a Unified ISSU	423
	Table 30: Routing Engine Status Before Upgrading	431
	Table 31: Routing Engine Status After Upgrading and Rebooting Both Routing Engines	434
	Table 32: Routing Engine Status After Upgrading, Rebooting, and Switching Mastership	435
	Table 33: Routing Engine Status Before Upgrading and Manually Rebooting the Backup Routing Engine	438
	Table 34: Routing Engine Status After Upgrading and Before Manually Rebooting the Backup Routing Engine	441
	Table 35: Routing Engine Status After Upgrading, Manually Rebooting, and Before Switching Mastership	443
	Table 36: Routing Engine Status After Upgrading, Manually Rebooting, and Switching Mastership	444
	Table 37: Routing Engine Status Before Upgrading and Rebooting One Routing Engine	445

	Table 38: Routing Engine Status After Upgrading One Routing Engine and Before Upgrading the Other Routing Engine	448
	Table 39: Routing Engine Status After Upgrading, Manually Rebooting, and Switching Mastership	451
Part 12	Performing Nonstop Software Upgrade (NSSU)	
Chapter 31	Getting Started with NSSU and Understanding How NSSU Works	471
	Table 40: Platform and Release Support for NSSU	477
Part 13	Configuration Statements and Operational Commands	
Chapter 46	Operational Commands	661
	Table 41: show bgp neighbor Output Fields	664
	Table 42: show ospf overview Output Fields	688
	Table 43: show chassis nonstop-upgrade Output Fields	743
	Table 44: show chassis nonstop-upgrade node-group Output Fields	745
	Table 45: show chassis power-budget-statistics Output Fields	746
	Table 46: show chassis redundant-power-system Output Fields	750
	Table 47: show redundant-power-system led Output Fields	752
	Table 48: show redundant-power-system power-supply Output Fields	756
	Table 49: show redundant-power-system status Output Fields	758
	Table 50: show redundant-power-system upgrade Output Fields	760
	Table 51: show redundant-power-system version Output Fields	761
	Table 52: show chassis ssb Output Fields	763
	Table 53: show nonstop-routing Output Fields	765
	Table 54: show pfe ssb Output Fields	768
	Table 55: show system switchover Output Fields	776
	Table 56: show task replication Output Fields	780
	Table 57: show vrrp Output Fields	783
	Table 58: show vrrp track Output Fields	794

About the Documentation

- Documentation and Release Notes on page xxv
- Using the Examples in This Manual on page xxv
- Documentation Conventions on page xxvii
- Documentation Feedback on page xxix
- Requesting Technical Support on page xxix

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>.

Using the Examples in This Manual

If you want to use the examples in this manual, you can use the **load merge** or the **load merge relative** command. These commands cause the software to merge the incoming configuration into the current candidate configuration. The example does not become active until you commit the candidate configuration.

If the example configuration contains the top level of the hierarchy (or multiple hierarchies), the example is a *full example*. In this case, use the **load merge** command.

If the example configuration does not start at the top level of the hierarchy, the example is a *snippet*. In this case, use the **load merge relative** command. These procedures are described in the following sections.

Merging a Full Example

To merge a full example, follow these steps:

1. From the HTML or PDF version of the manual, copy a configuration example into a text file, save the file with a name, and copy the file to a directory on your routing platform.

For example, copy the following configuration to a file and name the file **ex-script.conf**. Copy the **ex-script.conf** file to the **/var/tmp** directory on your routing platform.

```
system {
  scripts {
    commit {
      file ex-script.xml;
    }
  }
}
interfaces {
  fxp0 {
    disable;
    unit 0 {
      family inet {
        address 10.0.0.1/24;
      }
    }
  }
}
```

2. Merge the contents of the file into your routing platform configuration by issuing the **load merge** configuration mode command:

```
[edit]
user@host# load merge /var/tmp/ex-script.conf
load complete
```

Merging a Snippet

To merge a snippet, follow these steps:

1. From the HTML or PDF version of the manual, copy a configuration snippet into a text file, save the file with a name, and copy the file to a directory on your routing platform.

For example, copy the following snippet to a file and name the file **ex-script-snippet.conf**. Copy the **ex-script-snippet.conf** file to the **/var/tmp** directory on your routing platform.

```
commit {
  file ex-script-snippet.xml; }
```

2. Move to the hierarchy level that is relevant for this snippet by issuing the following configuration mode command:

```
[edit]
user@host# edit system scripts
[edit system scripts]
```

3. Merge the contents of the file into your routing platform configuration by issuing the **load merge relative** configuration mode command:

```
[edit system scripts]
user@host# load merge relative /var/tmp/ex-script-snippet.conf
load complete
```

For more information about the **load** command, see [CLI Explorer](#).

Documentation Conventions

Table 1 on page xxvii 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
Bold text like this	Represents text that you type.	To enter configuration mode, type the configure command: user@host> configure
Fixed-width text like this	Represents output that appears on the terminal screen.	user@host> show chassis alarms No alarms currently active
<i>Italic text like this</i>	<ul style="list-style-type: none"> Introduces or emphasizes important new terms. Identifies guide names. Identifies RFC and Internet draft titles. 	<ul style="list-style-type: none"> A policy <i>term</i> is a named structure that defines match conditions and actions. <i>Junos OS CLI User Guide</i> RFC 1997, <i>BGP Communities Attribute</i>
<i>Italic text like this</i>	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# set system domain-name <i>domain-name</i>
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	<ul style="list-style-type: none"> To configure a stub area, include the stub statement at the [edit protocols ospf area area-id] hierarchy level. The console port is labeled CONSOLE.
< > (angle brackets)	Encloses optional keywords or variables.	stub <default-metric <i>metric</i> >;
(pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	broadcast multicast (<i>string1</i> <i>string2</i> <i>string3</i>)
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	rsvp { # Required for dynamic MPLS only
[] (square brackets)	Encloses a variable for which you can substitute one or more values.	community name members [<i>community-ids</i>]
Indentation and braces ({ })	Identifies a level in the configuration hierarchy.	[edit] routing-options { static { route default { nexthop <i>address</i> ; retain; } } }
;(semicolon)	Identifies a leaf statement at a configuration hierarchy level.	

GUI Conventions


Table 2: Text and Syntax Conventions (continued)

Convention	Description	Examples
Bold text like this	Represents graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none"> In the Logical Interfaces box, select All Interfaces. To cancel the configuration, click Cancel.
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select Protocols>Ospf .

Documentation Feedback

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

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



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

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active Juniper Care or Partner Support Services support contract, or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

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

Self-Help Online Tools and Resources

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

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

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

Creating a Service Request with JTAC

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

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

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

PART 1

Overview

- [High Availability Overview on page 3](#)

CHAPTER 1

High Availability Overview

- [Understanding High Availability Features on Juniper Networks Routers on page 3](#)
- [High Availability-Related Features in Junos OS on page 9](#)
- [High Availability Features for EX Series Switches Overview on page 9](#)

Understanding High Availability Features on Juniper Networks Routers

For Juniper Networks routing platforms running the Junos operating system (Junos OS), *high availability* refers to the hardware and software components that provide redundancy and reliability for packet-based communications. This topic provides brief overviews of the following high availability features:

- [Routing Engine Redundancy on page 3](#)
- [Graceful Routing Engine Switchover on page 3](#)
- [Nonstop Bridging on page 4](#)
- [Nonstop Active Routing on page 4](#)
- [Graceful Restart on page 5](#)
- [Nonstop Active Routing Versus Graceful Restart on page 6](#)
- [Effects of a Routing Engine Switchover on page 7](#)
- [VRRP on page 7](#)
- [Unified ISSU on page 7](#)
- [Interchassis Redundancy for MX Series Routers Using Virtual Chassis on page 8](#)

Routing Engine Redundancy

Redundant Routing Engines are two Routing Engines that are installed in the same routing platform. One functions as the master, while the other stands by as a backup should the master Routing Engine fail. On routing platforms with dual Routing Engines, network reconvergence takes place more quickly than on routing platforms with a single Routing Engine.

Graceful Routing Engine Switchover

Graceful Routing Engine switchover (GRES) enables a routing platform with redundant Routing Engines to continue forwarding packets, even if one Routing Engine fails. Graceful Routing Engine switchover preserves interface and kernel information. Traffic is not

interrupted. However, graceful Routing Engine switchover does not preserve the control plane. Neighboring routers detect that the router has experienced a restart and react to the event in a manner prescribed by individual routing protocol specifications.



NOTE: To preserve routing during a switchover, graceful Routing Engine switchover must be combined with either graceful restart protocol extensions or nonstop active routing. For more information, see [“Understanding Graceful Routing Engine Switchover” on page 135](#) and [“Nonstop Active Routing Concepts” on page 199](#).



NOTE: In T Series routers, TX Matrix routers, and TX Matrix Plus routers, the control plane is preserved in case of GRES with NSR, and 75% of line rate worth of traffic per Packet Forwarding Engine remains uninterrupted during GRES.

Nonstop Bridging

Nonstop bridging enables an MX Series 5G Universal Routing Platform with redundant Routing Engines to switch from a primary Routing Engine to a backup Routing Engine without losing Layer 2 Control Protocol (L2CP) information. Nonstop bridging uses the same infrastructure as graceful Routing Engine switchover to preserve interface and kernel information. However, nonstop bridging also saves L2CP information by running the Layer 2 Control Protocol process (l2cpd) on the backup Routing Engine.



NOTE: To use nonstop bridging, you must first enable graceful Routing Engine switchover.

Nonstop bridging is supported for the following Layer 2 control protocols:

- Spanning Tree Protocol (STP)
- Rapid Spanning Tree Protocol (RSTP)
- Multiple Spanning Tree Protocol (MSTP)
- VLAN Spanning Tree Protocol (VSTP)

For more information, see [“Nonstop Bridging Concepts” on page 185](#).

Nonstop Active Routing

Nonstop active routing (NSR) enables a routing platform with redundant Routing Engines to switch from a primary Routing Engine to a backup Routing Engine without alerting peer nodes that a change has occurred. Nonstop active routing uses the same infrastructure as graceful Routing Engine switchover to preserve interface and kernel information. However, nonstop active routing also preserves routing information and

protocol sessions by running the routing protocol process (rpd) on both Routing Engines. In addition, nonstop active routing preserves TCP connections maintained in the kernel.



NOTE: To use nonstop active routing, you must also configure graceful Routing Engine switchover.

For a list of protocols and features supported by nonstop active routing, see [“Nonstop Active Routing Protocol and Feature Support” on page 206](#).

For more information about nonstop active routing, see [“Nonstop Active Routing Concepts” on page 199](#).

Graceful Restart

With routing protocols, any service interruption requires an affected router to recalculate adjacencies with neighboring routers, restore routing table entries, and update other protocol-specific information. An unprotected restart of a router can result in forwarding delays, route flapping, wait times stemming from protocol reconvergence, and even dropped packets. To alleviate this situation, graceful restart provides extensions to routing protocols. These protocol extensions define two roles for a router—*restarting* and *helper*. The extensions signal neighboring routers about a router undergoing a restart and prevent the neighbors from propagating the change in state to the network during a graceful restart wait interval. The main benefits of graceful restart are uninterrupted packet forwarding and temporary suppression of all routing protocol updates. Graceful restart enables a router to pass through intermediate convergence states that are hidden from the rest of the network.

When a router is running graceful restart and the router stops sending and replying to protocol liveness messages (hellos), the adjacencies assume a graceful restart and begin running a timer to monitor the restarting router. During this interval, helper routers do not process an adjacency change for the router that they assume is restarting, but continue active routing with the rest of the network. The helper routers assume that the router can continue stateful forwarding based on the last preserved routing state during the restart.

If the router was actually restarting and is back up before the graceful timer period expires in all of the helper routers, the helper routers provide the router with the routing table, topology table, or label table (depending on the protocol), exit the graceful period, and return to normal network routing.

If the router does not complete its negotiation with helper routers before the graceful timer period expires in all of the helper routers, the helper routers process the router's change in state and send routing updates, so that convergence occurs across the network. If a helper router detects a link failure from the router, the topology change causes the helper router to exit the graceful wait period and to send routing updates, so that network convergence occurs.

To enable a router to undergo a graceful restart, you must include the **graceful-restart** statement at the global **[edit routing-options]** or **[edit routing-instances instance-name routing-options]** hierarchy level. You can, optionally, modify the global settings at the

individual protocol level. When a routing session is started, a router that is configured with graceful restart must negotiate with its neighbors to support it when it undergoes a graceful restart. A neighboring router will accept the negotiation and support helper mode without requiring graceful restart to be configured on the neighboring router.



NOTE: A Routing Engine switchover event on a helper router that is in graceful wait state causes the router to drop the wait state and to propagate the adjacency's state change to the network.

Graceful restart is supported for the following protocols and applications:

- BGP
- ES-IS
- IS-IS
- OSPF/OSPFv3
- PIM sparse mode
- RIP/RIPng
- MPLS-related protocols, including:
 - Label Distribution Protocol (LDP)
 - Resource Reservation Protocol (RSVP)
 - Circuit cross-connect (CCC)
 - Translational cross-connect (TCC)
- Layer 2 and Layer 3 virtual private networks (VPNs)

For more information, see [“Graceful Restart Concepts” on page 233](#).

Nonstop Active Routing Versus Graceful Restart

Nonstop active routing and graceful restart are two different methods of maintaining high availability. Graceful restart requires a router restart. A router undergoing a graceful restart relies on its neighbors (or helpers) to restore its routing protocol information. The restart is the mechanism by which helpers are signaled to exit the wait interval and start providing routing information to the restarting router. For more information, see [“Graceful Restart Concepts” on page 233](#).

In contrast, nonstop active routing does not involve a router restart. Both the master and backup Routing Engines are running the routing protocol process (rpd) and exchanging updates with neighbors. When one Routing Engine fails, the router simply switches to the active Routing Engine to exchange routing information with neighbors. Because of these feature differences, nonstop routing and graceful restart are mutually exclusive. Nonstop active routing cannot be enabled when the router is configured as a graceful restarting router. If you include the **graceful-restart** statement at any hierarchy level and the **nonstop-routing** statement at the **[edit routing-options]** hierarchy level and try to

commit the configuration, the commit request fails. For more information, see [“Nonstop Active Routing Concepts” on page 199](#).

Effects of a Routing Engine Switchover

[“Effects of a Routing Engine Switchover” on page 7](#) describes the effects of a Routing Engine switchover when no high availability features are enabled and when graceful Routing Engine switchover, graceful restart, and nonstop active routing features are enabled.

VRRP

The Virtual Router Redundancy Protocol (VRRP) enables hosts on a LAN to make use of redundant routing platforms (master and backup pairs) on the LAN, requiring only the static configuration of a single default route on the hosts.

The VRRP routing platform pairs share the IP address corresponding to the default route configured on the hosts. At any time, one of the VRRP routing platforms is the master (active) and the others are backups. If the master fails, one of the backup routers or switches becomes the new master router.

VRRP has advantages in ease of administration and network throughput and reliability:

- It provides a virtual default routing platform.
- It enables traffic on the LAN to be routed without a single point of failure.
- A virtual backup router can take over a failed default router:
 - Within a few seconds.
 - With a minimum of VRRP traffic.
 - Without any interaction with the hosts.

Devices running VRRP dynamically elect master and backup routers. You can also force assignment of master and backup routers using priorities from 1 through 255, with 255 being the highest priority.

In VRRP operation, the default master router sends advertisements to backup routers at regular intervals (default 1 second). If a backup router does not receive an advertisement for a set period, the backup router with the next highest priority takes over as master and begins forwarding packets.

As of Junos OS Release 13.2, VRRP nonstop active routing (NSR) is enabled only when you configure the **nonstop-routing** statement at the **[edit routing-options]** or **[edit logical-system *logical-system-name* routing-options]** hierarchy level.

For more information, see [“Understanding VRRP” on page 315](#).

Unified ISSU

A unified in-service software upgrade (unified ISSU) enables you to upgrade between two different Junos OS Releases with no disruption on the control plane and with minimal disruption of traffic. Unified ISSU is only supported by dual Routing Engine platforms. In

addition, graceful Routing Engine switchover (GRES) and nonstop active routing (NSR) must be enabled.

With a unified ISSU, you can eliminate network downtime, reduce operating costs, and deliver higher services levels. For more information, see [“Getting Started with Unified In-Service Software Upgrade”](#) on page 387.

Interchassis Redundancy for MX Series Routers Using Virtual Chassis

Interchassis redundancy is a high availability feature that can span equipment located across multiple geographies to prevent network outages and protect routers against access link failures, uplink failures, and wholesale chassis failures without visibly disrupting the attached subscribers or increasing the network management burden for service providers. As more high-priority voice and video traffic is carried on the network, interchassis redundancy has become a requirement for providing stateful redundancy on broadband subscriber management equipment such as broadband services routers, broadband network gateways, and broadband remote access servers. Interchassis redundancy support enables service providers to fulfill strict service-level agreements (SLAs) and avoid unplanned network outages to better meet the needs of their customers.

To provide a stateful interchassis redundancy solution for MX Series 5G Universal Routing Platforms, you can configure a Virtual Chassis. A *Virtual Chassis* configuration interconnects two MX Series routers into a logical system that you can manage as a single network element. The member routers in a Virtual Chassis are designated as the *master router* (also known as the *protocol master*) and the *backup router* (also known as the *protocol backup*). The member routers are interconnected by means of dedicated *Virtual Chassis ports* that you configure on Trio Modular Port Concentrator/Modular Interface Card (MPC/MIC) interfaces.

An MX Series Virtual Chassis is managed by the *Virtual Chassis Control Protocol (VCCP)*, which is a dedicated control protocol based on IS-IS. VCCP runs on the Virtual Chassis port interfaces and is responsible for building the Virtual Chassis topology, electing the Virtual Chassis master router, and establishing the interchassis routing table to route traffic within the Virtual Chassis.

Starting with Junos OS Release 11.2, Virtual Chassis configurations are supported on MX240, MX480, and MX960 Universal Routing Platforms with Trio MPC/MIC interfaces and dual Routing Engines. In addition, graceful Routing Engine switchover (GRES) and nonstop active routing (NSR) must be enabled on both member routers in the Virtual Chassis.

Related Documentation

- [High Availability-Related Features in Junos OS on page 9](#)

High Availability-Related Features in Junos OS

Related redundancy and reliability features include:

- Redundant power supplies, host modules, host subsystems, and forwarding boards. For more information, see the *Junos OS Administration Library* and the *Junos OS Hardware Network Operations Guide*.
- Additional link-layer redundancy, including Automatic Protection Switching (APS) for SONET interfaces, Multiplex Section Protection (MSP) for SDH interfaces, and DLSw redundancy for Ethernet interfaces. For more information, see the *Junos OS Network Interfaces Library for Routing Devices*.
- Bidirectional Forwarding Detection (BFD) works with other routing protocols to detect failures rapidly. For more information, see the *Junos OS Routing Protocols Library*.
- Redirection of Multiprotocol Label Switching (MPLS) label-switched path (LSP) traffic—Mechanisms such as link protection, node-link protection, and fast reroute recognize link and node failures, allowing MPLS LSPs to select a bypass LSP to circumvent failed links or devices. For more information, see the *MPLS Applications Feature Guide*.

Related Documentation

- [Understanding High Availability Features on Juniper Networks Routers on page 3](#)

High Availability Features for EX Series Switches Overview

High availability refers to the hardware and software components that provide redundancy and reliability for network communications. This topic covers the following high availability features of Juniper Networks EX Series Ethernet Switches:

- [VRRP on page 9](#)
- [Graceful Protocol Restart on page 10](#)
- [Redundant Routing Engines on page 10](#)
- [Virtual Chassis on page 11](#)
- [Graceful Routing Engine Switchover on page 11](#)
- [Link Aggregation on page 12](#)
- [Nonstop Active Routing and Nonstop Bridging on page 12](#)
- [Nonstop Software Upgrade on page 12](#)
- [Redundant Power System on page 13](#)

VRRP

You can configure Virtual Router Redundancy Protocol (VRRP) for IP and IPv6 on most switch interfaces, including Gigabit Ethernet interfaces, high-speed Gigabit Ethernet uplink interfaces, and logical interfaces. When VRRP is configured, the switches act as virtual routing platforms. VRRP enables hosts on a LAN to make use of redundant routing

platforms on that LAN without requiring more than the static configuration of a single default route on the hosts. The VRRP routing platforms share the IP address corresponding to the default route configured on the hosts. At any time, one of the VRRP routing platforms is the master (active) and the others are backups. If the master routing platform fails, one of the backup routing platforms becomes the new master, providing a virtual default routing platform and enabling traffic on the LAN to be routed without relying on a single routing platform. Using VRRP, a backup switch can take over a failed default switch within a few seconds. This is done with minimum loss of VRRP traffic and without any interaction with the hosts.

Graceful Protocol Restart

With standard implementations of routing protocols, any service interruption requires an affected switch to recalculate adjacencies with neighboring switches, restore routing table entries, and update other protocol-specific information. An unprotected restart of a switch can result in forwarding delays, route flapping, wait times stemming from protocol reconvergence, and even dropped packets. Graceful protocol restart enables a restarting switch and its neighbors to continue forwarding packets without disrupting network performance. Because neighboring switches assist in the restart (these neighbors are called helper switches), the restarting switch can quickly resume full operation without recalculating algorithms from scratch.

On the switches, graceful protocol restart can be applied to aggregate and static routes and for routing protocols (BGP, IS-IS, OSPF, and RIP).

Graceful protocol restart works similarly for the different routing protocols. The main benefits of graceful protocol restart are uninterrupted packet forwarding and temporary suppression of all routing protocol updates. Graceful protocol restart thus allows a switch to pass through intermediate convergence states that are hidden from the rest of the network. Most graceful restart implementations define two types of switches—the restarting switch and the helper switch. The restarting switch requires rapid restoration of forwarding state information so that it can resume the forwarding of network traffic. The helper switch assists the restarting switch in this process. Individual graceful restart configuration statements typically apply to either the restarting switch or the helper switch.

Redundant Routing Engines

Redundant Routing Engines are two Routing Engines that are installed in a switch or a Virtual Chassis. When a switch has two Routing Engines, one functions as the master, while the other stands by as a backup in case the master Routing Engine fails. When a Virtual Chassis has two Routing Engines, the switch in the master role functions as the master Routing Engine and the switch in the backup role functions as the backup Routing Engine. Redundant Routing Engines are supported on Juniper Networks EX6200 Ethernet Switches, Juniper Networks EX8200 Ethernet Switches, and on all EX Series Virtual Chassis configurations.

The master Routing Engine receives and transmits routing information, builds and maintains routing tables, communicates with interfaces and Packet Forwarding Engine components of the switch, and has full control over the control plane of the switch.

The backup Routing Engine stays in sync with the master Routing Engine in terms of protocol states, forwarding tables, and so forth. If the master becomes unavailable, the backup Routing Engine takes over the functions that the master Routing Engine performs.

Network reconvergence takes place more quickly on switches and on Virtual Chassis with redundant Routing Engines than on switches and on Virtual Chassis with a single Routing Engine.

Virtual Chassis

A Virtual Chassis is multiple switches connected together that operate as a single network entity. The advantages of connecting multiple switches into a Virtual Chassis include better-managed bandwidth at a network layer, simplified configuration and maintenance because multiple devices can be managed as a single device, a simplified Layer 2 network topology that minimizes or eliminates the need for loop prevention protocols such as Spanning Tree Protocol (STP), and improved fault tolerance and high availability. A Virtual Chassis improves high availability for the following reasons:

- **Dual Routing Engine support.** A Virtual Chassis automatically has two Routing Engines—the switches in the master and backup **routing-engine** roles—and, therefore, provides more high availability options than standalone switches. Many high availability features, including graceful protocol restart, graceful Routing Engine switchover (GRES), nonstop software upgrade (NSSU), nonstop active routing (NSR), and nonstop bridging (NSB), are available for an EX Series Virtual Chassis that are not available on standalone EX Series switches.
- **Increased fault tolerance.** You increase your fault tolerance options when you configure your EX Series switches into a Virtual Chassis. You can, for instance, configure interfaces into a link aggregation group (LAG) with member interfaces on different member switches in the same Virtual Chassis to ensure network traffic is received by a Virtual Chassis even when a switch or physical interface in the Virtual Chassis fails.

Juniper Networks EX2200 Ethernet Switches, Juniper Networks EX3300 Ethernet Switches, Juniper Networks EX4200 Ethernet Switches, Juniper Networks EX4300 Ethernet Switches, Juniper Networks EX4500 Ethernet Switches, Juniper Networks EX4550 Ethernet Switches, or Juniper Networks EX8200 Ethernet Switches can form a Virtual Chassis. EX4200, EX4500, and EX4550 switches can be interconnected together to form a mixed Virtual Chassis.

Graceful Routing Engine Switchover

You can configure graceful Routing Engine switchover (GRES) on a switch with redundant Routing Engines or on a Virtual Chassis, allowing control to switch from the master Routing Engine to the backup Routing Engine with minimal interruption to network communications. When you configure GRES, the backup Routing Engine automatically synchronizes with the master Routing Engine to preserve kernel state information and forwarding state. Any updates to the master Routing Engine are replicated to the backup Routing Engine as soon as they occur. If the kernel on the master Routing Engine stops operating, the master Routing Engine experiences a hardware failure, or the administrator initiates a manual switchover, mastership switches to the backup Routing Engine.

When the backup Routing Engine assumes mastership in a redundant failover configuration (that is, when GRES is not enabled), the Packet Forwarding Engines initialize their state to the boot state before they connect to the new master Routing Engine. In contrast, in a GRES configuration, the Packet Forwarding Engines do not reinitialize their state, but resynchronize their state to that of the new master Routing Engine. The interruption to traffic is minimal.

Link Aggregation

You can combine multiple physical Ethernet ports to form a logical point-to-point link, known as a link aggregation group (LAG) or bundle. A LAG provides more bandwidth than a single Ethernet link can provide. Additionally, link aggregation provides network redundancy by load-balancing traffic across all available links. If one of the links should fail, the system automatically load-balances traffic across all remaining links. In a Virtual Chassis, LAGs can be used to load-balance network traffic between member switches, which increases high availability by ensuring that network traffic is received by the Virtual Chassis even if a single interface fails for any reason.

The number of Ethernet interfaces you can include in a LAG and the number of LAGs you can configure on a switch depend on the switch model.

Nonstop Active Routing and Nonstop Bridging

Nonstop active routing (NSR) provides high availability in a switch with redundant Routing Engines by enabling transparent switchover of the Routing Engines without requiring restart of supported Layer 3 routing protocols. Both Routing Engines are fully active in processing protocol sessions, and so each can take over for the other. The switchover is transparent to neighbor routing devices, which do not detect that a change has occurred.

Nonstop bridging (NSB) provides the same mechanism for Layer 2 protocols. NSB provides high availability in a switch with redundant Routing Engines by enabling transparent switchover of the Routing Engines without requiring restart of supported Layer 2 protocols. Both Routing Engines are fully active in processing protocol sessions, and so each can take over for the other. The switchover is transparent to neighbor switching devices, which do not detect that a change has occurred.

To use NSR or NSB, you must also configure GRES.

Nonstop Software Upgrade

Nonstop software upgrade (NSSU) allows you to upgrade the software on a switch with dual Routing Engines or on a Virtual Chassis in an automated manner with minimal traffic disruption. NSSU takes advantage of GRES and NSR to enable upgrading the Junos OS version with no disruption to the control plane. In addition, NSSU minimizes traffic disruption by:

- Upgrading line cards one at a time in an EX6200 switch, EX8200 switch, or EX8200 Virtual Chassis, permitting traffic to continue to flow through the line cards that are not being upgraded.
- Upgrading member switches one at a time in all other Virtual Chassis, permitting traffic to continue to flow through the members that are not being upgraded.

By configuring LAGs such that the member links reside on different line cards or Virtual Chassis members, you can achieve minimal traffic disruption when performing an NSSU.

Redundant Power System

Most Juniper Networks Ethernet Switches have a built-in capability for redundant power supplies—therefore if one power supply fails on those switches, the other power supply takes over. However, EX2200 switches and EX3300 switches have only one internal fixed power supply. If an EX2200 switch or EX3300 switch is deployed in a critical situation, we recommend that you connect a Redundant Power System (RPS) to that switch to supply backup power if the internal power supply fails. RPS is not a primary power supply—it only provides backup power to switches when the single dedicated power supply fails. An RPS operates in parallel with the single dedicated power supplies of the switches connected to it and provides all connected switches enough power to support either Power over Ethernet (PoE) or non-PoE devices. For more information about RPS, see [“EX Series Redundant Power System Hardware Overview” on page 305](#).

Related Documentation

- For more information about high availability features, see the [Junos OS High Availability Configuration Guide](#).
- *Understanding EX Series Virtual Chassis*
- *EX8200 Virtual Chassis Overview*
- [Understanding Nonstop Active Routing on EX Series Switches on page 203](#)
- [Understanding Nonstop Software Upgrade on EX Series Switches on page 471](#)
- *EX Series Redundant Power System Hardware Guide*

PART 2

Configuring Switching Control Board Redundancy

- [Understanding How Switching Control Board Redundancy Prevents Network Failures on page 17](#)
- [Configuring Switching Control Board Redundancy on page 23](#)

CHAPTER 2

Understanding How Switching Control Board Redundancy Prevents Network Failures

- [Understanding Switching Control Board Redundancy on page 17](#)

Understanding Switching Control Board Redundancy

This section describes the following redundant switching control boards:



NOTE: A failover from a master switching control board to a backup switching control board occurs automatically when the master experiences a hardware failure or when you have configured the software to support a change in mastership based on specific conditions. You can also manually switch mastership by issuing specific `request chassis` commands. In this section, the term *failover* refers to an automatic event, whereas *switchover* refers to either an automatic or a manual event.

- [Redundant CFEBs on the M10i Router on page 17](#)
- [Redundant FEBs on the M120 Router on page 18](#)
- [Redundant SSBs on the M20 Router on page 20](#)
- [Redundant SFMs on the M40e and M160 Routers on page 20](#)

Redundant CFEBs on the M10i Router

On the M10i router, the CFEB performs the following functions:

- Route lookups—Performs route lookups using the forwarding table stored in synchronous SRAM (SSRAM).
- Management of shared memory—Uniformly allocates incoming data packets throughout the router's shared memory.

- Transfer of outgoing data packets—Passes data packets to the destination Fixed Interface Card (FIC) or Physical Interface Card (PIC) when the data is ready to be transmitted.
- Transfer of exception and control packets—Passes exception packets to the microprocessor on the CFEB, which processes almost all of them. The remainder are sent to the Routing Engine for further processing. Any errors originating in the Packet Forwarding Engine and detected by the CFEB are sent to the Routing Engine using system log messages.

The M10i router has two CFEBs, one that is configured to act as the master and the other that serves as a backup in case the master fails. You can initiate a manual switchover by issuing the **request chassis cfeb master switch** command. For more information, see the *Junos OS Administration Library*.

Redundant FEBs on the M120 Router

The M120 router supports up to six Forwarding Engine Boards (FEBs). Flexible PIC Concentrator (FPCs), which host PICs, are separate from the FEBs, which handle packet forwarding. FPCs are located on the front of the chassis and provide power and management to PICs through the midplane. FEBs are located on the back of the chassis and receive signals from the midplane, which the FEBs process for packet forwarding. The midplane allows any FEB to carry traffic for any FPC.

To configure the mapping of FPCs to FEBs, use the **fpc-feb-connectivity** statement as described in the *Junos OS Administration Library*. You cannot specify a connection between an FPC and a FEB configured as a backup. If an FPC is not specified to connect to a FEB, the FPC is assigned automatically to the FEB with the same slot number. For example, the FPC in slot 1 is assigned to the FEB in slot 1.

You can configure one FEB as a backup for one or more FEBs by configuring a FEB redundancy group. When a FEB fails, the backup FEB can quickly take over packet forwarding. A redundancy group must contain exactly one backup FEB and can optionally contain one primary FEB and multiple other FEBs. A FEB can belong to only one group. A group can provide backup on a one-to-one basis (primary-to-backup), a many-to-one basis (two or more other-FEBs-to-backup), or a combination of both (one primary-to-backup and one or more other-FEBs-to-backup).

When you configure a primary FEB in a redundancy group, the backup FEB mirrors the exact forwarding state of the primary FEB. If switchover occurs from a primary FEB, the backup FEB does not reboot. A manual switchover from the primary FEB to the backup FEB results in less than 1 second of traffic loss. Failover from the primary FEB to the backup FEB results in less than 10 seconds of traffic loss.

If a failover occurs from the other FEB and a primary FEB is specified for the group, the backup FEB reboots so that the forwarding state from the other FEB can be downloaded to the backup FEB and forwarding can continue. Automatic failover from a FEB that is not specified as a primary FEB results in higher packet loss. The duration of packet loss depends on the number of interfaces and on the size of the routing table, but it can be minutes.

If a failover from a FEB occurs when no primary FEB is specified in the redundancy group, the backup FEB does not reboot and the interfaces on the FPC connected to the previously active FEB remain online. The backup FEB must obtain the entire forwarding state from the Routing Engine after a switchover, and this update may take a few minutes. If you do not want the interfaces to remain online during the switchover for the other FEB, configure a primary FEB for the redundancy group.

Failover to a backup FEB occurs automatically if a FEB in a redundancy group fails. You can disable automatic failover for any redundancy group by including the **no-auto-failover** statement at the **[edit chassis redundancy feb redundancy-group group-name]** hierarchy level.

You can also initiate a manual switchover by issuing the **request chassis redundancy feb slot slot-number switch-to-backup** command, where **slot-number** is the number of the active FEB. For more information, see the [CLI Explorer](#).

The following conditions result in failover as long as the backup FEB in a redundancy group is available:

- The FEB is absent.
- The FEB experienced a hard error while coming online.
- A software failure on the FEB resulted in a crash.
- Ethernet connectivity from a FEB to a Routing Engine failed.
- A hard error on the FEB, such as a power failure, occurred.
- The FEB was disabled when the offline button for the FEB was pressed.
- The software watchdog timer on the FEB expired.
- Errors occurred on the links between all the active fabric planes and the FEB. This situation results in failover to the backup FEB if it has at least one valid fabric link.
- Errors occurred on the link between the FEB and all of the FPCs connected to it.

After a switchover occurs, a backup FEB is no longer available for the redundancy group. You can revert from the backup FEB to the previously active FEB by issuing the operational mode command **request chassis redundancy feb slot slot-number revert-from-backup**, where **slot-number** is the number of the previously active FEB. For more information, see the [CLI Explorer](#).

When you revert from the backup FEB, it becomes available again for a switchover. If the redundancy group does not have a primary FEB, the backup FEB reboots after you revert back to the previously active FEB. If the FEB to which you revert back is not a primary FEB, the backup FEB is rebooted so that it can align with the state of the primary FEB.

If you modify the configuration for an existing redundancy group so that a FEB connects to a different FPC, the FEB is rebooted unless the FEB was already connected to one or two Type 1 FPCs and the change only resulted in the FEB being connected either to one additional or one fewer Type 1 FPC. For more information about how to map a connection between an FPC and a FEB, see the *Junos OS Administration Library*. If you change the primary FEB in a redundancy group, the backup FEB is rebooted. The FEB is also rebooted

if you change a backup FEB to a nonbackup FEB or change an active FEB to a backup FEB.

To view the status of configured FEB redundancy groups, issue the **show chassis redundancy feb** operational mode command. For more information, see the [CLI Explorer](#).

Redundant SSBs on the M20 Router

The System and Switch Board (SSB) on the M20 router performs the following major functions:

- Shared memory management on the FPCs—The Distributed Buffer Manager ASIC on the SSB uniformly allocates incoming data packets throughout shared memory on the FPCs.
- Outgoing data cell transfer to the FPCs—A second Distributed Buffer Manager ASIC on the SSB passes data cells to the FPCs for packet reassembly when the data is ready to be transmitted.
- Route lookups—The Internet Processor ASIC on the SSB performs route lookups using the forwarding table stored in SSRAM. After performing the lookup, the Internet Processor ASIC informs the midplane of the forwarding decision, and the midplane forwards the decision to the appropriate outgoing interface.
- System component monitoring—The SSB monitors other system components for failure and alarm conditions. It collects statistics from all sensors in the system and relays them to the Routing Engine, which sets the appropriate alarm. For example, if a temperature sensor exceeds the first internally defined threshold, the Routing Engine issues a “high temp” alarm. If the sensor exceeds the second threshold, the Routing Engine initiates a system shutdown.
- Exception and control packet transfer—The Internet Processor ASIC passes exception packets to a microprocessor on the SSB, which processes almost all of them. The remaining packets are sent to the Routing Engine for further processing. Any errors that originate in the Packet Forwarding Engine and are detected by the SSB are sent to the Routing Engine using system log messages.
- FPC reset control—The SSB monitors the operation of the FPCs. If it detects errors in an FPC, the SSB attempts to reset the FPC. After three unsuccessful resets, the SSB takes the FPC offline and informs the Routing Engine. Other FPCs are unaffected, and normal system operation continues.

The M20 router holds up to two SSBs. One SSB is configured to act as the master and the other is configured to serve as a backup in case the master fails. You can initiate a manual switchover by issuing the **request chassis ssb master switch** command. For more information, see the [CLI Explorer](#).

Redundant SFMs on the M40e and M160 Routers

The M40e and M160 routers have redundant Switching and Forwarding Modules (SFMs). The SFMs contain the Internet Processor II ASIC and two Distributed Buffer Manager ASICs. SFMs ensure that all traffic leaving the FPCs is handled properly. SFMs provide route lookup, filtering, and switching.

The M40e router holds up to two SFMs, one that is configured to act as the master and the other configured to serve as a backup in case the master fails. Removing the standby SFM has no effect on router function. If the active SFM fails or is removed from the chassis, forwarding halts until the standby SFM boots and becomes active. It takes approximately 1 minute for the new SFM to become active. Synchronizing router configuration information can take additional time, depending on the complexity of the configuration.

The M160 router holds up to four SFMs. All SFMs are active at the same time. A failure or taking an SFM offline has no effect on router function. Forwarding continues uninterrupted.

You can initiate a manual switchover by issuing the **request chassis sfm master switch** command. For more information, see the [CLI Explorer](#).

**Related
Documentation**

- [Understanding High Availability Features on Juniper Networks Routers on page 3](#)
- [Understanding Routing Engine Redundancy on Juniper Networks Routers on page 111](#)
- [Configuring CFEB Redundancy on the M10i Router on page 23](#)
- [Configuring FEB Redundancy on the M120 Router on page 24](#)
- [Configuring SFM Redundancy on M40e and M160 Routers on page 26](#)
- [Configuring SSB Redundancy on the M20 Router on page 27](#)
- *show chassis redundancy feb*
- *request chassis cb*

CHAPTER 3

Configuring Switching Control Board Redundancy

- [Configuring CFEB Redundancy on the M10i Router on page 23](#)
- [Configuring FEB Redundancy on the M120 Router on page 24](#)
- [Example: Configuring FEB Redundancy on M120 Routers on page 25](#)
- [Configuring SFM Redundancy on M40e and M160 Routers on page 26](#)
- [Configuring SSB Redundancy on the M20 Router on page 27](#)
- [Configuring the Junos OS to Support Redundancy on Routers Having Multiple Routing Engines or Switching Boards on page 27](#)

Configuring CFEB Redundancy on the M10i Router

The Compact Forwarding Engine Board (CFEB) on the M10i router provides route lookup, filtering, and switching on incoming data packets, and then directs outbound packets to the appropriate interface for transmission to the network. The CFEB communicates with the Routing Engine using a dedicated 100-Mbps Fast Ethernet link that transfers routing table data from the Routing Engine to the forwarding table in the integrated ASIC. The link is also used to transfer from the CFEB to the Routing Engine routing link-state updates and other packets destined for the router that have been received through the router interfaces.

To configure a CFEB redundancy group, include the following statements at the **[edit chassis redundancy]** hierarchy level:

```
[edit chassis redundancy]  
  cfeb slot-number (always | preferred);
```

slot-number can be 0 or 1.

always defines the CFEB as the sole device.

preferred defines a preferred CFEB.

To manually switch CFEB mastership, issue the **request chassis cfeb master switch** command. To view CFEB status, issue the **show chassis cfeb** command.

- Related Documentation
- [Understanding Switching Control Board Redundancy on page 17](#)

Configuring FEB Redundancy on the M120 Router

To configure a FEB redundancy group for the M120 router, include the following statements at the `[edit chassis redundancy feb]` hierarchy level:

```
[edit chassis redundancy feb]
redundancy-group group-name {
  description description;
  feb slot-number (backup | primary);
  no-auto-failover;
}
```

group-name is the unique name for the redundancy group. The maximum length is 39 alphanumeric characters.

slot-number is the slot number of each FEB you want to include in the redundancy group. The range is from 0 through 5. You must specify exactly one FEB as a backup FEB per redundancy group. Include the **backup** keyword when configuring the backup FEB and make sure that the FEB is not connected to an FPC.

Include the **primary** keyword to optionally specify one primary FEB per redundancy group. When the **primary** keyword is specified for a particular FEB, that FEB is configured for 1:1 redundancy. With 1:1 redundancy, the backup FEB contains the same forwarding state as the primary FEB. When no FEB in the redundancy group is configured as a primary FEB, the redundancy group is configured for *n*:1 redundancy. In this case, the backup FEB has no forwarding state. When a FEB fails, the forwarding state must be downloaded from the Routing Engine to the backup FEB before forwarding continues.

A combination of 1:1 and *n*:1 redundancy is possible when more than two FEBs are present in a group. The backup FEB contains the same forwarding state as the primary FEB, so that when the primary FEB fails, 1:1 failover is in effect. When a nonprimary FEB fails, the backup FEB must be rebooted so that the forwarding state from the nonprimary FEB is installed on the backup FEB before it can continue forwarding.

You can optionally include the **description** statement to describe a redundancy group.

Automatic failover is enabled by default. To disable automatic failover, include the **no-auto-failover** statement. If you disable automatic failover, you can perform only a manual switchover using the operational command **request chassis redundancy feb slot slot-number switch-to-backup**.

To view FEB status, issue the **show chassis feb** command. For more information, see the [CLI Explorer](#).

- Related Documentation
- [Understanding Switching Control Board Redundancy on page 17](#)
 - [Example: Configuring FEB Redundancy on M120 Routers on page 25](#)

Example: Configuring FEB Redundancy on M120 Routers

In the following configuration, two FEB redundancy groups are created:

- A FEB redundancy group named **group0** with the following properties:
 - Contains three FEBs (0 through 2).
 - Has a primary FEB (2).
 - Has a unique backup FEB (0).
 - Automatic failover is disabled.

When an active FEB in **group0** fails, automatic failover to the backup FEB does not occur. For **group0**, you can only perform a manual switchover.

- A FEB redundancy group named **group1** with the following properties:
 - Two FEBs (3 and 5). There is no primary FEB.
 - A unique backup FEB (5).
 - Automatic failover is enabled by default.

When **feb 3** in **group1** fails, an automatic failover occurs.

Because you must explicitly configure an FPC *not* to connect to the backup FEB, connectivity is set to none between **fpc 0** and **feb 0** and between **fpc 5** and **feb 5**.



NOTE: For information about the **fpc-feb-connectivity** statement, see the *Junos OS Administration Library*.

FPC to primary FEB connectivity is not explicitly configured, so by default, the software automatically assigns connectivity based on the numerical order of the FPCs.

```
[edit]
chassis {
  fpc-feb-connectivity {
    fpc 0 feb none;
    fpc 5 feb none;
  }
  redundancy feb {
    redundancy-group group0 {
      description "Interfaces to Customer X";
      feb 2 primary;
      feb 1;
      feb 0 backup;
      no-auto-failover;
    }
    redundancy-group group1 {
      feb 3;
      feb 5 backup;
    }
  }
}
```

**Related
Documentation**

- [Understanding Switching Control Board Redundancy on page 17](#)
- [Configuring FEB Redundancy on the M120 Router on page 24](#)

Configuring SFM Redundancy on M40e and M160 Routers

By default, the Switching and Forwarding Module (SFM) in slot 0 is the master and the SFM in slot 1 is the backup. To modify the default configuration, include the **sfm** statement at the **[edit chassis redundancy]** hierarchy level:

```
[edit chassis redundancy]
sfm slot-number (always | preferred);
```


On the M40e router, **slot-number** is 0 or 1. On the M160 router, **slot-number** is 0 through 3.

always defines the SFM as the sole device.

preferred defines a preferred SFM.

To manually switch mastership between SFMs, issue the **request chassis sfm master switch** command. To view SFM status, issue the **show chassis sfm** command. For more information, see the [CLI Explorer](#).

Related Documentation

- [Understanding Switching Control Board Redundancy on page 17](#)

Configuring SSB Redundancy on the M20 Router

For M20 routers with two System and Switch Boards (SSBs), you can configure which SSB is the master and which is the backup. By default, the SSB in slot 0 is the master and the SSB in slot 1 is the backup. To modify the default configuration, include the **ssb** statement at the **[edit chassis redundancy]** hierarchy level:

```
[edit chassis redundancy]
  ssb slot-number (always | preferred);
```

slot-number is 0 or 1.

always defines the SSB as the sole device.

preferred defines a preferred SSB.

To manually switch mastership between SSBs, issue the **request chassis ssb master switch** command.

To display SSB status information, issue the **show chassis ssb** command. The command output displays the number of times the mastership has changed, the SSB slot number, and the current state of the SSB: master, backup, or empty. For more information, see the [CLI Explorer](#).

Related Documentation

- [Understanding Switching Control Board Redundancy on page 17](#)

Configuring the Junos OS to Support Redundancy on Routers Having Multiple Routing Engines or Switching Boards

For routers that have multiple Routing Engines or these multiple switching control boards: Switching and Forwarding Modules (SFMs), System and Switch Boards (SSBs), Forwarding Engine Boards (FEBs), or Compact Forwarding Engine Boards (CFEBs), you can configure redundancy properties.

To configure redundancy, include the following redundancy statements at the **[edit chassis]** hierarchy level:

```
redundancy {  
  cfeb slot (always | preferred);  
  failover {  
    on-disk-failure  
    on-loss-of-keepalives;  
  }  
  feb {  
    redundancy-group group-name {  
      feb slot-number (backup | primary);  
      description description;  
      no-auto-failover;  
    }  
  }  
  graceful-switchover;  
  keepalive-time seconds;  
  routing-engine slot-number (master | backup | disabled);  
  sfm slot-number (always | preferred);  
  ssb slot-number (always | preferred);  
}
```

Related Documentation

- [Understanding Routing Engine Redundancy on Juniper Networks Routers on page 111](#)

PART 3

Configuring Bidirectional Forwarding Detection (BFD)

- [Understanding How BFD Detects Network Failures on page 31](#)
- [Configuring BFD on page 53](#)

CHAPTER 4

Understanding How BFD Detects Network Failures

- [Understanding BFD for Static Routes for Faster Network Failure Detection on page 31](#)
- [Understanding BFD for BGP on page 36](#)
- [Understanding BFD for OSPF on page 38](#)
- [Understanding BFD for IS-IS on page 41](#)
- [Understanding BFD for RIP on page 43](#)
- [Understanding Independent Micro BFD Sessions for LAG on page 44](#)
- [Understanding Distributed BFD on page 47](#)
- [Understanding Static Route State When BFD is in Admin Down State on page 52](#)

Understanding BFD for Static Routes for Faster Network Failure Detection

The Bidirectional Forwarding Detection (BFD) protocol is a simple hello mechanism that detects failures in a network. BFD works with a wide variety of network environments and topologies. A pair of routing devices exchanges BFD packets. Hello packets are sent at a specified, regular interval. A neighbor failure is detected when the routing device stops receiving a reply after a specified interval. The BFD failure detection timers have shorter time limits than the static route failure detection mechanisms, so they provide faster detection.

The BFD failure detection timers can be adjusted to be faster or slower. The lower the BFD failure detection timer value, the faster the failure detection and vice versa. For example, the timers can adapt to a higher value if the adjacency fails (that is, the timer detects failures more slowly). Or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (Rx) interval by two if the local BFD instance is the reason for the session flap. The transmission (Tx) interval is increased by two if the remote BFD instance is the reason for the session flap. You can use the **clear bfd adaptation** command to return BFD interval timers to their configured values. The **clear bfd adaptation** command is hitless, meaning that the command does not affect traffic flow on the routing device.

By default, BFD is supported on single-hop static routes.



NOTE: On MX Series devices, multihop BFD is not supported on a static route if the static route is configured with more than one next hop. It is recommended that you avoid using multiple next hops when a multihop BFD is required for a static route.

To enable failure detection, include the **bfd-liveness-detection** statement in the static route configuration.



NOTE: Starting with Junos OS Release 15.1X49-D70 and Junos OS Release 17.3R1, the **bfd-liveness-detection** command includes the **description** field. The **description** is an attribute under the **bfd-liveness-detection** object and it is supported only on SRX Series devices. This field is applicable only for the static routes.

In Junos OS Release 9.1 and later, the BFD protocol is supported for IPv6 static routes. Global unicast and link-local IPv6 addresses are supported for static routes. The BFD protocol is not supported on multicast or anycast IPv6 addresses. For IPv6, the BFD protocol supports only static routes and only in Junos OS Release 9.3 and later. IPv6 for BFD is also supported for the eBGP protocol.



NOTE: Inline BFD is supported on PTX5000 routers with third-generation FPCs starting in Junos OS Release 15.1F3 and 16.1R2. Inline BFD is supported on PTX3000 routers with third-generation FPCs starting in Junos OS Release 15.1F6 and 16.1R2.



NOTE:

There are three types of BFD sessions based on the source from which BFD packets are sent to the neighbors. Different types of BFD sessions and their descriptions are:

Type of BFD session	Description
Non-distributed BFD	BFD sessions running completely on the Routing Engine.
Distributed BFD	BFD sessions running completely on the Packet Forwarding Engine.

Type of BFD session	Description
Inline BFD NOTE: Starting in Junos OS Release 13.3, inline BFD is supported only on static MX Series routers with MPCs/MICs that have configured enhanced-ip . NOTE: Starting in Junos OS Release 16.1R1, the inline BFD sessions are supported on integrated routing and bridging (IRB) interfaces.	BFD sessions running on the FPC hardware.

To configure the BFD protocol for IPv6 static routes, include the **bfd-liveness-detection** statement at the **[edit routing-options rib inet6.0 static route *destination-prefix*]** hierarchy level.

In Junos OS Release 8.5 and later, you can configure a hold-down interval to specify how long the BFD session must remain up before a state change notification is sent.

To specify the hold-down interval, include the **holddown-interval** statement in the BFD configuration.

You can configure a number in the range from 0 through 255,000 milliseconds. The default is 0. If the BFD session goes down and then comes back up during the hold-down interval, the timer is restarted.



NOTE: If a single BFD session includes multiple static routes, the hold-down interval with the highest value is used.

To specify the minimum transmit and receive intervals for failure detection, include the **minimum-interval** statement in the BFD configuration.

This value represents both the minimum interval after which the local routing device transmits hello packets and the minimum interval after which the routing device expects to receive a reply from the neighbor with which it has established a BFD session. You can configure a number in the range from 1 through 255,000 milliseconds. Optionally, instead of using this statement, you can configure the minimum transmit and receive intervals separately using the **transmit-interval**, **minimum-interval**, and **minimum-receive-interval** statements.



NOTE: BFD is an intensive protocol that consumes system resources. Specifying a minimum interval for BFD of less than 100 ms for Routing Engine-based sessions and 10 ms for distributed BFD sessions can cause undesired BFD flapping.

Depending on your network environment, these additional recommendations might apply:

- For large-scale network deployments with a large number of BFD sessions, specify a minimum interval of 300 ms for Routing Engine-based sessions and 100 ms for distributed BFD sessions.
- For very large-scale network deployments with a large number of BFD sessions, contact Juniper Networks customer support for more information.
- For BFD sessions to remain up during a Routing Engine switchover event when nonstop active routing (NSR) is configured, specify a minimum interval of 2500 ms for Routing Engine-based sessions. For distributed BFD sessions with NSR configured, the minimum interval recommendations are unchanged and depend only on your network deployment.

To specify the minimum receive interval for failure detection, include the **minimum-receive-interval** statement in the BFD configuration. This value represents the minimum interval after which the routing device expects to receive a reply from a neighbor with which it has established a BFD session. You can configure a number in the range from 1 through 255,000 milliseconds. Optionally, instead of using this statement, you can configure the minimum receive interval using the **minimum-interval** statement at the **[edit routing-options static route destination-prefix bfd-liveness-detection]** hierarchy level.

To specify the number of hello packets not received by the neighbor that causes the originating interface to be declared down, include the **multiplier** statement in the BFD configuration.

The default value is 3. You can configure a number in the range from 1 through 255.

To specify a threshold for detecting the adaptation of the detection time, include the **threshold** statement in the BFD configuration.

When the BFD session detection time adapts to a value equal to or higher than the threshold, a single trap and a system log message are sent. The detection time is based on the multiplier of the **minimum-interval** or the **minimum-receive-interval** value. The threshold must be a higher value than the multiplier for either of these configured values. For example if the **minimum-receive-interval** is 300 ms and the **multiplier** is 3, the total detection time is 900 ms. Therefore, the detection time threshold must have a value higher than 900.

To specify the minimum transmit interval for failure detection, include the **transmit-interval** **minimum-interval** statement in the BFD configuration.

This value represents the minimum interval after which the local routing device transmits hello packets to the neighbor with which it has established a BFD session. You can configure a value in the range from 1 through 255,000 milliseconds. Optionally, instead of using this statement, you can configure the minimum transmit interval using the **minimum-interval** statement at the **[edit routing-options static route destination-prefix bfd-liveness-detection]** hierarchy level.

To specify the threshold for the adaptation of the transmit interval, include the **transmit-interval threshold** statement in the BFD configuration.

The threshold value must be greater than the transmit interval. When the BFD session transmit time adapts to a value greater than the threshold, a single trap and a system log message are sent. The detection time is based on the multiplier of the value for the **minimum-interval** or the **minimum-receive-interval** statement at the **[edit routing-options static route destination-prefix bfd-liveness-detection]** hierarchy level. The threshold must be a higher value than the multiplier for either of these configured values.

To specify the BFD version, include the **version** statement in the BFD configuration. The default is to have the version detected automatically.

To include an IP address for the next hop of the BFD session, include the **neighbor** statement in the BFD configuration.



NOTE: You must configure the **neighbor** statement if the next hop specified is an interface name. If you specify an IP address as the next hop, that address is used as the neighbor address for the BFD session.

In Junos OS Release 9.0 and later, you can configure BFD sessions not to adapt to changing network conditions.

To disable BFD adaptation, include the **no-adaptation** statement in the BFD configuration.



NOTE: We recommend that you not disable BFD adaptation unless it is preferable *not* to have BFD adaptation in your network.



NOTE: If BFD is configured only on one end of a static route, the route is removed from the routing table. BFD establishes a session when BFD is configured on both ends of the static route.

BFD is not supported on ISO address families in static routes. BFD does support IS-IS.

If you configure graceful Routing Engine switchover (GRES) at the same time as BFD, GRES does not preserve the BFD state information during a failover.

Release History Table

Release	Description
16.1R1	Starting in Junos OS Release 16.1R1, the inline BFD sessions are supported on integrated routing and bridging (IRB) interfaces.
15.1X49-D70	Starting with Junos OS Release 15.1X49-D70 and Junos OS Release 17.3R1, the bfd-liveness-detection command includes the description field. The description is an attribute under the bfd-liveness-detection object and it is supported only on SRX Series devices. This field is applicable only for the static routes.
15.1F6	Inline BFD is supported on PTX3000 routers with third-generation FPCs starting in Junos OS Release 15.1F6 and 16.1R2.
15.1F3	Inline BFD is supported on PTX5000 routers with third-generation FPCs starting in Junos OS Release 15.1F3 and 16.1R2.
13.3	Starting in Junos OS Release 13.3, inline BFD is supported only on static MX Series routers with MPCs/MICs that have configured enhanced-ip .

Related Documentation

- [Enabling Dedicated and Real-Time BFD on page 106](#)
- [Example: Configuring BFD for Static Routes for Faster Network Failure Detection on page 53](#)
- [Example: Enabling BFD on Qualified Next Hops in Static Routes for Route Selection](#)

Understanding BFD for BGP

The Bidirectional Forwarding Detection (BFD) protocol is a simple hello mechanism that detects failures in a network. Hello packets are sent at a specified, regular interval. A neighbor failure is detected when the routing device stops receiving a reply after a specified interval. BFD works with a wide variety of network environments and topologies. The failure detection timers for BFD have shorter time limits than default failure detection mechanisms for BGP, so they provide faster detection.



NOTE: Configuring both BFD and graceful restart for BGP on the same device is counterproductive. When an interface goes down, BFD detects this instantly, stops traffic forwarding and the BGP session goes down whereas graceful restart forwards traffic despite the interface failure, this behavior might cause network issues. Hence we do not recommend configuring both BFD and graceful restart on the same device.



NOTE: QFX5000 Series switches and EX4600 switches do not support minimum interval values of less than 1 second.

The BFD failure detection timers can be adjusted to be faster or slower. The lower the BFD failure detection timer value, the faster the failure detection and vice versa. For example, the timers can adapt to a higher value if the adjacency fails (that is, the timer detects failures more slowly). Or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds (15000 milliseconds). A back-off algorithm increases the receive (Rx) interval by two if the local BFD instance is the reason for the session flap. The transmission (Tx) interval is increased by two if the remote BFD instance is the reason for the session flap. You can use the **clear bfd adaptation** command to return BFD interval timers to their configured values. The **clear bfd adaptation** command is hitless, meaning that the command does not affect traffic flow on the routing device.



NOTE: On all SRX Series devices, high CPU utilization triggered for reasons such as CPU intensive commands and SNMP walks causes the BFD protocol to flap while processing large BGP updates. (Platform support depends on the Junos OS release in your installation.)

Starting with Junos OS Release 15.1X49-D100, SRX340, SRX345, and SRX1500 devices support dedicated BFD.

Starting with Junos OS Release 15.1X49-D100, SRX300 and SRX320 devices support real-time BFD.

Starting with Junos OS Release 15.1X49-D110, SRX550M devices support dedicated BFD.

In Junos OS Release 8.3 and later, BFD is supported on internal BGP (IBGP) and multihop external BGP (EBGP) sessions as well as on single-hop EBGP sessions. In Junos OS Release 9.1 through Junos OS Release 11.1, BFD supports IPv6 interfaces in static routes only. In Junos OS Release 11.2 and later, BFD supports IPv6 interfaces with BGP.

Release History Table

Release	Description
15.1X49-D100	Starting with Junos OS Release 15.1X49-D100, SRX340, SRX345, and SRX1500 devices support dedicated BFD.
15.1X49-D100	Starting with Junos OS Release 15.1X49-D100, SRX300 and SRX320 devices support real-time BFD.
11.2	In Junos OS Release 11.2 and later, BFD supports IPv6 interfaces with BGP.
9.1	In Junos OS Release 9.1 through Junos OS Release 11.1, BFD supports IPv6 interfaces in static routes only.
8.3	In Junos OS Release 8.3 and later, BFD is supported on internal BGP (IBGP) and multihop external BGP (EBGP) sessions as well as on single-hop EBGP sessions.

- Related Documentation**
- [Enabling Dedicated and Real-Time BFD on page 106](#)

Understanding BFD for OSPF

The Bidirectional Forwarding Detection (BFD) protocol is a simple hello mechanism that detects failures in a network. BFD works with a wide variety of network environments and topologies. A pair of routing devices exchange BFD packets. Hello packets are sent at a specified, regular interval. A neighbor failure is detected when the routing device stops receiving a reply after a specified interval. The BFD failure detection timers have shorter time limits than the OSPF failure detection mechanisms, so they provide faster detection.

The BFD failure detection timers are adaptive and can be adjusted to be faster or slower. The lower the BFD failure detection timer value, the faster the failure detection and vice versa. For example, the timers can adapt to a higher value if the adjacency fails (that is, the timer detects failures more slowly). Or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (Rx) interval by two if the local BFD instance is the reason for the session flap. The transmission (Tx) interval is increased by two if the remote BFD instance is the reason for the session flap. You can use the **clear bfd adaptation** command to return BFD interval timers to their configured values. The **clear bfd adaptation** command is hitless, meaning that the command does not affect traffic flow on the routing device.



NOTE: QFX5000 Series switches and EX4600 switches do not support minimum interval values of less than 1 second.



NOTE: BFD is supported for OSPFv3 in Junos OS Release 9.3 and later.

You can configure the following BFD protocol settings:

- **detection-time threshold**—Threshold for the adaptation of the detection time. When the BFD session detection time adapts to a value equal to or greater than the configured threshold, a single trap and a single system log message are sent.
- **full-neighbors-only**—Ability to establish BFD sessions only for OSPF neighbors with full neighbor adjacency. The default behavior is to establish BFD sessions for all OSPF neighbors. This setting is available in Junos OS Release 9.5 and later.
- **minimum-interval**—Minimum transmit and receive interval for failure detection. This setting configures both the minimum interval after which the local routing device transmits hello packets and the minimum interval after which the routing device expects to receive a reply from the neighbor with which it has established a BFD session. Both intervals are in milliseconds. You can also specify the minimum transmit and receive intervals separately using the **transmit-interval**, **minimum-interval**, and **minimum-receive-interval** statements.



NOTE: BFD is an intensive protocol that consumes system resources. Specifying a minimum interval for BFD of less than 100 ms for Routing Engine-based sessions and 10 ms for distributed BFD sessions can cause undesired BFD flapping.

Depending on your network environment, these additional recommendations might apply:

- For large-scale network deployments with a large number of BFD sessions, specify a minimum interval of no less than 500 ms. An interval of 1000 ms is recommended to avoid any instability issues.
- For very large-scale network deployments with a large number of BFD sessions, contact Juniper Networks customer support for more information.
- For BFD sessions to remain up during a Routing Engine switchover event when nonstop active routing (NSR) is configured, specify a minimum interval of 2500 ms for Routing Engine-based sessions. Without NSR, Routing Engine-based sessions can have a minimum interval of 100 ms. In OSPFv3, BFD is always based in the Routing Engine, meaning that BFD is not distributed. For distributed BFD sessions with NSR configured, the minimum interval recommendations are unchanged and depend only on your network deployment.
- On a single QFX5100 switch, when you add a QFX-EM-4Q expansion module, specify a minimum interval higher than 1000 ms.
- **minimum-receive-interval**—Minimum receive interval for failure detection. This setting configures the minimum receive interval, in milliseconds, after which the routing device

expects to receive a hello packet from a neighbor with which it has established a BFD session. You can also specify the minimum receive interval using the **minimum-interval** statement.

- **multiplier**—Multiplier for hello packets. This setting configures the number of hello packets that are not received by a neighbor, which causes the originating interface to be declared down. By default, three missed hello packets cause the originating interface to be declared down.
- **no-adaptation**—Disables BFD adaptation. This setting disables BFD sessions from adapting to changing network conditions. This setting is available in Junos OS Release 9.0 and later.



NOTE: We recommend that you do not disable BFD adaptation unless it is preferable not to have BFD adaptation in your network.

- **transmit-interval minimum-interval**—Minimum transmit interval for failure detection. This setting configures the minimum transmit interval, in milliseconds, at which the local routing device transmits hello packets to the neighbor with which it has established a BFD session. You can also specify the minimum transmit interval using the **minimum-interval** statement.
- **transmit-interval threshold**—Threshold for the adaptation of the BFD session transmit interval. When the transmit interval adapts to a value greater than the threshold, a single trap and a single system log message are sent. The threshold value must be greater than the minimum transmit interval. If you attempt to commit a configuration with a threshold value less than the minimum transmit interval, the routing device displays an error and does not accept the configuration.
- **version**—BFD version. This setting configures the BFD version used for detection. You can explicitly configure BFD version 1, or the routing device can automatically detect the BFD version. By default, the routing device automatically detects the BFD version automatically, which is either 0 or 1.

You can also trace BFD operations for troubleshooting purposes.

**Related
Documentation**

- [Example: Configuring BFD for OSPF on page 70](#)
- *bfd-liveness-detection*

Understanding BFD for IS-IS

The Bidirectional Forwarding Detection (BFD) protocol is a simple hello mechanism that detects failures in a network. Hello packets are sent at a specified, regular interval. A neighbor failure is detected when the routing device stops receiving a reply after a specified interval. BFD works with a wide variety of network environments and topologies. The failure detection timers for BFD have shorter time limits than the failure detection mechanisms of IS-IS, providing faster detection.

The BFD failure detection timers are adaptive and can be adjusted to be faster or slower. For example, the timers can adapt to a higher value if the adjacency fails, or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (RX) interval by two if the local BFD instance is the reason for the session flap. The transmission (TX) interval is increased by two if the remote BFD instance is the reason for the session flap.

You can use the **clear bfd adaptation** command to return BFD interval timers to their configured values. The **clear bfd adaptation** command is hitless, meaning that the command does not affect traffic flow on the routing device.



NOTE: Starting with Junos OS Release 16.1R1, you can configure IS-IS BFD sessions for IPv6 by including the `bfd-liveness-detection` statement at the `[edit protocols isis interface interface-name family inet|inet6]` hierarchy level.

- For interfaces that support both IPv4 and IPv6 routing, the `bfd-liveness-detection` statement must be configured separately for each inet family.
- BFD over IPv6 link local address is currently not distributed because IS-IS uses link local addresses for forming adjacencies.
- BFD sessions over IPv6 must not have the same aggressive detection intervals as IPv4 sessions.
- BFD IPv6 sessions with detection intervals less than 2.5 seconds are currently not supported when nonstop active routing (NSR) is enabled.



NOTE: QFX5000 Series switches and EX4600 switches do not support minimum interval values of less than 1 second.

To detect failures in the network, the set of statements in [Table 3 on page 42](#) are used in the configuration.

Table 3: Configuring BFD for IS-IS

Statement	Description
bfd-liveness-detection	Enable failure detection.
minimum-interval milliseconds	<p>Specify the minimum transmit and receive intervals for failure detection.</p> <p>This value represents the minimum interval at which the local router transmits hellos packets as well as the minimum interval at which the router expects to receive a reply from a neighbor with which it has established a BFD session. You can configure a number from 1 through 255,000 milliseconds. You can also specify the minimum transmit and receive intervals separately.</p> <p>NOTE: BFD is an intensive protocol that consumes system resources. Specifying a minimum interval for BFD less than 100 ms for Routing Engine-based sessions and 10 ms for distributed BFD sessions can cause undesired BFD flapping.</p> <p>Depending on your network environment, these additional recommendations might apply:</p> <ul style="list-style-type: none"> For large-scale network deployments with a large number of BFD sessions, specify a minimum interval of 300 ms for Routing Engine-based sessions and 100 ms for distributed BFD sessions. For very large-scale network deployments with a large number of BFD sessions, please contact Juniper Networks customer support for more information. For BFD sessions to remain up during a Routing Engine switchover event when nonstop active routing (NSR) is configured, specify a minimum interval of 2500 ms for Routing Engine-based sessions. For distributed BFD sessions with nonstop active routing configured, the minimum interval recommendations are unchanged and depend only on your network deployment.
minimum-receive-interval milliseconds	<p>Specify only the minimum receive interval for failure detection.</p> <p>This value represents the minimum interval at which the local router expects to receive a reply from a neighbor with which it has established a BFD session. You can configure a number from 1 through 255,000 milliseconds.</p>
multiplier number	<p>Specify the number of hello packets not received by the neighbor that causes the originating interface to be declared down.</p> <p>The default is 3, and you can configure a value from 1 through 225.</p>
no-adaptation	<p>Disable BFD adaptation.</p> <p>In Junos OS Release 9.0 and later, you can specify that the BFD sessions not adapt to changing network conditions.</p> <p>NOTE: We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.</p>
threshold	<p>Specify the threshold for the following:</p> <ul style="list-style-type: none"> Adaptation of the detection time <p>When the BFD session detection time adapts to a value equal to or greater than the threshold, a single trap and a system log message are sent.</p> Transmit interval <p>NOTE: The threshold value must be greater than the minimum transmit interval multiplied by the multiplier number.</p>

Table 3: Configuring BFD for IS-IS (continued)

Statement	Description
<code>transmit-interval</code> <code>minimum-interval</code>	Specify the minimum transmit interval for failure detection. This value represents the minimum interval at which the local routing device transmits hello packets to the neighbor with which it has established a BFD session. You can configure a value from 1 through 255,000 milliseconds.
<code>version</code>	Specify the BFD version used for detection. The default is to have the version detected automatically.



NOTE: You can trace BFD operations by including the `traceoptions` statement at the `[edit protocols bfd]` hierarchy level.

For a list of hierarchy levels at which you can include these statements, see the statement summary sections for these statements.

**Related
Documentation**

- [Example: Configuring BFD for IS-IS on page 74](#)
- [Understanding BFD Authentication for IS-IS](#)

Understanding BFD for RIP

The Bidirectional Forwarding Detection (BFD) Protocol is a simple hello mechanism that detects failures in a network. Hello packets are sent at a specified, regular interval. A neighbor failure is detected when the routing device stops receiving a reply after a specified interval. BFD works with a wide variety of network environments and topologies. BFD failure detection times are shorter than RIP detection times, providing faster reaction times to various kinds of failures in the network. Instead of waiting for the routing protocol neighbor timeout, BFD provides rapid detection of link failures. BFD timers are adaptive and can be adjusted to be more or less aggressive. For example, a timer can adapt to a higher value if the adjacency fails, or a neighbor can negotiate a higher value for a timer than the one configured. Note that the functionality of configuring BFD for RIP described in this topic is not supported in Junos OS Releases 15.1X49, 15.1X49-D30, or 15.1X49-D40.

BFD enables quick failover between a primary and a secondary routed path. The protocol tests the operational status of the interface multiple times per second. BFD provides for configuration timers and thresholds for failure detection. For example, if the minimum interval is set for 50 milliseconds and the threshold uses the default value of three missed messages, a failure is detected on an interface within 200 milliseconds of the failure.

Intervening devices (for example, an Ethernet LAN switch) hide link-layer failures from routing protocol peers, such as when two routers are connected by way of a LAN switch, where the local interface status remains up even when a physical fault happens on the remote link. Link-layer failure detection times vary, depending on the physical media and

the Layer 2 encapsulation. BFD can provide fast failure detection times for all media types, encapsulations, topologies, and routing protocols.

To enable BFD for RIP, both sides of the connection must receive an update message from the peer. By default, RIP does not export any routes. Therefore, you must enable update messages to be sent by configuring an export policy for routes before a BFD session is triggered.

Release History Table

Release	Description
15.1X49	Note that the functionality of configuring BFD for RIP described in this topic is not supported in Junos OS Releases 15.1X49, 15.1X49-D30, or 15.1X49-D40.

Understanding Independent Micro BFD Sessions for LAG

Starting with Junos OS Release 13.3, this feature is supported on the following PIC/FPC types:

- PC-1XGE-XENPAK (Type 3 FPC)
- PD-4XGE-XFP (Type 4 FPC)
- PD-5-10XGE-SFP (Type 4 FPC)
- 24x10GE (LAN/WAN) SFP, 12x10GE (LAN/WAN) SFP, 1x100GE Type 5 PICs
- All MPCs on MX Series with Ethernet MICs
- FPC-PTX-P1-A on PTX5000 with 10-Gigabit Ethernet interfaces
- FPC2-PTX-P1A on PTX5000 with 10-Gigabit Ethernet interfaces in Junos OS Release 14.1 and later
- All FPCs on PTX Series with Ethernet interfaces in Junos OS Release 14.1R3 and later 14.1 releases, and Junos 14.2 and later



TIP: See *PTX Series PIC/FPC Compatibility* for a list of PICs that are supported on each PTX Series FPC.



NOTE: Micro-BFD configuration with interface addresses is not supported on PTX routers on FPC3 and QFX10000 line of switches.

The Bidirectional Forwarding Detection (BFD) protocol is a simple detection protocol that quickly detects failures in the forwarding paths. A link aggregation group (LAG) combines multiple links between devices that are in point-to-point connections, thereby increasing bandwidth, providing reliability, and allowing load balancing. To run a BFD session on LAG interfaces, configure an independent, asynchronous mode BFD session on every LAG member link in a LAG bundle. Instead of a single BFD session monitoring

the status of the UDP port, independent micro BFD sessions monitor the status of individual member links.

The individual BFD sessions determine the Layer 2 and Layer 3 connectivity of each member link in the LAG. Once a BFD session is established on a particular link, the member links are attached to the LAG and the load balancer either by a static configuration or by the Link Aggregation Control Protocol (LACP). If the member links are attached to the LAG by a static configuration, the device control process acts as the client to the micro BFD session. When member links are attached to the LAG by the LACP, the LACP acts as the client to the micro BFD session.

When the micro BFD session is up, a LAG link is established and data is transmitted over that LAG link. If the micro BFD session on a member link is down, that particular member link is removed from the load balancer, and the LAG managers stop directing traffic to that link. These micro BFD sessions are independent of each other despite having a single client that manages the LAG interface.



NOTE:

- Starting with Junos OS Release 13.3, IANA has allocated 01-00-5E-90-00-01 as the dedicated MAC address for micro BFD. Dedicated MAC mode is used by default for micro BFD sessions, in accordance with the latest draft for BFD over LAG.
- In Junos OS, MicroBFD control packets are always untagged by default. For L2 aggregated interfaces, the configuration must include vlan-tagging or flexible-vlan-tagging in the Aggregated Ethernet with BFD. Otherwise, the system will throw error while committing the configuration.
- When you enable MicroBFD on an aggregated Ethernet Interface, the aggregated Interface can receive MicroBFD packets. Starting with Junos OS Release 19.3 and later, for MPC10E and MPC11E MPCs, you cannot apply firewall filters on the MicroBFD packets received on the aggregated Ethernet Interface. For MPC1E through MPC9E, you can apply firewall filters on the MicroBFD packets received on the aggregated Ethernet Interface only if the aggregated Ethernet Interface is configured as an untagged Interface.

Micro BFD sessions run in the following modes:

- Distribution Mode—Micro BFD sessions are distributed by default at Layer 3.
- Non-Distribution Mode—You can configure the BFD session to run in this mode by including the **no-delegate-processing** statement under periodic packet management (PPM). In this mode, the packets are being sent or received by the Routing Engine at Layer 2.

A pair of routing devices in a LAG exchange BFD packets at a specified, regular interval. The routing device detects a neighbor failure when it stops receiving a reply after a specified interval. This allows the quick verification of member link connectivity with or without LACP. A UDP port distinguishes BFD over LAG packets from BFD over single-hop IP.



NOTE: IANA has allocated 6784 as the UDP destination port for micro BFD.

To enable failure detection for LAG networks for aggregated Ethernet interfaces:

- Include the **bfd-liveness-detection** statement in the configuration.
- Specify a hold-down interval value to set the minimum time that the BFD session must remain up before a state change notification is sent to the other members in the LAG network.
- Specify the minimum interval that indicates the time interval for transmitting and receiving data.
- Starting with Junos OS Release 14.1, specify the neighbor in a BFD session. In releases prior to Junos OS Release 16.1, you must configure the loopback address of the remote destination as the neighbor address. Beginning with Junos OS Release 16.1, you can also configure this feature on MX series routers with aggregated Ethernet interface address of the remote destination as the neighbor address.



NOTE: On T1600 and T4000 routers, you cannot configure the local aggregated Ethernet interface address of the remote destination as the neighbor address.



CAUTION: Deactivate **bfd-liveness-detection** at the [edit interfaces aex aggregated-ether-options] hierarchy level or deactivate the aggregated Ethernet interface before changing the neighbor address from loopback IP address to aggregated Ethernet interface IP address. Modifying the local and neighbor address without deactivating **bfd-liveness-detection** or the aggregated Ethernet interface first might cause micro BFD sessions failure.



NOTE: Beginning with Release 16.1R2, Junos OS checks and validates the configured micro BFD local-address against the interface or loopback IP address before the configuration commit. Junos OS performs this check on both IPv4 and IPv6 micro BFD address configurations, and if they do not match, the commit fails.



NOTE: This feature works only when both the devices support BFD. If BFD is configured at one end of the LAG, this feature does not work.

For the IPv6 address family, disable duplicate address detection before configuring this feature with AE interface addresses. To disable duplicate address detection, include the **dad-disable** statement at the [edit interface aex unit y family inet6] hierarchy level.

Release History Table

Release	Description
19.3	Starting with Junos OS Release 19.3 and later, for MPC10E and MPC11E MPCs, you cannot apply firewall filters on the MicroBFD packets received on the aggregated Ethernet Interface. For MPC1E through MPC9E, you can apply firewall filters on the MicroBFD packets received on the aggregated Ethernet Interface only if the aggregated Ethernet Interface is configured as an untagged Interface.
16.1	Beginning with Junos OS Release 16.1, you can also configure this feature on MX series routers with aggregated Ethernet interface address of the remote destination as the neighbor address.
16.1	Beginning with Release 16.1R2, Junos OS checks and validates the configured micro BFD local-address against the interface or loopback IP address before the configuration commit.
14.1	Starting with Junos OS Release 14.1, specify the neighbor in a BFD session. In releases prior to Junos OS Release 16.1, you must configure the loopback address of the remote destination as the neighbor address.
13.3	Starting with Junos OS Release 13.3, IANA has allocated 01-00-5E-90-00-01 as the dedicated MAC address for micro BFD.

Related Documentation

- [authentication on page 506](#)
- [bfd-liveness-detection on page 507](#)
- [detection-time on page 509](#)
- [transmit-interval on page 512](#)

Understanding Distributed BFD

Bidirectional Forwarding Detection (BFD) is a protocol to verify the liveness of data path.

The terms *nondistributed BFD* and *centralized BFD* refer to BFD that runs on the Routing Engine. The term *distributed BFD* refers to BFD that runs on the Packet Forwarding Engine.



NOTE: By default, SRX Series devices operate in centralized BFD mode.

- **Single-hop BFD**—Single-hop BFD in Junos OS runs in distributed mode by default. The exceptions are OSPFv3 BFD and PIMv6 BFD, for which only nondistributed BFD is supported. Single-hop BFD control packets use UDP port 3784.
- **Multihop BFD**—One desirable application of BFD is to detect connectivity to routing devices that span multiple network hops and follow unpredictable paths. This is known as a multihop session. Prior to Junos OS Release 12.3, multihop BFD is nondistributed and runs on the Routing Engine. Starting in Junos OS Release 12.3, multihop BFD runs in distributed mode by default. Multihop BFD control packets use UDP port 4784.



NOTE: In a multichassis link aggregation group setup, Inter-Chassis Control Protocol (ICCP) uses BFD in multihop mode. Multihop BFD runs in centralized mode in this kind of setup prior to Junos OS Release 12.3 and continues to do so as of Junos OS Release 12.3 and later.

For both single-hop BFD and multihop BFD, the BFD session can be made to run on the Routing Engine (in nondistributed mode) by configuring **set routing-options ppm no-delegate-processing** and then running the **clear bfd session** command.

The benefits of distributed BFD are mainly in the scaling and performance areas.

The benefits are as follows:

- Allows for the creation of a larger number of BFD sessions.
- Runs BFD sessions with a shorter transfer/receive timer interval, which can in turn be used to bring down the overall detection time.
- Separates the functionality of BFD from that of the Routing Engine. This means that a BFD session can stay up during graceful restart, even with an aggressive interval. The minimum interval for Routing Engine-based BFD sessions to survive graceful Routing Engine switchover is 2500 ms. This is improved to sub-second times with distribution.
- Offloads the processing to the FPC CPU. This frees up the Routing Engine CPU, resulting in improved scaling and performance for Routing Engine-based applications.
- Starting with Junos OS Release 15.1X49-D100, dedicated BFD is supported on SRX340, SRX345, and SRX1500 devices.

Starting with Junos OS Release 15.1X49-D100, real-time BFD is supported on SRX300 and SRX320 devices.

Starting with Junos OS Release 15.1X49-D110, dedicated BFD is supported on SRX550M devices.

Starting with Junos OS Release 12.3X48-D60, dedicated BFD is supported on SRX240, SRX550, and SRX650 devices.

Starting with Junos OS Release 12.3X48-D60, real-time BFD is supported on SRX100, SRX110, SRX210, and SRX220 devices.

- To enable dedicated BFD on SRX100, SRX110, SRX210, SRX220, SRX240, SRX300, SRX320, SRX340, SRX345, SRX550, SRX550M, SRX650, and SRX1500 devices, use the **set chassis dedicated-ukern-cpu** command.

Enabling dedicated BFD impacts traffic throughput as one CPU core is removed from data plane processing.

- To enable real-time BFD on SRX100, SRX110, SRX210, SRX220, SRX240, SRX300, SRX320, SRX340, SRX345, SRX550, SRX550M, and SRX650 devices, use the **set chassis realtime-ukern-thread** command.

Enabling real-time BFD does not impact data plane performance. Higher priority is given to the pfe process handling BFD in distributed mode. This is suitable for scenarios where the number of BFD sessions are less.

Table 4 on page 49 lists the BFD modes supported on SRX Series devices.

Table 4: BFD Modes Supported on SRX Series Devices

SRX Series Device	Default BFD Mode	Distributed BFD	Real-Time BFD	Dedicated Core
SRX100	Centralized	Configuration	Configuration (Optional)	Not supported
SRX110	Centralized	Configuration	Configuration (Optional)	Not supported
SRX210	Centralized	Configuration	Configuration (Optional)	Not supported
SRX220	Centralized	Configuration	Configuration (Optional)	Not supported
SRX240	Centralized	Configuration	Configuration	Configuration (Optional)
SRX300	Centralized	Configuration	Configuration (Optional)	Not supported
SRX320	Centralized	Configuration	Configuration (Optional)	Not supported
SRX340	Centralized	Configuration	Configuration	Configuration (Optional)
SRX345	Centralized	Configuration	Configuration	Configuration (Optional)
SRX550	Centralized	Configuration	Configuration	Configuration (Optional)
SRX550M	Centralized	Configuration	Configuration	Configuration (Optional)
SRX650	Centralized	Configuration	Configuration	Configuration (Optional)
SRX1500	Centralized	Configuration	Not supported	Configuration (Optional)
SRX4100	Centralized	Not supported	Not supported	Not supported
SRX4200	Centralized	Not supported	Not supported	Not supported
SRX5400	Centralized	Not supported	Not supported	Not supported
SRX5600	Centralized	Not supported	Not supported	Not supported
SRX5800	Centralized	Not supported	Not supported	Not supported

To determine if a BFD peer is running distributed BFD, run the **show bfd sessions extensive** command and look for **Remote is control-plane independent** in the command output.

For distributed BFD to work, you need to configure the lo0 interface with unit 0 and the appropriate family.

```
# set interfaces lo0 unit 0 family inet
# set interfaces lo0 unit 0 family inet6
# set interfaces lo0 unit 0 family mpls
```

This is true for the following types of BFD sessions:

- BFD over ae logical interfaces, both IPv4 and IPv6
- Multihop BFD, both IPv4 and IPv6
- BFD over VLAN interfaces in EX Series switches, both IPv4 and IPv6
- Virtual Circuit Connectivity Verification (VCCV) BFD (Layer 2 circuit, Layer 3 VPN, and VPLS) (MPLS)



NOTE: Starting in Junos OS Release 13.3R5, if you apply a firewall filter on a loopback interface for a multihop BFD session with a delegated anchor FPC, Junos OS does not execute this filter, because there is an implicit filter on all ingress FPCs to forward packets to the anchor FPC. Therefore, the firewall filter on the loopback interface is not applied on these packets. If you do not want these packets to be forwarded to the anchor FPC, you can configure the `no-delegate-processing` option.

For information about troubleshooting BFD, see [Juniper Networks Knowledge Base article 26746](#).



NOTE: Starting in Junos OS Release 13.3, the distribution of adjacency entry (the IP addresses of adjacent routers) and transmit entry (the IP address of transmitting routers) for a BFD session is asymmetric. This is because an adjacency entry that requires rules might or might not be distributed based on the redirect rule, and the distribution of transmit entries is *not* dependent on the redirect rule.

The term *redirect rule* here denotes the capability of an interface to send protocol redirect messages. See *Disabling the Transmission of Redirect Messages on an Interface*.

Release History Table

Release	Description
15.1X49-D100	Starting with Junos OS Release 15.1X49-D100, dedicated BFD is supported on SRX340, SRX345, and SRX1500 devices.
15.1X49-D100	Starting with Junos OS Release 15.1X49-D100, real-time BFD is supported on SRX300 and SRX320 devices.
13.3R5	Starting in Junos OS Release 13.3R5, if you apply a firewall filter on a loopback interface for a multihop BFD session with a delegated anchor FPC, Junos OS does not execute this filter, because there is an implicit filter on all ingress FPCs to forward packets to the anchor FPC.
13.3	Starting in Junos OS Release 13.3, the distribution of adjacency entry (the IP addresses of adjacent routers) and transmit entry (the IP address of transmitting routers) for a BFD session is asymmetric.

**Related
Documentation**

- *show bfd session*
- [Understanding BFD for RIP on page 43](#)
- [Understanding BFD for Static Routes for Faster Network Failure Detection on page 31](#)
- [Understanding BFD for BGP on page 36](#)
- [Understanding BFD for IS-IS on page 41](#)
- [Understanding BFD for OSPF on page 38](#)
- *Understanding EBGP Multihop*

Understanding Static Route State When BFD is in Admin Down State

The Bidirectional Forwarding Detection (BFD) Admin Down state is used to bring down a BFD session administratively (applicable for normal BFD session and micro BFD session), to protect client applications from BFD configuration removal, license issues, and clearing of BFD sessions.

When BFD enters the Admin Down state, BFD notifies the new state to its peer for a failure detection time and after the time expires, the client stops transmitting packets.

For the Admin Down state to work, the peer, which receives the Admin Down state notification, must have the capability to distinguish between administratively down state and real link failure.

A BFD session moves to the Admin Down state under the following conditions:

- If BFD configuration is removed for the last client tied to a BFD session, BFD moves to Admin Down state and communicates the change to the peer, to enable the client protocols without going down.
- If BFD license is removed on the client, BFD moves to Admin Down state and communicates the change to the remote system to enable the client protocols without going down.
- When **clear bfd session** command is executed, the BFD sessions move to Admin Down state before restarting. This **clear bfd session** command also ensures that the client applications are not impacted.

Starting from Junos OS 16.1R1 release, you can set the state of static route in BFD Admin Down state by configuring one of the following commands:

- **set routing-options static static-route bfd-admin-down active**—BFD Admin Down state pulls down the static route.
- **set routing-options static static-route bfd-admin-down passive**—BFD Admin Down state does not pull down the static route.

Related Documentation

- [Understanding BFD for Static Routes for Faster Network Failure Detection on page 31](#)
- [Example: Configuring BFD for Static Routes for Faster Network Failure Detection on page 53](#)

CHAPTER 5

Configuring BFD

- [Example: Configuring BFD for Static Routes for Faster Network Failure Detection on page 53](#)
- [Example: Configuring BFD on Internal BGP Peer Sessions on page 59](#)
- [Example: Configuring BFD for OSPF on page 70](#)
- [Example: Configuring BFD for IS-IS on page 74](#)
- [Example: Configuring BFD for RIP on page 81](#)
- [Configuring Independent Micro BFD Sessions for LAG on page 87](#)
- [Example: Configuring Independent Micro BFD Sessions for LAG on page 93](#)
- [Configuring BFD for PIM on page 103](#)
- [Enabling Dedicated and Real-Time BFD on page 106](#)

Example: Configuring BFD for Static Routes for Faster Network Failure Detection

This example shows how to configure Bidirectional Forwarding Detection (BFD) for static routes.

- [Requirements on page 53](#)
- [Overview on page 53](#)
- [Configuration on page 54](#)
- [Verification on page 58](#)

Requirements

In this example, no special configuration beyond device initialization is required.

Overview

There are many practical applications for static routes. Static routing is often used at the network edge to support attachment to stub networks, which, given their single point of entry and egress, are well suited to the simplicity of a static route. In Junos OS, static routes have a global preference of 5. Static routes are activated if the specified next hop is reachable.


In this example, you configure the static route 192.168.47.0/24 from the provider network to the customer network, using the next-hop address of 172.16.1.2. You also configure a

static default route of 0.0.0.0/0 from the customer network to the provider network, using a next-hop address of 172.16.1.1.

For demonstration purposes, some loopback interfaces are configured on Device B and Device D. These loopback interfaces provide addresses to ping and thus verify that the static routes are working.

Figure 1 on page 54 shows the sample network.

Figure 1: Customer Routes Connected to a Service Provider



Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device B

```
set interfaces ge-1/2/0 unit 0 description B->D
set interfaces ge-1/2/0 unit 0 family inet address 172.16.1.1/24
set interfaces lo0 unit 57 family inet address 10.0.0.1/32
set interfaces lo0 unit 57 family inet address 10.0.0.2/32
set routing-options static route 192.168.47.0/24 next-hop 172.16.1.2
set routing-options static route 192.168.47.0/24 bfd-liveness-detection minimum-interval 1000
set routing-options static route 192.168.47.0/24 bfd-liveness-detection description Site-xxx
set protocols bfd traceoptions file bfd-trace
set protocols bfd traceoptions flag all
```

Device D

```
set interfaces ge-1/2/0 unit 1 description D->B
set interfaces ge-1/2/0 unit 1 family inet address 172.16.1.2/24
```

```

set interfaces lo0 unit 2 family inet address 192.168.47.5/32
set interfaces lo0 unit 2 family inet address 192.168.47.6/32
set routing-options static route 0.0.0.0/0 next-hop 172.16.1.1
set routing-options static route 0.0.0.0/0 bfd-liveness-detection minimum-interval 1000
set protocols bfd traceoptions file bfd-trace
set protocols bfd traceoptions flag all

```

Step-by-Step Procedure The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure BFD for static routes:

1. On Device B, configure the interfaces.

```

[edit interfaces]
user@B# set ge-1/2/0 unit 0 description B->D
user@B# set ge-1/2/0 unit 0 family inet address 172.16.1.1/24
user@B# set lo0 unit 57 family inet address 10.0.0.1/32
user@B# set lo0 unit 57 family inet address 10.0.0.2/32

```

2. On Device B, create a static route and set the next-hop address.

```

[edit routing-options]
user@B# set static route 192.168.47.0/24 next-hop 172.16.1.2

```

3. On Device B, configure BFD for the static route.

```

[edit routing-options]
user@B# set static route 192.168.47.0/24 bfd-liveness-detection minimum-interval
1000
set routing-options static route 192.168.47.0/24 bfd-liveness-detection description
Site-xxx

```

4. On Device B, configure tracing operations for BFD.

```

[edit protocols]
user@B# set bfd traceoptions file bfd-trace
user@B# set bfd traceoptions flag all

```

5. If you are done configuring Device B, commit the configuration.

```

[edit]
user@B# commit

```

6. On Device D, configure the interfaces.

```

[edit interfaces]

```

```
user@D# set ge-1/2/0 unit 1 description D->B
user@D# set ge-1/2/0 unit 1 family inet address 172.16.1.2/24
user@D# set lo0 unit 2 family inet address 192.168.47.5/32
user@D# set lo0 unit 2 family inet address 192.168.47.6/32
```

7. On Device D, create a static route and set the next-hop address.

```
[edit routing-options]
user@D# set static route 0.0.0.0/0 next-hop 172.16.1.1
```

8. On Device D, configure BFD for the static route.

```
[edit routing-options]
user@D# set static route 0.0.0.0/0 bfd-liveness-detection minimum-interval 1000
```

9. On Device D, configure tracing operations for BFD.

```
[edit protocols]
user@D# set bfd traceoptions file bfd-trace
user@D# set bfd traceoptions flag all
```

10. If you are done configuring Device D, commit the configuration.

```
[edit]
user@D# commit
```

Results

Confirm your configuration by issuing the **show interfaces**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
Device B user@B# show interfaces
ge-1/2/0 {
  unit 0 {
    description B->D;
    family inet {
      address 172.16.1.1/24;
    }
  }
}
lo0 {
  unit 57 {
    family inet {
      address 10.0.0.1/32;
      address 10.0.0.2/32;
    }
  }
}
```

```

    }
}

```

```

user@D# show protocols
bfd {
  traceoptions {
    file bfd-trace;
    flag all;
  }
}

```

```

user@B# show routing-options
static {
  route 192.168.47.0/24 {
    next-hop 172.16.1.2;
    bfd-liveness-detection {
      description Site- xxx;
      minimum-interval 1000;
    }
  }
}

```

Device D

```

user@D# show interfaces
ge-1/2/0 {
  unit 1 {
    description D->B;
    family inet {
      address 172.16.1.2/24;
    }
  }
}
lo0 {
  unit 2 {
    family inet {
      address 192.168.47.5/32;
      address 192.168.47.6/32;
    }
  }
}

```

```

user@D# show routing-options
static {
  route 0.0.0.0/0 {
    next-hop 172.16.1.1;
    bfd-liveness-detection {
      description Site - xxx;
      minimum-interval 1000;
    }
  }
}

```

Verification

Confirm that the configuration is working properly.

- [Verifying That BFD Sessions Are Up on page 58](#)
- [Viewing Detailed BFD Events on page 59](#)

Verifying That BFD Sessions Are Up

Purpose Verify that the BFD sessions are up, and view details about the BFD sessions.

Action From operational mode, enter the **show bfd session extensive** command.

```
user@B> show bfd session extensive
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
172.16.1.2	Up	lt-1/2/0.0	3.000	1.000	3

Client Static, description Site-xxx, TX interval 1.000, RX interval 1.000
 Session up time 00:14:30
 Local diagnostic None, remote diagnostic None
 Remote state Up, version 1
 Replicated, routing table index 172
 Min async interval 1.000, min slow interval 1.000
 Adaptive async TX interval 1.000, RX interval 1.000
 Local min TX interval 1.000, minimum RX interval 1.000, multiplier 3
 Remote min TX interval 1.000, min RX interval 1.000, multiplier 3
 Local discriminator 2, remote discriminator 1
 Echo mode disabled/inactive

1 sessions, 1 clients
 Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps



NOTE: The description Site- <xxx> is supported only on the SRX Series devices.

If each client has more than one description field, then it displays "and more" along with the first description field.

```
user@D> show bfd session extensive
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
172.16.1.1	Up	lt-1/2/0.1	3.000	1.000	3

Client Static, TX interval 1.000, RX interval 1.000
 Session up time 00:14:35
 Local diagnostic None, remote diagnostic None
 Remote state Up, version 1
 Replicated, routing table index 170
 Min async interval 1.000, min slow interval 1.000
 Adaptive async TX interval 1.000, RX interval 1.000
 Local min TX interval 1.000, minimum RX interval 1.000, multiplier 3


```
Remote min TX interval 1.000, min RX interval 1.000, multiplier 3
Local discriminator 1, remote discriminator 2
Echo mode disabled/inactive
```

```
1 sessions, 1 clients
Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps
```

Meaning The TX interval 1.000, RX interval 1.000 output represents the setting configured with the **minimum-interval** statement. All of the other output represents the default settings for BFD. To modify the default settings, include the optional statements under the **bfd-liveness-detection** statement.

Viewing Detailed BFD Events

Purpose View the contents of the BFD trace file to assist in troubleshooting, if needed.

Action From operational mode, enter the **file show /var/log/bfd-trace** command.

```
user@B> file show /var/log/bfd-trace
```

```
Nov 23 14:26:55 Data (9) len 35: (hex) 42 46 44 20 70 65 72 69 6f 64 69 63 20
78 6d 69 74 20 72
Nov 23 14:26:55 PPM Trace: BFD periodic xmit rt tbl index 172
Nov 23 14:26:55 Received Downstream TraceMsg (22) len 108:
Nov 23 14:26:55 IfIndex (3) len 4: 0
Nov 23 14:26:55 Protocol (1) len 1: BFD
Nov 23 14:26:55 Data (9) len 83: (hex) 70 70 6d 64 5f 62 66 64 5f 73 65 6e 64
6d 73 67 20 3a 20
Nov 23 14:26:55 PPM Trace: pmd_bfd_sendmsg : socket 12 len 24, ifl 78 src
172.16.1.1 dst 172.16.1.2 errno 65
Nov 23 14:26:55 Received Downstream TraceMsg (22) len 93:
Nov 23 14:26:55 IfIndex (3) len 4: 0
Nov 23 14:26:55 Protocol (1) len 1: BFD
Nov 23 14:26:55 Data (9) len 68: (hex) 42 46 44 20 70 65 72 69 6f 64 69 63 20
78 6d 69 74 20 74
```

Meaning BFD messages are being written to the trace file.

Related Documentation

- [Understanding BFD for Static Routes for Faster Network Failure Detection on page 31](#)

Example: Configuring BFD on Internal BGP Peer Sessions

This example shows how to configure internal BGP (IBGP) peer sessions with the Bidirectional Forwarding Detection (BFD) protocol to detect failures in a network.

- [Requirements on page 60](#)
- [Overview on page 60](#)

- [Configuration on page 62](#)
- [Verification on page 66](#)

Requirements

No special configuration beyond device initialization is required before you configure this example.

Overview

The minimum configuration to enable BFD on IBGP sessions is to include the **bfd-liveness-detection minimum-interval** statement in the BGP configuration of all neighbors participating in the BFD session. The **minimum-interval** statement specifies the minimum transmit and receive intervals for failure detection. Specifically, this value represents the minimum interval after which the local routing device transmits hello packets as well as the minimum interval that the routing device expects to receive a reply from a neighbor with which it has established a BFD session. You can configure a value from 1 through 255,000 milliseconds.

Optionally, you can specify the minimum transmit and receive intervals separately using the **transmit-interval minimum-interval** and **minimum-receive-interval** statements. For information about these and other optional BFD configuration statements, see **bfd-liveness-detection**.



NOTE: BFD is an intensive protocol that consumes system resources. Specifying a minimum interval for BFD less than 100 milliseconds for Routing Engine-based sessions and less than 10 milliseconds for distributed BFD sessions can cause undesired BFD flapping.


Depending on your network environment, these additional recommendations might apply:

- To prevent BFD flapping during the general Routing Engine switchover event, specify a minimum interval of 5000 milliseconds for Routing Engine-based sessions. This minimum value is required because, during the general Routing Engine switchover event, processes such as RPD, MIBD, and SNMPD utilize CPU resources for more than the specified threshold value. Hence, BFD processing and scheduling is affected because of this lack of CPU resources.
- For BFD sessions to remain up during the dual chassis cluster control link scenario, when the first control link fails, specify the minimum interval of 6000 milliseconds to prevent the LACP from flapping on the secondary node for Routing Engine-based sessions.
- For large-scale network deployments with a large number of BFD sessions, specify a minimum interval of 300 milliseconds for Routing Engine-based sessions and 100 milliseconds for distributed BFD sessions.
- For very large-scale network deployments with a large number of BFD sessions, contact Juniper Networks customer support for more information.
- For BFD sessions to remain up during a Routing Engine switchover event when nonstop active routing (NSR) is configured, specify a minimum interval of 2500 milliseconds for Routing Engine-based sessions. For distributed BFD sessions with NSR configured, the minimum interval recommendations are unchanged and depend only on your network deployment.

BFD is supported on the default routing instance (the main router), routing instances, and logical systems. This example shows BFD on logical systems.

Figure 2 on page 62 shows a typical network with internal peer sessions.

Figure 2: Typical Network with IBGP Sessions



Configuration

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device A

```

set logical-systems A interfaces lt-1/2/0 unit 1 description to-B
set logical-systems A interfaces lt-1/2/0 unit 1 encapsulation ethernet
set logical-systems A interfaces lt-1/2/0 unit 1 peer-unit 2
set logical-systems A interfaces lt-1/2/0 unit 1 family inet address 10.10.10.1/30
set logical-systems A interfaces lo0 unit 1 family inet address 192.168.6.5/32
set logical-systems A protocols bgp group internal-peers type internal
set logical-systems A protocols bgp group internal-peers traceoptions file bgp-bfd
set logical-systems A protocols bgp group internal-peers traceoptions flag bfd detail
set logical-systems A protocols bgp group internal-peers local-address 192.168.6.5
set logical-systems A protocols bgp group internal-peers export send-direct
set logical-systems A protocols bgp group internal-peers bfd-liveness-detection
  minimum-interval 1000
set logical-systems A protocols bgp group internal-peers neighbor 192.163.6.4
set logical-systems A protocols bgp group internal-peers neighbor 192.168.40.4
set logical-systems A protocols ospf area 0.0.0.0 interface lo0.1 passive
set logical-systems A protocols ospf area 0.0.0.0 interface lt-1/2/0.1
set logical-systems A policy-options policy-statement send-direct term 2 from protocol
  direct
set logical-systems A policy-options policy-statement send-direct term 2 then accept
set logical-systems A routing-options router-id 192.168.6.5
set logical-systems A routing-options autonomous-system 17

```

Device B

```

set logical-systems B interfaces lt-1/2/0 unit 2 description to-A
set logical-systems B interfaces lt-1/2/0 unit 2 encapsulation ethernet
set logical-systems B interfaces lt-1/2/0 unit 2 peer-unit 1
set logical-systems B interfaces lt-1/2/0 unit 2 family inet address 10.10.10.2/30
set logical-systems B interfaces lt-1/2/0 unit 5 description to-C
set logical-systems B interfaces lt-1/2/0 unit 5 encapsulation ethernet

```

```

set logical-systems B interfaces lt-1/2/0 unit 5 peer-unit 6
set logical-systems B interfaces lt-1/2/0 unit 5 family inet address 10.10.10.5/30
set logical-systems B interfaces lo0 unit 2 family inet address 192.163.6.4/32
set logical-systems B protocols bgp group internal-peers type internal
set logical-systems B protocols bgp group internal-peers local-address 192.163.6.4
set logical-systems B protocols bgp group internal-peers export send-direct
set logical-systems B protocols bgp group internal-peers bfd-liveness-detection
    minimum-interval 1000
set logical-systems B protocols bgp group internal-peers neighbor 192.168.40.4
set logical-systems B protocols bgp group internal-peers neighbor 192.168.6.5
set logical-systems B protocols ospf area 0.0.0.0 interface lo0.2 passive
set logical-systems B protocols ospf area 0.0.0.0 interface lt-1/2/0.2
set logical-systems B protocols ospf area 0.0.0.0 interface lt-1/2/0.5
set logical-systems B policy-options policy-statement send-direct term 2 from protocol
    direct
set logical-systems B policy-options policy-statement send-direct term 2 then accept
set logical-systems B routing-options router-id 192.163.6.4
set logical-systems B routing-options autonomous-system 17

```

Device C

```

set logical-systems C interfaces lt-1/2/0 unit 6 description to-B
set logical-systems C interfaces lt-1/2/0 unit 6 encapsulation ethernet
set logical-systems C interfaces lt-1/2/0 unit 6 peer-unit 5
set logical-systems C interfaces lt-1/2/0 unit 6 family inet address 10.10.10.6/30
set logical-systems C interfaces lo0 unit 3 family inet address 192.168.40.4/32
set logical-systems C protocols bgp group internal-peers type internal
set logical-systems C protocols bgp group internal-peers local-address 192.168.40.4
set logical-systems C protocols bgp group internal-peers export send-direct
set logical-systems C protocols bgp group internal-peers bfd-liveness-detection
    minimum-interval 1000
set logical-systems C protocols bgp group internal-peers neighbor 192.163.6.4
set logical-systems C protocols bgp group internal-peers neighbor 192.168.6.5
set logical-systems C protocols ospf area 0.0.0.0 interface lo0.3 passive
set logical-systems C protocols ospf area 0.0.0.0 interface lt-1/2/0.6
set logical-systems C policy-options policy-statement send-direct term 2 from protocol
    direct
set logical-systems C policy-options policy-statement send-direct term 2 then accept
set logical-systems C routing-options router-id 192.168.40.4
set logical-systems C routing-options autonomous-system 17

```

Configuring Device A

Step-by-Step Procedure The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device A:

1. Set the CLI to Logical System A.

```
user@host> set cli logical-system A
```

2. Configure the interfaces.

```
[edit interfaces lt-1/2/0 unit 1]
user@host:A# set description to-B
user@host:A# set encapsulation ethernet
user@host:A# set peer-unit 2
user@host:A# set family inet address 10.10.1/30
[edit interfaces lo0 unit 1]
user@host:A# set family inet address 192.168.6.5/32
```

3. Configure BGP.

The **neighbor** statements are included for both Device B and Device C, even though Device A is not directly connected to Device C.

```
[edit protocols bgp group internal-peers]
user@host:A# set type internal
user@host:A# set local-address 192.168.6.5
user@host:A# set export send-direct
user@host:A# set neighbor 192.163.6.4
user@host:A# set neighbor 192.168.40.4
```

4. Configure BFD.

```
[edit protocols bgp group internal-peers]
user@host:A# set bfd-liveness-detection minimum-interval 1000
```

You must configure the same minimum interval on the connecting peer.

5. (Optional) Configure BFD tracing.

```
[edit protocols bgp group internal-peers]
user@host:A# set traceoptions file bgp-bfd
user@host:A# set traceoptions flag bfd detail
```

6. Configure OSPF.

```
[edit protocols ospf area 0.0.0.0]
user@host:A# set interface lo0.1 passive
user@host:A# set interface lt-1/2/0.1
```

7. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit policy-options policy-statement send-direct term 2]
user@host:A# set from protocol direct
user@host:A# set then accept
```

8. Configure the router ID and the autonomous system (AS) number.

```
[edit routing-options]
user@host:A# set router-id 192.168.6.5
user@host:A# set autonomous-system 17
```

9. If you are done configuring the device, enter **commit** from configuration mode. Repeat these steps to configure Device B and Device C.

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host:A# show interfaces
lt-1/2/0 {
  unit 1 {
    description to-B;
    encapsulation ethernet;
    peer-unit 2;
    family inet {
      address 10.10.10.1/30;
    }
  }
}
lo0 {
  unit 1 {
    family inet {
      address 192.168.6.5/32;
    }
  }
}
```

```
user@host:A# show policy-options
policy-statement send-direct {
  term 2 {
    from protocol direct;
    then accept;
  }
}
```

```
user@host:A# show protocols
bgp {
  group internal-peers {
    type internal;
    traceoptions {
      file bgp-bfd;
      flag bfd detail;
    }
    local-address 192.168.6.5;
```

```
export send-direct;
bfd-liveness-detection {
    minimum-interval 1000;
}
neighbor 192.163.6.4;
neighbor 192.168.40.4;
}
}
ospf {
    area 0.0.0.0 {
        interface lo0.1 {
            passive;
        }
        interface lt-1/2/0.1;
    }
}
```

```
user@host:A# show routing-options
router-id 192.168.6.5;
autonomous-system 17;
```

Verification

Confirm that the configuration is working properly.

- [Verifying That BFD Is Enabled on page 66](#)
- [Verifying That BFD Sessions Are Up on page 67](#)
- [Viewing Detailed BFD Events on page 68](#)
- [Viewing Detailed BFD Events After Deactivating and Reactivating a Loopback Interface on page 69](#)

Verifying That BFD Is Enabled

Purpose Verify that BFD is enabled between the IBGP peers.

Action From operational mode, enter the **show bgp neighbor** command. You can use the **| match bfd** filter to narrow the output.

```
user@host:A> show bgp neighbor | match bfd
Options: <BfdEnabled>
BFD: enabled, up
Trace file: /var/log/A/bgp-bfd size 131072 files 10
Options: <BfdEnabled>
BFD: enabled, up
Trace file: /var/log/A/bgp-bfd size 131072 files 10
```

Meaning The output shows that Logical System A has two neighbors with BFD enabled. When BFD is not enabled, the output displays **BFD: disabled, down**, and the **<BfdEnabled>** option is absent. If BFD is enabled and the session is down, the output displays **BFD: enabled,**

down. The output also shows that BFD-related events are being written to a log file because trace operations are configured.

Verifying That BFD Sessions Are Up

Purpose Verify that the BFD sessions are up, and view details about the BFD sessions.

Action From operational mode, enter the **show bfd session extensive** command.

```
user@host:A> show bfd session extensive
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
192.163.6.4	Up		3.000	1.000	3
Client BGP, TX interval 1.000, RX interval 1.000					
Session up time 00:54:40					
Local diagnostic None, remote diagnostic None					
Remote state Up, version 1					
Logical system 12, routing table index 25					
Min async interval 1.000, min slow interval 1.000					
Adaptive async TX interval 1.000, RX interval 1.000					
Local min TX interval 1.000, minimum RX interval 1.000, multiplier 3					
Remote min TX interval 1.000, min RX interval 1.000, multiplier 3					
Local discriminator 10, remote discriminator 9					
Echo mode disabled/inactive					
Multi-hop route table 25, local-address 192.168.6.5					
192.168.40.4	Up		3.000	1.000	3
Client BGP, TX interval 1.000, RX interval 1.000					
Session up time 00:48:03					
Local diagnostic None, remote diagnostic None					
Remote state Up, version 1					
Logical system 12, routing table index 25					
Min async interval 1.000, min slow interval 1.000					
Adaptive async TX interval 1.000, RX interval 1.000					
Local min TX interval 1.000, minimum RX interval 1.000, multiplier 3					
Remote min TX interval 1.000, min RX interval 1.000, multiplier 3					
Local discriminator 14, remote discriminator 13					
Echo mode disabled/inactive					
Multi-hop route table 25, local-address 192.168.6.5					

2 sessions, 2 clients
Cumulative transmit rate 2.0 pps, cumulative receive rate 2.0 pps

Meaning The **TX interval 1.000, RX interval 1.000** output represents the setting configured with the **minimum-interval** statement. All of the other output represents the default settings for BFD. To modify the default settings, include the optional statements under the **bfd-liveness-detection** statement.

Viewing Detailed BFD Events

Purpose View the contents of the BFD trace file to assist in troubleshooting, if needed.

Action From operational mode, enter the `file show /var/log/A/bgp-bfd` command.

```
user@host:A> file show /var/log/A/bgp-bfd
```

```
Aug 15 17:07:25 trace_on: Tracing to "/var/log/A/bgp-bfd" started
Aug 15 17:07:26.492190 bgp_peer_init: BGP peer 192.163.6.4 (Internal AS 17) local
address 192.168.6.5 not found. Leaving peer idled
Aug 15 17:07:26.493176 bgp_peer_init: BGP peer 192.168.40.4 (Internal AS 17) local
address 192.168.6.5 not found. Leaving peer idled
Aug 15 17:07:32.597979 task_connect: task BGP_17.192.163.6.4+179 addr
192.163.6.4+179: No route to host
Aug 15 17:07:32.599623 bgp_connect_start: connect 192.163.6.4 (Internal AS 17):
No route to host
Aug 15 17:07:36.869394 task_connect: task BGP_17.192.168.40.4+179 addr
192.168.40.4+179: No route to host
Aug 15 17:07:36.870624 bgp_connect_start: connect 192.168.40.4 (Internal AS 17):
No route to host
Aug 15 17:08:04.599220 task_connect: task BGP_17.192.163.6.4+179 addr
192.163.6.4+179: No route to host
Aug 15 17:08:04.601135 bgp_connect_start: connect 192.163.6.4 (Internal AS 17):
No route to host
Aug 15 17:08:08.869717 task_connect: task BGP_17.192.168.40.4+179 addr
192.168.40.4+179: No route to host
Aug 15 17:08:08.869934 bgp_connect_start: connect 192.168.40.4 (Internal AS 17):
No route to host
Aug 15 17:08:36.603544 advertising receiving-speaker only capability to neighbor
192.163.6.4 (Internal AS 17)
Aug 15 17:08:36.606726 bgp_read_message: 192.163.6.4 (Internal AS 17): 0 bytes
buffered
Aug 15 17:08:36.609119 Initiated BFD session to peer 192.163.6.4 (Internal AS
17): address=192.163.6.4 ifindex=0 ifname=(none) txivl=1000 rxivl=1000 mult=3
ver=255
Aug 15 17:08:36.734033 advertising receiving-speaker only capability to neighbor
192.168.40.4 (Internal AS 17)
Aug 15 17:08:36.738436 Initiated BFD session to peer 192.168.40.4 (Internal AS
17): address=192.168.40.4 ifindex=0 ifname=(none) txivl=1000 rxivl=1000 mult=3
ver=255
Aug 15 17:08:40.537552 BFD session to peer 192.163.6.4 (Internal AS 17) up
Aug 15 17:08:40.694410 BFD session to peer 192.168.40.4 (Internal AS 17) up
```

Meaning Before the routes are established, the **No route to host** message appears in the output. After the routes are established, the last two lines show that both BFD sessions come up.

Viewing Detailed BFD Events After Deactivating and Reactivating a Loopback Interface

Purpose Check to see what happens after bringing down a router or switch and then bringing it back up. To simulate bringing down a router or switch, deactivate the loopback interface on Logical System B.

Action 1. From configuration mode, enter the **deactivate logical-systems B interfaces lo0 unit 2 family inet** command.

```
user@host:A# deactivate logical-systems B interfaces lo0 unit 2 family inet
user@host:A# commit
```

2. From operational mode, enter the **file show /var/log/A/bgp-bfd** command.

```
user@host:A> file show /var/log/A/bgp-bfd
...
Aug 15 17:20:55.995648 bgp_read_v4_message:9747: NOTIFICATION received from
192.163.6.4 (Internal AS 17): code 6 (Cease) subcode 6 (Other Configuration
Change)
Aug 15 17:20:56.004508 Terminated BFD session to peer 192.163.6.4 (Internal
AS 17)
Aug 15 17:21:28.007755 task_connect: task BGP_17.192.163.6.4+179 addr
192.163.6.4+179: No route to host
Aug 15 17:21:28.008597 bgp_connect_start: connect 192.163.6.4 (Internal AS
17): No route to host
```

3. From configuration mode, enter the **activate logical-systems B interfaces lo0 unit 2 family inet** command.

```
user@host:A# activate logical-systems B interfaces lo0 unit 2 family inet
user@host:A# commit
```

4. From operational mode, enter the **file show /var/log/A/bgp-bfd** command.

```
user@host:A> file show /var/log/A/bgp-bfd
...
Aug 15 17:25:53.623743 advertising receiving-speaker only capability to neighbor
192.163.6.4 (Internal AS 17)
Aug 15 17:25:53.631314 Initiated BFD session to peer 192.163.6.4 (Internal AS
17): address=192.163.6.4 ifindex=0 ifname=(none) txivl=1000 rxivl=1000 mult=3
ver=255
Aug 15 17:25:57.570932 BFD session to peer 192.163.6.4 (Internal AS 17) up
```

Related Documentation

- *Example: Configuring BFD Authentication for BGP*

Example: Configuring BFD for OSPF

This example shows how to configure the Bidirectional Forwarding Detection (BFD) protocol for OSPF.

- [Requirements on page 70](#)
- [Overview on page 70](#)
- [Configuration on page 72](#)
- [Verification on page 73](#)

Requirements

Before you begin:

- Configure the device interfaces. See the *Junos OS Network Interfaces Library for Routing Devices*.
- Configure the router identifiers for the devices in your OSPF network. See *Example: Configuring an OSPF Router Identifier*.
- Control OSPF designated router election. See *Example: Controlling OSPF Designated Router Election*.
- Configure a single-area OSPF network. See *Example: Configuring a Single-Area OSPF Network*.
- Configure a multiarea OSPF network. See *Example: Configuring a Multiarea OSPF Network*.
- Configure a multiarea OSPF network. See *Example: Configuring a Multiarea OSPF Network*.

Overview

An alternative to adjusting the OSPF hello interval and dead interval settings to increase route convergence is to configure BFD. The BFD protocol is a simple hello mechanism that detects failures in a network. The BFD failure detection timers have shorter timer limits than the OSPF failure detection mechanisms, thereby providing faster detection.

BFD is useful on interfaces that are unable to detect failure quickly, such as Ethernet interfaces. Other interfaces, such as SONET interfaces, already have built-in failure detection. Configuring BFD on those interfaces is unnecessary.

You configure BFD on a pair of neighboring OSPF interfaces. Unlike the OSPF hello interval and dead interval settings, you do not have to enable BFD on all interfaces in an OSPF area.

In this example, you enable failure detection by including the **bfd-liveness-detection** statement on the neighbor OSPF interface **fe-0/1/0** in area 0.0.0.0 and configure the BFD packet exchange interval to 300 milliseconds, configure 4 as the number of missed hello packets that causes the originating interface to be declared down, and configure BFD sessions only for OSPF neighbors with full neighbor adjacency by including the following settings:

- **full-neighbors-only**—In Junos OS Release 9.5 and later, configures the BFD protocol to establish BFD sessions only for OSPF neighbors with full neighbor adjacency. The default behavior is to establish BFD sessions for all OSPF neighbors.
- **minimum-interval**—Configures the minimum interval, in milliseconds, after which the local routing device transmits hello packets as well as the minimum interval after which the routing device expects to receive a reply from the neighbor with which it has established a BFD session. You can configure a number in the range from 1 through 255,000 milliseconds. You can also specify the minimum transmit and receive intervals separately using the **transmit-interval**, **minimum-interval**, and **minimum-receive-interval** statements.



NOTE: BFD is an intensive protocol that consumes system resources. Specifying a minimum interval for BFD of less than 100 ms for Routing Engine-based sessions and 10 ms for distributed BFD sessions can cause undesired BFD flapping.

Depending on your network environment, these additional recommendations might apply:

- For large-scale network deployments with a large number of BFD sessions, specify a minimum interval of no less than 500 ms. An interval of 1000 ms is recommended to avoid any instability issues.



NOTE:

- For the **bfdd** process, the detection time interval set is lower than 300 ms. If there is a high priority process such as **ppmd** running on the system, the CPU might spend time on the **ppmd** process rather than the **bfdd** process.
- For branch SRX Series devices, we recommend 1000 ms as the minimum keepalive time interval for BFD packets.
- For very large-scale network deployments with a large number of BFD sessions, contact Juniper Networks customer support for more information.
- For BFD sessions to remain up during a Routing Engine switchover event when nonstop active routing (NSR) is configured, specify a minimum interval of 2500 ms for Routing Engine-based sessions. For distributed BFD sessions with NSR configured, the minimum interval recommendations are unchanged and depend only on your network deployment.

- **multiplier**—Configures the number of hello packets not received by a neighbor that causes the originating interface to be declared down. By default, three missed hello packets cause the originating interface to be declared down. You can configure a value in the range from 1 through 255.

Configuration

CLI Quick Configuration To quickly configure the BFD protocol for OSPF, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter **commit** from configuration mode.

```
[edit]
set protocols ospf area 0.0.0.0 interface fe-0/0/1 bfd-liveness-detection minimum-interval
  300
set protocols ospf area 0.0.0.0 interface fe-0/0/1 bfd-liveness-detection multiplier 4
set protocols ospf area 0.0.0.0 interface fe-0/0/1 bfd-liveness-detection full-neighbors-only
```

Step-by-Step Procedure To configure the BFD protocol for OSPF on one neighboring interface:

1. Create an OSPF area.



NOTE: To specify OSPFv3, include the **ospf3** statement at the [edit protocols] hierarchy level.

```
[edit]
user@host# edit protocols ospf area 0.0.0.0
```

2. Specify the interface.

```
[edit protocols ospf area 0.0.0.0]
user@host# set interface fe-0/0/1
```

3. Specify the minimum transmit and receive intervals.

```
[edit protocols ospf area 0.0.0.0 ]
user@host# set interface fe-0/0/1 bfd-liveness-detection minimum-interval 300
```

4. Configure the number of missed hello packets that cause the originating interface to be declared down.

```
[edit protocols ospf area 0.0.0.0 ]
user@host# set interface fe-0/0/1 bfd-liveness-detection multiplier 4
```

5. Configure BFD sessions only for OSPF neighbors with full neighbor adjacency.

```
[edit protocols ospf area 0.0.0.0 ]
user@host# set interface fe-0/0/1 bfd-liveness-detection full-neighbors-only
```

6. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.0 ]
user@host# commit
```



NOTE: Repeat this entire configuration on the other neighboring interface.

Results Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
area 0.0.0.0 {
  interface fe-0/0/1.0 {
    bfd-liveness-detection {
      minimum-interval 300;
      multiplier 4;
      full-neighbors-only;
    }
  }
}
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

Verification

Confirm that the configuration is working properly.

Verifying the BFD Sessions

Purpose Verify that the OSPF interfaces have active BFD sessions, and that session components have been configured correctly.

Action From operational mode, enter the **show bfd session detail** command.

Meaning The output displays information about the BFD sessions.

- The Address field displays the IP address of the neighbor.
- The Interface field displays the interface you configured for BFD.

- The State field displays the state of the neighbor and should show Full to reflect the full neighbor adjacency that you configured.
- The Transmit Interval field displays the time interval you configured to send BFD packets.
- The Multiplier field displays the multiplier you configured.

Related Documentation

- [Understanding BFD for OSPF on page 38](#)
- [Understanding BFD Authentication for OSPF](#)
- [Example: Configuring BFD Authentication for OSPF](#)

Example: Configuring BFD for IS-IS

This example describes how to configure the Bidirectional Forwarding Detection (BFD) protocol to detect failures in an IS-IS network.



NOTE: BFD is not supported with ISIS for IPV6 on QFX10000 series switches.

-
- [Requirements on page 74](#)
 - [Overview on page 74](#)
 - [Configuration on page 75](#)
 - [Verification on page 79](#)

Requirements

Before you begin, configure IS-IS on both routers. See *Example: Configuring IS-IS* for information about the required IS-IS configuration.

This example uses the following hardware and software components:


- Junos OS Release 7.3 or later
- M Series, MX Series, and T Series routers

Overview

This example shows two routers connected to each other. A loopback interface is configured on each router. IS-IS and BFD protocols are configured on both routers.

[Figure 3 on page 75](#) shows the sample network.

Figure 3: Configuring BFD for IS-IS



Configuration

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Router R1

```
set protocols isis interface so-0/0/0 bfd-liveness-detection detection-time threshold 5
set protocols isis interface so-0/0/0 bfd-liveness-detection minimum-interval 2
set protocols isis interface so-0/0/0 bfd-liveness-detection minimum-receive-interval 1
set protocols isis interface so-0/0/0 bfd-liveness-detection no-adaptation
set protocols isis interface so-0/0/0 bfd-liveness-detection transmit-interval threshold 3
set protocols isis interface so-0/0/0 bfd-liveness-detection transmit-interval
  minimum-interval 1
set protocols isis interface so-0/0/0 bfd-liveness-detection multiplier 2
set protocols isis interface so-0/0/0 bfd-liveness-detection version automatic
```

Router R2

```
set protocols isis interface so-0/0/0 bfd-liveness-detection detection-time threshold 6
set protocols isis interface so-0/0/0 bfd-liveness-detection minimum-interval 3
set protocols isis interface so-0/0/0 bfd-liveness-detection minimum-receive-interval 1
set protocols isis interface so-0/0/0 bfd-liveness-detection no-adaptation
set protocols isis interface so-0/0/0 bfd-liveness-detection transmit-interval threshold 4
set protocols isis interface so-0/0/0 bfd-liveness-detection transmit-interval
  minimum-interval 1
set protocols isis interface so-0/0/0 bfd-liveness-detection multiplier 2
set protocols isis interface so-0/0/0 bfd-liveness-detection version automatic
```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode*.



NOTE: To simply configure BFD for IS-IS, only the `minimum-interval` statement is required. The BFD protocol selects default parameters for all the other configuration statements when you use the `bfd-liveness-detection` statement without specifying any parameters.



NOTE: You can change parameters at any time without stopping or restarting the existing session. BFD automatically adjusts to the new parameter value. However, no changes to BFD parameters take place until the values resynchronize with each BFD peer.

To configure BFD for IS-IS on Routers R1 and R2:

1. Enable BFD failure detection for IS-IS.

```
[edit protocols isis]  
user@R1# set interface so-0/0/0 bfd-liveness-detection
```

```
[edit protocols isis]  
user@R2# set interface so-0/0/0 bfd-liveness-detection
```

2. Configure the threshold for the adaptation of the detection time, which must be greater than the multiplier number multiplied by the minimum interval.

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]  
user@R1# set detection-time threshold 5
```

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]  
user@R2# set detection-time threshold 6
```

3. Configure the minimum transmit and receive intervals for failure detection.

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]  
user@R1# set minimum-interval 2
```

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]  
user@R2# set minimum-interval 3
```

4. Configure only the minimum receive interval for failure detection.

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R1# set minimum-receive-interval 1
```

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R2# set minimum-receive-interval 1
```

5. Disable BFD adaptation.

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R1# set no-adaptation
```

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R2# set no-adaptation
```

6. Configure the threshold for the transmit interval, which must be greater than the minimum transmit interval.

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R1# set transmit-interval threshold 3
```

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R2# set transmit-interval threshold 4
```

7. Configure the minimum transmit interval for failure detection.

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R1# set transmit-interval minimum-interval 1
```

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R2# set transmit-interval minimum-interval 1
```

8. Configure the multiplier number, which is the number of hello packets not received by the neighbor that causes the originating interface to be declared down.

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R1# set multiplier 2
```

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R2# set multiplier 2
```

9. Configure the BFD version used for detection.

The default is to have the version detected automatically.

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R1# set version automatic
```

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R2# set version automatic
```

Results

From configuration mode, confirm your configuration by issuing the **show protocols isis interface** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show protocols isis interface so-0/0/0
```

```
bfd-liveness-detection {
  version automatic;
  minimum-interval 2;
  minimum-receive-interval 1;
  multiplier 2;
  no-adaptation;
  transmit-interval {
    minimum-interval 1;
    threshold 3;
  }
  detection-time {
    threshold 5;
  }
}
...
```

```
user@R2# show protocols isis interface so-0/0/0
```

```
bfd-liveness-detection {
  version automatic;
  minimum-interval 3;
  minimum-receive-interval 1;
  multiplier 2;
  no-adaptation;
  transmit-interval {
    minimum-interval 1;
    threshold 4;
  }
  detection-time {
    threshold 6;
  }
}
...
```

Verification

Confirm that the configuration is working properly.

- [Verifying the Connection Between Routers R1 and R2 on page 79](#)
- [Verifying That IS-IS Is Configured on page 79](#)
- [Verifying That BFD Is configured on page 80](#)

Verifying the Connection Between Routers R1 and R2

Purpose Make sure that Routers R1 and R2 are connected to each other.

Action Ping the other router to check the connectivity between the two routers as per the network topology.

```
user@R1> ping 10.0.0.2
```

```
PING 10.0.0.2 (10.0.0.2): 56 data bytes
64 bytes from 10.0.0.2: icmp_seq=0 ttl=64 time=1.367 ms
64 bytes from 10.0.0.2: icmp_seq=1 ttl=64 time=1.662 ms
64 bytes from 10.0.0.2: icmp_seq=2 ttl=64 time=1.291 ms
^C
--- 10.0.0.2 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.291/1.440/1.662/0.160 ms
```

```
user@R2> ping 10.0.0.1
```

```
PING 10.0.0.1 (10.0.0.1): 56 data bytes
64 bytes from 10.0.0.1: icmp_seq=0 ttl=64 time=1.287 ms
64 bytes from 10.0.0.1: icmp_seq=1 ttl=64 time=1.310 ms
64 bytes from 10.0.0.1: icmp_seq=2 ttl=64 time=1.289 ms
^C
--- 10.0.0.1 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.287/1.295/1.310/0.010 ms
```

Meaning Routers R1 and R2 are connected to each other.

Verifying That IS-IS Is Configured

Purpose Make sure that the IS-IS instance is running on both routers.

Action Use the **show isis database** statement to check if the IS-IS instance is running on both routers, R1 and R2.

```
user@R1> show isis database
```

```
IS-IS level 1 link-state database:
LSP ID      Sequence Checksum Lifetime Attributes
R1.00-00    0x4a571  0x30c5    1195 L1 L2
R2.00-00    0x4a586  0x4b7e    1195 L1 L2
R2.02-00    0x330ca1 0x3492    1196 L1 L2
  3 LSPs
```

```
IS-IS level 2 link-state database:
LSP ID      Sequence Checksum Lifetime Attributes
R1.00-00    0x4a856  0x5db0    1194 L1 L2
R2.00-00    0x4a89d  0x149b    1194 L1 L2
R2.02-00    0x1fb2ff 0xd302    1194 L1 L2
  3 LSPs
```

```
user@R2> show isis database
```

```
IS-IS level 1 link-state database:
LSP ID      Sequence Checksum Lifetime Attributes
R1.00-00    0x4b707  0xcc80    1195 L1 L2
R2.00-00    0x4b71b  0xeb37    1198 L1 L2
R2.02-00    0x33c2ce 0xb52d    1198 L1 L2
  3 LSPs
```

```
IS-IS level 2 link-state database:
LSP ID      Sequence Checksum Lifetime Attributes
R1.00-00    0x4b9f2  0xee70    1192 L1 L2
R2.00-00    0x4ba41  0x9862    1197 L1 L2
R2.02-00    0x3      0x6242    1198 L1 L2
  3 LSPs
```

Meaning IS-IS is configured on both routers, R1 and R2.

Verifying That BFD Is configured

Purpose Make sure that the BFD instance is running on both routers, R1 and R2.

Action Use the `show bfd session detail` statement to check if BFD instance is running on the routers.

```
user@R1> show bfd session detail
```

```
Address          State      Interface    Detect   Transmit
10.0.0.2         Up         so-0/0/0     Time    Interval Multiplier
                2.000     1.000        2
Client ISIS R2, TX interval 0.001, RX interval 0.001
Client ISIS R1, TX interval 0.001, RX interval 0.001
Session down time 00:00:00, previous up time 00:00:15
Local diagnostic NbrSignal, remote diagnostic NbrSignal
Remote state AdminDown, version 1
Router 3, routing table index 17
```

```
1 sessions, 2 clients
Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps
```

```
user@R2> show bfd session detail
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
10.0.0.1	Up	so-0/0/0	2.000	1.000	2

```
Client ISIS R2, TX interval 0.001, RX interval 0.001
Session down time 00:00:00, previous up time 00:00:05
Local diagnostic NbrSignal, remote diagnostic NbrSignal
Remote state AdminDown, version 1
Router 2, routing table index 15
```

```
1 sessions, 1 clients
Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps
```

Meaning BFD is configured on Routers R1 and R2 for detecting failures in the IS-IS network.

Related Documentation

- [Understanding BFD for IS-IS on page 41](#)

Example: Configuring BFD for RIP

This example shows how to configure Bidirectional Forwarding Detection (BFD) for a RIP network.

- [Requirements on page 81](#)
- [Overview on page 81](#)
- [Configuration on page 83](#)
- [Verification on page 86](#)

Requirements

No special configuration beyond device initialization is required before configuring this example.

Overview

To enable failure detection, include the **bfd-liveness-detection** statement:

```
bfd-liveness-detection {
  detection-time {
    threshold milliseconds;
  }
  minimum-interval milliseconds;
  minimum-receive-interval milliseconds;
  multiplier number;
  no-adaptation;
  transmit-interval {
```

```
    threshold milliseconds;  
    minimum-interval milliseconds;  
  }  
  version (1 | automatic);  
}
```

Optionally, you can specify the threshold for the adaptation of the detection time by including the **threshold** statement. When the BFD session detection time adapts to a value equal to or greater than the threshold, a single trap and a system log message are sent.

To specify the minimum transmit and receive interval for failure detection, include the **minimum-interval** statement. This value represents the minimum interval at which the local routing device transmits hello packets as well as the minimum interval at which the routing device expects to receive a reply from a neighbor with which it has established a BFD session. You can configure a value in the range from 1 through 255,000 milliseconds. This examples sets a minimum interval of 600 milliseconds.



NOTE: BFD is an intensive protocol that consumes system resources. Specifying a minimum interval for BFD of less than 100 ms for Routing Engine-based sessions and 10 ms for distributed BFD sessions can cause undesired BFD flapping.

Depending on your network environment, these additional recommendations might apply:

- For large-scale network deployments with a large number of BFD sessions, specify a minimum interval of 300 ms for Routing Engine-based sessions and 100 ms for distributed BFD sessions.
- For very large-scale network deployments with a large number of BFD sessions, contact Juniper Networks customer support for more information.
- For BFD sessions to remain up during a Routing Engine switchover event when nonstop active routing (NSR) is configured, specify a minimum interval of 2500 ms for Routing Engine-based sessions. For distributed BFD sessions with nonstop active routing configured, the minimum interval recommendations are unchanged and depend only on your network deployment.

You can optionally specify the minimum transmit and receive intervals separately.

To specify only the minimum receive interval for failure detection, include the **minimum-receive-interval** statement. This value represents the minimum interval at which the local routing device expects to receive a reply from a neighbor with which it has established a BFD session. You can configure a value in the range from 1 through 255,00 milliseconds.

To specify only the minimum transmit interval for failure detection, include the **transmit-interval minimum-interval** statement. This value represents the minimum interval

at which the local routing device transmits hello packets to the neighbor with which it has established a BFD session. You can configure a value in the range from 1 through 255,000 milliseconds.

To specify the number of hello packets not received by a neighbor that causes the originating interface to be declared down, include the **multiplier** statement. The default is 3, and you can configure a value in the range from 1 through 255.

To specify the threshold for detecting the adaptation of the transmit interval, include the **transmit-interval threshold** statement. The threshold value must be greater than the transmit interval.


To specify the BFD version used for detection, include the **version** statement. The default is to have the version detected automatically.

You can trace BFD operations by including the **traceoptions** statement at the **[edit protocols bfd]** hierarchy level.

In Junos OS Release 9.0 and later, you can configure BFD sessions not to adapt to changing network conditions. To disable BFD adaptation, include the **no-adaptation** statement. We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

Figure 4 on page 83 shows the topology used in this example.

Figure 4: RIP BFD Network Topology



“CLI Quick Configuration” on page 83 shows the configuration for all of the devices in Figure 4 on page 83. The section “Step-by-Step Procedure” on page 84 describes the steps on Device R1.

Configuration

CLI Quick Configuration	To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level.
Device R1	<pre> set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30 set protocols bfd traceoptions file bfd-trace set protocols bfd traceoptions flag all set protocols rip group rip-group export advertise-routes-through-rip set protocols rip group rip-group neighbor fe-1/2/0.1 </pre>

```
set protocols rip group rip-group bfd-liveness-detection minimum-interval 600
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

Device R2

```
set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces fe-1/2/1 unit 5 family inet address 10.0.0.5/30
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.2
set protocols rip group rip-group neighbor fe-1/2/1.5
set protocols rip group rip-group bfd-liveness-detection minimum-interval 600
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

Device R3

```
set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.6
set protocols rip group rip-group bfd-liveness-detection minimum-interval 600
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a BFD for a RIP network:

1. Configure the network interfaces.

```
[edit interfaces]
user@R1# set fe-1/2/0 unit 1 family inet address 10.0.0.1/30
```

2. Create the RIP group and add the interface.

To configure RIP in Junos OS, you must configure a group that contains the interfaces on which RIP is enabled. You do not need to enable RIP on the loopback interface.

```
[edit protocols rip group rip-group]
user@R1# set neighbor fe-1/2/0.1
```

3. Create the routing policy to advertise both direct and RIP-learned routes.

```
[edit policy-options policy-statement advertise-routes-through-rip term 1]
user@R1# set from protocol direct
user@R1# set from protocol rip
user@R1# set then accept
```

4. Apply the routing policy.

In Junos OS, you can only apply RIP export policies at the group level.

```
[edit protocols rip group rip-group]
user@R1# set export advertise-routes-through-rip
```

5. Enable BFD.

```
[edit protocols rip group rip-group]
user@R1# set bfd-liveness-detection minimum-interval 600
```

6. Configure tracing operations to track BFD messages.

```
[edit protocols bfd traceoptions]
user@R1# set file bfd-trace
user@R1# set flag all
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R1# show interfaces
fe-1/2/0 {
  unit 1 {
    family inet {
      address 10.0.0.1/30;
    }
  }
}
```

```
user@R1# show protocols
bfd {
  traceoptions {
    file bfd-trace;
    flag all;
  }
}
rip {
  group rip-group {
    export advertise-routes-through-rip;
    bfd-liveness-detection {
```

```
        minimum-interval 600;
    }
    neighbor fe-1/2/0.1;
}

user@R1# show policy-options
policy-statement advertise-routes-through-rip {
    term 1 {
        from protocol [ direct rip ];
        then accept;
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Verifying That the BFD Sessions Are Up on page 86](#)
- [Checking the BFD Trace File on page 86](#)

Verifying That the BFD Sessions Are Up

Purpose Make sure that the BFD sessions are operating.

Action From operational mode, enter the **show bfd session** command.

```
user@R1> show bfd session
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
10.0.0.2	Up	fe-1/2/0.1	1.800	0.600	3

1 sessions, 1 clients
Cumulative transmit rate 1.7 pps, cumulative receive rate 1.7 pps

Meaning The output shows that there are no authentication failures.

Checking the BFD Trace File

Purpose Use tracing operations to verify that BFD packets are being exchanged.

Action From operational mode, enter the **show log** command.

```
user@R1> show log bfd-trace
```

```
Feb 16 10:26:32 PPM Trace: BFD periodic xmit to 10.0.0.2 (IFL 124, rtbl 53,
single-hop port)
```

```

Feb 16 10:26:32 Received Downstream TraceMsg (24) len 86:
Feb 16 10:26:32   IfIndex (3) len 4: 0
Feb 16 10:26:32   Protocol (1) len 1: BFD
Feb 16 10:26:32   Data (9) len 61: (hex) 42 46 44 20 70 61 63 6b 65 74 20 66 72
    6f 6d 20 31 30 2e
Feb 16 10:26:32 PPM Trace: BFD packet from 10.0.0.1 (IFL 73, rtbl 56, ttl 255)
absorbed
Feb 16 10:26:32 Received Downstream TraceMsg (24) len 60:
Feb 16 10:26:32   IfIndex (3) len 4: 0
Feb 16 10:26:32   Protocol (1) len 1: BFD
Feb 16 10:26:32   Data (9) len 35: (hex) 42 46 44 20 70 65 72 69 6f 64 69 63 20
    78 6d 69 74 20 6f
...

```

Meaning The output shows the normal functioning of BFD.

Configuring Independent Micro BFD Sessions for LAG

The Bidirectional Forwarding Detection (BFD) protocol is a simple detection protocol that quickly detects failures in the forwarding paths. A link aggregation group (LAG) combines multiple links between devices that are in point-to-point connections, thereby increasing bandwidth, providing reliability, and allowing load balancing. To run a BFD session on LAG interfaces, configure an independent, asynchronous mode BFD session on every LAG member link in a LAG bundle. Instead of a single BFD session monitoring the status of the UDP port, independent micro BFD sessions monitor the status of individual member links.

To enable failure detection for aggregated Ethernet interfaces:

1. Include the following statement in the configuration at the **[edit interfaces aex aggregated-ether-options]** hierarchy level:

```

bfd-liveness-detection {
  authentication {
    algorithm algorithm-name;
    key-chain key-chain-name;
    loose-check;
  }
  detection-time {
    threshold milliseconds;
  }
  holddown-interval milliseconds;
  local-address bfd-local-address;
  minimum-interval milliseconds;
  minimum-receive-interval milliseconds;
  multiplier number;
  neighbor bfd-neighbor-address;
  no-adaptation;
  transmit-interval {
    minimum-interval milliseconds;
    threshold milliseconds;
  }
}

```

```
    }  
    version (1 | automatic);  
}
```

2. Configure the authentication criteria of the BFD session for LAG.

To specify the authentication criteria, include the **authentication** statement:

```
bfd-liveness-detection {  
  authentication {  
    algorithm algorithm-name;  
    key-chain key-chain-name;  
    loose-check;  
  }  
}
```

- Specify the algorithm to be used to authenticate the BFD session. You can use one of the following algorithms for authentication:
 - keyed-md5
 - keyed-sha-1
 - meticulous-keyed-md5
 - meticulous-keyed-sha-1
 - simple-password
- To configure the key chain, specify the name that is associated with the security key for the BFD session. The name you specify must match one of the key chains configured in the **authentication-key-chains** *key-chain* statement at the **[edit security]** hierarchy level.
- Configure loose authentication checking on the BFD session. Use only for transitional periods when authentication might not be configured at both ends of the BFD session.

3. Configure BFD timers for aggregated Ethernet interfaces.

To specify the BFD timers, include the **detection-time** statement:

```
bfd-liveness-detection {  
  detection-time {  
    threshold milliseconds;  
  }  
}
```

Specify the threshold value. This is the maximum time interval for detecting a BFD neighbor. If the transmit interval is greater than this value, the device triggers a trap.

4. Configure a hold-down interval value to set the minimum time that the BFD session must remain up before a state change notification is sent to the other members in the LAG network.

To specify the hold-down interval, include the **holddown-interval** statement:

```
bfd-liveness-detection {
    holddown-interval milliseconds;
}
```

You can configure a number in the range from 0 through 255,000 milliseconds, and the default is 0. If the BFD session goes down and then comes back up during the hold-down interval, the timer is restarted.

This value represents the minimum interval at which the local routing device transmits BFD packets, as well as the minimum interval in which the routing device expects to receive a reply from a neighbor with which it has established a BFD session. You can configure a number in the range from 1 through 255,000 milliseconds. You can also specify the minimum transmit and receive intervals separately.

5. Configure the source address for the BFD session.

To specify a local address, include the **local-address** statement:

```
bfd-liveness-detection {
    local-address bfd-local-address;
}
```

The BFD local address is the loopback address of the source of the BFD session.



NOTE: Beginning with Junos OS Release 16.1, you can also configure this feature with the AE interface address as the local address in a micro BFD session. For the IPv6 address family, disable duplicate address detection before configuring this feature with the AE interface address. To disable duplicate address detection, include the **dad-disable** statement at the [edit interface *aex* unit *y* family inet6] hierarchy level.

Beginning with Release 16.1R2, Junos OS checks and validates the configured micro BFD **local-address** against the interface or loopback IP address before the configuration commit. Junos OS performs this check on both IPv4 and IPv6 micro BFD address configurations, and if they do not match, the commit fails.

6. Specify the minimum interval that indicates the time interval for transmitting and receiving data.

This value represents the minimum interval at which the local routing device transmits BFD packets, as well as the minimum interval in which the routing device expects to receive a reply from a neighbor with which it has established a BFD session. You can configure a number in the range from 1 through 255,000 milliseconds. You can also specify the minimum transmit and receive intervals separately.

To specify the minimum transmit and receive intervals for failure detection, include the **minimum-interval** statement:

```
bfd-liveness-detection {  
  minimum-interval milliseconds;  
}
```



NOTE: BFD is an intensive protocol that consumes system resources. Specifying a minimum interval for BFD less than 100 ms for Routing Engine-based sessions and 10 ms for distributed BFD sessions can cause undesired BFD flapping.

Depending on your network environment, these additional recommendations might apply:

- For large-scale network deployments with a large number of BFD sessions, specify a minimum interval of 300 ms for Routing Engine-based sessions and 100 ms for distributed BFD sessions.
- For very large-scale network deployments with a large number of BFD sessions, contact Juniper Networks customer support for more information.
- For BFD sessions to remain up during a Routing Engine switchover event when nonstop active routing is configured, specify a minimum interval of 2500 ms for Routing Engine-based sessions. For distributed BFD sessions with nonstop active routing configured, the minimum interval recommendations are unchanged and depend only on your network deployment.

7. Specify only the minimum receive interval for failure detection by including the **minimum-receive-interval** statement:

```
bfd-liveness-detection {  
  minimum-receive-interval milliseconds;  
}
```

This value represents the minimum interval in which the local routing device expects to receive a reply from a neighbor with which it has established a BFD session. You can configure a number in the range from 1 through 255,000 milliseconds.

8. Specify the number of BFD packets that were not received by the neighbor that causes the originating interface to be declared down by including the **multiplier** statement:

```
bfd-liveness-detection {  
  multiplier number;  
}
```

The default value is 3. You can configure a number in the range from 1 through 255.

9. Configure the neighbor in a BFD session.

The neighbor address can be either an IPv4 or an IPv6 address.

To specify the next hop of the BFD session, include the **neighbor** statement:

```
bfd-liveness-detection {
  neighbor bfd-neighbor-address;
}
```

The BFD neighbor address is the loopback address of the remote destination of the BFD session.



NOTE: Beginning with Junos OS Release 16.1, you can also configure the AE interface address of the remote destination as the BFD neighbor address in a micro BFD session.

10. (Optional) Configure BFD sessions not to adapt to changing network conditions.

To disable BFD adaptation, include the **no-adaptation** statement:

```
bfd-liveness-detection {
  no-adaptation;
}
```



NOTE: We recommend that you do not disable BFD adaptation unless it is preferable not to have BFD adaptation in your network.

11. Specify a threshold for detecting the adaptation of the detection time by including the **threshold** statement:

```
bfd-liveness-detection {
  detection-time {
    threshold milliseconds;
  }
}
```

When the BFD session detection time adapts to a value equal to or greater than the threshold, a single trap and a system log message are sent. The detection time is based on the multiplier of the minimum-interval or the minimum-receive-interval value. The threshold must be a higher value than the multiplier for either of these configured values. For example, if the minimum-receive-interval is 300 ms and the multiplier is 3, the total detection time is 900 ms. Therefore, the detection time threshold must have a value greater than 900.

12. Specify only the minimum transmit interval for failure detection by including the **transmit-interval minimum-interval** statement:

```
bfd-liveness-detection {
```

```
transmit-interval {  
    minimum-interval milliseconds;  
}  
}
```

This value represents the minimum interval at which the local routing device transmits BFD packets to the neighbor with which it has established a BFD session. You can configure a value in the range from 1 through 255,000 milliseconds.

13. Specify the transmit threshold for detecting the adaptation of the transmit interval by including the **transmit-interval threshold** statement:

```
bfd-liveness-detection {  
    transmit-interval {  
        threshold milliseconds;  
    }  
}
```

The threshold value must be greater than the transmit interval. When the BFD session detection time adapts to a value greater than the threshold, a single trap and a system log message are sent. The detection time is based on the multiplier of the minimum-interval or the minimum-receive-interval value. The threshold must be a higher value than the multiplier for either of these configured values.

14. Specify the BFD version by including the **version** statement:

```
bfd-liveness-detection {  
    version (1 | automatic);  
}
```

The default is to have the version detected automatically.



NOTE: This feature works when both the devices support BFD. If BFD is configured at only one end of the LAG, this feature does not work.

**Related
Documentation**

- [authentication on page 506](#)
- [bfd-liveness-detection on page 507](#)
- [detection-time on page 509](#)
- [Example: Configuring Independent Micro BFD Sessions for LAG on page 93](#)

Example: Configuring Independent Micro BFD Sessions for LAG

This example shows how to configure an independent micro BFD session for aggregated Ethernet interfaces.

- [Requirements on page 93](#)
- [Overview on page 93](#)
- [Configuration on page 94](#)
- [Verification on page 99](#)

Requirements

This example uses the following hardware and software components:

- MX Series routers with Junos Trio chipset
- T Series routers with Type 4 FPC or Type 5 FPC
 - BFD for LAG is supported on the following PIC types on T-Series:
 - PC-1XGE-XENPAK (Type 3 FPC),
 - PD-4XGE-XFP (Type 4 FPC),
 - PD-5-10XGE-SFP (Type 4 FPC),
 - 24x10GE (LAN/WAN) SFPP, 12x10GE (LAN/WAN) SFPP, 1X100GE Type 5 PICs
- PTX Series routers with 24X10GE (LAN/WAN) SFPP
- Junos OS Release 13.3 or later running on all devices

Overview

The example includes two routers that are directly connected. Configure two aggregated Ethernet interfaces, AE0 for IPv4 connectivity and AE1 for IPv6 connectivity. Configure micro BFD session on the AE0 bundle using IPv4 addresses as local and neighbor endpoints on both routers. Configure micro BFD session on the AE1 bundle using IPv6 addresses as local and neighbor endpoints on both routers. This example verifies that independent micro BFD sessions are active in the output.

Topology

Figure 5 on page 93 shows the sample topology.

Figure 5: Configuring an Independent Micro BFD Session for LAG



Configuration

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Router R0

```
set interfaces ge-1/0/1 unit 0 family inet address 20.20.20.1/30
set interfaces ge-1/0/1 unit 0 family inet6 address 3ffe::1:1/126
set interfaces xe-4/0/0 gigether-options 802.3ad ae0
set interfaces xe-4/0/1 gigether-options 802.3ad ae0
set interfaces xe-4/1/0 gigether-options 802.3ad ae1
set interfaces xe-4/1/1 gigether-options 802.3ad ae1
set interfaces lo0 unit 0 family inet address 10.255.106.107/32
set interfaces lo0 unit 0 family inet6 address 201:DB8:251::aa:aa:1/126
set interfaces ae0 aggregated-ether-options bfd-liveness-detection minimum-interval 100
set interfaces ae0 aggregated-ether-options bfd-liveness-detection neighbor 10.255.106.102
set interfaces ae0 aggregated-ether-options bfd-liveness-detection local-address 10.255.106.107
set interfaces ae0 aggregated-ether-options minimum-links 1
set interfaces ae0 aggregated-ether-options link-speed 10g
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 unit 0 family inet address 10.0.0.1/30
set interfaces ae1 aggregated-ether-options bfd-liveness-detection minimum-interval 100
set interfaces ae1 aggregated-ether-options bfd-liveness-detection multiplier 3
set interfaces ae1 aggregated-ether-options bfd-liveness-detection neighbor 201:DB8:251::bb:bb:1
set interfaces ae1 aggregated-ether-options bfd-liveness-detection local-address 201:DB8:251::aa:aa:1
set interfaces ae1 aggregated-ether-options minimum-links 1
set interfaces ae1 aggregated-ether-options link-speed 10g
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 unit 0 family inet6 address 5555::1/126
set interface ae1 unit 0 family inet6 dad-disable
set routing-options nonstop-routing
set routing-options static route 30.30.30.0/30 next-hop 10.0.0.2
set routing-options rib inet6.0 static route 3ffe::1:2/126 next-hop 5555::2
set protocols bfd traceoptions file bfd
set protocols bfd traceoptions file size 100m
set protocols bfd traceoptions file files 10
set protocols bfd traceoptions flag all
```

Router R1

```
set interfaces ge-1/1/8 unit 0 family inet address 30.30.30.1/30
set interfaces ge-1/1/8 unit 0 family inet6 address 3ffe::1:2/126
set interfaces xe-0/0/0 gigether-options 802.3ad ae0
set interfaces xe-0/0/1 gigether-options 802.3ad ae0
set interfaces xe-0/0/2 gigether-options 802.3ad ae1
set interfaces xe-0/0/3 gigether-options 802.3ad ae1
set interfaces lo0 unit 0 family inet address 10.255.106.102/32
```

```

set interfaces lo0 unit 0 family inet6 address 201:DB8:251::bb:bb:1/126
set interfaces ae0 aggregated-ether-options bfd-liveness-detection minimum-interval
  150
set interfaces ae0 aggregated-ether-options bfd-liveness-detection multiplier 3
set interfaces ae0 aggregated-ether-options bfd-liveness-detection neighbor
  10.255.106.107
set interfaces ae0 aggregated-ether-options bfd-liveness-detection local-address
  10.255.106.102
set interfaces ae0 aggregated-ether-options minimum-links 1
set interfaces ae0 aggregated-ether-options link-speed 10g
set interfaces ae0 aggregated-ether-options lacp passive
set interfaces ae0 unit 0 family inet address 10.0.0.2/30
set interfaces ae1 aggregated-ether-options bfd-liveness-detection minimum-interval
  200
set interfaces ae1 aggregated-ether-options bfd-liveness-detection multiplier 3
set interfaces ae1 aggregated-ether-options bfd-liveness-detection neighbor
  201:DB8:251::aa:aa:1
set interfaces ae1 aggregated-ether-options bfd-liveness-detection local-address
  201:DB8:251::bb:bb:1
set interfaces ae1 aggregated-ether-options minimum-links 1
set interfaces ae1 aggregated-ether-options link-speed 10g
set interfaces ae1 aggregated-ether-options lacp passive
set interfaces ae1 unit 0 family inet6 address 5555::2/126
set routing-options static route 20.20.20.0/30 next-hop 10.0.0.1
set routing-options rib inet6.0 static route 3ffe::1/126 next-hop 5555::1

```

Configuring a Micro BFD Session for Aggregated Ethernet Interfaces

Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see “*Using the CLI Editor in Configuration Mode*” in the *CLI User Guide*.



NOTE: Repeat this procedure for Router R1, modifying the appropriate interface names, addresses, and any other parameters for each router.

To configure a micro BFD session for aggregated Ethernet interfaces on Router R0:

1. Configure the physical interfaces.

```

[edit interfaces]
user@R0# set ge-1/0/1 unit 0 family inet address 20.20.20.1/30
user@R0# set ge-1/0/1 unit 0 family inet6 address 3ffe::1/126
user@R0# set xe-4/0/0 gigether-options 802.3ad ae0
user@R0# set xe-4/0/1 gigether-options 802.3ad ae0
user@R0# set xe-4/1/0 gigether-options 802.3ad ae1
user@R0# set xe-4/1/1 gigether-options 802.3ad ae1

```

2. Configure the loopback interface.

```
[edit interfaces]
user@R0# set lo0 unit 0 family inet address 10.255.106.107/32
user@R0# set lo0 unit 0 family inet6 address 201:DB8:251::aa:aa:1/128
```

3. Configure an IP address on the aggregated Ethernet interface ae0 with either IPv4 or IPv6 addresses, as per your network requirements.

```
[edit interfaces]
user@R0# set ae0 unit 0 family inet address 10.0.0.1/30
```

4. Set the routing option, create a static route, and set the next-hop address.



NOTE: You can configure either an IPv4 or IPv6 static route, depending on your network requirements.

```
[edit routing-options]
user@R0# set nonstop-routing
user@R0# set static route 30.30.30.0/30 next-hop 10.0.0.2
user@R0# set rib inet6.0 static route 3ffe::1:2/126 next-hop 5555::2
```

5. Configure the Link Aggregation Control Protocol (LACP).

```
[edit interfaces]
user@R0# set ae0 aggregated-ether-options lacp active
```

6. Configure BFD for the aggregated Ethernet interface ae0, and specify the minimum interval, local IP address, and the neighbor IP address.

```
[edit interfaces]
user@R0# set ae0 aggregated-ether-options bfd-liveness-detection
  minimum-interval 100
user@R0# set ae0 aggregated-ether-options bfd-liveness-detection multiplier 3
user@R0# set ae0 aggregated-ether-options bfd-liveness-detection neighbor
  10.255.106.102
user@R0# set ae0 aggregated-ether-options bfd-liveness-detection local-address
  10.255.106.107
user@R0# set ae0 aggregated-ether-options minimum-links 1
user@R0# set ae0 aggregated-ether-options link-speed 10g
```

7. Configure an IP address on the aggregated Ethernet interface ae1.
You can assign either IPv4 or IPv6 addresses as per your network requirements.

```
[edit interfaces]
user@R0# set ae1 unit 0 family inet6 address 5555::1/126
```

8. Configure BFD for the aggregated Ethernet interface ae1.

```
[edit interfaces]
user@R0# set ae1 aggregated-ether-options bfd-liveness-detection
  minimum-interval 100
user@R0# set ae1 aggregated-ether-options bfd-liveness-detection multiplier 3
user@R0# set ae1 aggregated-ether-options bfd-liveness-detection neighbor
  201:DB8:251::bb:bb:1
user@R0# set ae1 aggregated-ether-options bfd-liveness-detection local-address
  201:DB8:251::aa:aa:1
user@R0# set ae1 aggregated-ether-options minimum-links 1
user@R0# set ae1 aggregated-ether-options link-speed 10g
```



NOTE: Beginning with Junos OS Release 16.1, you can also configure this feature with the AE interface address as the local address in a micro BFD session.

Beginning with Release 16.1R2, Junos OS checks and validates the configured micro BFD local-address against the interface or loopback IP address before the configuration commit. Junos OS performs this check on both IPv4 and IPv6 micro BFD address configurations, and if they do not match, the commit fails.

9. Configure tracing options for BFD for troubleshooting.

```
[edit protocols]
user@R0# set bfd traceoptions file bfd
user@R0# set bfd traceoptions file size 100m
user@R0# set bfd traceoptions file files 10
user@R0# set bfd traceoptions flag all
```

Results

From configuration mode, enter the **show interfaces**, **show protocols**, and **show routing-options** commands and confirm your configuration. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R0> show interfaces
traceoptions {
  flag bfd-events;
}
ge-1/0/1 {
  unit 0 {
    family inet {
      address 20.20.20.1/30;
    }
    family inet6 {
```

```
        address 3ffe::1:1/126;
    }
}
xe-4/0/0 {
    enable;
    gigether-options {
        802.3ad ae0;
    }
}
xe-4/0/1 {
    gigether-options {
        802.3ad ae0;
    }
}
xe-4/1/0 {
    enable;
    gigether-options {
        802.3ad ae1;
    }
}
xe-4/1/1 {
    gigether-options {
        802.3ad ae1;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 10.255.106.107/32;
        }
        family inet6 {
            address 201:DB8:251::aa:aa:1/128;
        }
    }
}
ae0 {
    aggregated-ether-options {
        bfd-liveness-detection {
            minimum-interval 100;
            neighbor 10.255.106.102;
            local-address 10.255.106.107;
        }
        minimum-links 1;
        link-speed 10g;
        lacp {
            active;
        }
    }
    unit 0 {
        family inet {
            address 10.0.0.1/30;
        }
    }
}
```



```

ae1 {
  aggregated-ether-options {
    bfd-liveness-detection {
      minimum-interval 100;
      multiplier 3;
      neighbor 201:DB8:251::bb:bb:1;
      local-address 201:DB8:251::aa:aa:1;
    }
    minimum-links 1
    link-speed 10g;
  }
  unit 0 {
    family inet6 {
      address 5555::1/126;
    }
  }
}

```

```

user@R0> show protocols
bfd {
  traceoptions {
    file bfd size 100m files 10;
    flag all;
  }
}

```

```

user@R0> show routing-options
nonstop-routing ;
rib inet6.0 {
  static {
    route 3ffe:1:2/126 {
      next-hop 5555::2;
    }
  }
}
static {
  route 30.30.30.0/30 {
    next-hop 10.0.0.2;
  }
}

```

If you are done configuring the device, commit the configuration.

```
user@R0# commit
```

Verification

Confirm that the configuration is working properly.

- [Verifying That the Independent BFD Sessions Are Up on page 100](#)
- [Viewing Detailed BFD Events on page 102](#)

Verifying That the Independent BFD Sessions Are Up

Purpose Verify that the micro BFD sessions are up, and view details about the BFD sessions.

Action From operational mode, enter the **show bfd session extensive** command.

```
user@R0> show bfd session extensive
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
10.255.106.102	Up	xe-4/0/0	9.000	3.000	3

```
Client LACPD, TX interval 0.100, RX interval 0.100
Session up time 4d 23:13, previous down time 00:00:06
Local diagnostic None, remote diagnostic None
Remote heard, hears us, version 1
Replicated
Session type: Micro BFD
Min async interval 0.100, min slow interval 1.000
Adaptive async TX interval 0.100, RX interval 0.100
Local min TX interval 0.100, minimum RX interval 0.100, multiplier 3
Remote min TX interval 3.000, min RX interval 3.000, multiplier 3
Local discriminator 21, remote discriminator 75
Echo mode disabled/inactive
Remote is control-plane independent
Session ID: 0x0
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
10.255.106.102	Up	xe-4/0/1	9.000	3.000	3

```
Client LACPD, TX interval 0.100, RX interval 0.100
Session up time 4d 23:13, previous down time 00:00:07
Local diagnostic None, remote diagnostic None
Remote heard, hears us, version 1
Replicated
Session type: Micro BFD
Min async interval 0.100, min slow interval 1.000
Adaptive async TX interval 0.100, RX interval 0.100
Local min TX interval 0.100, minimum RX interval 0.100, multiplier 3
Remote min TX interval 3.000, min RX interval 3.000, multiplier 3
Local discriminator 19, remote discriminator 74
Echo mode disabled/inactive
Remote is control-plane independent
Session ID: 0x0
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
201:DB8:251::bb:bb:1 3	Up	xe-4/1/1	9.000	3.000	3

```
Client LACPD, TX interval 0.100, RX interval 0.100
Session up time 4d 23:13
Local diagnostic None, remote diagnostic None
Remote not heard, hears us, version 1
Replicated
Session type: Micro BFD
Min async interval 0.100, min slow interval 1.000
Adaptive async TX interval 0.100, RX interval 0.100
Local min TX interval 1.000, minimum RX interval 0.100, multiplier 3
Remote min TX interval 3.000, min RX interval 3.000, multiplier 3
Local discriminator 17, remote discriminator 67
Echo mode disabled/inactive, no-absorb, no-refresh
Remote is control-plane independent
Session ID: 0x0
```

```

Address          State      Interface    Detect    Transmit
201:DB8:251::bb:bb:1  UP        xe-4/1/0     Time     Interval Multiplier
3
Client LACPD, TX interval 0.100, RX interval 0.100
Session up time 4d 23:13
Local diagnostic None, remote diagnostic None
Remote not heard, hears us, version 1
Replicated
Session type: Micro BFD
Min async interval 0.100, min slow interval 1.000
Adaptive async TX interval 0.100, RX interval 0.100
Local min TX interval 1.000, minimum RX interval 0.100, multiplier 3
Remote min TX interval 3.000, min RX interval 3.000, multiplier 3
Local discriminator 16, remote discriminator 66
Echo mode disabled/inactive, no-absorb, no-refresh
Remote is control-plane independent
Session ID: 0x0

4 sessions, 4 clients
Cumulative transmit rate 2.0 pps, cumulative receive rate 1.7 pps

```

Meaning The Micro BFD field represents the independent micro BFD sessions running on the links in a LAG. The TX interval *item*, RX interval *item* output represents the setting configured with the **minimum-interval** statement. All of the other output represents the default settings for BFD. To modify the default settings, include the optional statements under **bfd-liveness-detection** statement.

Viewing Detailed BFD Events

Purpose View the contents of the BFD trace file to assist in troubleshooting, if required.

Action From operational mode, enter the **file show /var/log/bfd** command.

```

user@R0> file show /var/log/bfd
Jun  5 00:48:59 Protocol (1) len 1: BFD
Jun  5 00:48:59 Data (9) len 41: (hex) 42 46 44 20 6e 65 69 67 68 62 6f 72 20
31 30 2e 30 2e 30
Jun  5 00:48:59 PPM Trace: BFD neighbor 10.255.106.102 (IFL 349) set, 9 0
Jun  5 00:48:59 Received Downstream RcvPkt (19) len 108:
Jun  5 00:48:59 IfIndex (3) len 4: 329
Jun  5 00:48:59 Protocol (1) len 1: BFD
Jun  5 00:48:59 SrcAddr (5) len 8: 10.255.106.102
Jun  5 00:48:59 Data (9) len 24: (hex) 00 88 03 18 00 00 00 4b 00 00 00 15 00
2d c6 c0 00 2d c6
Jun  5 00:48:59 PktError (26) len 4: 0
Jun  5 00:48:59 RtbIdx (24) len 4: 0
Jun  5 00:48:59 MultiHop (64) len 1: (hex) 00
Jun  5 00:48:59 Unknown (168) len 1: (hex) 01
Jun  5 00:48:59 Unknown (171) len 2: (hex) 02 3d
Jun  5 00:48:59 Unknown (172) len 6: (hex) 80 71 1f c7 81 c0
Jun  5 00:48:59 Authenticated (121) len 1: (hex) 01

```

```

Jun  5 00:48:59 BFD packet from 10.0.0.2 (IFL 329), len 24
Jun  5 00:48:59     Ver 0, diag 0, mult 3, len 24
Jun  5 00:48:59     Flags: IHU Fate
Jun  5 00:48:59     My discr 0x0000004b, your discr 0x00000015
Jun  5 00:48:59     Tx ivl 3000000, rx ivl 3000000, echo rx ivl 0
Jun  5 00:48:59 [THROTTLE]bfd_rate_limit_can_accept_pkt: session 10.255.106.102
    is up or already in program thread
Jun  5 00:48:59 Replicate: marked session (discr 21) for update

```

Meaning BFD messages are being written to the specified trace file.

- Related Documentation**
- [authentication on page 506](#)
 - [bfd-liveness-detection on page 507](#)
 - [detection-time on page 509](#)
 - [Configuring Independent Micro BFD Sessions for LAG on page 87](#)

Configuring BFD for PIM

The Bidirectional Forwarding Detection (BFD) Protocol is a simple hello mechanism that detects failures in a network. BFD works with a wide variety of network environments and topologies. A pair of routing devices exchanges BFD packets. Hello packets are sent at a specified, regular interval. A neighbor failure is detected when the routing device stops receiving a reply after a specified interval. The BFD failure detection timers have shorter time limits than the Protocol Independent Multicast (PIM) hello hold time, so they provide faster detection.

The BFD failure detection timers are adaptive and can be adjusted to be faster or slower. The lower the BFD failure detection timer value, the faster the failure detection and vice versa. For example, the timers can adapt to a higher value if the adjacency fails (that is, the timer detects failures more slowly). Or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (Rx) interval by two if the local BFD instance is the reason for the session flap. The transmission (Tx) interval is increased by two if the remote BFD instance is the reason for the session flap. You can use the **clear bfd adaptation** command to return BFD interval timers to their configured values. The **clear bfd adaptation** command is hitless, meaning that the command does not affect traffic flow on the routing device.

You must specify the minimum transmit and minimum receive intervals to enable BFD on PIM.

To enable failure detection:

1. Configure the interface globally or in a routing instance.

This example shows the global configuration.

```
[edit protocols pim]
user@host# edit interface fe-1/0/0.0 family inet bfd-liveness-detection
```

2. Configure the minimum transmit interval.

This is the minimum interval after which the routing device transmits hello packets to a neighbor with which it has established a BFD session. Specifying an interval smaller than 300 ms can cause undesired BFD flapping.

```
[edit protocols pim interface fe-1/0/0.0 family inet bfd-liveness-detection]
user@host# set transmit-interval 350
```

3. Configure the minimum interval after which the routing device expects to receive a reply from a neighbor with which it has established a BFD session.

Specifying an interval smaller than 300 ms can cause undesired BFD flapping.

```
[edit protocols pim interface fe-1/0/0.0 family inet bfd-liveness-detection]
user@host# set minimum-receive-interval 350
```

4. (Optional) Configure other BFD settings.

As an alternative to setting the receive and transmit intervals separately, configure one interval for both.

```
[edit protocols pim interface fe-1/0/0.0 family inet bfd-liveness-detection]
user@host# set minimum-interval 350
```

5. Configure the threshold for the adaptation of the BFD session detection time.

When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

```
[edit protocols pim interface fe-1/0/0.0 family inet bfd-liveness-detection]
user@host# set detection-time threshold 800
```

6. Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.

```
[edit protocols pim interface fe-1/0/0.0 family inet bfd-liveness-detection]
user@host# set multiplier 50
```

7. Configure the BFD version.

```
[edit protocols pim interface fe-1/0/0.0 family inet bfd-liveness-detection]
user@host# set version 1
```

8. Specify that BFD sessions should not adapt to changing network conditions.

We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

```
[edit protocols pim interface fe-1/0/0.0 family inet bfd-liveness-detection]  
user@host# set no-adaptation
```

9. Verify the configuration by checking the output of the **show bfd session** command.

Related Documentation

- *show bfd session*

Enabling Dedicated and Real-Time BFD

Starting with Junos OS Release 15.1X49-D100, SRX340, SRX345, and SRX1500 devices support dedicated BFD.

Starting with Junos OS Release 15.1X49-D100, SRX300 and SRX320 devices support real-time BFD.

Starting with Junos OS Release 15.1X49-D110, SRX550M devices support distributed BFD.

Starting with Junos OS Release 12.3X48-D60, dedicated BFD is supported on SRX240, SRX550, and SRX650 devices.

Starting with Junos OS Release 12.3X48-D60, real-time BFD is supported on SRX100, SRX110, SRX210, and SRX220 devices.

- SRX Series devices support the following maximum number of BFD sessions supported are:
 - Up to four sessions on SRX100, SRX110, SRX210, SRX220, SRX300, and SRX320 devices
 - Up to 50 sessions on SRX240, SRX340, SRX345, SRX550, SRX550M, and SRX650 devices
 - Up to 120 sessions on SRX1500 devices
 - The supported failure detection interval has improved

To enable dedicated BFD on the SRX100, SRX110, SRX210, SRX220, SRX240, SRX300, SRX320,, SRX340, SRX345, SRX550, SRX550M, SRX650, and SRX1500 devices:

1. Include the **dedicated-ukern-cpu** statement at the **[edit chassis]** hierarchy level and then commit the configuration.

[edit]

```
user@host# set chassis dedicated-ukern-cpu
```

```
user@host# commit
```

The following warning message to reboot the system displays when you commit the configuration:

warning: Packet processing throughput may be impacted in dedicated-ukernel-cpu mode.

warning: A reboot is required for dedicated-ukernel-cpu mode to be enabled. Please use "request system reboot" to reboot the system.

commit complete

2. Reboot the device to enable the configuration:

```
user@host> request system reboot
```


3. Verify that dedicated BFD is enabled.

```
user@host> show chassis dedicated-ukern-cpu
```

Dedicated Ukern CPU Status: Enabled

To enable real-time BFD on the SRX100, SRX110, SRX210, SRX220, SRX240, SRX300, SRX320, SRX340, SRX345, SRX550, SRX550M, and SRX650 devices:

1. Include the **realtime-ukern-thread** statement at the **[edit chassis]** hierarchy level and then commit the configuration.

```
[edit]
```

```
user@host# set chassis realtime-ukern-thread
```

```
user@host# commit
```

The following warning message to reboot the system displays when you commit the configuration:

WARNING: realtime-ukern-thread is enable. Please use the command request system reboot.

2. Reboot the device to enable the configuration:

```
user@host> request system reboot
```

3. Verify that real-time BFD is enabled.

```
user@host> show chassis realtime-ukern-thread
```

realtime Ukern thread Status: Enabled

Related Documentation

- [Understanding BFD for BGP on page 36](#)
- [Understanding Distributed BFD on page 47](#)
- [show chassis dedicated-ukern-cpu on page 693](#)
- [show chassis realtime-ukern-thread on page 694](#)

PART 4

Configuring Routing Engine Redundancy

- [Understanding How Routing Engine Redundancy Prevents Network Failures on page 111](#)
- [Configuring Routing Engine Redundancy on page 119](#)

CHAPTER 6

Understanding How Routing Engine Redundancy Prevents Network Failures

- [Understanding Routing Engine Redundancy on Juniper Networks Routers on page 111](#)

Understanding Routing Engine Redundancy on Juniper Networks Routers

This topic contains the following sections:

- [Routing Engine Redundancy Overview on page 111](#)
- [Conditions That Trigger a Routing Engine Failover on page 112](#)
- [Default Routing Engine Redundancy Behavior on page 113](#)
- [Routing Engine Redundancy on a TX Matrix Router on page 114](#)
- [Routing Engine Redundancy on a TX Matrix Plus Router on page 115](#)
- [Situations That Require You to Halt Routing Engines on page 116](#)

Routing Engine Redundancy Overview

Redundant Routing Engines are two Routing Engines that are installed in the same routing platform. One functions as the master, while the other stands by as a backup should the master Routing Engine fail. On routing platforms with dual Routing Engines, network reconvergence takes place more quickly than on routing platforms with a single Routing Engine.

When a Routing Engine is configured as master, it has full functionality. It receives and transmits routing information, builds and maintains routing tables, communicates with interfaces and Packet Forwarding Engine components, and has full control over the chassis. When a Routing Engine is configured to be the backup, it does not communicate with the Packet Forwarding Engine or chassis components.



NOTE: On devices running Junos OS Release 8.4 or later, both Routing Engines cannot be configured to be master at the same time. This configuration causes the commit check to fail.

A failover from the master Routing Engine to the backup Routing Engine occurs automatically when the master Routing Engine experiences a hardware failure or when

you have configured the software to support a change in mastership based on specific conditions. You can also manually switch Routing Engine mastership by issuing one of the **request chassis routing-engine** commands. In this topic, the term *failover* refers to an automatic event, whereas *switchover* refers to either an automatic or a manual event.

When a failover or a switchover occurs, the backup Routing Engine takes control of the system as the new master Routing Engine.

- If graceful Routing Engine switchover is not configured, when the backup Routing Engine becomes master, it resets the switch plane and downloads its own version of the microkernel to the Packet Forwarding Engine components. Traffic is interrupted while the Packet Forwarding Engine is reinitialized. All kernel and forwarding processes are restarted.
- If graceful Routing Engine switchover is configured, interface and kernel information is preserved. The switchover is faster because the Packet Forwarding Engines are not restarted. The new master Routing Engine restarts the routing protocol process (rpd). All hardware and interfaces are acquired by a process that is similar to a warm restart. For more information about graceful Routing Engine switchover, see [“Understanding Graceful Routing Engine Switchover” on page 135](#).
- If graceful Routing Engine switchover and nonstop active routing (NSR) are configured, traffic is not interrupted during the switchover. Interface, kernel, and routing protocol information is preserved. For more information about nonstop active routing, see [“Nonstop Active Routing Concepts” on page 199](#).
- If graceful Routing Engine switchover and graceful restart are configured, traffic is not interrupted during the switchover. Interface and kernel information is preserved. Graceful restart protocol extensions quickly collect and restore routing information from the neighboring routers. For more information about graceful restart, see [“Graceful Restart Concepts” on page 233](#).

Conditions That Trigger a Routing Engine Failover

The following events can result in an automatic change in Routing Engine mastership, depending on your configuration:

- The routing platform experiences a hardware failure. A change in Routing Engine mastership occurs if either the Routing Engine or the associated host module or subsystem is abruptly powered off. You can also configure the backup Routing Engine to take mastership if it detects a hard disk error on the master Routing Engine. To enable this feature, include the **failover on-disk-failure** statement at the **[edit chassis redundancy]** hierarchy level.
- The routing platform experiences a software failure, such as a kernel crash or a CPU lock. You must configure the backup Routing Engine to take mastership when it detects a loss of keepalive signal. To enable this failover method, include the **failover on-loss-of-keepalives** statement at the **[edit chassis redundancy]** hierarchy level.
- The routing platform experiences an em0 interface failure on the master Routing Engine. You must configure the backup Routing Engine to take mastership when it

detects the em0 interface failure. To enable this failover method, include the **on-re-to-fpc-stale** statement at the **[edit chassis redundancy failover]** hierarchy level.

- A specific software process fails. You can configure the backup Routing Engine to take mastership when one or more specified processes fail at least four times within 30 seconds. Include the **failover other-routing-engine** statement at the **[edit system processes process-name]** hierarchy level.

If any of these conditions is met, a message is logged and the backup Routing Engine attempts to take mastership. By default, an alarm is generated when the backup Routing Engine becomes active. After the backup Routing Engine takes mastership, it continues to function as master even after the originally configured master Routing Engine has successfully resumed operation. You must manually restore it to its previous backup status. (However, if at any time one of the Routing Engines is not present, the other Routing Engine becomes master automatically, regardless of how redundancy is configured.)

Default Routing Engine Redundancy Behavior

By default, Junos OS uses **re0** as the master Routing Engine and **re1** as the backup Routing Engine. Unless otherwise specified in the configuration, **re0** always becomes master when the acting master Routing Engine is rebooted.



NOTE: A single Routing Engine in the chassis always becomes the master Routing Engine even if it was previously the backup Routing Engine.

Perform the following steps to see how the default Routing Engine redundancy setting works:

1. Ensure that **re0** is the master Routing Engine.
2. Manually switch the state of Routing Engine mastership by issuing the **request chassis routing-engine master switch** command from the master Routing Engine. **re0** is now the backup Routing Engine and **re1** is the master Routing Engine.



NOTE: On the next reboot of the master Routing Engine, Junos OS returns the router to the default state because you have not configured the Routing Engines to maintain this state after a reboot.

3. Reboot the master Routing Engine **re1**.

The Routing Engine boots up and reads the configuration. Because you have not specified in the configuration which Routing Engine is the master, **re1** uses the default configuration as the backup. Now both **re0** and **re1** are in a backup state. Junos OS detects this conflict and, to prevent a no-master state, reverts to the default configuration to direct **re0** to become master.

Routing Engine Redundancy on a TX Matrix Router

In a routing matrix, all master Routing Engines in the TX Matrix router and connected T640 routers must run the same Junos OS release. Likewise, all backup Routing Engines in a routing matrix must run the same Junos OS release. When you run the same Junos OS release on all master and backup Routing Engines in a routing matrix, a change in mastership to any backup Routing Engine in the routing matrix does not cause a change in mastership in any other chassis in the routing matrix.



CAUTION: (Routing matrix based on the TX Matrix or TX Matrix Plus routers only) Within the routing matrix, we recommend that all Routing Engines run the same Junos OS release. If you run different releases on the Routing Engines and a change in mastership occurs on any backup Routing Engine in the routing matrix based on TX Matrix router or TX Matrix Plus router, one or all routers might become logically disconnected from the TX Matrix router or the TX Matrix Plus router and cause data loss.

If the same Junos OS release is not running on all master and backup Routing Engines in the routing matrix, the following consequences occur when the **failover on-loss-of-keepalives** statement is included at the **[edit chassis redundancy]** hierarchy level:

- When the **failover on-loss-of-keepalives** statement is included at the **[edit chassis redundancy]** hierarchy level and you or a host subsystem initiates a change in mastership to the backup Routing Engine in the TX Matrix router, the master Routing Engines in the T640 routers detect a software release mismatch with the new master Routing Engine in the TX Matrix router and switch mastership to their backup Routing Engines.
- When you manually change mastership to a backup Routing Engine in a T640 router using the **request chassis routing-engine master** command, the new master Routing Engine in the T640 router detects a software release mismatch with the master Routing Engine in the TX Matrix router and relinquishes mastership to the original master Routing Engine. (Routing Engine mastership in the TX Matrix router does not switch in this case.)
- When a host subsystem initiates a change in mastership to a backup Routing Engine in a T640 router because the master Routing Engine has failed, the T640 router is logically disconnected from the TX Matrix router. To reconnect the T640 router, initiate a change in mastership to the backup Routing Engine in the TX Matrix router, or replace the failed Routing Engine in the T640 router and switch mastership to it. The replacement Routing Engine must be running the same software release as the master Routing Engine in the TX Matrix router.

If the same Junos OS release is not running on all master and backup Routing Engines in the routing matrix, the following consequences occur when the **failover on-loss-of-keepalives** statement *is not* included at the **[edit chassis redundancy]** hierarchy level:

- If you initiate a change in mastership to the backup Routing Engine in the TX Matrix router, all T640 routers are logically disconnected from the TX Matrix router. To reconnect the T640 routers, switch mastership of all master Routing Engines in the T640 routers to their backup Routing Engines.
- If you initiate a change in mastership to a backup Routing Engine in a T640 router, the T640 router is logically disconnected from the TX Matrix router. To reconnect the T640 router, switch mastership of the new master Routing Engine in the T640 router back to the original master Routing Engine.

Routing Engine Redundancy on a TX Matrix Plus Router

In a routing matrix, all master Routing Engines in the TX Matrix Plus router and the connected LCC must run the same Junos OS release. Likewise, all backup Routing Engines in a routing matrix must run the same Junos OS release. When you run the same Junos OS release on all master and backup Routing Engines in the routing matrix, a change in mastership to any backup Routing Engine in the routing matrix does not cause a change in mastership in any other chassis in the routing matrix.



CAUTION: (Routing matrix based on the TX Matrix or TX Matrix Plus routers only) Within the routing matrix, we recommend that all Routing Engines run the same Junos OS release. If you run different releases on the Routing Engines and a change in mastership occurs on any backup Routing Engine in the routing matrix based on a TX Matrix router or a TX Matrix Plus router, one or all routers might become logically disconnected from the TX Matrix router or the TX Matrix Plus router and cause data loss.

If the same Junos OS release is not running on all master and backup Routing Engines in the routing matrix, the following scenarios occur when the **failover on-loss-of-keepalives** statement *is* included at the **[edit chassis redundancy]** hierarchy level:

- When the **failover on-loss-of-keepalives** statement is included at the **[edit chassis redundancy]** hierarchy level and you or a host subsystem initiates a change in mastership to the backup Routing Engine in the TX Matrix Plus router, the master Routing Engines in the connected LCC detect a software release mismatch with the new master Routing Engine in the TX Matrix Plus router and switch mastership to their backup Routing Engines.
- When you manually change mastership to a backup Routing Engine in a connected LCC by using the **request chassis routing-engine master** command, the new master Routing Engine in the connected LCC detects a software release mismatch with the master Routing Engine in the TX Matrix Plus router and relinquishes mastership to the original master Routing Engine. (Routing Engine mastership in the TX Matrix Plus router does not switch in this case.)

- When a host subsystem initiates a change in mastership to a backup Routing Engine in a connected LCC because the master Routing Engine has failed, the connected LCC is logically disconnected from the TX Matrix Plus router. To reconnect the connected LCC, initiate a change in mastership to the backup Routing Engine in the TX Matrix Plus router, or replace the failed Routing Engine in the connected LCC and switch mastership to it. The replacement Routing Engine must be running the same software release as the master Routing Engine in the TX Matrix Plus router.

If the same Junos OS release is not running on all master and backup Routing Engines in the routing matrix, the following scenarios occur when the **failover on-loss-of-keepalives** statement *is not* included at the **[edit chassis redundancy]** hierarchy level:

- If you initiate a change in mastership to the backup Routing Engine in the TX Matrix Plus router, all connected LCCs are logically disconnected from the TX Matrix Plus router. To reconnect the connected LCC, switch mastership of all master Routing Engines in the connected LCC to their backup Routing Engines.
- If you initiate a change in mastership to a backup Routing Engine in a connected LCC, the connected LCC is logically disconnected from the TX Matrix Plus router. To reconnect the connected LCC, switch mastership of the new master Routing Engine in the connected LCC back to the original master Routing Engine.

Situations That Require You to Halt Routing Engines

Before you shut the power off to a routing platform that has two Routing Engines or before you remove the master Routing Engine, you must first halt the backup Routing Engine and then halt the master Routing Engine. Otherwise, you might need to reinstall Junos OS. You can use the **request system halt both-routing-engines** command on the master Routing Engine, which first shuts down the master Routing Engine and then shuts down the backup Routing Engine. To shut down only the backup Routing Engine, issue the **request system halt** command on the backup Routing Engine.

If you halt the master Routing Engine and do not power it off or remove it, the backup Routing Engine remains inactive unless you have configured it to become the master when it detects a loss of keepalive signal from the master Routing Engine.



NOTE: To restart the router, you must log in to the console port (rather than the Ethernet management port) of the Routing Engine. When you log in to the console port of the master Routing Engine, the system automatically reboots. After you log in to the console port of the backup Routing Engine, press Enter to reboot it.



NOTE: If you have upgraded the backup Routing Engine, first reboot it and then reboot the master Routing Engine.

Related Documentation

- [Understanding High Availability Features on Juniper Networks Routers on page 3](#)

- [Understanding Switching Control Board Redundancy on page 17](#)
- [Configuring Routing Engine Redundancy on page 119](#)

CHAPTER 7

Configuring Routing Engine Redundancy

- [Configuring Routing Engine Redundancy on page 119](#)
- [Initial Routing Engine Configuration Example on page 124](#)
- [Copying a Configuration File from One Routing Engine to the Other on page 125](#)
- [Loading a Software Package from the Other Routing Engine on page 127](#)

Configuring Routing Engine Redundancy

The following sections describe how to configure Routing Engine redundancy:



NOTE: To complete the tasks in the following sections, `re0` and `re1` configuration groups must be defined. For more information about configuration groups, see the *CLI User Guide*.

- [Modifying the Default Routing Engine Mastership on page 119](#)
- [Configuring Automatic Failover to the Backup Routing Engine on page 120](#)
- [Manually Switching Routing Engine Mastership on page 122](#)
- [Verifying Routing Engine Redundancy Status on page 123](#)

Modifying the Default Routing Engine Mastership

For routers with two Routing Engines, you can configure which Routing Engine is the master and which is the backup. By default, the Routing Engine in slot 0 is the master (`re0`) and the one in slot 1 is the backup (`re1`).



NOTE: In systems with two Routing Engines, both Routing Engines cannot be configured to be master at the same time. This configuration causes the commit check to fail.

To modify the default configuration, include the `routing-engine` statement at the `[edit chassis redundancy]` hierarchy level:

```
[edit chassis redundancy]  
routing-engine slot-number (master | backup | disabled);
```

slot-number can be 0 or 1. To configure the Routing Engine to be the master, specify the **master** option. To configure it to be the backup, specify the **backup** option. To disable a Routing Engine, specify the **disabled** option.



NOTE: To switch between the master and the backup Routing Engines, see [“Manually Switching Routing Engine Mastership” on page 122](#).

Configuring Automatic Failover to the Backup Routing Engine

The following sections describe how to configure automatic failover to the backup Routing Engine when certain failures occur on the master Routing Engine.

- [Without Interruption to Packet Forwarding on page 120](#)
- [On Detection of a Hard Disk Error on the Master Routing Engine on page 120](#)
- [On Detection of a Broken LCMD Connectivity Between the VM and RE on page 120](#)
- [On Detection of a Loss of Keepalive Signal from the Master Routing Engine on page 121](#)
- [On Detection of the em0 Interface Failure on the Master Routing Engine on page 122](#)
- [When a Software Process Fails on page 122](#)

Without Interruption to Packet Forwarding

For routers with two Routing Engines, you can configure graceful Routing Engine switchover (GRES). When graceful switchover is configured, socket reconnection occurs seamlessly without interruption to packet forwarding. For information about how to configure graceful Routing Engine switchover, see [“Configuring Graceful Routing Engine Switchover” on page 148](#).

On Detection of a Hard Disk Error on the Master Routing Engine

After you configure a backup Routing Engine, you can direct it to take mastership automatically if it detects a hard disk error from the master Routing Engine. To enable this feature, include the **on-disk-failure** statement at the **[edit chassis redundancy failover]** hierarchy level.

```
[edit chassis redundancy failover]  
on-disk-failure;
```

On Detection of a Broken LCMD Connectivity Between the VM and RE

Set the following configuration that will result in an automatic RE switchover when the LCMD connectivity between VM and RE is broken. To enable this feature, include the **on-loss-of-vm-host-connection** statement at the **[edit chassis redundancy failover]** hierarchy level.

```
[edit chassis redundancy failover]  
on-loss-of-vm-host-connection;
```

If the LCMD process is crashing on the master, the system will switchover after one minute provided the backup RE LCMD connection is stable. The system will not switchover under

the following conditions: if the backup RE LCMD connection is unstable or if the current master just gained mastership. When the master has just gained mastership, the switchover happens only after four minutes.

On Detection of a Loss of Keepalive Signal from the Master Routing Engine

After you configure a backup Routing Engine, you can direct it to take mastership automatically if it detects a loss of keepalive signal from the master Routing Engine.

To enable failover on receiving a loss of keepalive signal, include the **on-loss-of-keepalives** statement at the **[edit chassis redundancy failover]** hierarchy level:

```
[edit chassis redundancy failover]
on-loss-of-keepalives;
```

When graceful Routing Engine switchover is not configured, by default, failover occurs after 300 seconds (5 minutes). You can configure a shorter or longer time interval.



NOTE: The keepalive time period is reset to 360 seconds when the master Routing Engine has been manually rebooted or halted.

To change the keepalive time period, include the **keepalive-time** statement at the **[edit chassis redundancy]** hierarchy level:

```
[edit chassis redundancy]
keepalive-time seconds;
```

The range for **keepalive-time** is 2 through 10,000 seconds.

The following example describes the sequence of events if you configure the backup Routing Engine to detect a loss of keepalive signal in the master Routing Engine:

1. Manually configure a **keepalive-time** of 25 seconds.
2. After the Packet Forwarding Engine connection to the primary Routing Engine is lost and the keepalive timer expires, packet forwarding is interrupted.
3. After 25 seconds of keepalive loss, a message is logged, and the backup Routing Engine attempts to take mastership. An alarm is generated when the backup Routing Engine becomes active, and the display is updated with the current status of the Routing Engine.
4. After the backup Routing Engine takes mastership, it continues to function as master.



NOTE: When graceful Routing Engine switchover is configured, the keepalive signal is automatically enabled and the failover time is set to 2 seconds (4 seconds on M20 routers). You cannot manually reset the keepalive time.



NOTE: When you halt or reboot the master Routing Engine, Junos OS resets the keepalive time to 360 seconds, and the backup Routing Engine does not take over mastership until the 360-second keepalive time period expires.

A former master Routing Engine becomes a backup Routing Engine if it returns to service after a failover to the backup Routing Engine. To restore master status to the former master Routing Engine, you can use the **request chassis routing-engine master switch** operational mode command.

If at any time one of the Routing Engines is not present, the remaining Routing Engine becomes master automatically, regardless of how redundancy is configured.

On Detection of the em0 Interface Failure on the Master Routing Engine

After you configure a backup Routing Engine, you instruct it to take mastership automatically if the em0 interface fails on the master Routing Engine. To enable this feature, include the **on-re-to-fpc-stale** statement at the **[edit chassis redundancy failover]** hierarchy level.

```
[edit chassis redundancy failover]
on-re-to-fpc-stale;
```

When a Software Process Fails

To configure automatic switchover to the backup Routing Engine if a software process fails, include the **failover other-routing-engine** statement at the **[edit system processes process-name]** hierarchy level:

```
[edit system processes process-name]
failover other-routing-engine;
```

process-name is one of the valid process names. If this statement is configured for a process, and that process fails four times within 30 seconds, the router reboots from the other Routing Engine. Another statement available at the **[edit system processes]** hierarchy level is **failover alternate-media**. For information about the alternate media option, see the *Junos OS Administration Library*.

Manually Switching Routing Engine Mastership

To manually switch Routing Engine mastership, use one of the following commands:

- On the backup Routing Engine, request that the backup Routing Engine take mastership by issuing the **request chassis routing-engine master acquire** command.
- On the master Routing Engine, request that the backup Routing Engine take mastership by using the **request chassis routing-engine master release** command.
- On either Routing Engine, switch mastership by issuing the **request chassis routing-engine master switch** command.

Verifying Routing Engine Redundancy Status

A separate log file is provided for redundancy logging at `/var/log/mastership`. To view the log, use the **file show /var/log/mastership** command. [Table 5 on page 123](#) lists the mastership log event codes and descriptions.

Table 5: Routing Engine Mastership Log

Event Code	Description
E_NULL = 0	The event is a null event.
E_CFG_M	The Routing Engine is configured as master.
E_CFG_B	The Routing Engine is configured as backup.
E_CFG_D	The Routing Engine is configured as disabled.
E_MAXTRY	The maximum number of tries to acquire or release mastership was exceeded.
E_REQ_C	A claim mastership request was sent.
E_ACK_C	A claim mastership acknowledgement was received.
E_NAK_C	A claim mastership request was not acknowledged.
E_REQ_Y	Confirmation of mastership is requested.
E_ACK_Y	Mastership is acknowledged.
E_NAK_Y	Mastership is not acknowledged.
E_REQ_G	A release mastership request was sent by a Routing Engine.
E_ACK_G	The Routing Engine acknowledged release of mastership.
E_CMD_A	The command request chassis routing-engine master acquire was issued from the backup Routing Engine.
E_CMD_F	The command request chassis routing-engine master acquire force was issued from the backup Routing Engine.
E_CMD_R	The command request chassis routing-engine master release was issued from the master Routing Engine.
E_CMD_S	The command request chassis routing-engine master switch was issued from a Routing Engine.
E_NO_ORE	No other Routing Engine is detected.
E_TMOUT	A request timed out.

Table 5: Routing Engine Mastership Log (continued)

Event Code	Description
E_NO_IPC	Routing Engine connection was lost.
E_ORE_M	Other Routing Engine state was changed to master.
E_ORE_B	Other Routing Engine state was changed to backup.
E_ORE_D	Other Routing Engine state was changed to disabled.

- Related Documentation**
- [Understanding Routing Engine Redundancy on Juniper Networks Routers on page 111](#)
 - [Understanding Switching Control Board Redundancy on page 17](#)

Initial Routing Engine Configuration Example

You can use configuration groups to ensure that the correct IP addresses are used for each Routing Engine and to maintain a single configuration file for both Routing Engines.

The following example defines configuration groups **re0** and **re1** with separate IP addresses. These well-known configuration group names take effect only on the appropriate Routing Engine.

```
groups {
  re0 {
    system {
      host-name my-re0;
    }
    interfaces {
      fxp0 {
        description "10/100 Management interface";
        unit 0 {
          family inet {
            address 10.255.2.40/24;
          }
        }
      }
    }
  }
  re1 {
    system {
      host-name my-re1;
    }
    interfaces {
      fxp0 {
        description "10/100 Management interface";
        unit 0 {
          family inet {
            address 10.255.2.41/24;
          }
        }
      }
    }
  }
}
```

```

    }
  }
}

```

You can assign an additional IP address to the management Ethernet interface (**fxp0** in this example) on both Routing Engines. The assigned address uses the **master-only** keyword and is identical for both Routing Engines, ensuring that the IP address for the master Routing Engine can be accessed at any time. The address is active only on the master Routing Engine's management Ethernet interface. During a Routing Engine switchover, the address moves over to the new master Routing Engine.

For example, on **re0**, the configuration is:

```

[edit groups re0 interfaces fxp0]
unit 0 {
  family inet {
    address 10.17.40.131/25 {
      master-only;
    }
    address 10.17.40.132/25;
  }
}

```

On **re1**, the configuration is:

```

[edit groups re1 interfaces fxp0]
unit 0 {
  family inet {
    address 10.17.40.131/25 {
      master-only;
    }
    address 10.17.40.133/25;
  }
}

```

For more information about the initial configuration of dual Routing Engines, see the *Software Installation and Upgrade Guide*. For more information about assigning an additional IP address to the management Ethernet interface with the **master-only** keyword on both Routing Engines, see the *CLI User Guide*.

Related Documentation

- [Understanding Routing Engine Redundancy on Juniper Networks Routers on page 111](#)
- [Understanding Switching Control Board Redundancy on page 17](#)

Copying a Configuration File from One Routing Engine to the Other

You can use either the console port or the management Ethernet port to establish connectivity between the two Routing Engines. You can then copy or use FTP to transfer

the configuration from the master to the backup, and load the file and commit it in the normal way.

To connect to the other Routing Engine using the management Ethernet port, issue the following command:

```
user@host> request routing-engine login (other-routing-engine | re0 | re1)
```

On a TX Matrix router, to make connections to the other Routing Engine using the management Ethernet port, issue the following command:

```
user@host> request routing-engine login (backup | lcc number | master |  
other-routing-engine | re0 | re1)
```

For more information about the **request routing-engine login** command, see the [CLI Explorer](#).

To copy a configuration file from one Routing Engine to the other, issue the **file copy** command:

```
user@host> file copy source destination
```

In this case, **source** is the name of the configuration file. These files are stored in the directory **/config**. The active configuration is **/config/juniper.conf**, and older configurations are in **/config/juniper.conf {1...9}**. The **destination** is a file on the other Routing Engine.

The following example copies a configuration file from Routing Engine 0 to Routing Engine 1:

```
user@host> file copy /config/juniper.conf re1:/var/tmp/copied-juniper.conf
```

The following example copies a configuration file from Routing Engine 0 to Routing Engine 1 on a TX Matrix router:

```
user@host> file copy /config/juniper.conf scc-re1:/var/tmp/copied-juniper.conf
```

To load the configuration file, enter the **load replace** command at the **[edit]** hierarchy level:

```
user@host> load replace /var/tmp/copied-juniper.conf
```



CAUTION: Make sure you change any IP addresses specified in the management Ethernet interface configuration on Routing Engine 0 to addresses appropriate for Routing Engine 1.

Related Documentation

- [Understanding Routing Engine Redundancy on Juniper Networks Routers on page 111](#)
- [Understanding Switching Control Board Redundancy on page 17](#)
- [Loading a Software Package from the Other Routing Engine on page 127](#)

Loading a Software Package from the Other Routing Engine

You can load a package from the other Routing Engine onto the local Routing Engine using the existing **request system software add *package-name*** command:

```
user@host> request system software add re(0|1):/filename
```

In the **re** portion of the URL, specify the number of the other Routing Engine. In the ***filename*** portion of the URL, specify the path to the package. Packages are typically in the directory **/var/sw/pkg**.

Related Documentation

- [Understanding Routing Engine Redundancy on Juniper Networks Routers on page 111](#)
- [Understanding Switching Control Board Redundancy on page 17](#)
- [Copying a Configuration File from One Routing Engine to the Other on page 125](#)

CHAPTER 8

Configuring Load Balancing

- [Understanding Aggregated Ethernet Load Balancing on page 129](#)
- [Configuring Adaptive Load Balancing on page 131](#)

Understanding Aggregated Ethernet Load Balancing

The link aggregation feature is used to bundle several physical aggregated Ethernet interfaces to form one logical interface. One or more links are aggregated to form a virtual link or link aggregation group (LAG). The MAC client treats this virtual link as if it were a single link. Link aggregation increases bandwidth, provides graceful degradation as failure occurs, and increases availability.

In addition to these benefits, an aggregated Ethernet bundle is enhanced to provide load-balancing capabilities that ensure that the link utilization among the member links of the aggregated Ethernet bundle are fully and efficiently utilized.

The load-balancing feature allows a device to divide incoming and outgoing traffic along multiple paths or interfaces in order to reduce congestion in the network. Load balancing improves the utilization of various network paths and provides more effective network bandwidth.

Typically, the applications that use load balancing include:

- Aggregated Interfaces (Layer 2)

Aggregated Interfaces (also called AE for aggregated Ethernet, and AS for aggregated SONET) are a Layer 2 mechanism for load-balancing across multiple interfaces between two devices. Because this is a Layer 2 load-balancing mechanism, all of the individual component links must be between the same two devices on each end. Junos OS supports a non-signaled (static) configuration for Ethernet and SONET, as well as the 802.3ad standardized LACP protocol for negotiation over Ethernet links.

- Equal-Cost Multipath (ECMP) (Layer 3)

By default, when there are multiple equal-cost paths to the same destination for the active route, Junos OS uses a hash algorithm to choose one of the next-hop addresses to install in the forwarding table. Whenever the set of next hops for a destination changes in any way, the next-hop address is rechosen using the hash algorithm. There is also an option that allows multiple next-hop addresses to be installed in the forwarding table, known as per-packet load balancing.

ECMP load balancing can be:

- Across BGP paths (BGP multipath)
- Within a BGP path, across multiple LSPs

In complex Ethernet topologies, traffic imbalances occur due to increased traffic flow, and load balancing becomes challenging for some of the following reasons:

- Incorrect load balancing by aggregate next hops
- Incorrect packet hash computation
- Insufficient variance in the packet flow
- Incorrect pattern selection

As a result of traffic imbalance, the load is not well distributed causing congestion in certain links, whereas some other links are not efficiently utilized.

To overcome these challenges, Junos OS provides the following solutions for resolving the genuine traffic imbalance on aggregated Ethernet bundles (IEEE 802.3ad).

- Adaptive Load Balancing

Adaptive load balancing uses a feedback mechanism to correct a genuine traffic imbalance. To correct the imbalance weights, the bandwidth and packet stream of links are adapted to achieve efficient traffic distribution across the links in an AE bundle.

To configure adaptive load balancing, include the **adaptive** statement at the **[edit interfaces aex aggregated-ether-options load-balance]** hierarchy level.



NOTE: Adaptive load balancing is not supported if the VLAN ID is configured on the aggregated Ethernet interface. This limitation affects the PTX Series Packet Transport Routers and QFX10000 switches only.

To configure the tolerance value as a percentage, include the **tolerance** optional keyword at the **[edit interfaces aex aggregated-ether-options load-balance adaptive]** hierarchy level.

To configure adaptive load balancing based on packets per second (instead of the default bits per second setting), include the **pps** optional keyword at the **[edit interfaces aex aggregated-ether-options load-balance adaptive]** hierarchy level.

To configure the scan interval for the hash value based on the sample rate for the last two seconds, include the **scan-interval** optional keyword at the **[edit interfaces aex aggregated-ether-options load-balance adaptive]** hierarchy level.



NOTE: The **pps** and **scan-interval** optional keywords are supported on PTX Series Packet Transport Routers only.

- Per-Packet Random Spray Load Balancing

When the adaptive load-balancing option fails, per-packet random spray load balancing serves as a last resort. It ensures that the members of an AE bundle are equally loaded without taking bandwidth into consideration. Per packet causes packet reordering and hence is recommended only if the applications absorb reordering. Per-packet random spray eliminates traffic imbalance that occurs as a result of software errors, except for packet hash.

To configure per-packet random spray load balancing, include the **per-packet** statement at the **[edit interfaces aex aggregated-ether-options load-balance]** hierarchy level.



NOTE: The Per-Packet option for load balancing is not supported on PTX Series Packet Transport Routers.

The aggregated Ethernet load-balancing solutions are mutually exclusive. When more than one of the load-balancing solutions is configured, the solution that is configured last overrides the previously configured one. You can verify the load-balancing solution being used by issuing the **show interfaces aex aggregated-ether-options load-balance** command.

Related Documentation

- *per-packet*

Configuring Adaptive Load Balancing

This topic describes how to configure adaptive load balancing. Adaptive load balancing maintains efficient utilization of member link bandwidth for an aggregated Ethernet (AE) bundle. Adaptive load balancing uses a feedback mechanism to correct traffic load imbalance by adjusting the bandwidth and packet streams on links within an AE bundle.

Before you begin:

- Configure a set of interfaces with a protocol family and IP address. These interfaces can make up the membership for the AE bundle.
- Create an AE bundle by configuring a set of router interfaces as aggregated Ethernet and with a specific AE group identifier.

To configure adaptive load balancing for an AE bundles:

1. Enable adaptive load balancing on the AE bundle:

```
[edit interfaces ae-x aggregated-ether-options load-balance]
user@router# set adaptive
```



NOTE: To configure adaptive load balancing on aggregated Ethernet bundles with mixed link speeds, use the following statement:

```
user@router# set interfaces ae0 aggregated-ether-options link-speed
mixed load-balance adaptive
```

2. Configure the scan interval value for adaptive load balancing on the AE bundle. The scan interval value determines the length of the traffic scan by multiplying the integer value with a 30-second time period:

```
[edit interfaces ae-x aggregated-ether-options load-balance adaptive]
user@router# set scan-interval multiplier
```

3. Configure the tolerance percentage value. The tolerance value determines the allowed deviation in the traffic rates among the members of the AE bundle before the router triggers an adaptive load balancing update:

```
[edit interfaces ae-x aggregated-ether-options load-balance adaptive]
user@router# set tolerance percentage
```

4. (Optional) Enable packet-per-second-based adaptive load balancing on the AE bundle:

```
[edit interfaces ae-x aggregated-ether-options load-balance adaptive]
user@router# set pps
```

Related Documentation

- [adaptive on page 500](#)

PART 5

Configuring Graceful Routing Engine Switchover (GRES)

- [Understanding How GRES Enables Uninterrupted Packet Forwarding During a Routing Engine Switchover on page 135](#)
- [GRES System Requirements on page 143](#)
- [Configuring GRES on page 147](#)
- [Configuring Ethernet Automatic Protection Switching for High Availability on page 157](#)

CHAPTER 9

Understanding How GRES Enables Uninterrupted Packet Forwarding During a Routing Engine Switchover

- [Understanding Graceful Routing Engine Switchover on page 135](#)

Understanding Graceful Routing Engine Switchover

This topic contains the following sections:

- [Graceful Routing Engine Switchover Concepts on page 135](#)
- [Effects of a Routing Engine Switchover on page 140](#)
- [Graceful Routing Engine Switchover on Aggregated Services interfaces on page 141](#)

Graceful Routing Engine Switchover Concepts

The graceful Routing Engine switchover (GRES) feature in Junos OS enables a router with redundant Routing Engines to continue forwarding packets, even if one Routing Engine fails. GRES preserves interface and kernel information. Traffic is not interrupted. However, GRES does not preserve the control plane.



NOTE: On T Series routers, TX Matrix routers, and TX Matrix Plus routers, the control plane is preserved in case of GRES with nonstop active routing (NSR), and nearly 75 percent of line rate worth of traffic per Packet Forwarding Engine remains uninterrupted during GRES.

Neighboring routers detect that the router has experienced a restart and react to the event in a manner prescribed by individual routing protocol specifications.

To preserve routing during a switchover, GRES must be combined with either:

- Graceful restart protocol extensions
- Nonstop active routing (NSR)

Any updates to the master Routing Engine are replicated to the backup Routing Engine as soon as they occur.



NOTE: Because of its synchronization requirements and logic, NSR/GRES performance is limited by the slowest Routing Engine in the system.

Mastership switches to the backup Routing Engine if:

- The master Routing Engine kernel stops operating.
- The master Routing Engine experiences a hardware failure.
- The administrator initiates a manual switchover.



NOTE: To quickly restore or to preserve routing protocol state information during a switchover, GRES must be combined with either graceful restart or nonstop active routing, respectively. For more information about graceful restart, see [“Graceful Restart Concepts” on page 233](#). For more information about nonstop active routing, see [“Nonstop Active Routing Concepts” on page 199](#).

If the backup Routing Engine does not receive a keepalive from the master Routing Engine after 2 seconds (4 seconds on M20 routers), it determines that the master Routing Engine has failed; and assumes mastership.

The Packet Forwarding Engine:

- Seamlessly disconnects from the old master Routing Engine
- Reconnects to the new master Routing Engine
- Does not reboot
- Does not interrupt traffic

The new master Routing Engine and the Packet Forwarding Engine then become synchronized. If the new master Routing Engine detects that the Packet Forwarding Engine state is not up to date, it resends state update messages.



NOTE: Starting with Junos OS Release 12.2, if adjacencies between the restarting router and the neighboring peer 'helper' routers time out, graceful restart protocol extensions are unable to notify the peer 'helper' routers about the impending restart. Graceful restart can then stop and cause interruptions in traffic.

To ensure that these adjacencies are maintained, change the *hold-time* for IS-IS protocols from the default of 27 seconds to a value higher than 40 seconds.



NOTE: Successive Routing Engine switchover events must be a minimum of 240 seconds (4 minutes) apart after both Routing Engines have come up.

If the router or switch displays a warning message similar to **Standby Routing Engine is not ready for graceful switchover. Packet Forwarding Engines that are not ready for graceful switchover might be reset**, do not attempt switchover. If you choose to proceed with switchover, only the Packet Forwarding Engines that were not ready for graceful switchover are reset. None of the FPCs should spontaneously restart. We recommend that you wait until the warning no longer appears and then proceed with the switchover.



NOTE: Starting from Junos OS Release 14.2, when you perform GRES on MX Series routers, you must execute the `clear synchronous-ethernet wait-to-restore` operational mode command on the new master Routing Engine to clear the wait-to-restore timer on it. This is because the `clear synchronous-ethernet wait-to-restore` operational mode command clears the wait-to-restore timer only on the local Routing Engine.



NOTE: In a routing matrix with TX Matrix Plus router with 3D SIBs, for successive Routing Engine switchover, events must be a minimum of 900 seconds (15 minutes) apart after both Routing Engines have come up.

GRES must be performed on one line-card chassis (LCC) (of a TX Matrix router with 3D SIBs) at a time to avoid synchronization issues.



NOTE:


- We do *not* recommend performing a commit operation on the backup Routing Engine when GRES is enabled on the router or switch.
- We do *not* recommend enabling GRES on the backup Routing Engine in *any* scenario.



NOTE: On QFX10000 switches, we strongly recommend that you configure the `nsr-phantom-holdtime seconds` statement at the [edit routing-options] hierarchy level when nonstop routing is enabled with GRES. Doing so helps to prevent traffic loss. When you configure this statement, phantom IP addresses remain in the kernel during a switchover until the specified hold-time interval expires. After the interval expires, these routes are added to the appropriate routing tables. In an Ethernet VPN (EVPN)/VXLAN environment, we recommend that you specify a hold-time value of 300 seconds (5 minutes).

Figure 6 on page 138 shows the system architecture of graceful Routing Engine switchover and the process a routing platform follows to prepare for a switchover.

Figure 6: Preparing for a Graceful Routing Engine Switchover



NOTE: Check GRES readiness by executing both:


- The **request chassis routing-engine master switch check** command from the master Routing Engine
- The **show system switchover** command from the Backup Routing Engine

The switchover preparation process for GRES is as follows:

1. The master Routing Engine starts.
2. The routing platform processes (such as the chassis process [chassisd]) start.
3. The Packet Forwarding Engine starts and connects to the master Routing Engine.
4. All state information is updated in the system.
5. The backup Routing Engine starts.
6. The system determines whether GRES has been enabled.
7. The kernel synchronization process (ksyncd) synchronizes the backup Routing Engine with the master Routing Engine.
8. After ksyncd completes the synchronization, all state information and the forwarding table are updated.

Figure 7 on page 139 shows the effects of a switchover on the routing (or switching) platform.

Figure 7: Graceful Routing Engine Switchover Process



A switchover process comprises the following steps:

1. When keepalives from the master Routing Engine are lost, the system switches over gracefully to the backup Routing Engine.
2. The Packet Forwarding Engine connects to the backup Routing Engine, which becomes the new master.
3. Routing platform processes that are not part of GRES (such as the routing protocol process rpd) restart.
4. State information learned from the point of the switchover is updated in the system.
5. If configured, graceful restart protocol extensions collect and restore routing information from neighboring peer *helper* routers.



NOTE: For MX Series routers using enhanced subscriber management, the new backup Routing Engine (the former master Routing Engine) will reboot when a graceful Routing Engine switchover is performed. This cold restart resynchronizes the backup Routing Engine state with that of the new master Routing Engine, preventing discrepancies in state that might have occurred during the switchover.



NOTE: During GRES on T Series and M320 routers during GRES, the Switch Interface Boards (SIBs) are taken offline and restarted one by one. This is done to provide the Switch Processor Mezzanine Board (SPMB) that manages the SIB enough time to populate state information for its associated SIB. However, on a fully populated chassis where all FPCs are sending traffic at full line rate, there might be momentary packet loss during the switchover.



NOTE: When GRES is configured and the `restart chassis-control` command is executed on a TX Matrix Plus router with 3D SIBs, you cannot ascertain which Routing Engine becomes the master. This is because the `chassisd` process restarts with the execution of the `restart chassis-control` command. The `chassisd` process is responsible for maintaining and retaining mastership and when it is restarted, the new `chassisd` is processed based on the router or switch load. As a result, any one of the Routing Engines is made the master.

Effects of a Routing Engine Switchover

Table 6 on page 140 describes the effects of a Routing Engine switchover when different features are enabled:

- No high availability features
- Graceful Routing Engine switchover
- Graceful restart
- Nonstop active routing

Table 6: Effects of a Routing Engine Switchover

Feature	Benefits	Considerations
Dual Routing Engines only (no features enabled)	<ul style="list-style-type: none"> • When the switchover to the new master Routing Engine is complete, routing convergence takes place and traffic is resumed. 	<ul style="list-style-type: none"> • All physical interfaces are taken offline. • Packet Forwarding Engines restart. • The backup Routing Engine restarts the routing protocol process (<code>rpd</code>). • All hardware and interfaces are discovered by the new master Routing Engine. • The switchover takes several minutes. • All of the router's adjacencies are aware of the physical (interface alarms) and routing (topology) changes.
GRES enabled	<ul style="list-style-type: none"> • During the switchover, interface and kernel information is preserved. • The switchover is faster because the Packet Forwarding Engines are not restarted. 	<ul style="list-style-type: none"> • The new master Routing Engine restarts the routing protocol process (<code>rpd</code>). • All hardware and interfaces are acquired by a process that is similar to a warm restart. • All adjacencies are aware of the router's change in state.
GRES and NSR enabled	<ul style="list-style-type: none"> • Traffic is not interrupted during the switchover. • Interface and kernel information are preserved. 	<ul style="list-style-type: none"> • Unsupported protocols must be refreshed using the normal recovery mechanisms inherent in each protocol.

Table 6: Effects of a Routing Engine Switchover (continued)

Feature	Benefits	Considerations
GRES and graceful restart enabled	<ul style="list-style-type: none"> Traffic is not interrupted during the switchover. Interface and kernel information are preserved. Graceful restart protocol extensions quickly collect and restore routing information from the neighboring routers. 	<ul style="list-style-type: none"> Neighbors are required to support graceful restart, and a wait interval is required. The routing protocol process (rpd) restarts. For certain protocols, a significant change in the network can cause graceful restart to stop. Starting with Junos OS Release 12.2, if adjacencies between the restarting router and the neighboring peer 'helper' routers time out, graceful restart can stop and cause interruptions in traffic.

Graceful Routing Engine Switchover on Aggregated Services interfaces

If a graceful Routing Engine switchover (GRES) is triggered by an operational mode command, the state of aggregated services interfaces (ASIs) are not preserved. For example:

```
request interface <switchover | revert> asi-interface
```

However, if GRES is triggered by a CLI commit or FPC restart or crash, the backup Routing Engine updates the ASI state. For example:

```
set interface si-x/y/z disable
commit
```

Or:

```
request chassis fpc restart
```

Release History Table

Release	Description
14.2	Starting from Junos OS Release 14.2, when you perform GRES on MX Series routers, you must execute the clear synchronous-ethernet wait-to-restore operational mode command on the new master Routing Engine to clear the wait-to-restore timer on it.
12.2	Starting with Junos OS Release 12.2, if adjacencies between the restarting router and the neighboring peer 'helper' routers time out, graceful restart protocol extensions are unable to notify the peer 'helper' routers about the impending restart.
12.2	Starting with Junos OS Release 12.2, if adjacencies between the restarting router and the neighboring peer 'helper' routers time out, graceful restart can stop and cause interruptions in traffic.

Related Documentation

- [Understanding High Availability Features on Juniper Networks Routers on page 3](#)
- [Graceful Routing Engine Switchover System Requirements on page 143](#)
- [Configuring Graceful Routing Engine Switchover on page 148](#)
- [Configuring Graceful Routing Engine Switchover in a Virtual Chassis on page 151](#)
- [Configuring Graceful Routing Engine Switchover in a Virtual Chassis on page 151](#)
- [Requirements for Routers with a Backup Router Configuration on page 147](#)
- [Example: Configuring IS-IS for GRES with Graceful Restart on page 153](#)
- *hold-time*

GRES System Requirements

- [Graceful Routing Engine Switchover System Requirements on page 143](#)

Graceful Routing Engine Switchover System Requirements

Graceful Routing Engine switchover is supported on all routing (or switching) platforms that contain dual Routing Engines. All Routing Engines configured for graceful Routing Engine switchover must run the same Junos OS release. Hardware and software support for graceful Routing Engine switchover is described in the following sections:

- [Graceful Routing Engine Switchover Platform Support on page 143](#)
- [Graceful Routing Engine Switchover Feature Support on page 144](#)
- [Graceful Routing Engine Switchover DPC Support on page 145](#)
- [Graceful Routing Engine Switchover and Subscriber Access on page 145](#)
- [Graceful Routing Engine Switchover PIC Support on page 146](#)

Graceful Routing Engine Switchover Platform Support

To enable graceful Routing Engine switchover, your system must meet these minimum requirements:

- M20 and M40e routers—Junos OS Release 5.7 or later
- M10i router—Junos OS Release 6.1 or later
- M320 router—Junos OS Release 6.2 or later
- T320 router, T640 router, and TX Matrix router—Junos OS Release 7.0 or later
- M120 router—Junos OS Release 8.2 or later
- MX960 router—Junos OS Release 8.3 or later
- MX480 router—Junos OS Release 8.4 or later (8.4R2 recommended)
- MX240 router—Junos OS Release 9.0 or later
- PTX5000 router—Junos OS Release 12.1X48 or later
- Standalone T1600 router—Junos OS Release 8.5 or later
- Standalone T4000 router—Junos OS Release 12.1R2 or later

- TX Matrix Plus router—Junos OS Release 9.6 or later
- TX Matrix Plus router with 3D SIBs—Junos Release 13.1 or later
- EX Series switches with dual Routing Engines or in a Virtual Chassis — Junos OS Release 9.2 or later for EX Series switches
- QFX Series switches in a Virtual Chassis —Junos OS Release 13.2 or later for the QFX Series
- EX Series or QFX Series switches in a Virtual Chassis Fabric —Junos OS Release 13.2X51-D20 or later for the EX Series and QFX Series switches

For more information about support for graceful Routing Engine switchover, see the sections that follow.

Graceful Routing Engine Switchover Feature Support

Graceful Routing Engine switchover supports most Junos OS features in Release 5.7 and later. Particular Junos OS features require specific versions of Junos OS. See [Table 7 on page 144](#).

Table 7: Graceful Routing Engine Switchover Feature Support

Application	Junos OS Release
Aggregated Ethernet interfaces with Link Aggregation Control Protocol (LACP) and aggregated SONET interfaces	6.2
Asynchronous Transfer Mode (ATM) virtual circuits (VCs)	6.2
Logical systems	6.3
NOTE: In Junos OS Release 9.3 and later, the logical router feature is renamed to logical system.	
Multicast	6.4 (7.0 for TX Matrix router)
Multilink Point-to-Point Protocol (MLPPP) and Multilink Frame Relay (MLFR)	7.0
Automatic Protection Switching (APS)—The current active interface (either the designated working or the designated protect interface) remains the active interface during a Routing Engine switchover.	7.4
Point-to-multipoint Multiprotocol Label Switching MPLS LSPs (transit only)	7.4
Compressed Real-Time Transport Protocol (CRTP)	7.6
Virtual private LAN service (VPLS)	8.2
Ethernet Operation, Administration, and Management (OAM) as defined by IEEE 802.3ah	8.5

Table 7: Graceful Routing Engine Switchover Feature Support (continued)

Application	Junos OS Release
Extended DHCP relay agent	8.5
Ethernet OAM as defined by IEEE 802.1ag	9.0
Packet Gateway Control Protocol (PGCP) process (pgcpd) on Multiservices 500 PICs on T640 routers.	9.0
Subscriber access	9.4
Layer 2 Circuit and LDP-based VPLS pseudowire redundant configuration	9.6

The following constraints apply to graceful Routing Engine switchover feature support:

- When graceful Routing Engine switchover and aggregated Ethernet interfaces are configured in the same system, the aggregated Ethernet interfaces must not be configured for fast-polling LACP. When fast polling is configured, the LACP polls time out at the remote end during the Routing Engine mastership switchover. When LACP polling times out, the aggregated link and interface are disabled. The Routing Engine mastership change is fast enough that standard and slow LACP polling do not time out during the procedure. However, note that this restriction does not apply to MX Series Routers that are running Junos OS Release 9.4 or later and have distributed periodic packet management (PPM) enabled—which is the default configuration—on them. In such cases, you can configure graceful Routing Engine switchover and have aggregated Ethernet interfaces configured for fast-polling LACP on the same device.



NOTE: MACSec sessions will flap upon Graceful Routing Engine switchover.

Starting with Junos OS Release 13.2, when a graceful Routing Engine switchover occurs, the VRRP state does not change. VRRP is supported by graceful Routing Engine switchover only in the case that PPM delegation is enabled (which the default).

Graceful Routing Engine Switchover DPC Support

Graceful Routing Engine switchover supports all Dense Port Concentrators (DPCs) on the MX Series 5G Universal Routing Platforms running the appropriate version of Junos OS as shown in “[Graceful Routing Engine Switchover Platform Support](#)” on page 143. For more information about DPCs, see the *MX Series DPC Guide*.

Graceful Routing Engine Switchover and Subscriber Access

Graceful Routing Engine switchover currently supports most of the features directly associated with dynamic DHCP and dynamic PPPoE subscriber access. Graceful Routing Engine switchover also supports the unified in-service software upgrade (ISSU) for the DHCP access model and the PPPoE access model used by subscriber access.



NOTE: When graceful Routing Engine switchover is enabled for subscriber management, all Routing Engines in the router must have the same amount of DRAM for stable operation.

Graceful Routing Engine Switchover PIC Support

Graceful Routing Engine switchover is supported on most PICs, except for the services PICs listed in this section. The PIC must be on a supported routing platform running the appropriate version of Junos OS. For information about FPC types, FPC/PIC compatibility, and the initial Junos OS Release in which an FPC supported a particular PIC, see the PIC guide for your router platform.

The following constraints apply to graceful Routing Engine switchover support for services PICs:

- You can include the **graceful-switchover** statement at the **[edit chassis redundancy]** hierarchy level on a router with Adaptive Services, Multiservices, and Tunnel Services PICs configured on it and successfully commit the configuration. However, all services on these PICs—except the Layer 2 service packages and extension-provider and SDK applications on Multiservices PICs—are reset during a switchover.
- Graceful Routing Engine switchover is not supported on any Monitoring Services PICs or Multilink Services PICs. If you include the **graceful-switchover** statement at the **[edit chassis redundancy]** hierarchy level on a router with either of these PIC types configured on it and issue the **commit** command, the commit fails.
- Graceful Routing Engine switchover is not supported on Multiservices 400 PICs configured for monitoring services applications. If you include the **graceful-switchover** statement, the commit fails.



NOTE: When an unsupported PIC is online, you cannot enable graceful Routing Engine switchover. If graceful Routing Engine switchover is already enabled, an unsupported PIC cannot come online.

Release History Table

Release	Description
13.2	Starting with Junos OS Release 13.2, when a graceful Routing Engine switchover occurs, the VRRP state does not change.

Related Documentation

- [Understanding High Availability Features on Juniper Networks Routers on page 3](#)
- [Understanding Graceful Routing Engine Switchover on page 135](#)
- [Configuring Graceful Routing Engine Switchover on page 148](#)
- [Configuring Graceful Routing Engine Switchover in a Virtual Chassis on page 151](#)
- [Requirements for Routers with a Backup Router Configuration on page 147](#)

CHAPTER 11

Configuring GRES

- [Requirements for Routers with a Backup Router Configuration on page 147](#)
- [Configuring Graceful Routing Engine Switchover on page 148](#)
- [Configuring Graceful Routing Engine Switchover in a Virtual Chassis on page 151](#)
- [Preventing Graceful Routing Engine Switchover in the Case of Slow Disks on page 152](#)
- [Resetting Local Statistics on page 152](#)
- [Example: Configuring IS-IS for GRES with Graceful Restart on page 153](#)

Requirements for Routers with a Backup Router Configuration

If your Routing Engine configuration includes a **backup-router** statement or an **inet6-backup-router** statement, you can also use the **destination** statement to specify a subnet address or multiple subnet addresses for the backup router. Include destination subnets for the backup Routing Engine at the **[edit system (backup-router | inet6-backup-router) address]** hierarchy level. This requirement also applies to any T640 router connected to a TX Matrix router that includes a **backup-router** or **inet6-backup-router** statement.



NOTE: If you have a backup router configuration in which multiple static routes point to a gateway from the management Ethernet interface, you must configure prefixes that are more specific than the static routes or include the retain flag at the **[edit routing-options static route]** hierarchy level.

For example, if you configure the static route 172.16.0.0/12 from the management Ethernet interface for management purposes, you must specify the backup router configuration as follows:

```
backup-router 172.29.201.62 destination [172.16.0.0/13 172.16.128.0/13]
```

Related Documentation

- [Understanding Graceful Routing Engine Switchover on page 135](#)
- [Graceful Routing Engine Switchover System Requirements on page 143](#)

Configuring Graceful Routing Engine Switchover

This section contains the following topics:

- [Enabling Graceful Routing Engine Switchover on page 148](#)
- [Configuring Graceful Routing Engine Switchover with Graceful Restart on page 148](#)
- [Synchronizing the Routing Engine Configuration on page 148](#)
- [Verifying Graceful Routing Engine Switchover Operation on page 150](#)

Enabling Graceful Routing Engine Switchover

By default, graceful Routing Engine switchover (GRES) is disabled. To configure GRES, include the **graceful-switchover** statement at the **[edit chassis redundancy]** hierarchy level.

```
[edit chassis redundancy]  
graceful-switchover;
```

When you enable GRES, the command-line interface (CLI) indicates which Routing Engine you are using. For example:

```
{master} [edit]  
user@host#
```

To disable GRES, delete the **graceful-switchover** statement from the **[edit chassis redundancy]** hierarchy level.

Configuring Graceful Routing Engine Switchover with Graceful Restart

When using GRES with Graceful Restart, if adjacencies between the Routing Engine and the neighboring peer 'helper' routers time out, graceful restart protocol extensions are unable to notify the peer 'helper' routers about the impending restart. Graceful restart can then stop and cause interruptions in traffic.

To ensure that these adjacencies are kept, change the *hold-time* for IS-IS protocols from the default of 27 seconds to a value higher than 40 seconds.

Synchronizing the Routing Engine Configuration



.....
NOTE: A newly inserted backup Routing Engine automatically synchronizes its configuration with the master Routing Engine configuration.

When you configure GRES, you can bring the backup Routing Engine online after the master Routing Engine is already running. There is no requirement to start the two Routing Engines simultaneously.

.....
Only when you enable the graceful Routing Engine switchover, you can copy the running Junos OS version of the master Routing Engine to the backup Routing Engine.



NOTE: If the system is in ISSU state, you cannot copy the running Junos OS version of the master Router Engine.

Starting in Junos OS release 14.1, you can enable automatic synchronization of the master Routing Engine configuration with the backup Routing Engine by including the events CHASSISD_SNMP_TRAP7 statement at the [edit event-options policy *policy-name*] hierarchy level.

CHASSISD_SNMP_TRAP7 is a system event logging message that the chassis process (chassisd) generates a Simple Network Management Protocol (SNMP) trap with the seven indicated argument-value pairs. An example of an event script to trigger automatic synchronization of master to the backup Routing Engine is as follows:

```
[edit event-options]
policy UPGRADE-BACKUPRE {
events CHASSISD_SNMP_TRAP7;
attributes-match {
CHASSISD_SNMP_TRAP7.value5 matches "Routing Engine";
CHASSISD_SNMP_TRAP7.trap matches "Fru Online";
CHASSISD_SNMP_TRAP7.argument5 matches jnxFruName;
}
then {
event-script auto-image-upgrade.slax {
arguments {
trap "${$.trap}";
value5 "${$.value5}";
argument5 "${$.argument5}";
}
}
}
event-script {
file auto-image-upgrade.slax;
}
```

After receiving this event, the event policy on the master Router Engine is triggered and the image available in the `/var/sw/pkg` path is pushed to the backup Router Engine upgrade. During script execution, the image is copied to the backup Routing Engine's `/var/sw/pkg` path.



NOTE: If the image is not available in the `/var/sw/pkg` path, the script is terminated with an appropriate syslog message.

If the Routing Engine is running at the Junos OS Release 13.2 or later, the Junos automation scripts is synchronized automatically.

After the master Router Engine is rebooted, the event script available at the `/usr/libexec/scripts/event/auto-image-upgrade.slax` must be copied to the `/var/db/scripts/event` path.



NOTE: For MX Series routers using enhanced subscriber management, the new backup Routing Engine (the former master Routing Engine) will reboot when a graceful Routing Engine switchover is performed. This cold restart resynchronizes the backup Routing Engine state with that of the new master Routing Engine, preventing discrepancies in state that might have occurred during the switchover.

Verifying Graceful Routing Engine Switchover Operation

To verify whether GRES is enabled on the backup Routing Engine, issue the **show system switchover** command. When the output of the command indicates that the **Graceful switchover** field is set to **On**, GRES is operational. The status of the kernel database and configuration database synchronization between Routing Engines is also provided. For example:

```
Graceful switchover: On
Configuration database: Ready
Kernel database: Ready
Peer state: Steady state
```



NOTE: You must issue the **show system switchover** command on the backup Routing Engine. This command is not supported on the master Routing Engine.

For more information about the **show system switchover** command, see the [CLI Explorer](#).

Release History Table

Release	Description
14.1	Starting in Junos OS release 14.1, you can enable automatic synchronization of the master Routing Engine configuration with the backup Routing Engine by including the events CHASSISD_SNMP_TRAP7 statement at the [edit event-options policy <i>policy-name</i>] hierarchy level.

Related Documentation

- [Understanding Graceful Routing Engine Switchover on page 135](#)
- [Graceful Routing Engine Switchover System Requirements on page 143](#)
- [Requirements for Routers with a Backup Router Configuration on page 147](#)
- [Resetting Local Statistics on page 152](#)
- [graceful-switchover on page 513](#)
- [graceful-switchover on page 514](#)
- [Example: Configuring IS-IS for GRES with Graceful Restart on page 153](#)
- *hold-time*

Configuring Graceful Routing Engine Switchover in a Virtual Chassis

In a Virtual Chassis, one member switch is assigned the master role and has the master Routing Engine. Another member switch is assigned the backup role and has the backup Routing Engine. Graceful Routing Engine switchover (GRES) enables the master and backup Routing Engines in a Virtual Chassis configuration to switch from the master to backup without interruption to packet forwarding as a hitless failover solution. When you configure graceful Routing Engine switchover, the backup Routing Engine automatically synchronizes with the master Routing Engine to preserve kernel state information and the forwarding state.

To set up the Virtual Chassis configuration to use graceful Routing Engine switchover (GRES):

1. Set up a minimum of two switches in a Virtual Chassis configuration with mastership priority of 255:

```
[edit]
user@switch# set virtual-chassis member 0 mastership-priority 255
[edit]
user@switch# set virtual-chassis member 1 mastership-priority 255
```

2. Set up graceful Routing Engine switchover:

```
[edit]
user@switch# set chassis redundancy graceful-switchover
```

Commit the configuration.



NOTE: We recommend that you use the `commit synchronize` command to save any configuration changes that you make to a multimember Virtual Chassis.

Related Documentation

- *Example: Configuring an EX4200 Virtual Chassis with a Master and Backup in a Single Wiring Closet*
- [High Availability Features for EX Series Switches Overview on page 9](#)
- *Understanding EX Series Virtual Chassis*
- *Understanding QFX Series Virtual Chassis*

Preventing Graceful Routing Engine Switchover in the Case of Slow Disks

Unexpected slow disk access can happen for various reasons—a faulty or bad sector, for example—causing a hold up of the normal operation of processes such as the routing process (rpd). Eventually, the router's performance will be impacted. Under these circumstances, it may take longer for the typical failover mechanism to be triggered.

Juniper Networks has introduced a disk monitoring daemon to solve this dilemma. The daemon detects slow disk access and initiates failover. Failover can minimize the traffic impact and ease the load on the original master Routing Engine for its backlog clean up.

However, there are instances when you might not want failover to occur. You might commit a large set of changes or even minor changes that might lead to a series of updates on the routing topology. Such activity could lead to extensive disk access delay and, therefore, trigger failover. For expected disk access delays like this, where you do not want to trigger failover, you can choose to not have failover occur by setting the **chassis redundancy failover not-on-disk-underperform** configuration command. Another way is to disable the disk monitoring daemon completely by setting the **system processes gstatd disable** command.

To prevent failovers in the case of slow disks in the Routing Engine:

- Set the option for preventing gstatd from initiating failovers in response to slow disks at the **[edit chassis redundancy failover]** hierarchy level.

[edit]

```
user@host# set chassis redundancy failover not-on-disk-underperform
```

Related Documentation

- [not-on-disk-underperform on page 536](#)
- [Understanding Graceful Routing Engine Switchover on page 135](#)

Resetting Local Statistics

When you enable graceful Routing Engine switchover, the master Routing Engine configuration is copied and loaded to the backup Routing Engine. User files, accounting information, and trace options information are not replicated to the backup Routing Engine.

When a graceful Routing Engine switchover occurs, local statistics such as process statistics and networking statistics are displayed as a cumulative value from the time the process first came online. Because processes on the master Routing Engine can start at different times from the processes on the backup Routing Engine, the statistics on the two Routing Engines for the same process might differ. After a graceful Routing Engine switchover, we recommend that you issue the **clear interface statistics (interface-name | all)** command to reset the cumulative values for local statistics. Forwarding statistics are not affected by graceful Routing Engine switchover.

For information about how to use the **clear** command to clear statistics and protocol database information, see the [CLI Explorer](#).



NOTE: The **clear firewall** command cannot be used to clear the Routing Engine filter counters on a backup Routing Engine that is enabled for graceful Routing Engine switchover.

Related Documentation

- [Understanding Graceful Routing Engine Switchover on page 135](#)
- [Configuring Graceful Routing Engine Switchover on page 148](#)

Example: Configuring IS-IS for GRES with Graceful Restart

This example shows how to configure the Routing Engine's graceful restart protocol extensions using the intermediate system to intermediate system (IS-IS) interior gateway protocol (IGP) to successfully enable graceful Routing Engine switchover (GRES) with graceful restart.

- [Requirements on page 153](#)
- [Overview on page 153](#)
- [Configuration on page 154](#)
- [Verification on page 155](#)

Requirements

GRES prevents interruptions in network traffic if the master Routing Engine fails when combined with either:

- Graceful restart
- Nonstop active routing (NSR)

Before you follow the directions here to configure graceful restart, be sure you have enabled GRES, which is disabled by default. See [“Configuring Graceful Routing Engine Switchover” on page 148](#) for more information.

Overview

If adjacencies between the Routing Engine and the neighboring peer 'helper' routers time out, graceful restart protocol extensions are unable to notify the peer 'helper' routers about the impending restart. Graceful restart can then stop and cause interruptions in traffic.

To ensure that these adjacencies are kept, change the hold-time for IS-IS protocols from the default of 27 seconds to a value higher than 40 seconds.

If your system uses the open shortest pathway first (OSPF) protocol instead of IS-IS, see [Example: Configuring OSPF Timers](#) for configuration information.

Configuration

- [Configuring the IS-IS Protocol Hold Time for Graceful Restart on page 154](#)
- [Results on page 155](#)

CLI Quick Configuration

To quickly configure the hold-time, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the different hierarchy levels shown.

Each interface must be set individually, with a value for each level that the routing device operates on. The minimum recommended value of 41 seconds is used in this example, your system may require a higher value based on size and traffic.

Level 1 and level 2 can be set to different values.

[edit protocols]	set protocols isis interface ge-1/2/0 level 1 hold-time 41 set protocols isis interface ge-1/2/0 level 2 hold-time 41
[edit logical-systems logical-system-name]	set protocols isis interface ge-1/2/0 level 1 hold-time 41 set protocols isis interface ge-1/2/0 level 2 hold-time 41
[edit logical-systems logical-system-name routing-instances routing-instance-name]	set protocols isis interface ge-1/2/0 level 1 hold-time 41 set protocols isis interface ge-1/2/0 level 2 hold-time 41
[edit routing-instances routing-instance-name]	set protocols isis interface ge-1/2/0 level 1 hold-time 41 set protocols isis interface ge-1/2/0 level 2 hold-time 41

Configuring the IS-IS Protocol Hold Time for Graceful Restart

Step-by-Step Procedure

To configure the IS-IS hold-time for graceful restart:

1. Locate or set the interfaces.

```
set protocols isis interface interface-name
```

2. Set the network level and the hold-time in seconds for that level.

```
set protocols isis interface interface-name level 1 hold-time 41
```

3. If the routing device functions on more than one level, set the value for the other level.


```
set protocols isis interface interface-name level 2 hold-time 41
```

4. If you are done configuring the routing device, commit the configuration.



NOTE: Repeat the entire configuration on all routing devices in a shared network.

Results

Verification

- [Verifying the IS-IS Protocol Hold Time for Graceful Restart on page 155](#)

Verifying the IS-IS Protocol Hold Time for Graceful Restart

Purpose	Verify that the IS-IS protocol hold-time is set to 41 seconds or greater to ensure that graceful restart is enabled.
Action	Confirm your configuration by entering the show isis adjacency brief command from operational mode. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.
Meaning	A high enough IS-IS protocol hold-time value allows your system configuration to restart and ensures that even if a Routing Engine fails, traffic continues.
Related Documentation	<ul style="list-style-type: none"> • Understanding Graceful Routing Engine Switchover on page 135 • Configuring Graceful Routing Engine Switchover on page 148 • <i>Example: Configuring IS-IS</i> • <i>Example: Configuring OSPF Timers</i> • <i>interface</i> • <i>level</i> • <i>hold-time</i>

Configuring Ethernet Automatic Protection Switching for High Availability

- [Ethernet Automatic Protection Switching Overview on page 157](#)
- [Mapping of CCM Defects to APS Events on page 160](#)
- [Example: Configuring Protection Switching Between Psuedowires on page 161](#)

Ethernet Automatic Protection Switching Overview

Ethernet automatic protection switching (APS) is a linear protection scheme designed to protect VLAN based Ethernet networks.

With Ethernet APS, a protected domain is configured with two paths, a working path and a protection path. Both working and protection paths can be monitored using an Operations Administration Management (OAM) protocol like Connectivity Fault Management (CFM). Normally, traffic is carried on the working path (that is, the working path is the active path), and the protection path is disabled. If the working path fails, its protection status is marked as degraded (DG) and APS switches the traffic to the protection path, then the protection path becomes the active path.

APS uses two modes of operation, linear 1+1 protection switching architecture and linear 1:1 protection switching architecture. The linear 1+1 protection switching architecture operates with either unidirectional or bidirectional switching. The linear 1:1 protection switching architecture operates with bidirectional switching.

In the linear 1+1 protection switching architecture, the normal traffic is copied and fed to both working and protection paths with a permanent bridge at the source of the protected domain. The traffic on the working and protection transport entities is transmitted simultaneously to the sink of the protected domain, where a selection between the working and protection transport entities is made.

In the linear 1:1 protection switching architecture, the normal traffic is transported on either the working path or on the protection path using a selector bridge at the source of the protection domain. The selector at the sink of the protected domain selects the entity that carries the normal traffic.

Unidirectional and Bidirectional Switching

Unidirectional switching utilizes fully independent selectors at each end of the protected domain. Bidirectional switching attempts to configure the two end points with the same bridge and selector settings, even for a unidirectional failure. Unidirectional switching can protect two unidirectional failures in opposite directions on different entities.

Selective and Merging Selectors

In the linear 1:1 protection switching architecture, where traffic is sent only on the active path, there are two different ways in which the egress direction (the direction out of the protected segment) data forwarding can act: selective selectors and merging selectors. A selective selector forwards only traffic that is received from both the paths regardless of which one is currently active. In other words, with a merging selector the selection of the currently active path only affects the ingress direction. Merging selectors minimize the traffic loss during a protection switch, but they do not guarantee the delivery of the data packets in order.

Revertive and Nonrevertive Switching

For revertive switching, traffic is restored to the working path after the conditions causing the switch have cleared.


For nonrevertive switching, traffic is allowed to remain on the protection path even after the conditions causing the switch have cleared.



NOTE: The configuration on both the provider edge (PE) routers have to be either in revertive mode or non-revertive mode.

Protection Switching Between VPWS Pseudowires


Figure 8: Connections Terminating on Single PE



In the scenario diagrammed in Figure 8 on page 158, a Virtual Private Wire Service (VPWS) is provisioned between customer sites A and B using a single pseudowire (layer 2 circuit) in the core network, and two Multiprotocol Label Switching (MPLS) Label Switched Paths (LSPs) are provisioned, one for the working path and the other one for the protection

path. CFM CCM will be used to monitor the status of each LSP. Provider edge routers PE1 and PE2 run G.8031 Ethernet APS to select one of the LSPs as the active path. Once the active path is elected at the source end of the protection group, PE1 forwards traffic from site A to the elected active path. At the sink end of the protection group, PE2 implements a merging selector, meaning it forwards the traffic coming from both the LSPs to the customer site B.

Figure 9: Connections Terminating on a Different PE



In the scenario represented in Figure 9 on page 159, a VPWS is provisioned between customer sites A and B using two pseudowires (layer 2 circuit) in the core network, one for the working path and the other for the protection path. CFM CCM will be used to monitor the status of each pseudowire.

Provider edge router PE1 and MTU run G.8031 Ethernet APS to select one of the pseudowires as the active path. Once the active path is elected at the source end of the protection group, PE1 forwards the traffic from site A to the elected active path. At the sink end of the protection group, MTU implements a merging selector, meaning it forwards the traffic coming from both the pseudowires to customer site B.

CLI Configuration Statements

```
[edit protocols protection-group]
ethernet-aps profile1{
  protocol g8031;
  revert-time seconds;
  hold-time 0-10000ms;
  local-request lockout;
}
```

revert-time- By default, protection logic restores the use of the working path once it recovers. The revert-time statement specifies how much time should elapse before the path for data should be switched from Protection to Working once recovery for Working has occurred. A revert-time of zero indicates no reversion. It will default to 300 sec (5 minutes) if not configured.

hold-time- Once a failure is detected, APS waits until this timer expires before initiating the protection switch. The range of the hold-time timer is 0 to 10,000 milliseconds. It will default to zero if not configured.

local-request- Configuring this value to lockout or force-switch will trigger lockout or force-switch operation on the protection groups using this profile.

- Related Documentation**
- [Mapping of CCM Defects to APS Events on page 160](#)
 - [Example: Configuring Protection Switching Between Psuedowires on page 161](#)

Mapping of CCM Defects to APS Events

The continuity check message (CCM) engine marks the status of working and protected transport entities as either Down, Degraded, or Up.

Down—The monitored path is declared down if any of the following Multiple End Point (MEP) defects occur:

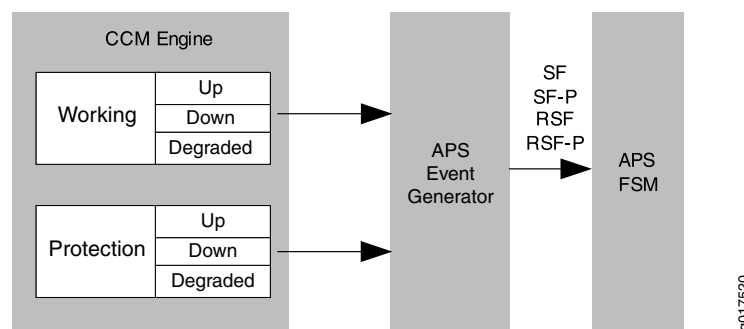
- Interface down
- CCM expiry
- RDI indicating signal failure

Degraded—The monitored path is declared degraded if any of the following MEP defects occur:

- FRR on
- FRR-ACK on

Up—The monitored path is declared up in the absence of any of the above events.

Figure 10: Understanding APS Events



As show in [Figure 10 on page 160](#), the APS event generator generates the following APS events based on the status of the working and protection paths:

- **SF**—Signal failure on working path
- **RSF**—Working path recovers from signal failure
- **SF-P**—Signal failure on protection path
- **RSF-P**—Protection path recovers from signal failure

- Related Documentation**
- [Ethernet Automatic Protection Switching Overview on page 157](#)
 - [Example: Configuring Protection Switching Between Psuedowires on page 161](#)

Example: Configuring Protection Switching Between Psuedowires

- [Requirements on page 161](#)
- [Overview and Topology on page 161](#)
- [Configuration on page 161](#)

Requirements


This example uses the following hardware and software components:

- Junos OS Release 11.2 or later
- 2 MX Series PE routers

Overview and Topology

The physical topology of the protection switching between psuedowires example is shown in [Figure 11 on page 161](#).

Figure 11: Topology of a Network Using VPWS Psuedowires



The following definitions describe the meaning of the device abbreviations used in [Figure 11 on page 161](#).

- Customer edge (CE) device—A device at the customer site that provides access to the service provider's VPN over a data link to one or more provider edge (PE) routers.
- Provider edge (PE) device—A device, or set of devices, at the edge of the provider network that presents the provider's view of the customer site.

Configuration

- Step-by-Step Procedure**
- To configure protection switching between psuedowires, perform these tasks:
1. Configure automatic protection switching.

```
protocols {
  protection-group {
    ethernet-aps {
      profile-1 {
        protocol g8031;
        hold-time 1000s;
        revert-time 5m;
      }
    }
  }
}
```

2. Configure the connectivity fault management.

```
ethernet {
  oam {
    connectivity-fault-management {
      maintenance-domain mdl {
        level 5;
      }
    }
  }
}
```

3. Configure the continuity check message for the working path.

```
maintenance-association W {
  protect maintenance-association P {
    aps-profile profile-1;
  }
  continuity-check {
    interval 1s;
  }
  mep 100 {
    interface ge-1/0/0.0 working;
    direction down;
    auto-discovery;
  }
}
```

4. Configure the continuity check message for the protection path.

```
maintenance-association P {
  continuity-check {
    interval 1s;
  }
  mep 100 {
    interface ge-1/0/0.0 protect;
    direction down;
    auto-discovery;
  }
}
```

Results Check the results of the configuration:


```

protocols {
  protection-group {
    ethernet-aps {
      profile-1 {
        protocol g8031;
        hold-time 1000s;
        revert-time 5m;
      }
    }
  }
}
ethernet {
  oam {
    connectivity-fault-management {
      maintenance-domain md1 {
        level 5;
        maintenance-association W {
          protect maintenance-association P {
            aps-profile profile-1;
          }
          continuity-check {
            interval 1s;
          }
          mep 100 {
            interface ge-1/0/0.0 working;
            direction down;
            auto-discovery;
          }
        }
        maintenance-association P {
          continuity-check {
            interval 1s;
          }
          mep 100 {
            interface ge-1/0/0.0 protect;
            direction down;
            auto-discovery;
          }
        }
      }
    }
  }
}

```

- Related Documentation**
- [Ethernet Automatic Protection Switching Overview on page 157](#)
 - [Mapping of CCM Defects to APS Events on page 160](#)

Configuring Ethernet Ring Protection Switching for High Availability

- [Ethernet Ring Protection Switching Overview on page 165](#)
- [Understanding Ethernet Ring Protection Switching Functionality on page 166](#)
- [Configuring Ethernet Ring Protection Switching on page 173](#)
- [Example: Ethernet Ring Protection Switching Configuration on MX Routers on page 174](#)

Ethernet Ring Protection Switching Overview

Ethernet ring protection switching (ERPS) helps achieve high reliability and network stability. Links in the ring will never form loops that fatally affect the network operation and services availability. The basic idea of an Ethernet ring is to use one specific link to protect the whole ring. This special link is called a *ring protection link (RPL)*. If no failure happens in other links of the ring, the RPL blocks the traffic and is not used. The RPL is controlled by a special node called an *RPL owner*. There is only one RPL owner in a ring. The RPL owner is responsible for blocking traffic over the RPL. Under ring failure conditions, the RPL owner is responsible for unblocking traffic over the RPL. A ring failure results in protection switching of the RPL traffic. An automatic protection switching (APS) protocol is used to coordinate the protection actions over the ring. Protection switching blocks traffic on the failed link and unblocks the traffic on the RPL. When the failure clears, revertive protection switching blocks traffic over the RPL and unblocks traffic on the link on which the failure is cleared.



NOTE: ERPS on AE interfaces is not supported on ACX Series routers except on ACX5000 Series routers.

The following standards provide detailed information on Ethernet ring protection switching:

- ITU-T Recommendation G.8032/Y.1344 version 1 and 2, *Ethernet Ring protection switching*. G.8032v1 supports a single ring topology and G.8032v2 supports multiple rings and ladder topology.



NOTE: EX2300 and EX3400 switches support G.8032v1 only.

- [ITU-T Y.1731, OAM functions and mechanisms for Ethernet-based networks](#)

For additional information on configuring Ethernet ring protection switching on EX Series switches, see *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*.

For additional information on configuring Ethernet ring protection switching on MX Series routers, see the *Layer 2 Configuration Guide* for a complete example of Ethernet rings and information about STP loop avoidance and prevention.

**Related
Documentation**

- [Understanding Ethernet Ring Protection Switching Functionality on page 166](#)
- [Configuring Ethernet Ring Protection Switching on page 173](#)
- [Example: Ethernet Ring Protection Switching Configuration on MX Routers on page 174](#)
- [Example: Configuring Ethernet Ring Protection Switching on EX Series Switches](#)
- [Ethernet Interfaces Feature Guide for Routing Devices](#)

Understanding Ethernet Ring Protection Switching Functionality

- [Acronyms on page 167](#)
- [Ring Nodes on page 167](#)
- [Ring Node States on page 167](#)
- [Default Logging of Basic State Transitions on EX Series Switches on page 168](#)
- [Logical Ring on page 168](#)
- [FDB Flush on page 168](#)
- [Traffic Blocking and Forwarding on page 169](#)
- [RPL Neighbor Node on page 169](#)
- [RAPS Message Blocking and Forwarding on page 169](#)
- [Dedicated Signaling Control Channel on page 171](#)
- [RAPS Message Termination on page 171](#)
- [Revertive and Non-revertive Modes on page 171](#)
- [Multiple Rings on page 171](#)
- [Node ID on page 171](#)
- [Ring ID on page 172](#)
- [Bridge Domains with the Ring Port \(MX Series Routers Only\) on page 172](#)
- [Wait-to-Block Timer on page 172](#)
- [Adding and Removing a Node on page 173](#)

Acronyms

The following acronyms are used in the discussion about Ethernet ring protection switching (ERPS):

- MA—Maintenance association
- MEP—Maintenance association end point
- OAM—Operations, administration, and management (Ethernet ring protection switching uses connectivity fault management daemon)
- FDB—MAC forwarding database
- STP—Spanning Tree Protocol
- RAPS—Ring automatic protection switching
- WTB—Wait to block. Note that WTB is always disabled on EX2300 and EX3400 switches because it is not supported in ERPSv1. Any configuration you make to the WTB setting on EX2300 and EX3400 switches has no effect. The output from the CLI command 'show protection-group ethernet-ring node-state detail' lists a WTB setting but that setting has no effect on EX2300 and EX3400 switches.
- WTR—Wait to restore. Note that on EX2300 and EX3400 switches only, the WTR configuration must be 5-12 minutes.
- RPL—Ring protection link

Ring Nodes

Multiple nodes are used to form a ring. There are two different node types:

- Normal node—The node has no special role on the ring.
- RPL owner node—The node owns the RPL and blocks or unblocks traffic over the RPL.

Ring Node States

The following are the different states for each node of a specific ring:

- init—Not a participant of a specific ring.
- idle—No failure on the ring; the node is performing normally. For a normal node, traffic is unblocked on both ring ports. For the RPL owner or RPL neighbor, traffic is blocked on the ring port that connects to the RPL and unblocked on the other ring port.
- protection—A failure occurred on the ring. For a normal node, traffic is blocked on the ring port that connects to the failing link and unblocked on working ring ports. For the RPL owner, traffic is unblocked on both ring ports if they connect to non-failure links.
- pending—The node is recovering from failure or its state after a **clear** command is used to remove the previous manual command. When a protection group is configured, the node enters the pending state. When a node is in pending state, the WTR or WTB timer will be running. All nodes are in pending state till WTR or WTB timer expiry.

- **force switch**—A force switch is issued. When a force switch is issued on a node in the ring all nodes in the ring will move into the force switch state.



NOTE: EX2300 and EX3400 switches do not support force switch.

- **manual switch**—A manual switch is issued. When a manual switch is issued on a node in the ring all nodes in the ring will move into the manual switch state.



NOTE: EX2300 and EX3400 switches do not support manual switch.

There can be only one RPL owner for each ring. The user configuration must guarantee this, because the APS protocol cannot check this.

Default Logging of Basic State Transitions on EX Series Switches

Starting with Junos OS Release 14.1X53-D15, EX Series switches automatically log basic state transitions for the ERPS protocol. Starting with Junos OS Release 18.2R1, EX2300 and EX3400 switches automatically log basic state transitions for the ERPS protocol. No configuration is required to initiate this logging. Basic state transitions include ERPS interface transitions from up to down, and down to up; and ERPS state transitions from idle to protection, and protection to idle.

The basic state transitions are logged in a single file named **erp-default**, which resides in the **/var/log** directory of the switch. The maximum size of this file is 15 MB.

Default logging for ERPS can capture initial ERPS interface and state transitions, which can help you troubleshoot issues that occur early in the ERPS protocol startup process. However, if more robust logging is needed, you can enable traceoptions for ERPS by entering the **traceoptions** statement in the **[edit protocols protection-group]** hierarchy.

Be aware that for ERPS, only default logging or traceoptions can be active at a time on the switch. That is, default logging for ERPS is automatically enabled and if you enable traceoptions for ERPS, the switch automatically disables default logging. Conversely, if you disable traceoptions for ERPS, the switch automatically enables default logging.

Logical Ring

You can define multiple logical-ring instances on the same physical ring. The logical ring feature currently supports only the physical ring, which means that two adjacent nodes of a ring must be physically connected and the ring must operate on the physical interface, not the VLAN. Multiple ring instances are usually defined with trunk mode ring interfaces.

FDB Flush

When ring protection switching occurs, normally an *FDB flush* is executed. The Ethernet ring control module uses the same mechanism as the STP to trigger the FDB flush. The Ethernet ring control module controls the ring port physical interface's default STP index to execute the FDB flush.



NOTE: Optimized flushing is not supported on EX2300 and EX3400 switches.

Starting with Junos OS Release 14.2, the FDB flush depends on the RAPS messages received on the both the ports of the ring node.

Traffic Blocking and Forwarding

Ethernet ring control uses the same mechanism as the STP to control forwarding or discarding of user traffic. The Ethernet ring control module sets the ring port physical interface default STP index state to forwarding or discarding in order to control user traffic.

RPL Neighbor Node

Starting with Junos OS Release 14.2, ring protection link neighbor nodes are supported. An RPL neighbor node is adjacent to the RPL and is not the RPL owner. If a node is configured with one interface as the protection-link-end and no protection-link-owner is present in its configuration, the node is an RPL neighbor node.



NOTE: RPL neighbor node is not supported on EX2300 and EX3400 switches.

RAPS Message Blocking and Forwarding

The router or switch treats the ring automatic protection switching (RAPS) message the same as it treats user traffic for forwarding RAPS messages between two ring ports. The ring port physical interface default STP index state also controls forwarding RAPS messages between the two ring ports. Other than forwarding RAPS messages between the two ring ports, as shown in [Figure 12 on page 169](#), the system also needs to forward the RAPS message between the CPU (Ethernet ring control module) and the ring port. This type of forwarding does not depend on the ring port physical interfaces' STP index state. The RAPS message is always sent by the router or switch through the ring ports, as shown in [Figure 13 on page 169](#). A RAPS message received from a discarding ring port is sent to the Ethernet ring control module, but is not sent to the other ring port.

Figure 12: Protocol Packets from the Network to the Router





Figure 13: Protocol Packets from the Router or Switch to the Network



Juniper Networks switches and Juniper Networks routers use different methods to achieve these routes.

The switches use forwarding database entries to direct the RAPS messages. The forwarding database entry (keyed by the RAPS multicast address and VLAN) has a composite next hop associated with it—the composite next hop associates the two ring interfaces with the forwarding database entry and uses the split horizon feature to prevent sending the packet out on the interface that it is received on. This is an example of the forwarding database entry relating to the RAPS multicast MAC (a result of the **show ethernet-switching table detail** command):

```
VLAN: v1, Tag: 101, MAC: 01:19:a7:00:00:01, Interface: ERP
Interfaces:                               ge-0/0/9.0, ge-0/0/3.0
Type: Static
Action: Mirror
Nexthop index: 1333
```

The routers use an implicit filter to achieve ERP routes. Each implicit filter binds to a bridge domain. Therefore, the east ring port control channel and the west ring port control channel of a particular ring instance must be configured to the same bridge domain. For each ring port control channel, a filter term is generated to control RAPS message forwarding. The filter number is the same as the number of bridge domains that contain the ring control channels. If a bridge domain contains control channels from multiple rings, the filter related to this bridge domain will have multiple terms and each term will relate to a control channel. The filter has command parts and control-channel related parts, as follows:

- Common terms:

- term 1: if [Ethernet type is not OAM Ethernet type (0x8902)]
 { accept packet }
- term 2: if [source MAC address belongs to this bridge]
 { drop packet, our packet loop through the ring and come back to home }
- term 3: if [destination is the RAPS PDU multicast address(0x01,0x19,0xa7,0x00,0x00,0x01) AND[ring port STP status is DISCARDING]
 { send to CPU }

- Control channel related terms:

- if [destination is the RAPS PDU multicast address(0x01,0x19,0xa7,0x00,0x00,0x01) AND[ring port STP status is FORWARDING] AND [Incoming interface IFL equal to control channel IFL]
 { send packet to CPU and send to the other ring port }
 default term: accept packet.

Dedicated Signaling Control Channel

For each ring port, a dedicated signaling control channel with a dedicated VLAN ID must be configured. In Ethernet ring configuration, only this control logical interface is configured and the underlying physical interface is the physical ring port. Each ring requires that two control physical interfaces be configured. These two logical interfaces must be configured in a bridge domain for routers (or the same VLAN for switches) in order to forward RAPS protocol data units (PDUs) between the two ring control physical interfaces. If the router control channel logical interface is not a trunk port, only control logical interfaces will be configured in ring port configuration. If this router control channel logical interface is a trunk port, in addition to the control channel logical interfaces, a dedicated VLAN ID must be configured for routers. For switches, always specify either a VLAN name or VLAN ID for all links.

RAPS Message Termination

The RAPS message starts from the originating node, travels through the entire ring, and terminates in the originating node unless a failure is present in the ring. The originating node must drop the RAPS message if the source MAC address in the RAPS message belongs to itself. The source MAC address is the node's node ID.

Revertive and Non-revertive Modes

In revertive operation, once the condition causing a switch has cleared, traffic is blocked on the RPL and restored to the working transport entity. In nonrevertive operation, traffic is allowed to use the RPL if it has not failed, even after a switch condition has cleared.



NOTE: Non-revertive mode is not supported on EX2300 and EX3400 switches.

Multiple Rings

The Ethernet ring control module supports multiple rings in each node (two logical interfaces are part of each ring). The ring control module also supports the interconnection of multiple rings. Interconnection of two rings means that two rings might share the same link or share the same node. Ring interconnection is supported only using non-virtual-channel mode. Ring interconnection using virtual channel mode is not supported.



NOTE: Interconnection of multiple rings is not supported on EX2300 and EX3400 switches.

Node ID

For each node in the ring, a unique *node ID* identifies each node. The node ID is the node's MAC address.

For routers only, you can configure this node ID when configuring the ring on the node or automatically select an ID like STP does. In most cases, you will not configure this and the router will select a node ID, like STP does. It should be the manufacturing MAC address. The ring node ID should not be changed, even if you change the manufacturing MAC address. Any MAC address can be used if you make sure each node in the ring has a different node ID. The node ID on switches is selected automatically and is not configurable.

Ring ID

The ring ID is used to determine the value of the last octet of the MAC destination address field of the RAPS protocol data units (PDUs) generated by the ERP control process. The ring ID is also used to discard any RAPS PDU, received by this ERP control process with a non-matching ring ID. Ring ID values 1 through 239 are supported.

Bridge Domains with the Ring Port (MX Series Routers Only)

On the routers, the protection group is seen as an abstract logical port that can be configured to any bridge domain. Therefore, if you configure one ring port or its logical interface in a bridge domain, you must configure the other related ring port or its logical interface to the same bridge domain. The bridge domain that includes the ring port acts as any other bridge domain and supports the IRB Layer 3 interface.

Wait-to-Block Timer

The RPL owner node uses a delay timer before initiating an RPL block in revertive mode of operation or before reverting to IDLE state after clearing manual commands. The Wait-to-Block (WTB) timer is used when clearing **force switch** and **manual switch** commands. As multiple **force switch** commands are allowed to coexist in an Ethernet ring, the WTB timer ensures that clearing of a single **force switch** command does not trigger the re-blocking of the RPL. When clearing a **manual switch** command, the WTB timer prevents the formation of a closed loop due to a possible timing anomaly where the RPL Owner Node receives an outdated remote **manual switch** request during the recovery process.

When recovering from a **manual switch** command, the delay timer must be long enough to receive any latent remote **force switch**, signal failure, or **manual switch** commands. This delay timer is called the WTB timer and is defined to be 5 seconds longer than the guard timer. This delay timer is activated on the RPL Owner Node. When the WTB timer expires, the RPL Owner Node initiates the reversion process by transmitting an RAPS (NR, RB) message. The WTB timer is deactivated when any higher-priority request preempts it.



NOTE: The Wait To Block Timer (WTB) is always disabled on EX2300 and EX3400 switches because it is not supported in ERPSv1. Any configuration you make to the WTB setting has no effect. The output from the CLI command 'show protection-group ethernet-ring node-state detail' lists a WTB setting but that setting has no effect.

Adding and Removing a Node

Starting with Junos OS Release 14.2, you can add or remove a node between two nodes in an Ethernet ring. Nodes are added or removed using the **force switch** command.



NOTE: EX2300 and EX3400 switches do not support force switch.

Release History Table

Release	Description
18.2R1	Starting with Junos OS Release 18.2R1, EX2300 and EX3400 switches automatically log basic state transitions for the ERPS protocol.
14.2	Starting with Junos OS Release 14.2, the FDB flush depends on the RAPS messages received on the both the ports of the ring node.
14.2	Starting with Junos OS Release 14.2, ring protection link neighbor nodes are supported.
14.2	Starting with Junos OS Release 14.2, you can add or remove a node between two nodes in an Ethernet ring.
14.1X53-D15	Starting with Junos OS Release 14.1X53-D15, EX Series switches automatically log basic state transitions for the ERPS protocol.

Related Documentation

- [Ethernet Ring Protection Switching Overview on page 165](#)
- [Configuring Ethernet Ring Protection Switching on page 173](#)
- [Example: Ethernet Ring Protection Switching Configuration on MX Routers on page 174](#)
- [Example: Configuring Ethernet Ring Protection Switching on EX Series Switches](#)
- [Configuring Ethernet Ring Protection Switching on Switches \(CLI Procedure\)](#)

Configuring Ethernet Ring Protection Switching

The inheritance model follows:

```

protection-group {
  ethernet-ring ring-name (
    node-id mac-address;
    ring-protection-link-owner;
    east-interface {
      control-channel channel-name {
        ring-protection-link-end;
      }
    }
    west-interface {
      node-id mac-address;
      control-channel channel-name {

```

```
        ring-protection-link-end;  
    }  
    data-channel {  
        vlan number;  
    }  
    guard-interval number;  
    restore-interval number;  
  }  
}
```

For each ring, a protection group must be configured. There may be several rings in each node, so there should be multiple protection groups corresponding to the related Ethernet rings.

Three interval parameters (**restore-interval**, **guard-interval**, and **hold-interval**) can be configured at the protection group level. These configurations are global configurations and apply to all Ethernet rings if the Ethernet ring doesn't have a more specific configuration for these values. If no parameter is configured at the protection group level, the global configuration of this parameter uses the default value.

Related Documentation

- [Ethernet Ring Protection Switching Overview on page 165](#)
- [Understanding Ethernet Ring Protection Switching Functionality on page 166](#)
- [Example: Ethernet Ring Protection Switching Configuration on MX Routers on page 174](#)
- [Example: Configuring Ethernet Ring Protection Switching on EX Series Switches](#)
- [Ethernet Interfaces Feature Guide for Routing Devices](#)

Example: Ethernet Ring Protection Switching Configuration on MX Routers

This example describes how to configure Ethernet ring protection switching on an MX Series router:

- [Requirements on page 174](#)
- [Ethernet Ring Overview and Topology on page 174](#)
- [Configuring a Three-Node Ring on page 175](#)

Requirements

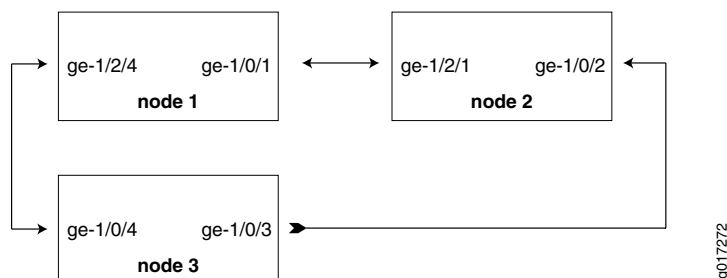
This example uses the following hardware and software components:

- Router node 1 running Junos OS with two Gigabit Ethernet interfaces.
- Router node 2 running Junos OS with two Gigabit Ethernet interfaces.
- Router node 3 running Junos OS with two Gigabit Ethernet interfaces.

Ethernet Ring Overview and Topology

This section describes a configuration example for a three-node ring. The ring topology is shown in [Figure 14 on page 175](#).

Figure 14: Example of a Three-Node Ring Topology



The configuration in this section is only for the RAPS channel. The bridge domain for user traffic is the same as the normal bridge domain. The only exception is if a bridge domain includes a ring port, then it must also include the other ring port of the same ring.

Configuring a Three-Node Ring

To configure Ethernet Ring Protection Switching on a three-node ring, perform these tasks:

- [Configuring Ethernet Ring Protection Switching on a Three-Node Ring on page 175](#)

Configuring Ethernet Ring Protection Switching on a Three-Node Ring

Step-by-Step Procedure

```

1.  interfaces {
    ge-1/0/1 {
        vlan-tagging;
        encapsulation flexible-ethernet-services;
        unit 1 {
            encapsulation vlan-bridge;
            vlan-id 100;
        }
    }
    ge-1/2/4 {
        vlan-tagging;
        encapsulation flexible-ethernet-services;
        unit 1 {
            encapsulation vlan-bridge;
            vlan-id 100;
        }
    }
}
bridge-domains {
    bd1 {
        domain-type bridge;
        interface ge-1/2/4.1;
        interface ge-1/0/1.1;
    }
}
protocols {
    protection-group {
        ethernet-ring pg101 {
            node-id 00:01:01:00:00:01;
            ring-protection-link-owner;
            east-interface {
                control-channel ge-1/0/1.1;
                ring-protection-link-end;
            }
        }
    }
}
  
```

```

    }
    west-interface {
        control-channel ge-1/2/4.1;
    }
}
}
protocols {
    oam {
        ethernet {
            connectivity-fault-management {
                action-profile rmep-defaults {
                    default-action {
                        interface-down;
                    }
                }
                maintenance-domain d1 {
                    level 0;
                    maintenance-association 100 {
                        mep 1 {
                            interface ge-1/0/1;
                            remote-mep 2 {
                                action-profile rmep-defaults;
                            }
                        }
                    }
                }
                maintenance-domain d2 {
                    level 0;
                    maintenance-association 100 {
                        mep 1 {
                            interface ge-1/2/4;
                            remote-mep 2 {
                                action-profile rmep-defaults;
                            }
                        }
                    }
                }
            }
        }
    }
}

```

```
2. interfaces {
    ge-1/0/2 {
        vlan-tagging;
        encapsulation flexible-ethernet-services;
        unit 1 {
            encapsulation vlan-bridge;
            vlan-id 100;
        }
    }

    ge-1/2/1 {
        vlan-tagging;
        encapsulation flexible-ethernet-services;
```

```

        unit 1 {
            encapsulation vlan-bridge;
            vlan-id 100;
        }
    }
}

bridge-domains {
    bd1 {
        domain-type bridge;
        interface ge-1/2/1.1;
        interface ge-1/0/2.1;
    }
}

protocols {
    protection-group {
        ethernet-ring pg102 {
            east-interface {
                control-channel ge-1/0/2.1;
            }
            west-interface {
                control-channel ge-1/2/1.1;
            }
        }
    }
}

protocols {
    oam {
        ethernet {
            connectivity-fault-management {
                action-profile rmep-defaults {
                    default-action {
                        interface-down;
                    }
                }
                maintenance-domain d1 {
                    level 0;
                    maintenance-association 100 {
                        mep 2 {
                            interface ge-1/2/1;
                            remote-mep 1 {
                                action-profile rmep-defaults;
                            }
                        }
                    }
                }
                maintenance-domain d3 {
                    level 0;
                    maintenance-association 100 {
                        mep 1 {
                            interface ge-1/0/2;
                            remote-mep 2 {
                                action-profile rmep-defaults;
                            }
                        }
                    }
                }
            }
        }
    }
}

```

```

    }
  }
}

```

```

3.  interfaces {
    ge-1/0/4 {
      vlan-tagging;
      encapsulation flexible-ethernet-services;
      unit 1 {
        encapsulation vlan-bridge;
        vlan-id 100;
      }
    }

    ge-1/0/3 {
      vlan-tagging;
      encapsulation flexible-ethernet-services;
      unit 1 {
        encapsulation vlan-bridge;
        vlan-id 100;
      }
    }
  }

  bridge-domains {
    bd1 {
      domain-type bridge;
      interface ge-1/0/4.1;
      interface ge-1/0/3.1;
    }
  }

  protocols {
    protection-group {
      ethernet-ring pg103 {
        east-interface {
          control-channel ge-1/0/3.1;
        }
        west-interface {
          control-channel ge-1/0/4.1;
        }
      }
    }
  }

  protocols {
    oam {
      ethernet {
        connectivity-fault-management {
          action-profile rmep-defaults {
            default-action {
              interface-down;
            }
          }
          maintenance-domain d2 {

```



```
level 0;
maintenance-association 100 {
    mep 2 {
        interface ge-1/0/4;
        remote-mep 1 {
            action-profile rmep-defaults;
        }
    }
}
maintenance-domain d3 {
    level 0;
    maintenance-association 100 {
        mep 2 {
            interface ge-1/0/3;
            remote-mep 1 {
                action-profile rmep-defaults;
            }
        }
    }
}
```

Examples: Ethernet RPS Output

This section provides output examples based on the configuration shown in “[Example: Ethernet Ring Protection Switching Configuration on MX Routers](#)” on page 174. The show commands used in these examples can help verify configuration and correct operation.

Normal Situation—RPL Owner Node

If the ring has no failure, the **show** command will have the following output for Node 1:

```
user@node1> show protection-group ethernet-ring aps
```

Ethernet Ring Name	Request/state	No Flush	Ring Protection	Link Blocked
pg101	NR	No	Yes	

```
Originator Remote Node ID
Yes
```

```
user@node1> show protection-group ethernet-ring interface
```

Ethernet ring port parameters for protection group pg101

Interface	Control Channel	Forward State	Ring Protection Link End
ge-1/0/1	ge-1/0/1.1	discarding	Yes
ge-1/2/4	ge-1/2/4.1	forwarding	No

Signal	Failure	Admin State
Clear		IFF ready
Clear		IFF ready

```
user@node1> show protection-group ethernet-ring node-state
```

Ethernet ring pg101	APS State idle	Event NR-RB	Ring Protection Link Owner Yes
------------------------	-------------------	----------------	-----------------------------------

Restore Timer disabled	Quard Timer disabled	Operation state operational
---------------------------	-------------------------	--------------------------------

```
user@node1> show protection-group ethernet-ring statistics group-name pg101
```

```
Ethernet Ring statistics for PG pg101
RAPS sent                : 1
RAPS received             : 0
Local SF happened:       : 0
Remote SF happened:       : 0
NR event happened:        : 0
NR-RB event happened:     : 1
```

Normal Situation—Other Nodes

For Node 2 and Node 3, the outputs should be the same:

```
user@node2> show protection-group ethernet-ring aps
```

Ethernet Ring Name pg102	Request/state NR	No Flush No	Ring Protection Link Blocked Yes
-----------------------------	---------------------	----------------	-------------------------------------

Originator	Remote Node ID
No	00:01:01:00:00:01

```
user@node2> show protection-group ethernet-ring interface
```

```
Ethernet ring port parameters for protection group pg102
```

Interface	Control Channel	Forward State	Ring Protection Link End
ge-1/2/1	ge-1/2/1.1	forwarding	No
ge-1/0/2	ge-1/0/2.1	forwarding	No

Signal Failure	Admin State
Clear	IFF ready
Clear	IFF ready

```
user@node2> show protection-group ethernet-ring node-state
```

Ethernet ring pg102	APS State idle	Event NR-RB	Ring Protection Link Owner No
------------------------	-------------------	----------------	----------------------------------

Restore Timer disabled	Quard Timer disabled	Operation state operational
---------------------------	-------------------------	--------------------------------

```
user@node2> show protection-group ethernet-ring statistics group-name pg102
```

```
Ethernet Ring statistics for PG pg101
RAPS sent                : 0
RAPS received             : 1
Local SF happened:       : 0
Remote SF happened:       : 0
NR event happened:        : 0
NR-RB event happened:     : 1
```

Failure Situation—RPL Owner Node

If the ring has a link failure between Node 2 and Node 3, the **show** command will have the following outputs for Node 1:

```
user@node1> show protection-group ethernet-ring aps
```

Ethernet Ring Name	Request/state	No Flush	Ring Protection Link Blocked
pg101	SF	NO	No

Originator	Remote Node ID
No	00:01:02:00:00:01

```
user@node1> show protection-group ethernet-ring interface
```

Ethernet ring port parameters for protection group pg101

Interface	Control Channel	Forward State	Ring Protection Link End
ge-1/0/1	ge-1/0/1.1	forwarding	Yes
ge-1/2/4	ge-1/2/4.1	forwarding	No

Signal Failure	Admin State
Clear	IFF ready
Clear	IFF ready

```
user@node1> show protection-group ethernet-ring node-state
```

Ethernet ring	APS State	Event	Ring Protection Link Owner
pg101	protected	SF	Yes

Restore Timer	Quard Timer	Operation state
disabled	disabled	operational

```
user@node1> show protection-group ethernet-ring statistics group-name pg101
```

Ethernet Ring statistics for PG pg101

RAPS sent	: 1
RAPS received	: 1
Local SF happened:	: 0
Remote SF happened:	: 1
NR event happened:	: 0
NR-RB event happened:	: 1

Failure Situation—Other Nodes

For Node 2 and Node 3, the outputs should be the same:

```
user@node2> show protection-group ethernet-ring aps
```

Ethernet Ring Name	Request/state	No Flush	Ring Protection Link Blocked
pg102	SF	No	No

Originator	Remote Node ID
Yes	00:00:00:00:00:00

```
user@node2> show protection-group ethernet-ring interface
```

Ethernet ring port parameters for protection group pg102

Interface	Control Channel	Forward State	Ring Protection Link End
ge-1/2/1	ge-1/2/1.1	forwarding	No

```
ge-1/0/2      ge-1/0/2.1      discarding      No
```

```
Signal Failure Admin State  
Clear          IFF ready  
set            IFF ready
```

```
user@node2> show protection-group ethernet-ring node-state
```

```
Ethernet ring    APS State    Event        Ring Protection Link Owner  
pg102            idle        NR-RB        No
```

```
Restore Timer    Quard Timer  Operation state  
disabled         disabled    operational
```

```
user@node2> show protection-group ethernet-ring statistics group-name pg102
```

```
Ethernet Ring statistics for PG pg101  
RAPS sent                : 1  
RAPS received             : 1  
Local SF happened:        : 1  
Remote SF happened:       : 0  
NR event happened:        : 0  
NR-RB event happened:     : 1
```

**Related
Documentation**

- [Ethernet Ring Protection Switching Overview on page 165](#)
- [Understanding Ethernet Ring Protection Switching Functionality on page 166](#)
- [Configuring Ethernet Ring Protection Switching on page 173](#)
- *Ethernet Interfaces Feature Guide for Routing Devices*

PART 6

Configuring Nonstop Bridging

- [Understanding How Nonstop Bridging Preserves Layer 2 Protocol Information During a Routing Engine Switchover on page 185](#)
- [Nonstop Bridging System Requirements on page 189](#)
- [Configuring Nonstop Bridging on page 191](#)

Understanding How Nonstop Bridging Preserves Layer 2 Protocol Information During a Routing Engine Switchover

- [Nonstop Bridging Concepts on page 185](#)
- [Understanding Nonstop Bridging on EX Series Switches on page 187](#)

Nonstop Bridging Concepts


Nonstop bridging uses the same infrastructure as graceful Routing Engine switchover (GRES) to preserve interface and kernel information. However, nonstop bridging also saves Layer 2 Control Protocol (L2CP) information by running the Layer 2 Control Protocol process (l2cpd) on the backup Routing Engine.



NOTE: To use nonstop bridging, you must first enable graceful Routing Engine switchover on your routing (or switching) platform. For more information about graceful Routing Engine switchover, see [“Understanding Graceful Routing Engine Switchover” on page 135](#).

[Figure 15 on page 186](#) shows the system architecture of nonstop bridging and the process a routing (or switching) platform follows to prepare for a switchover.

Figure 15: Nonstop Bridging Switchover Preparation Process




The switchover preparation process for nonstop bridging follows these steps:

1. The master Routing Engine starts.
2. The routing platform processes on the master Routing Engine (such as the chassis process [chassisd] and the Layer 2 Control Protocol process [l2cpd]) start.
3. The Packet Forwarding Engine starts and connects to the master Routing Engine.
4. All state information is updated in the system.
5. The backup Routing Engine starts, including the chassis process (chassisd) and the Layer 2 Control Protocol process (l2cpd).
6. The system determines whether graceful Routing Engine switchover and nonstop bridging have been enabled.
7. The kernel synchronization process (ksyncd) synchronizes the backup Routing Engine with the master Routing Engine.
8. For supported protocols, state information is updated directly between the l2cpds on the master and backup Routing Engines.

Figure 16 on page 187 shows the effects of a switchover on the routing platform.

Figure 16: Nonstop Bridging During a Switchover



The switchover process follows these steps:

1. When keepalives from the master Routing Engine are lost, the system switches over gracefully to the backup Routing Engine.
2. The Packet Forwarding Engine connects to the backup Routing Engine, which becomes the new master. Because the Layer 2 Control Protocol process (l2cpd) and chassis process (chassisd) are already running, these processes do not need to restart.
3. State information learned from the point of the switchover is updated in the system. Forwarding and bridging are continued during the switchover, resulting in minimal packet loss.

Related Documentation

- [Understanding High Availability Features on Juniper Networks Routers on page 3](#)
- [Nonstop Bridging System Requirements on page 189](#)
- [Configuring Nonstop Bridging on page 191](#)
- [Configuring Nonstop Bridging on Switches \(CLI Procedure\) on page 193](#)

Understanding Nonstop Bridging on EX Series Switches

You can configure nonstop bridging (NSB) to provide resilience for Layer 2 protocol sessions on a Juniper Networks EX Series Ethernet Switch or on an EX Series Virtual Chassis with redundant Routing Engines.

NSB operates by synchronizing all protocol information for NSB-supported Layer 2 protocols between the master and backup Routing Engines. If the switch has a Routing Engine switchover, the NSB-supported Layer 2 protocol sessions remain active because all session information is already synchronized to the backup Routing Engine. Traffic disruption for the NSB-supported Layer 2 protocol is minimal or nonexistent as a result of the switchover. The Routing Engine switchover is transparent to neighbor devices,

which do not detect any changes related to the NSB-supported Layer 2 protocol sessions on the switch.

For a list of the EX Series switches and Layer 2 protocols that support NSB, see *EX Series Switch Software Features Overview* and *EX Series Virtual Chassis Software Features Overview*.



NOTE: Nonstop bridging provides a transparent switchover mechanism only for Layer 2 protocol sessions. Nonstop active routing (NSR) provides a similar mechanism for Layer 3 protocol sessions. See [“Understanding Nonstop Active Routing on EX Series Switches” on page 203](#).

**Related
Documentation**

- For information about configuring NSB on EX Series switches that do not support the Enhanced Layer 2 Software (ELS) CLI style, see [Configuring Nonstop Bridging on EX Series Switches \(CLI Procedure\) on page 195](#)
- For information about configuring NSB on EX Series switches that support ELS, see [Configuring Nonstop Bridging on Switches \(CLI Procedure\) on page 193](#)

Nonstop Bridging System Requirements

- [Nonstop Bridging System Requirements on page 189](#)

Nonstop Bridging System Requirements

This topic contains the following sections:

- [Platform Support on page 189](#)
- [Protocol Support on page 189](#)

Platform Support

Nonstop bridging is supported on MX Series 5G Universal Routing Platforms. Your system must be running Junos OS Release 8.4 or later.

Nonstop bridging is supported on EX Series switches with redundant Routing Engines in a Virtual Chassis or in a Virtual Chassis Fabric.

Nonstop bridging is supported on QFX Series switches in a Virtual Chassis or in a Virtual Chassis Fabric.

For a list of the EX Series switches and Layer 2 protocols that support nonstop bridging, see *EX Series Switch Software Features Overview*.



NOTE: All Routing Engines configured for nonstop bridging must be running the same Junos OS release.

Protocol Support

Nonstop bridging is supported for the following Layer 2 control protocols:

- Spanning Tree Protocol (STP)
- Rapid Spanning Tree Protocol (RSTP)
- Multiple Spanning Tree Protocol (MSTP)
- VLAN Spanning Tree Protocol (VSTP)

**Related
Documentation**

- [Nonstop Bridging Concepts on page 185](#)
- [Configuring Nonstop Bridging on page 191](#)
- [Configuring Nonstop Bridging on Switches \(CLI Procedure\) on page 193](#)

CHAPTER 16

Configuring Nonstop Bridging

- [Configuring Nonstop Bridging on page 191](#)
- [Configuring Nonstop Bridging on Switches \(CLI Procedure\) on page 193](#)
- [Configuring Nonstop Bridging on EX Series Switches \(CLI Procedure\) on page 195](#)

Configuring Nonstop Bridging

This section includes the following topics:

- [Enabling Nonstop Bridging on page 191](#)
- [Synchronizing the Routing Engine Configuration on page 191](#)
- [Verifying Nonstop Bridging Operation on page 192](#)

Enabling Nonstop Bridging

Nonstop bridging requires you to configure graceful Routing Engine switchover (GRES). To enable graceful Routing Engine switchover, include the **graceful-switchover** statement at the **[edit chassis redundancy]** hierarchy level:

```
[edit chassis redundancy]  
graceful-switchover;
```

By default, nonstop bridging is disabled. To enable nonstop bridging, include the **nonstop-bridging** statement at the **[edit protocols layer2-control]** hierarchy level:

```
[edit protocols layer2-control]  
nonstop-bridging;
```

To disable nonstop active routing, remove the **nonstop-bridging** statement from the **[edit protocols layer2-control]** hierarchy level.

Synchronizing the Routing Engine Configuration

When you configure nonstop bridging, you must also include the **commit synchronize** statement at the **[edit system]** hierarchy level so that, by default, when you issue the **commit** command, the configuration changes are synchronized on both Routing Engines. If you issue the **commit synchronize** command at the **[edit]** hierarchy level on the backup Routing Engine, the Junos OS displays a warning and commits the candidate configuration.



NOTE: A newly inserted backup Routing Engine automatically synchronizes its configuration with the master Routing Engine configuration.

When you configure nonstop bridging, you can bring the backup Routing Engine online after the master Routing Engine is already running. There is no requirement to start the two Routing Engines simultaneously.

Verifying Nonstop Bridging Operation

When you enable nonstop bridging, you can issue Layer 2 Control Protocol-related operational mode commands on the backup Routing Engine. However, the output of the commands might not match the output of the same commands issued on the master Routing Engine.

Related Documentation

- [Nonstop Bridging Concepts on page 185](#)
- [Nonstop Bridging System Requirements on page 189](#)
- [nonstop-bridging on page 553](#)
- [Configuring Nonstop Bridging on EX Series Switches \(CLI Procedure\) on page 195](#)

Configuring Nonstop Bridging on Switches (CLI Procedure)



NOTE: This task uses switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see “[Configuring Nonstop Bridging on EX Series Switches \(CLI Procedure\)](#)” on page 195. For ELS details, see *Using the Enhanced Layer 2 Software CLI*.

You can configure nonstop bridging (NSB) to provide resilience for Layer 2 protocol sessions on a Juniper Networks EX Series switch with multiple Routing Engines or an EX Series or QFX Series switch in a Virtual Chassis or Virtual Chassis Fabric configuration. Limited support for NSB is also provided on QFX5100 and EX4600 standalone switches, but NSB is enabled *only* during an ISSU.

NSB operates by synchronizing all protocol information for NSB-supported Layer 2 protocols between the master and backup Routing Engines. If the switch has a Routing Engine switchover, the NSB-supported Layer 2 protocol sessions remain active because they are already synchronized on the backup Routing Engine. The Routing Engine switchover is transparent to neighbor devices, which do not detect any changes related to the Layer 2 protocol sessions. The neighboring devices and other devices on the network do not, therefore, have to resynchronize their Layer 2 protocol states to respond to the downtime on the switch—a process that adds network overhead and risks disrupting network performance—when a Routing Engine switchover occurs when NSB is enabled.



NOTE: If you are using a QFX5100 or EX4600 standalone switch and you want to use ISSU, configure Graceful Routing Engine switchover (GRES), NSB and nonstop active routing (NSR). You must configure NSB, GRES, and NSR in order to run ISSU. However, GRES, NSB and NSR are enabled *only* during the upgrade. During an ISSU, the Junos OS runs in two separate virtual machines (VMs)—one VM is in the master role acting as the master Routing Engine, and the other VM is in the backup role acting as the backup Routing Engine. The Junos OS is upgraded on the backup VM. After a successful software upgrade, the backup VM then becomes the master VM, and the original master VM is no longer needed and is shut down.

To configure NSB:

1. Enable graceful Routing Engine switchover (GRES):

```
[edit chassis redundancy]
user@switch# set graceful-switchover
```

2. Enable NSB:

```
[edit protocols layer2-control]
user@switch# set nonstop-bridging
```

3. Synchronize configuration changes between the Routing Engines:

```
[edit system]
user@switch# set commit synchronize
```

If you try to commit a configuration that includes NSB without including the **commit synchronize** statement, the commit fails.



NOTE: There is no requirement to start the two Routing Engines simultaneously. If the backup Routing Engine is not up when you use the **commit synchronize** statement, the candidate configuration is committed in the master Routing Engine. When the backup Routing Engine comes online, its configuration is automatically synchronized with that of the master.



BEST PRACTICE: After a graceful Routing Engine switchover, we recommend that you issue the clear interface statistics (*interface-name* | all) command to reset the cumulative values for local statistics on the new master Routing Engine.

Related Documentation

- [Performing an In-Service Software Upgrade \(ISSU\) with Non-Stop Routing on page 452](#)
- [Understanding Nonstop Bridging on EX Series Switches on page 187](#)
- [Nonstop Bridging Concepts on page 185](#)
- [Understanding In-Service Software Upgrade \(ISSU\) on page 399](#)

Configuring Nonstop Bridging on EX Series Switches (CLI Procedure)



NOTE: This task uses Junos OS for EX Series switches that do not support the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that supports ELS, see “[Configuring Nonstop Bridging on Switches \(CLI Procedure\)](#)” on page 193.

You can configure nonstop bridging (NSB) to provide resilience for Layer 2 protocol sessions on an EX Series switch with redundant Routing Engines.

Nonstop bridging operates by synchronizing all protocol information for NSB-supported Layer 2 protocols between the master and backup Routing Engines. If the switch has a Routing Engine switchover, the NSB-supported Layer 2 protocol sessions remain active because they are already synchronized on the backup Routing Engine. The Routing Engine switchover is transparent to neighbor devices, which do not detect any changes related to the Layer 2 protocol sessions on the switch.

To configure nonstop bridging:

1. Enable graceful Routing Engine switchover (GRES):

```
[edit chassis redundancy]
user@switch# set graceful-switchover
```

2. Configure the switch to always synchronize configuration changes between the Routing Engines:

```
[edit system]
user@switch# set commit synchronize
```

If you try to commit a configuration in which nonstop bridging is configured but synchronization of configuration changes is not configured, the configuration is not committed.

3. Enable nonstop bridging:

```
[edit ethernet-switching-options]
user@switch# set nonstop-bridging
```



NOTE: There is no requirement to start both Routing Engines simultaneously. If the backup Routing Engine is not up when you commit the configuration, the candidate configuration is committed in the master Routing Engine. When the backup Routing Engine comes online, the configuration is automatically synchronized.

- Related Documentation**
- *Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches*
 - [Understanding Nonstop Bridging on EX Series Switches on page 187](#)

PART 7

Configuring Nonstop Active Routing (NSR)

- [Understanding How Nonstop Active Routing Preserves Routing Protocol Information During a Routing Engine Switchover on page 199](#)
- [Nonstop Active Routing System Requirements on page 205](#)
- [Configuring Nonstop Active Routing on page 217](#)

Understanding How Nonstop Active Routing Preserves Routing Protocol Information During a Routing Engine Switchover

- [Nonstop Active Routing Concepts on page 199](#)
- [Understanding Nonstop Active Routing on EX Series Switches on page 203](#)

Nonstop Active Routing Concepts

Nonstop active routing (NSR) uses the same infrastructure as graceful Routing Engine switchover (GRES) to preserve interface and kernel information. However, NSR also saves routing protocol information by running the routing protocol process (rpd) on the backup Routing Engine. By saving this additional information, NSR is self-contained and does not rely on helper routers (or switches) to assist the routing platform in restoring routing protocol information. NSR is advantageous in networks in which neighbor routers (or switches) do not support graceful restart protocol extensions. As a result of this enhanced functionality, NSR is a natural replacement for graceful restart.

Starting with Junos OS Release 15.1R1, if you have NSR configured, it is never valid to issue the **restart routing** command in any form on the NSR master Routing Engine. Doing so results in a loss of protocol adjacencies and neighbors and a drop in traffic.



NOTE: To use NSR, you must first enable GRES on your routing (or switching) platform. For more information about GRES, see [“Understanding Graceful Routing Engine Switchover” on page 135](#).



NOTE: Starting with Junos OS Release 12.3, because of its synchronization requirements and logic, NSR or GRES performance is limited by the slowest Routing Engine in the system.




.....

NOTE: If NSR is enabled, certain system log (syslog) messages are sent from the backup Routing Engine if the configured syslog host is reachable through the fxp0 interface.

.....

Figure 17 on page 201 shows the system architecture of nonstop active routing and the process a routing (or switching) platform follows to prepare for a switchover.

Figure 17: Nonstop Active Routing Switchover Preparation Process




The switchover preparation process for NSR comprises the following steps:

1. The master Routing Engine starts.
2. The routing (or switching) platform processes on the master Routing Engine (such as the chassis process [chassisd] and the routing protocol process [rpd]) start.
3. The Packet Forwarding Engine starts and connects to the master Routing Engine.
4. All state information is updated in the system.
5. The backup Routing Engine starts, including the chassis process (chassisd) and the routing protocol process (rpd).
6. The system determines whether GRES and NSR have been enabled.
7. The kernel synchronization process (ksyncd) synchronizes the backup Routing Engine with the master Routing Engine.
8. For supported protocols, state information is updated directly between the routing protocol processes on the master and backup Routing Engines.

Figure 18 on page 202 shows the effects of a switchover on the routing platform.

Figure 18: Nonstop Active Routing During a Switchover



gd16810

The switchover process comprises the following steps:

1. When keepalives from the master Routing Engine are lost, the system switches over gracefully to the backup Routing Engine.
2. The Packet Forwarding Engine connects to the backup Routing Engine, which becomes the new master. Because the routing protocol process (rpd) and chassis process (chassisd) are already running, these processes do not need to restart.
3. State information learned from the point of the switchover is updated in the system. Forwarding and routing are continued during the switchover, resulting in minimal packet loss.
4. Peer routers (or switches) continue to interact with the routing platform as if no change had occurred. Routing adjacencies and session state relying on underlying routing information are preserved and not reset.



CAUTION: We recommend that you do not restart the routing protocol process (rpd) on master Routing Engine after enabling NSR, as it disrupts the protocol adjacency/peering sessions, resulting in traffic loss.

Release History Table

Release	Description
15.1R1	Starting with Junos OS Release 15.1R1, if you have NSR configured, it is never valid to issue the restart routing command in any form on the NSR master Routing Engine.
12.3	Starting with Junos OS Release 12.3, because of its synchronization requirements and logic, NSR or GRES performance is limited by the slowest Routing Engine in the system.

Related Documentation

- [Understanding High Availability Features on Juniper Networks Routers on page 3](#)
- [Nonstop Active Routing System Requirements on page 205](#)
- [Configuring Nonstop Active Routing on page 217](#)
- [Configuring Nonstop Active Routing on Switches on page 220](#)

Understanding Nonstop Active Routing on EX Series Switches

You can configure nonstop active routing (NSR) on an EX Series switch with redundant Routing Engines or on an EX Series Virtual Chassis to enable the transparent switchover of the Routing Engines in the event that one of the Routing Engines goes down.

Nonstop active routing provides high availability for Routing Engines by enabling transparent switchover of the Routing Engines without requiring restart of supported routing protocols. Both Routing Engines are fully active in processing protocol sessions, and so each can take over for the other. The switchover is transparent to neighbor routing devices, which do not detect that a change has occurred.

Enable nonstop active routing when neighbor routing devices are not configured to support graceful restart of protocols or when you want to ensure graceful restart of protocols for which graceful restart is not supported—such as PIM.

You do not need to start the two Routing Engines simultaneously to synchronize them for nonstop active routing. If both Routing Engines are not present or not up when you issue a **commit synchronize** statement, the candidate configuration is committed in the master Routing Engine and when the backup Routing Engine is inserted or comes online, its configuration is automatically synchronized with that of the master.

Nonstop active routing uses the same infrastructure as graceful Routing Engine switchover (GRES) to preserve interface and kernel information. However, nonstop active routing also saves routing protocol information by running the routing protocol process (**rpd**) on the backup Routing Engine. By saving this additional information, nonstop active routing does not rely on other routing devices to assist in restoring routing protocol information.



NOTE: After a graceful Routing Engine switchover, we recommend that you issue the **clear interface statistics** (*interface-name* | all) command to reset the cumulative values for local statistics on the new master Routing Engine.

If you suspect a problem with the synchronization of Routing Engines when nonstop active routing is enabled, you can gather troubleshooting information using trace options. For example, if certain protocols lose connectivity with neighbors after a graceful Routing Engine switchover with NSR enabled, you can use trace options to help isolate the problem. See [“Tracing Nonstop Active Routing Synchronization Events” on page 224](#).



NOTE: Graceful restart and nonstop active routing are mutually exclusive. You will receive an error message upon commit if both are configured.



NOTE: Nonstop active routing provides a transparent switchover mechanism only for Layer 3 protocol sessions. Nonstop bridging (NSB) provides a similar mechanism for Layer 2 protocol sessions. See [“Understanding Nonstop Bridging on EX Series Switches” on page 187](#).

**Related
Documentation**

- [Configuring Nonstop Active Routing on Switches on page 220](#)
- [Example: Configuring Nonstop Active Routing on Switches on page 226](#)

CHAPTER 18

Nonstop Active Routing System Requirements

- [Nonstop Active Routing System Requirements on page 205](#)

Nonstop Active Routing System Requirements

This section contains the following topics:

- [Nonstop Active Routing Platform and Switching Platform Support on page 205](#)
- [Nonstop Active Routing Protocol and Feature Support on page 206](#)
- [Nonstop Active Routing BFD Support on page 209](#)
- [Nonstop Active Routing BGP Support on page 210](#)
- [Nonstop Active Routing Layer 2 Circuit and VPLS Support on page 211](#)
- [Nonstop Active Routing PIM Support on page 212](#)
- [Nonstop Active Routing MSDP Support on page 214](#)
- [Nonstop Active Routing Support for RSVP-TE LSPs on page 215](#)

Nonstop Active Routing Platform and Switching Platform Support

[Table 8 on page 205](#) lists the platforms that support nonstop active routing (NSR).

Table 8: Nonstop Active Routing Platform Support

Platform	Junos OS Release
M10i router	8.4 or later
M20 router	8.4 or later
M40e router	8.4 or later
M120 router	9.0 or later
M320 router	8.4 or later
MX Series routers	9.0 or later

Table 8: Nonstop Active Routing Platform Support (continued)

Platform	Junos OS Release
PTX Series Packet Transport Routers	12.1R4 or later
<p>NOTE:</p> <p>Nonstop active routing (NSR) switchover on PTX series is supported only for the following MPLS and VPN protocols and applications using chained composite next hops:</p> <ul style="list-style-type: none"> • Labeled BGP • Layer 2 VPNs excluding Layer 2 interworking (Layer 2 switching) • Layer 3 VPNs • LDP • RSVP 	
T320 router, T640 router, and TX Matrix router	8.4 or later
Standalone T1600 router	8.5 or later
Standalone T4000 router	12.1R2 or later
TX Plus Matrix router	10.0 or later
TX Plus Matrix router with 3D SIBs	13.1 or later
EX Series switch with dual Routing Engines or in a Virtual Chassis	10.4 or later for EX Series switches
EX Series or QFX Series switches in a Virtual Chassis Fabric	13.2X51-D20 or later for the EX Series and QFX Series switches



NOTE: All Routing Engines configured for nonstop active routing must be running the same Junos OS release.

Nonstop Active Routing Protocol and Feature Support

Table 9 on page 206 lists the protocols that are supported by nonstop active routing.

Table 9: Nonstop Active Routing Protocol and Feature Support

Protocol	Junos OS Release
Aggregated Ethernet interfaces with Link Aggregation Control Protocol (LACP)	9.4 or later
Bidirectional Forwarding Detection (BFD)	8.5 or later
For more information, see "Nonstop Active Routing BFD Support" on page 209.	

Table 9: Nonstop Active Routing Protocol and Feature Support (continued)

Protocol	Junos OS Release
BGP	8.4 or later
For more information, see “Nonstop Active Routing BGP Support” on page 210.	
EVPN <ul style="list-style-type: none"> • EVPN with ingress replication for BUM traffic • EVPN-ETREE • EVPN-VPWS • EVPN -VXLAN • PBB-EVPN • EVPN with P2MP mLDP replication for BUM traffic For more information, please see <i>NSR and Unified ISSU Support for EVPN</i> .	16.2R1 or later (for EVPN with ingress replication for BUM traffic) 17.2R1 or later (for (EVPN-ETREE, EVPN-VPWS, EVPN-VXLAN, and PBB-EVPN) 18.2R1 or later (for EVPN with P2MP mLDP replication for BUM traffic)
Labeled BGP (PTX Series Packet Transport Routers: only)	12.1R4 or later
IS-IS	8.4 or later
LDP	8.4 or later
LDP-based virtual private LAN service (VPLS)	9.3 or later
LDP OAM (operation, administration, and management) features	9.6 or later
LDP (PTX Series Packet Transport Routers only) Nonstop active routing support for LDP includes: <ul style="list-style-type: none"> • LDP unicast transit LSPs • LDP egress LSPs for labeled internal BGP (IBGP) and external BGP (EBGP) • LDP over RSVP transit LSPs • LDP transit LSPs with indexed next hops • LDP transit LSPs with unequal cost load balancing • LDP Point-to-Multipoint LSPs • LDP ingress LSPs 	12.3R4 or later (for LDP Point-to-Multipoint LSPs) 13.3R1 or later (for LDP ingress LSPs) 13.3R1 or later
Layer 2 circuits	(on LDP-based VPLS) 9.2 or later (on RSVP-TE LSP) 11.1 or later
Layer 2 VPNs	9.1 or later
Layer 2 VPNs (PTX Series Packet Transport Routers only)	12.1R4 or later
NOTE: Nonstop active routing is not supported for Layer 2 interworking (Layer 2 stitching).	

Table 9: Nonstop Active Routing Protocol and Feature Support (continued)

Protocol	Junos OS Release
Layer 3 VPNs (see the first Note after this table for restrictions) Nonstop active routing support for Layer 3 VPNs include: <ul style="list-style-type: none"> • IPv4 labeled-unicast (ingress or egress) • IPv4-vpn unicast (ingress or egress) • IPv6 labeled-unicast (ingress or egress) • IPv6-vpn unicast (ingress or egress) 	9.2 or later
Layer 3 VPNs (PTX Series Packet Transport Routers only)	12.1R4 or later
Logical System support (Nonstop active routing support for logical systems to preserve interface and kernel information).	13.3R1 or later
Multicast Source Discovery Protocol (MSDP) For more information, see “Nonstop Active Routing MSDP Support” on page 214.	12.1 or later
OSPF/OSPFv3	8.4 or later
Protocol Independent Multicast (PIM) For more information, see “Nonstop Active Routing PIM Support” on page 212.	(for IPv4) 9.3 or later (for IPv6) 10.4 or later
RIP and RIP next generation (RIPng)	9.0 or later
RSVP (PTX Series Packet Transport Routers only) Nonstop active routing support for RSVP includes: <ul style="list-style-type: none"> • Point-to-Multipoint LSPs <ul style="list-style-type: none"> • RSVP Point-to-Multipoint ingress, transit, and egress LSPs using existing non-chained next hop. • RSVP Point-to-Multipoint transit LSPs using composite next hops for Point-to-Multipoint label routes. • Point-to-Point LSPs <ul style="list-style-type: none"> • RSVP Point-to-Point ingress, transit, and egress LSPs using non-chained next hops. • RSVP Point-to-Point transit LSPs using chained composite next hops. 	12.1R4 or later
RSVP-TE LSP For more information, see “Nonstop Active Routing Support for RSVP-TE LSPs” on page 215.	9.5 or later

Table 9: Nonstop Active Routing Protocol and Feature Support (continued)

Protocol	Junos OS Release
VPLS	(LDP-based) 9.1 or later (RSVP-TE-based) 11.2 or later
VRRP	13.2 or later
VRRP	13.2 or later



NOTE: Layer 3 VPN support does not include dynamic GRE tunnels, multicast VPNs, or BGP flow routes.



NOTE: If you configure a protocol that is not supported by nonstop active routing, the protocol operates as usual. When a switchover occurs, the state information for the unsupported protocol is not preserved and must be refreshed using the normal recovery mechanisms inherent in the protocol.



NOTE: On routers that have logical systems configured on them, NSR is only supported in the main instance.



NOTE: Starting with Junos OS Release 13.3R5, on EX9214 switches, the VRRP master state might change during graceful Routing Engine switchover, even when nonstop active routing is enabled.

Nonstop Active Routing BFD Support

Nonstop active routing supports the Bidirectional Forwarding Detection (BFD) protocol, which uses the topology discovered by routing protocols to monitor neighbors. The BFD protocol is a simple hello mechanism that detects failures in a network. Because BFD is streamlined to be efficient at fast liveness detection, when it is used in conjunction with routing protocols, routing recovery times are improved. With nonstop active routing enabled, BFD session states are not restarted when a Routing Engine switchover occurs.



NOTE: BFD session states are saved only for clients using aggregate or static routes or for BGP, IS-IS, OSPF/OSPFv3, PIM, or RSVP.

When a BFD session is distributed to the Packet Forwarding Engine, BFD packets continue to be sent during a Routing Engine switchover. If nondistributed BFD sessions are to be kept alive during a switchover, you must ensure that the session failure detection time

is greater than the Routing Engine switchover time. The following BFD sessions are not distributed to the Packet Forwarding Engine: multihop sessions, tunnel-encapsulated sessions, and sessions over integrated routing and bridging (IRB) interfaces.



NOTE: BFD is an intensive protocol that consumes system resources. Specifying a minimum interval for BFD less than 100 ms for Routing Engine-based sessions and 10 ms for distributed BFD sessions can cause undesired BFD flapping. The minimum-interval configuration statement is a BFD liveness detection parameter.

Depending on your network environment, these additional recommendations might apply:

- For large-scale network deployments with a large number of BFD sessions, specify a minimum interval of 300 ms for Routing Engine-based sessions, and 100 ms for distributed BFD sessions.
- For very large-scale network deployments with a large number of BFD sessions, contact Juniper Networks customer support for more information.
- For BFD sessions to remain up during a Routing Engine switchover event when nonstop active routing is configured, specify a minimum interval of 10 seconds for Routing Engine-based sessions. For distributed BFD sessions with nonstop active routing configured, the minimum interval recommendations are unchanged and depend only on your network deployment.

Nonstop Active Routing BGP Support

Nonstop active routing BGP support is subject to the following conditions:

- You must include the **path-selection external-router-ID** statement at the **[edit protocols bgp]** hierarchy level to ensure consistent path selection between the master and backup Routing Engines during and after the nonstop active routing switchover.
- Starting with Junos OS Release 14.1, you must include the **advertise-from-main-vpn-tables** statement at the **[edit protocols bgp]** hierarchy level to prevent BGP sessions from going down when route reflector (RR) or autonomous system border router (ASBR) functionality is enabled or disabled on a routing device that has VPN address families configured.
- BGP session uptime and downtime statistics are not synchronized between the primary and backup Routing Engines during Nonstop Active Routing and ISSU. The backup Routing Engine maintains its own session uptime based on the time when the backup first becomes aware of the established sessions. For example, if the backup Routing Engine is rebooted (or if you run **restart routing** on the backup Routing Engine), the backup's uptime is a short duration, because the backup has just learned about the established sessions. If the backup is operating when the BGP sessions first come up on the primary, the uptime on the primary and the uptime on the backup are almost

the same duration. After a Routing Engine switchover, the new master continues from the time left on the backup Routing Engine.

- If the BGP peer in the master Routing Engine has negotiated address-family capabilities that are not supported for nonstop active routing, then the corresponding BGP neighbor state on the backup Routing Engine shows as idle. On switchover, the BGP session is reestablished from the new master Routing Engine.

Only the following address families are supported for nonstop active routing.



NOTE: Address families are supported only on the main instance of BGP. Only unicast is supported on VRF instances.

- inet labeled-unicast
- inet-mdt
- inet multicast
- inet-mvpn
- inet unicast
- inet-vpn unicast
- inet6 labeled-unicast
- inet6 multicast
- inet6-mvpn
- inet6 unicast
- inet6-vpn unicast
- iso-vpn
- l2vpn signaling
- route-target
- BGP route dampening does not work on the backup Routing Engine when nonstop active routing is enabled.

Nonstop Active Routing Layer 2 Circuit and VPLS Support

Nonstop active routing supports Layer 2 circuit and VPLS on both LDP-based and RSVP-TE-based networks. Nonstop active routing support enables the backup Routing Engine to track the label advertised by Layer 2 circuit and VPLS on the primary Routing Engine, and to use the same label after the Routing Engine switchover.

in Junos OS Release 9.6 and later, nonstop active routing support is extended to the Layer 2 circuit and LDP-based VPLS pseudowire redundant configurations.

Nonstop Active Routing PIM Support

Nonstop active routing supports Protocol Independent Multicast (PIM) with stateful replication on backup Routing Engines. State information replicated on the backup Routing Engine includes information about neighbor relationships, join and prune events, rendezvous point (RP) sets, synchronization between routes and next hops, multicast session states, and the forwarding state between the two Routing Engines.



NOTE: Nonstop active routing for PIM is supported for IPv4 on Junos OS Release 9.3 and later, and for IPv6 on Junos OS Release 10.4 and later. Starting with Release 11.1, Junos OS also supports nonstop active routing for PIM on devices that have both IPv4 and IPv6 configured on them.

To configure nonstop active routing for PIM, include the same statements in the configuration as for other protocols: the **nonstop-routing** statement at the [edit routing-options] hierarchy level and the **graceful-switchover** statement at the [edit chassis redundancy] hierarchy level. To trace PIM nonstop active routing events, include the **flag nsr-synchronization** statement at the [edit protocols pim traceoptions] hierarchy level.



NOTE: The **clear pim join**, **clear pim register**, and **clear pim statistics** operational mode commands are not supported on the backup Routing Engine when nonstop active routing is enabled.

Nonstop active routing support varies for different PIM features. The features fall into the following three categories: supported features, unsupported features, and incompatible features.

Supported features:

- Auto-RP



NOTE: Nonstop active routing PIM support on IPv6 does not support auto-RP because IPv6 does not support auto-RP.

- Bootstrap router (BSR)
- Static RPs
- Embedded RP on non-RP IPv6 routers
- Local RP



NOTE: RP set information synchronization is supported for local RP and BSR (on IPv4 and IPv6), autoRP (on IPv4), and embedded RP (on IPv6).

- BFD

- Dense mode
- Sparse mode
- Source-specific multicast (SSM)
- Draft Rosen multicast VPNs (MVPNs)
- Anycast RP (anycast RP set information synchronization and anycast RP register state synchronization on IPv4 and IPv6 configurations)
- Flow maps
- Unified ISSU
- Policy features such as neighbor policy, bootstrap router export and import policies, scope policy, flow maps, and reverse path forwarding (RPF) check policies
- Upstream assert synchronization
- PIM join load balancing

Starting with Release 12.2, Junos OS extends the nonstop active routing PIM support to draft Rosen MVPNs. Nonstop active routing PIM support for draft Rosen MVPNs enables nonstop active routing-enabled devices to preserve draft Rosen MPVN-related information—such as default and data multicast distribution tree (MDT) states—across switchovers. In releases earlier than 12.2, nonstop active routing PIM configuration was incompatible with draft Rosen MPVN configuration.

The backup Routing Engine sets up the default MDT based on the configuration and the information it receives from the master Routing Engine, and keeps updating the default MDT state information.

However, for data MDTs, the backup Routing Engine relies on the master Routing Engine to provide updates when data MDTs are created, updated, or deleted. The backup Routing Engine neither monitors data MDT flow rates nor triggers a data MDT switchover based on variations in flow rates. Similarly, the backup Routing Engine does not maintain the data MDT delay timer or timeout timer. It does not send MDT join TLV packets for the data MDTs until it takes over as the master Routing Engine. After the switchover, the new master Routing Engine starts sending MDT join TLV packets for each data MDT, and also resets the data MDT timers. Note that the expiration time for the timers might vary from the original values on the previous master Routing Engine.

Starting with Release 12.3, Junos OS extends the Protocol Independent Multicast (PIM) nonstop active routing support to IGMP-only interfaces.

In Junos OS releases earlier than 12.3, the PIM joins created on IGMP-only interfaces were not replicated on the backup Routing Engine. Thus, the corresponding multicast routes were marked as pruned (meaning discarded) on the backup Routing Engine. Because of this limitation, after a switchover, the new master Routing Engine had to wait for the IGMP module to come up and start receiving reports to create PIM joins and to install multicast routes. This caused traffic loss until the multicast joins and routes were reinstated.

However, in Junos OS Release 12.3 and later, the multicast joins on the IGMP-only interfaces are mapped to PIM states, and these states are replicated on the backup Routing Engine. If the corresponding PIM states are available on the backup, the multicast routes are marked as forwarding on the backup Routing Engine. This enables uninterrupted traffic flow after a switchover. This enhancement covers IGMPv2, IGMPv3, MLDv1, and MLDv2 reports and leaves.

Unsupported features: You can configure the following PIM features on a router along with nonstop active routing, but they function as if nonstop active routing is not enabled. In other words, during Routing Engine switchover and other outages, their state information is not preserved, and traffic loss is to be expected.

- Internet Group Management Protocol (IGMP) exclude mode
- IGMP snooping

Nonstop active routing is not supported for next-generation MVPNs with PIM provider tunnels. The commit operation fails if the configuration includes both nonstop active routing and next-generation MVPNs with PIM provider tunnels.

Junos OS provides a configuration statement that disables nonstop active routing for PIM only, so that you can activate incompatible PIM features and continue to use nonstop active routing for the other protocols on the router. Before activating an incompatible PIM feature, include the **nonstop-routing disable** statement at the **[edit protocols pim]** hierarchy level. Note that in this case, nonstop active routing is disabled for all PIM features, not just incompatible features.

Nonstop Active Routing MSDP Support

Starting with Release 12.1, Junos OS extends nonstop active routing support to the Multicast Source Discovery Protocol (MSDP).

Nonstop active routing support for MSDP preserves the following MSDP-related information across the switchover:

- MSDP configuration and peer information
- MSDP peer socket information
- Source-active and related information

However, note that the following restrictions or limitations apply to nonstop active routing MSDP support:

- Because the backup Routing Engine learns the active source information by processing the source-active messages from the network, synchronizing of source active information between the master and backup Routing Engines might take up to 60 seconds. So, no planned switchover is allowed within 60 seconds of the initial replication of the sockets.
- Similarly, Junos OS does not support two planned switchovers within 240 seconds of each other.

Junos OS enables you to trace MSDP nonstop active routing events by including the **flag nsr-synchronization** statement at the **[edit protocols mspd traceoptions]** hierarchy level.

Nonstop Active Routing Support for RSVP-TE LSPs

Junos OS extends nonstop active routing support to label-switching routers (LSRs) and Layer 2 Circuits that are part of an RSVP-TE LSP. Nonstop active routing support on LSRs ensures that the master to backup Routing Engine switchover on an LSR remains transparent to the network neighbors and that the LSP information remains unaltered during and after the switchover.

You can use the **show rsvp version** command to view the nonstop active routing mode and state on an LSR. Similarly, you can use the **show mpls lsp** and **show rsvp session** commands on the backup Routing Engine to view the state recreated on the backup Routing Engine.

The Junos OS nonstop active routing feature is also supported on RSVP point-to-multipoint LSPs. Nonstop active routing support for RSVP point-to-multipoint egress and transit LSPs was added in Junos OS Release 11.4, and for ingress LSPs in Release 12.1. During the switchover, the LSP comes up on the backup Routing Engine that shares and synchronizes the state information with the master Routing Engine before and after the switchover. Nonstop active routing support for point-to-multipoint transit and egress LSPs ensures that the switchover remains transparent to the network neighbors, and preserves the LSP information across the switchover.

Starting with Release 14.1R1, Junos OS extends nonstop active routing support to the next-generation multicast VPNs (MVPNs).

The **show rsvp session detail** command enables you to check the point-to-multipoint LSP remerge state information (**P2MP LSP re-merge**; possible values are **head**, **member**, and **none**).

Starting with Release 14.1R1, Junos OS extends nonstop active routing support for point-to-multipoint LSPs used by VPLS and MVPN.

However, Junos OS does not support nonstop active routing for the following features:

- Generalized Multiprotocol Label Switching (GMPLS) and LSP hierarchy
- Interdomain or loose-hop expansion LSPs
- BFD liveness detection
- Starting with Junos OS Release 14.2, Setup protection

Nonstop active routing support for RSVP-TE LSPs is subject to the following limitations and restrictions:

- Detour LSPs are not maintained across a switchover and so, detour LSPs might fail to come back online after the switchover.
- Control plane statistics corresponding to the **show rsvp statistics** and **show rsvp interface detail | extensive** commands are not maintained across Routing Engine switchovers.

- Statistics from the backup Routing Engine are not reported for **show mpls lsp statistics** and **monitor mpls label-switched-path** commands. However, if a switchover occurs, the backup Routing Engine, after taking over as the master, starts reporting statistics. Note that the **clear statistics** command issued on the old master Routing Engine does not have any effect on the new master Routing Engine, which reports statistics, including any uncleared statistics.
- State timeouts might take additional time during nonstop active routing switchover. For example, if a switchover occurs after a neighbor has missed sending two hello messages to the master, the new master Routing Engine waits for another three hello periods before timing out the neighbor.
- On the RSVP ingress router, if you configure auto-bandwidth functionality, the bandwidth adjustment timers are set in the new master after the switchover. This causes a one-time increase in the length of time required for the bandwidth adjustment after the switchover occurs.
- Backup LSPs —LSPs that are established between the point of local repair (PLR) and the merge point after a node or link failure—are not preserved during a Routing Engine switchover.
- When nonstop active routing is enabled, graceful restart is not supported. However, graceful restart helper mode is supported.

Release History Table

Release	Description
14.2	Starting with Junos OS Release 14.2, Setup protection
14.1	Starting with Junos OS Release 14.1, you must include the advertise-from-main-vpn-tables statement at the [edit protocols bgp] hierarchy level to prevent BGP sessions from going down when route reflector (RR) or autonomous system border router (ASBR) functionality is enabled or disabled on a routing device that has VPN address families configured.
14.1	Starting with Release 14.1R1, Junos OS extends nonstop active routing support to the next-generation multicast VPNs (MVPNs).
14.1	Starting with Release 14.1R1, Junos OS extends nonstop active routing support for point-to-multipoint LSPs used by VPLS and MVPN.
13.3R5	Starting with Junos OS Release 13.3R5, on EX9214 switches, the VRRP master state might change during graceful Routing Engine switchover, even when nonstop active routing is enabled.

Related Documentation

- [Nonstop Active Routing Concepts on page 199](#)
- [Configuring Nonstop Active Routing on page 217](#)
- [Configuring Nonstop Active Routing on Switches on page 220](#)
- [Example: Configuring Nonstop Active Routing on Switches on page 226](#)

CHAPTER 19

Configuring Nonstop Active Routing

- [Configuring Nonstop Active Routing on page 217](#)
- [Configuring Nonstop Active Routing on Switches on page 220](#)
- [Preventing Automatic Reestablishment of BGP Peer Sessions After NSR Switchovers on page 221](#)
- [Example: Configuring Nonstop Active Routing on page 222](#)
- [Tracing Nonstop Active Routing Synchronization Events on page 224](#)
- [Resetting Local Statistics on page 226](#)
- [Example: Configuring Nonstop Active Routing on Switches on page 226](#)

Configuring Nonstop Active Routing

This section includes the following topics:

- [Enabling Nonstop Active Routing on page 217](#)
- [Synchronizing the Routing Engine Configuration on page 218](#)
- [Verifying Nonstop Active Routing Operation on page 219](#)

Enabling Nonstop Active Routing

Nonstop active routing (NSR) requires you to configure graceful Routing Engine switchover (GRES). To enable graceful Routing Engine switchover, include the **graceful-switchover** statement at the **[edit chassis redundancy]** hierarchy level:

```
[edit chassis redundancy]  
graceful-switchover;
```

By default, nonstop active routing is disabled. To enable nonstop active routing, include the **nonstop-routing** statement at the **[edit routing-options]** hierarchy level:

```
[edit routing-options]  
nonstop-routing;
```

To disable nonstop active routing, remove the **nonstop-routing** statement from the **[edit routing-options]** hierarchy level.



NOTE: When you enable nonstop active routing, you cannot enable automatic route distinguishers for multicast VPN routing instances. Automatic route distinguishers are enabled by configuring the `route-distinguisher-id` statement at the `[edit routing-instances instance-name]` hierarchy level; for more information, see the *Junos OS VPNs Library for Routing Devices*.

If the routing protocol process (rpd) on the NSR master Routing Engine crashes, the master Routing Engine simply restarts rpd (with no Routing Engine switchover), which impacts routing protocol adjacencies and neighbors and results in traffic loss. To prevent this negative impact on traffic flow, configure the **switchover-on-routing-crash** statement at the `[edit system]` hierarchy level. This configuration forces an NSR Routing Engine switchover if rpd on the master Routing Engine crashes.

```
[edit system]
user@host# set switchover-on-routing-crash
```

To enable the routing platform to switch over to the backup Routing Engine when the routing protocol process (rpd) fails rapidly three times in succession, include the **other-routing-engine** statement at the `[edit system processes routing failover]` hierarchy level.

For more information about the **other-routing-engine** statement, see the *Junos OS Administration Library*.

Synchronizing the Routing Engine Configuration

When you configure nonstop active routing, you must also include the **commit synchronize** statement at the `[edit system]` hierarchy level so that configuration changes are synchronized on both Routing Engines:

```
[edit system]
commit synchronize;
```

If you try to commit the nonstop active routing configuration without including the **commit synchronize** statement, the commit fails.

If you configure the **commit synchronize** statement at the `[edit system]` hierarchy level and issue a commit in the master Routing Engine, the master configuration is automatically synchronized with the backup.

However, if the backup Routing Engine is down when you issue the commit, the Junos OS displays a warning and commits the candidate configuration in the master Routing Engine. When the backup Routing Engine comes up, its configuration will automatically be synchronized with the master.



NOTE: A newly inserted backup Routing Engine automatically synchronizes its configuration with the master Routing Engine configuration.

When you configure nonstop active routing, you can bring the backup Routing Engine online after the master Routing Engine is already running. There is no requirement to start the two Routing Engines simultaneously.



CAUTION: We recommend that you do not restart Routing Protocol Process (rpd) on master Routing Engine after enabling nonstop active routing, as it disrupts the protocol adjacency/peering sessions, resulting in traffic loss.

Verifying Nonstop Active Routing Operation

To see whether or not nonstop active routing is enabled, issue the **show task replication** command. For BGP nonstop active routing, you must also issue the **show bgp replication** command.



CAUTION: If BGP is configured, before attempting nonstop active routing switchover, check the output of **show bgp replication** to confirm that BGP routing table synchronization has completed on the backup Routing Engine. The **complete** status in the output of **show task replication** only indicates that the socket replication has completed and the BGP synchronization is in progress. To determine whether BGP synchronization is complete, you must check the **Protocol state** and **Synchronization state** fields in the output of **show bgp replication** on the master Routing Engine. The **Protocol state** must be **idle** and the **Synchronization state** must be **complete**. If you perform NSR switchover before the BGP synchronization has completed, the BGP session might flap.

For more information about these commands, see the [CLI Explorer](#).

When you enable nonstop active routing or graceful Routing Engine switchover and issue routing-related operational mode commands on the backup Routing Engine (such as **show route**, **show bgp neighbor**, **show ospf database**, and so on), the output might not match the output of the same commands issued on the master Routing Engine. For example, it is normal for the routing table on the backup Routing Engine to contain persistent phantom routes that are not present in the routing table on the master Routing Engine.

To display BFD state replication status, issue the **show bfd session** command. The **replicated** flag appears in the output for this command when a BFD session has been replicated to the backup Routing Engine. For more information, see the [CLI Explorer](#).

Related Documentation

- [Nonstop Active Routing Concepts on page 199](#)

- [Nonstop Active Routing System Requirements on page 205](#)
- [Tracing Nonstop Active Routing Synchronization Events on page 224](#)
- [Resetting Local Statistics on page 226](#)
- [Example: Configuring Nonstop Active Routing on page 222](#)
- [nonstop-routing on page 546](#)

Configuring Nonstop Active Routing on Switches

Nonstop active routing (NSR) provides a mechanism for transparent switchover of the Routing Engines without necessitating restart of supported routing protocols. Both Routing Engines are fully active in processing protocol sessions, and so each can take over for the other. The switchover is transparent to neighbors.

You can configure NSR on an on a Juniper Networks EX Series switch with multiple Routing Engines or an EX Series or QFX Series switch in a Virtual Chassis or Virtual Chassis Fabric configuration.

To configure nonstop active routing:

1. Enable graceful Routing Engine switchover (GRES):

```
[edit chassis redundancy]
user@switch# set graceful-switchover
```

2. Enable nonstop active routing (by default, nonstop active routing is disabled):

```
[edit routing-options]
user@switch# set nonstop-routing
```

3. Synchronize configuration changes between the Routing Engines:

```
[edit system]
user@switch# set commit synchronize
```

If you try to commit the nonstop active routing configuration without including the **commit synchronize** statement, the commit fails.



NOTE: There is no requirement to start the two Routing Engines simultaneously. If the backup Routing Engine is not up when you issue the **commit synchronize** command, the candidate configuration is committed in the master Routing Engine. When the backup Routing Engine is inserted or comes online, its configuration is automatically synchronized with that of the master.



BEST PRACTICE: After a graceful Routing Engine switchover, we recommend that you issue the clear interface statistics (*interface-name* | all) command to reset the cumulative values for local statistics on the new master Routing Engine.

To disable nonstop active routing:

```
[edit routing-options]
user@switch# delete nonstop-routing
```

Related Documentation

- [Example: Configuring Nonstop Active Routing on Switches on page 226](#)
- [Tracing Nonstop Active Routing Synchronization Events on page 224](#)
- [Understanding Nonstop Active Routing on EX Series Switches on page 203](#)
- [Nonstop Active Routing Concepts on page 199](#)

Preventing Automatic Reestablishment of BGP Peer Sessions After NSR Switchovers

It is useful to prevent a BGP peer session from automatically being reestablished after a nonstop active routing (NSR) switchover when you have applied routing policies configured in the dynamic database. When NSR is enabled, the dynamic database is not synchronized with the backup Routing Engine. Therefore, when a switchover occurs, import and export policies configured in the dynamic database might no longer be available. For more information about configuring dynamic routing policies, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide*.



NOTE: The BGP established timers are not maintained across switchovers.

You can configure the routing device not to reestablish a BGP peer session after an NSR switchover either for a specified period or until you manually reestablish the session. Include the **idle-after-switch-over** statement at the **[edit protocols bgp]** hierarchy level:

```
idle-after-switch-over (forever | seconds);
```

For a list of hierarchy levels at which you can configure this statement, see the configuration statement summary for this statement.

For **seconds**, specify a value from 1 through 4294967295. The BGP peer session is not reestablished until after the specified period. If you specify the **forever** option, the BGP peer session is not reestablished until you issue the **clear bgp neighbor** command.

Example: Configuring Nonstop Active Routing

The following example enables graceful Routing Engine switchover, nonstop active routing, and nonstop active routing trace options for BGP, IS-IS, and OSPF.

```
[edit]
system commit {
  synchronize;
}
chassis {
  redundancy {
    graceful-switchover; # This enables graceful Routing Engine switchover on
                        # the routing platform.
  }
}
interfaces {
  so-0/0/0 {
    unit 0 {
      family inet {
        address 10.0.1.1/30;
      }
      family iso;
    }
  }
  so-0/0/1 {
    unit 0 {
      family inet {
        address 10.1.1.1/30;
      }
      family iso;
    }
  }
  so-0/0/2 {
    unit 0 {
      family inet {
        address 10.2.1.1/30;
      }
      family iso;
    }
  }
  so-0/0/3 {
    unit 0 {
      family inet {
        address 10.3.1.1/30;
      }
      family iso;
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 192.168.2.1/32;
      }
      family iso {
```

```

        address 49.0004.1921.6800.2001.00;
    }
}
}
routing-options {
    nonstop-routing; # This enables nonstop active routing on the routing platform.
    router-id 192.168.2.1;
    autonomous-system 65432;
}
protocols {
    bgp {
        traceoptions {
            flag nsr-synchronization detail; # This logs nonstop active routing
            # events for BGP.
        }
        advertise-from-main-vpn-tables;
        local-address 192.168.2.1;
        group external-group {
            type external;
            export BGP_export;
            neighbor 192.168.1.1 {
                family inet {
                    unicast;
                }
            }
            peer-as 65103;
        }
        group internal-group {
            type internal;
            neighbor 192.168.10.1;
            neighbor 192.168.11.1;
            neighbor 192.168.12.1;
        }
    }
}
isis {
    traceoptions {
        flag nsr-synchronization detail; # This logs nonstop active routing events
        # for IS-IS.
    }
    interface all;
    interface fxp0.0 {
        disable;
    }
    interface lo0.0 {
        passive;
    }
}
ospf {
    traceoptions {
        flag nsr-synchronization detail; # This logs nonstop active routing events
        # for OSPF.
    }
    area 0.0.0.0 {
        interface all;
    }
}

```

```
        interface fxp0.0 {
            disable;
        }
        interface lo0.0 {
            passive;
        }
    }
}
policy-options {
    policy-statement BGP_export {
        term direct {
            from {
                protocol direct;
            }
            then accept;
        }
        term final {
            then reject;
        }
    }
}
```

- Related Documentation**
- [Configuring Nonstop Active Routing on page 217](#)
 - [Tracing Nonstop Active Routing Synchronization Events on page 224](#)

Tracing Nonstop Active Routing Synchronization Events

To track the progress of nonstop active routing synchronization between Routing Engines, you can configure nonstop active routing trace options flags for each supported protocol and for BFD sessions and record these operations to a log file.

To configure nonstop active routing trace options for supported routing protocols, include the **nsr-synchronization** statement at the **[edit protocols *protocol-name* traceoptions flag]** hierarchy level and optionally specify one or more of the **detail**, **disable**, **receive**, and **send** options:

```
[edit protocols]
bgp {
    traceoptions {
        flag nsr-synchronization <detail> <disable> <receive> <send>;
    }
}
isis {
    traceoptions {
        flag nsr-synchronization <detail> <disable> <receive> <send>;
    }
}
ldp {
    traceoptions {
        flag nsr-synchronization <detail> <disable> <receive> <send>;
    }
}
```

```

    }
  }
  mpls {
    traceoptions {
      flag nsr-synchronization;
      flag nsr-synchronization-detail;
    }
  }
  msdp {
    traceoptions {
      flag nsr-synchronization <detail> <disable> <receive> <send>;
    }
  }
  (ospf | ospf3) {
    traceoptions {
      flag nsr-synchronization <detail> <disable> <receive> <send>;
    }
  }
  rip {
    traceoptions {
      flag nsr-synchronization <detail> <disable> <receive> <send>;
    }
  }
  ripng {
    traceoptions {
      flag nsr-synchronization <detail> <disable> <receive> <send>;
    }
  }
  pim {
    traceoptions {
      flag nsr-synchronization <detail> <disable> <receive> <send>;
    }
  }
}

```

To configure nonstop active routing trace options for BFD sessions, include the **nsr-synchronization** and **nsr-packet** statements at the **[edit protocols bfd traceoptions flag]** hierarchy level.

```

[edit protocols]
bfd {
  traceoptions {
    flag nsr-synchronization;
    flag nsr-packet;
  }
}

```

To trace the Layer 2 VPN signaling state replicated from routes advertised by BGP, include the **nsr-synchronization** statement at the **[edit routing-options traceoptions flag]** hierarchy level. This flag also traces the label and logical interface association that VPLS receives from the kernel replication state.

```

[edit routing-options]
traceoptions {
  flag nsr-synchronization;
}

```

```
}
```

**Related
Documentation**

- [Configuring Nonstop Active Routing on page 217](#)
- [Configuring Nonstop Active Routing on Switches on page 220](#)
- [Example: Configuring Nonstop Active Routing on Switches on page 226](#)
- [Example: Configuring Nonstop Active Routing on page 222](#)

Resetting Local Statistics

After a graceful Routing Engine switchover, we recommend that you issue the **clear interface statistics** (*interface-name* | **all**) command to reset the cumulative values for local statistics on the new master Routing Engine.

**Related
Documentation**

- [Configuring Nonstop Active Routing on page 217](#)
- [Tracing Nonstop Active Routing Synchronization Events on page 224](#)

Example: Configuring Nonstop Active Routing on Switches

Nonstop active routing (NSR) provides high availability for Routing Engines by enabling transparent switchover of the Routing Engines without necessitating restart of supported routing protocols. Both Routing Engines are fully active in processing protocol sessions, and so each can take over for the other. The switchover is transparent to neighbors.

This example describes how to configure nonstop active routing on switches with multiple Routing Engines or on an EX Series or a QFX series switch in a Virtual Chassis or Virtual Chassis Fabric configuration.

- [Requirements on page 226](#)
- [Overview and Topology on page 227](#)
- [Configuration on page 227](#)
- [Verification on page 228](#)
- [Troubleshooting on page 229](#)

Requirements

This example uses the following hardware and software components:

- An EX Series with multiple Routing Engines or on an EX Series or a QFX series switch in a Virtual Chassis or Virtual Chassis Fabric configuration
- Junos OS Release 10.4 or later for EX Series switches
- Junos OS Release 13.2X51-D20 or later for QFX Series switches

Overview and Topology

Configure nonstop active routing on any EX Series with multiple Routing Engines or on an EX Series or a QFX series switch in a Virtual Chassis or Virtual Chassis Fabric configuration. Nonstop active routing is advantageous in networks where neighbor routing devices do not support graceful restart protocol extensions.

The topology used in this example consists of an EX8200 switch with redundant Routing Engines connected to neighbor routing devices that are not configured to support graceful restart of protocols.

Configuration

CLI Quick Configuration To quickly configure nonstop active routing, copy the following commands and paste them into the switch terminal window:

```
[edit]
set chassis redundancy graceful-switchover
set routing-options nonstop-routing
set system commit synchronize
```

Step-by-Step Procedure To configure nonstop active routing on a switch:

1. Enable graceful Routing Engine switchover (GRES):

```
[edit chassis redundancy]
user@switch# set graceful-switchover
```

2. Enable nonstop active routing (by default, nonstop active routing is disabled):

```
[edit routing-options]
user@switch# set nonstop-routing
```

3. Synchronize configuration changes between the Routing Engines:

```
[edit system]
user@switch# set commit synchronize
```

If you try to commit the nonstop active routing configuration without including the **commit synchronize** statement, the commit fails.



NOTE: If the backup Routing Engine is down when you issue the commit, a warning is displayed and the candidate configuration is committed in the master Routing Engine. When the backup Routing Engine comes up, its configuration is automatically synchronized with that of the master. If you subsequently insert or bring up a backup Routing Engine, it automatically synchronizes its configuration with the master Routing Engine configuration.

Results

Check the results of the configuration:

```
[edit]
user@switch# show
chassis {
  redundancy {
    graceful-switchover;
  }
}
routing-options {
  nonstop-routing;
}
system {
  commit synchronize;
}
```

Verification

To confirm that the configuration is working properly, perform these tasks:

- [Verifying That Nonstop Active Routing Is Working Correctly on the Switch on page 228](#)

[Verifying That Nonstop Active Routing Is Working Correctly on the Switch](#)

Purpose Verify that nonstop active routing is enabled.

Action Issue the `show task replication` command:

```
user@switch# show task replication
Stateful Replication: Enabled
RE mode: Master

Protocol                Synchronization Status
OSPF                    Complete
RIP                     Complete
PIM                     Complete
RSVP                    Complete
```

Meaning This output shows that nonstop active routing (Stateful Replication) is enabled on master routing engine. If nonstop routing is not enabled, instead of the output shown above:

- On the backup routing engine the following error message is displayed: **“error: the routing subsystem is not running.”**
- On the master routing engine, the following output is displayed if nonstop routing is not enabled:

```
Stateful Replication: Disabled
RE mode: Master
```

Troubleshooting

To troubleshoot nonstop active routing, perform these tasks:

- [Investigating Problems with Synchronization of Routing Engines When NSR Is Enabled on page 229](#)

[Investigating Problems with Synchronization of Routing Engines When NSR Is Enabled](#)

Problem A protocol loses connectivity with neighbors after a graceful Routing Engine switchover (GRES) occurs with nonstop active routing (NSR) enabled.

Solution Use trace options to help isolate the problem and gather troubleshooting information. Using the information gathered from trace options, you can confirm or eliminate the synchronization of the Routing Engines as the cause of the loss of connectivity for the protocol. See [“Tracing Nonstop Active Routing Synchronization Events” on page 224](#).

Related Documentation

- [Configuring Nonstop Active Routing on Switches on page 220](#)
- [Tracing Nonstop Active Routing Synchronization Events on page 224](#)
- [Understanding Nonstop Active Routing on EX Series Switches on page 203](#)
- [Nonstop Active Routing Concepts on page 199](#)

PART 8

Configuring Graceful Restart

- [Understanding How Graceful Restart Enables Uninterrupted Packet Forwarding When a Router Is Restarted on page 233](#)
- [Graceful Restart System Requirements on page 241](#)
- [Configuring Graceful Restart on page 243](#)

CHAPTER 20

Understanding How Graceful Restart Enables Uninterrupted Packet Forwarding When a Router Is Restarted

- [Graceful Restart Concepts on page 233](#)
- [Graceful Restart for Aggregate and Static Routes on page 234](#)
- [Graceful Restart and Routing Protocols on page 235](#)
- [Graceful Restart and MPLS-Related Protocols on page 237](#)
- [Understanding Restart Signaling-Based Helper Mode Support for OSPF Graceful Restart on page 238](#)
- [Graceful Restart and Layer 2 and Layer 3 VPNs on page 239](#)
- [Graceful Restart on Logical Systems on page 240](#)

Graceful Restart Concepts

With routing protocols, any service interruption requires that an affected router recalculate adjacencies with neighboring routers, restore routing table entries, and update other protocol-specific information. An unprotected restart of a router can result in forwarding delays, route flapping, wait times stemming from protocol reconvergence, and even dropped packets. The main benefits of graceful restart are uninterrupted packet forwarding and temporary suppression of all routing protocol updates. Graceful restart enables a router to pass through intermediate convergence states that are hidden from the rest of the network.

Three main types of graceful restart are available on Juniper Networks routing platforms:

- Graceful restart for aggregate and static routes and for routing protocols—Provides protection for aggregate and static routes and for Border Gateway Protocol (BGP), End System-to-Intermediate System (ES-IS), Intermediate System-to-Intermediate System (IS-IS), Open Shortest Path First (OSPF), Routing Information Protocol (RIP), next-generation RIP (RIPng), and Protocol Independent Multicast (PIM) sparse mode routing protocols.
- Graceful restart for MPLS-related protocols—Provides protection for Label Distribution Protocol (LDP), Resource Reservation Protocol (RSVP), circuit cross-connect (CCC), and translational cross-connect (TCC). (Not supported on OCX Series switches.)

- Graceful restart for virtual private networks (VPNs)—Provides protection for Layer 2 and Layer 3 VPNs.

Graceful restart works similarly for routing protocols and MPLS protocols and combines components of these protocol types to enable graceful restart in VPNs. The main benefits of graceful restart are uninterrupted packet forwarding and temporary suppression of all routing protocol updates. Graceful restart thus enables a router to pass through intermediate convergence states that are hidden from the rest of the network.

Most graceful restart implementations define two types of routers—the restarting router and the helper router. The restarting router requires rapid restoration of forwarding state information so it can resume the forwarding of network traffic. The helper router assists the restarting router in this process. Graceful restart configuration statements typically affect either the restarting router or the helper router.

Related Documentation

- [Understanding High Availability Features on Juniper Networks Routers on page 3](#)
- [Graceful Restart System Requirements on page 241](#)
- [Graceful Restart for Aggregate and Static Routes on page 234](#)
- [Graceful Restart and Routing Protocols on page 235](#)
- [Graceful Restart and MPLS-Related Protocols on page 237](#)
- [Graceful Restart and Layer 2 and Layer 3 VPNs on page 239](#)
- [Graceful Restart on Logical Systems on page 240](#)
- [Configuring Graceful Restart on page 244](#)
- [Configuring Graceful Restart for QFabric Systems on page 282](#)

Graceful Restart for Aggregate and Static Routes

When you include the **graceful-restart** statement at the **[edit routing-options]** hierarchy level, any static routes or aggregated routes that have been configured are protected. Because no helper router assists in the restart, these routes are retained in the forwarding table while the router restarts (rather than being discarded or refreshed).

Related Documentation

- [Graceful Restart Concepts on page 233](#)
- [Graceful Restart System Requirements on page 241](#)
- [Enabling Graceful Restart on page 243](#)
- [Verifying Graceful Restart Operation on page 290](#)
- [Configuring Graceful Restart on page 244](#)

Graceful Restart and Routing Protocols

This section covers the following topics:

- [BGP on page 235](#)
- [IS-IS on page 235](#)
- [OSPF and OSPFv3 on page 235](#)
- [PIM Sparse Mode on page 236](#)
- [RIP and RIPv6 on page 237](#)

BGP

When a router enabled for BGP graceful restart restarts, it retains BGP peer routes in its forwarding table and marks them as stale. However, it continues to forward traffic to other peers (or receiving peers) during the restart. To reestablish sessions, the restarting router sets the “restart state” bit in the BGP OPEN message and sends it to all participating peers. The receiving peers reply to the restarting router with messages containing end-of-routing-table markers. When the restarting router or switch receives all replies from the receiving peers, the restarting router performs route selection, the forwarding table is updated, and the routes previously marked as stale are discarded. At this point, all BGP sessions are reestablished and the restarting peer can receive and process BGP messages as usual.

While the restarting router does its processing, the receiving peers also temporarily retain routing information. When a receiving peer detects a TCP transport reset, it retains the routes received and marks the routes as stale. After the session is reestablished with the restarting router or switch, the stale routes are replaced with updated route information.

IS-IS

Normally, IS-IS routers move neighbor adjacencies to the down state when changes occur. However, a router enabled for IS-IS graceful restart sends out Hello messages with the Restart Request (RR) bit set in a restart type length value (TLV) message. This indicates to neighboring routers that a graceful restart is in progress and to leave the IS-IS adjacency intact. The neighboring routers must interpret and implement restart signaling themselves. Besides maintaining the adjacency, the neighbors send complete sequence number PDUs (CSNPs) to the restarting router and flood their entire database.

The restarting router never floods any of its own link-state PDUs (LSPs), including pseudonode LSPs, to IS-IS neighbors while undergoing graceful restart. This enables neighbors to reestablish their adjacencies without transitioning to the down state and enables the restarting router to reinitiate a smooth database synchronization.

OSPF and OSPFv3

When a router enabled for OSPF graceful restart restarts, it retains routes learned before the restart in its forwarding table. The router does not allow new OSPF link-state advertisements (LSAs) to update the routing table. This router continues to forward traffic to other OSPF neighbors (or helper routers), and sends only a limited number of

LSAs during the restart period. To reestablish OSPF adjacencies with neighbors, the restarting router must send a grace LSA to all neighbors. In response, the helper routers enter helper mode and send an acknowledgement back to the restarting router. If there are no topology changes, the helper routers continue to advertise LSAs as if the restarting router had remained in continuous OSPF operation.

When the restarting router receives replies from all the helper routers, the restarting router selects routes, updates the forwarding table, and discards the old routes. At this point, full OSPF adjacencies are reestablished and the restarting router receives and processes OSPF LSAs as usual. When the helper routers no longer receive grace LSAs from the restarting router or the topology of the network changes, the helper routers also resume normal operation.



NOTE: For more information about the standard helper mode implementation, see RFC 3623, *Graceful OSPF Restart*.

Starting with Release 11.3, Junos OS supports the restart signaling-based helper mode for OSPF graceful restart configurations. The helper modes, both standard and restart signaling-based, are enabled by default. In restart signaling-based helper mode implementations, the restarting router relays the restart status to its neighbors only after the restart is complete. When the restart is complete, the restarting router sends hello messages to its helper routers with the **restart signal (RS)** bit set in the hello packet header. When a helper router receives a hello packet with the **RS** bit set in the header, the helper router returns a hello message to the restarting router. The reply hello message from the helper router contains the **ResyncState** flag and the **ResyncTimeout** timer that enable the restarting router to keep track of the helper routers that are syncing up with it. When all helpers complete the synchronization, the restarting router exits the restart mode.



NOTE:

For more information about restart signaling-based graceful restart helper mode implementation, see RFC 4811, *OSPF Out-of-Band Link State Database (LSDB) Resynchronization*, RFC 4812, *OSPF Restart Signaling*, and RFC 4813, *OSPF Link-Local Signaling*.

Restart signaling-based graceful restart helper mode is not supported for OSPFv3 configurations.

PIM Sparse Mode

PIM sparse mode uses a mechanism called a *generation identifier* to indicate the need for graceful restart. Generation identifiers are included by default in PIM hello messages. An initial generation identifier is created by each PIM neighbor to establish device capabilities. When one of the PIM neighbors restarts, it sends a new generation identifier to its neighbors. All neighbors that support graceful restart and are connected by point-to-point links assist by sending multicast updates to the restarting neighbor.

The restart phase completes when either the PIM state becomes stable or when the restart interval timer expires. If the neighbors do not support graceful restart or connect to each other using multipoint interfaces, the restarting router uses the restart interval timer to define the restart period.

RIP and RIPng

When a router enabled for RIP graceful restart restarts, routes that have been configured are protected. Because no helper router assists in the restart, these routes are retained in the forwarding table while the router restarts (rather than being discarded or refreshed).

Related Documentation

- [Graceful Restart Concepts on page 233](#)
- [Graceful Restart System Requirements on page 241](#)
- [Configuring Routing Protocols Graceful Restart on page 270](#)
- [Verifying Graceful Restart Operation on page 290](#)
- [Configuring Graceful Restart on page 244](#)
- [Example: Configuring IS-IS for GRES with Graceful Restart on page 153](#)

Graceful Restart and MPLS-Related Protocols

This section contains the following topics:

- [LDP on page 237](#)
- [RSVP on page 238](#)
- [CCC and TCC on page 238](#)

LDP

LDP graceful restart enables a router whose LDP control plane is undergoing a restart to continue to forward traffic while recovering its state from neighboring routers. It also enables a router on which helper mode is enabled to assist a neighboring router that is attempting to restart LDP.

During session initialization, a router advertises its ability to perform LDP graceful restart or to take advantage of a neighbor performing LDP graceful restart by sending the graceful restart TLV. This TLV contains two fields relevant to LDP graceful restart: the reconnect time and the recovery time. The values of the reconnect and recovery times indicate the graceful restart capabilities supported by the router.

The reconnect time is configured in Junos OS as 60 seconds and is not user-configurable. The reconnect time is how long the helper router waits for the restarting router to establish a connection. If the connection is not established within the reconnect interval, graceful restart for the LDP session is terminated. The maximum reconnect time is 120 seconds and is not user-configurable. The maximum reconnect time is the maximum value that a helper router accepts from its restarting neighbor.

When a router discovers that a neighboring router is restarting, it waits until the end of the recovery time before attempting to reconnect. The recovery time is the length of time a router waits for LDP to restart gracefully. The recovery time period begins when an initialization message is sent or received. This time period is also typically the length of time that a neighboring router maintains its information about the restarting router, so it can continue to forward traffic.

You can configure LDP graceful restart both in the master instance for the LDP protocol and for a specific routing instance. You can disable graceful restart at the global level for all protocols, at the protocol level for LDP only, and for a specific routing instance only.

RSVP

RSVP graceful restart enables a router undergoing a restart to inform its adjacent neighbors of its condition. The restarting router requests a grace period from the neighbor or peer, which can then cooperate with the restarting router. The restarting router can still forward MPLS traffic during the restart period; convergence in the network is not disrupted. The restart is not visible to the rest of the network, and the restarting router is not removed from the network topology. RSVP graceful restart can be enabled on both transit routers and ingress routers. It is available for both point-to-point LSPs and point-to-multipoint LSPs.

CCC and TCC

CCC and TCC graceful restart enables Layer 2 connections between customer edge (CE) routers to restart gracefully. These Layer 2 connections are configured with the **remote-interface-switch** or **lsp-switch** statements. Because these CCC and TCC connections have an implicit dependency on RSVP LSPs, graceful restart for CCC and TCC uses the RSVP graceful restart capabilities.

RSVP graceful restart must be enabled on the provider edge (PE) routers and provider (P) routers to enable graceful restart for CCC and TCC. Also, because RSVP is used as the signaling protocol for signaling label information, the neighboring router must use helper mode to assist with the RSVP restart procedures.

Related Documentation

- [Graceful Restart Concepts on page 233](#)
- [Graceful Restart System Requirements on page 241](#)
- [Configuring Graceful Restart for MPLS-Related Protocols on page 277](#)
- [Configuring Graceful Restart on page 244](#)

Understanding Restart Signaling-Based Helper Mode Support for OSPF Graceful Restart

Starting with Release 11.4, Junos OS supports restart signaling-based helper mode for OSPF graceful restart configurations.

**NOTE:**

- Restart signaling-based graceful restart helper mode is not supported for OSPFv3 configurations.
- Junos OS releases prior to Release 11.4 and OSPFv3 configurations support only standard helper mode as defined in RFC 3623 . For more information about the standard helper mode implementation, see RFC 3623 and the *Junos OS High Availability Configuration Guide*.

Both standard and restart signaling-based helper modes are enabled by default, irrespective of the graceful-restart configuration status on the device.

In restart signaling-based helper mode implementations, the restarting router informs the restart status to its neighbors only after the restart is complete. When the restart is complete, the restarting router sends hello messages to its helper routers with the **restart signal (RS)** bit set in the hello packet header. When a helper router receives a hello packet with the **RS** bit set in the header, the helper router returns a hello message to the restarting router. The reply hello message from the helper router contains the **ResyncState** flag and the **ResyncTimeout** timer that enable the restarting router to keep track of the helper routers that are syncing up with it. When all helpers complete the synchronization, the restarting router exits the restart mode.

For more information about restart signaling-based graceful restart helper mode implementation, see RFC 4811, *OSPF Out-of-Band Link State Database (LSDB) Resynchronization*, RFC 4812, *OSPF Restart Signaling* and RFC 4813, *OSPF Link-Local Signaling*.

**Related
Documentation**

- [Example: Managing Helper Modes for OSPF Graceful Restart on page 286](#)
- [Tracing Restart Signaling-Based Helper Mode Events for OSPF Graceful Restart on page 289](#)

Graceful Restart and Layer 2 and Layer 3 VPNs

VPN graceful restart uses three types of restart functionality:

1. BGP graceful restart functionality is used on all PE-to-PE BGP sessions. This affects sessions carrying any service signaling data for network layer reachability information (NLRI), for example, an IPv4 VPN or Layer 2 VPN NLRI.
2. OSPF, IS-IS, LDP, or RSVP graceful restart functionality is used in all core routers. Routes added by these protocols are used to resolve Layer 2 and Layer 3 VPN NLRI.
3. Protocol restart functionality is used for any Layer 3 protocol (RIP, OSPF, LDP, and so on) used between the PE and customer edge (CE) routers. This does not apply to Layer 2 VPNs because Layer 2 protocols used between the CE and PE routers do not have graceful restart capabilities.

Before VPN graceful restart can work properly, all of the components must restart gracefully. In other words, the routers must preserve their forwarding states and request neighbors to continue forwarding to the router in case of a restart. If all of the conditions are satisfied, VPN graceful restart imposes the following rules on a restarting router:

- The router must wait to receive all BGP NLRI information from other PE routers before advertising routes to the CE routers.
- The router must wait for all protocols in all routing instances to converge (or complete the restart process) before it sends CE router information to other PE routers. In other words, the router must wait for all instance information (whether derived from local configuration or advertisements received from a remote peer) to be processed before it sends this information to other PE routers.
- The router must preserve all forwarding state in the **instance.mpls.0** tables until the new labels and transit routes are allocated and announced to other PE routers (and CE routers in a carrier-of-carriers scenario).

If any condition is not met, VPN graceful restart does not succeed in providing uninterrupted forwarding between CE routers across the VPN infrastructure.

**Related
Documentation**

- [Graceful Restart Concepts on page 233](#)
- [Graceful Restart System Requirements on page 241](#)
- [Configuring Logical System Graceful Restart on page 281](#)
- [Verifying Graceful Restart Operation on page 290](#)
- [Configuring Graceful Restart on page 244](#)

Graceful Restart on Logical Systems

Graceful restart for a logical system functions much as graceful restart does in the main router. The only difference is the location of the **graceful-restart** statement:

- For a logical system, include the **graceful-restart** statement at the **[edit logical-systems *logical-system-name* routing-options]** hierarchy level.
- For a routing instance inside a logical system, include the **graceful-restart** statement at both the **[edit logical-systems *logical-system-name* routing-options]** and **[edit logical-systems *logical-system-name* routing-instances *instance-name* routing-options]** hierarchy levels.

**Related
Documentation**

- [Graceful Restart Concepts on page 233](#)
- [Graceful Restart System Requirements on page 241](#)
- [Configuring Logical System Graceful Restart on page 281](#)
- [Verifying Graceful Restart Operation on page 290](#)
- [Configuring Graceful Restart on page 244](#)

Graceful Restart System Requirements

- [Graceful Restart System Requirements on page 241](#)

Graceful Restart System Requirements

Graceful restart is supported on all routing platforms. To implement graceful restart for particular features, your system must meet these minimum requirements:

- Junos OS Release 5.3 or later for aggregate route, BGP, IS-IS, OSPF, RIP, RIPng, or static route graceful restart.
- Junos OS Release 5.5 or later for RSVP on egress provider edge (PE) routers.
- Junos OS Release 5.5 or later for LDP graceful restart.
- Junos OS Release 5.6 or later for the CCC, TCC, Layer 2 VPN, or Layer 3 VPN implementations of graceful restart.
- Junos OS Release 6.1 or later for RSVP graceful restart on ingress PE routers.
- Junos OS Release 6.4 or later for PIM sparse mode graceful restart.
- Junos OS Release 7.4 or later for ES-IS graceful restart.
- Junos OS Release 8.5 or later for BFD session (helper mode only)—If a node is undergoing a graceful restart and its BFD sessions are distributed to the Packet Forwarding Engine, the peer node can help the peer with the graceful restart.
- Junos OS Release 9.2 or later for BGP to support helper mode without requiring that graceful restart be configured.

**Related
Documentation**

- [Graceful Restart Concepts on page 233](#)

Configuring Graceful Restart

- [Enabling Graceful Restart on page 243](#)
- [Configuring Graceful Restart on page 244](#)
- [Configuring Routing Protocols Graceful Restart on page 270](#)
- [Configuring Graceful Restart for MPLS-Related Protocols on page 277](#)
- [Configuring VPN Graceful Restart on page 280](#)
- [Configuring Logical System Graceful Restart on page 281](#)
- [Configuring Graceful Restart for QFabric Systems on page 282](#)
- [Example: Managing Helper Modes for OSPF Graceful Restart on page 286](#)
- [Tracing Restart Signaling-Based Helper Mode Events for OSPF Graceful Restart on page 289](#)
- [Verifying Graceful Restart Operation on page 290](#)

Enabling Graceful Restart

Graceful restart is disabled by default. You must configure graceful restart at the **[edit routing-options]** or **[edit routing-instances *instance-name* routing-options]** hierarchy level to enable the feature globally.

For example:

```
routing-options {  
  graceful-restart;  
}
```

You can, optionally, modify the global settings at the individual protocol level or, as of Junos OS 15.1, at the individual routing instance level.



NOTE: If you configure graceful restart after a BGP or LDP session has been established, the BGP or LDP session restarts and the peers negotiate graceful restart capabilities.

To disable graceful restart, include the **disable** statement. You can do this globally for all protocols by including the **disable** statement at the **[edit routing-options]** hierarchy

level, or you can disable graceful restart for a single protocol by including the `disable` statement at the `[edit protocols protocol graceful-restart]` hierarchy level. To configure a time period for complete restart, include the `restart-duration` statement. You can specify a number between 120 and 900.

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

When you include the `graceful-restart` statement at the `[edit routing-options]` hierarchy level, graceful restart is also enabled for aggregate and static routes.

Release History Table

Release	Description
15.1	You can, optionally, modify the global settings at the individual protocol level or, as of Junos OS 15.1, at the individual routing instance level.

Related Documentation

- [Graceful Restart Concepts on page 233](#)
- [Graceful Restart System Requirements on page 241](#)
- [Graceful Restart for Aggregate and Static Routes on page 234](#)
- [Configuring Graceful Restart on page 244](#)

Configuring Graceful Restart

To enable graceful restart, include the `graceful-restart` statement at the `[edit routing-instance instance-name routing-options]` or `[edit routing-options]` hierarchy level. This enables graceful restart globally for all routing protocols. You can, optionally, modify or supplement the global settings at the individual protocol level.



NOTE: When set protocols `bgp group group-name allow network` is configured to accept dynamic BGP sessions, `unconfigured-peer-graceful-restart` statement should be configured to avoid traffic drop during graceful restart or graceful Routing Engine switchover.

For example:

```
protocols {
  bgp {
    group ext {
      graceful-restart {
        restart-time 400;
      }
    }
  }
}
routing-options {
```

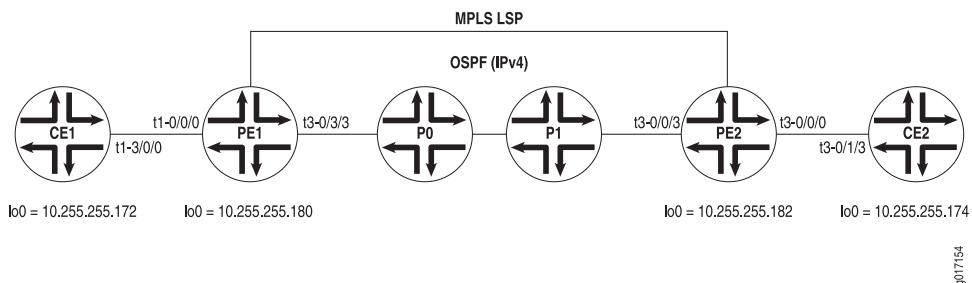
```

    graceful-restart;
}

```

Figure 19 on page 245 shows a standard MPLS VPN network. Routers CE1 and CE2 are customer edge routers, PE1 and PE2 are provider edge routers, and P0 is a provider core router. Several Layer 3 VPNs are configured across this network, as well as one Layer 2 VPN. Interfaces are shown in the diagram and are not included in the configuration example that follows.

Figure 19: Layer 3 VPN Graceful Restart Topology



Router CE1 On Router CE1, configure the following protocols on the logical interfaces of **t3-3/1/0**: OSPF on unit 101, RIP on unit 102, BGP on unit 103, and IS-IS on unit 512. Also configure graceful restart, BGP, IS-IS, OSPF, and RIP on the main instance to be able to connect to the routing instances on Router PE1.

```

[edit]
interfaces {
  t3-3/1/0 {
    encapsulation frame-relay;
    unit 100 {
      dlci 100;
      family inet {
        address 10.96.100.2/30;
      }
    }
    unit 101 {
      dlci 101;
      family inet {
        address 10.96.101.2/30;
      }
    }
    unit 102 {
      dlci 102;
      family inet {
        address 10.96.102.2/30;
      }
    }
    unit 103 {
      dlci 103;
      family inet {
        address 10.96.103.2/30;
      }
    }
  }
}

```

```
    unit 512 {
      dlci 512;
      family inet {
        address 10.96.252.1/30;
      }
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.245.14.172/32;
        primary;
      }
      address 10.96.110.1/32;
      address 10.96.111.1/32;
      address 10.96.112.1/32;
      address 10.96.113.1/32;
      address 10.96.116.1/32;
    }
    family iso {
      address 47.0005.80ff.f800.0000.0108.0001.0102.4501.4172.00;
    }
  }
}
routing-options {
  graceful-restart;
  autonomous-system 65100;
}
protocols {
  bgp {
    group CE-PE-INET {
      type external;
      export BGP_INET_LB_DIRECT;
      neighbor 10.96.103.1 {
        local-address 10.96.103.2;
        family inet {
          unicast;
        }
      }
      peer-as 65103;
    }
  }
  isis {
    export ISIS_L2VPN_LB_DIRECT;
    interface t3-3/1/0.512;
  }
  ospf {
    export OSPF_LB_DIRECT;
    area 0.0.0.0 {
      interface t3-3/1/0.101;
    }
  }
  rip {
    group RIP {
      export RIP_LB_DIRECT;
    }
  }
}
```

```
        neighbor t3-3/1/0.102;
    }
}
policy-options {
    policy-statement OSPF_LB_DIRECT {
        term direct {
            from {
                protocol direct;
                route-filter 10.96.101.0/30 exact;
                route-filter 10.96.111.1/32 exact;
            }
            then accept;
        }
        term final {
            then reject;
        }
    }
    policy-statement RIP_LB_DIRECT {
        term direct {
            from {
                protocol direct;
                route-filter 10.96.102.0/30 exact;
                route-filter 10.96.112.1/32 exact;
            }
            then accept;
        }
        term final {
            then reject;
        }
    }
    policy-statement BGP_INET_LB_DIRECT {
        term direct {
            from {
                protocol direct;
                route-filter 10.96.103.0/30 exact;
                route-filter 10.96.113.1/32 exact;
            }
            then accept;
        }
        term final {
            then reject;
        }
    }
    policy-statement ISIS_L2VPN_LB_DIRECT {
        term direct {
            from {
                protocol direct;
                route-filter 10.96.116.1/32 exact;
            }
            then accept;
        }
        term final {
            then reject;
        }
    }
}
```

```
}  
}
```

Router PE1 On Router PE1, configure graceful restart in the master instance, along with BGP, OSPF, MPLS, and LDP. Next, configure several protocol-specific instances of graceful restart. By including instances for BGP, OSPF, Layer 2 VPNs, RIP, and static routes, you can observe the wide range of options available when you implement graceful restart. Configure the following protocols in individual instances on the logical interfaces of **t3-0/0/0**: a static route on unit 100, OSPF on unit 101, RIP on unit 102, BGP on unit 103, and Frame Relay on unit 512 for the Layer 2 VPN instance.

```
[edit]  
interfaces {  
  t3-0/0/0 {  
    dce;  
    encapsulation frame-relay-ccc;  
    unit 100 {  
      dlci 100;  
      family inet {  
        address 10.96.100.1/30;  
      }  
      family mpls;  
    }  
    unit 101 {  
      dlci 101;  
      family inet {  
        address 10.96.101.1/30;  
      }  
      family mpls;  
    }  
    unit 102 {  
      dlci 102;  
      family inet {  
        address 10.96.102.1/30;  
      }  
      family mpls;  
    }  
    unit 103 {  
      dlci 103;  
      family inet {  
        address 10.96.103.1/30;  
      }  
      family mpls;  
    }  
    unit 512 {  
      encapsulation frame-relay-ccc;  
      dlci 512;  
    }  
  }  
}  
t1-0/1/0 {  
  unit 0 {  
    family inet {  
      address 10.96.0.2/30;  
    }  
  }  
}
```

```

    }
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 10.245.14.176/32;
    }
    family iso {
      address 47.0005.80ff.f800.0000.0108.0001.0102.4501.4176.00;
    }
  }
}
}
routing-options {
  graceful-restart;
  router-id 10.245.14.176;
  autonomous-system 69;
}
protocols {
  mpls {
    interface all;
  }
  bgp {
    group PEPE {
      type internal;
      neighbor 10.245.14.182 {
        local-address 10.245.14.176;
        family inet-vpn {
          unicast;
        }
        family l2vpn {
          unicast;
        }
      }
    }
  }
}
ospf {
  area 0.0.0.0 {
    interface t1-0/1/0.0;
    interface fxp0.0 {
      disable;
    }
    interface lo0.0 {
      passive;
    }
  }
}
}
ldp {
  interface all;
}
}
policy-options {
  policy-statement STATIC-import {

```

```
        from community STATIC;
        then accept;
    }
    policy-statement STATIC-export {
        then {
            community add STATIC;
            accept;
        }
    }
    policy-statement OSPF-import {
        from community OSPF;
        then accept;
    }
    policy-statement OSPF-export {
        then {
            community add OSPF;
            accept;
        }
    }
    policy-statement RIP-import {
        from community RIP;
        then accept;
    }
    policy-statement RIP-export {
        then {
            community add RIP;
            accept;
        }
    }
    policy-statement BGP-INET-import {
        from community BGP-INET;
        then accept;
    }
    policy-statement BGP-INET-export {
        then {
            community add BGP-INET;
            accept;
        }
    }
    policy-statement L2VPN-import {
        from community L2VPN;
        then accept;
    }
    policy-statement L2VPN-export {
        then {
            community add L2VPN;
            accept;
        }
    }
    community BGP-INET members target:69:103;
    community L2VPN members target:69:512;
    community OSPF members target:69:101;
    community RIP members target:69:102;
    community STATIC members target:69:100;
}
```



```

routing-instances {
  BGP-INET {
    instance-type vrf;
    interface t3-0/0/0.103;
    route-distinguisher 10.245.14.176:103;
    vrf-import BGP-INET-import;
    vrf-export BGP-INET-export;
    routing-options {
      graceful-restart;
      autonomous-system 65103;
    }
    protocols {
      bgp {
        group BGP-INET {
          type external;
          export BGP-INET-import;
          neighbor 10.96.103.2 {
            local-address 10.96.103.1;
            family inet {
              unicast;
            }
          }
          peer-as 65100;
        }
      }
    }
  }
  L2VPN {
    instance-type l2vpn;
    interface t3-0/0/0.512;
    route-distinguisher 10.245.14.176:512;
    vrf-import L2VPN-import;
    vrf-export L2VPN-export;
    protocols {# There is no graceful-restart statement for Layer 2 VPN instances.
      l2vpn {
        encapsulation-type frame-relay;
        site CE1-ISIS {
          site-identifier 512;
          interface t3-0/0/0.512 {
            remote-site-id 612;
          }
        }
      }
    }
  }
  OSPF {
    instance-type vrf;
    interface t3-0/0/0.101;
    route-distinguisher 10.245.14.176:101;
    vrf-import OSPF-import;
    vrf-export OSPF-export;
    routing-options {
      graceful-restart;
    }
    protocols {

```

```
        ospf {
            export OSPF-import;
            area 0.0.0.0 {
                interface all;
            }
        }
    }
}
RIP {
    instance-type vrf;
    interface t3-0/0/0.102;
    route-distinguisher 10.245.14.176:102;
    vrf-import RIP-import;
    vrf-export RIP-export;
    routing-options {
        graceful-restart;
    }
    protocols {
        rip {
            group RIP {
                export RIP-import;
                neighbor t3-0/0/0.102;
            }
        }
    }
}
STATIC {
    instance-type vrf;
    interface t3-0/0/0.100;
    route-distinguisher 10.245.14.176:100;
    vrf-import STATIC-import;
    vrf-export STATIC-export;
    routing-options {
        graceful-restart;
        static {
            route 10.96.110.1/32 next-hop t3-0/0/0.100;
        }
    }
}
}
```

Router P0 On Router P0, configure graceful restart in the main instance, along with OSPF, MPLS, and LDP. This allows the protocols on the PE routers to reach one another.

```
[edit]
interfaces {
    t3-0/1/3 {
        unit 0 {
            family inet {
                address 10.96.0.5/30;
            }
            family mpls;
        }
    }
}
```

```

}
t1-0/2/0 {
  unit 0 {
    family inet {
      address 10.96.0.1/30;
    }
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 10.245.14.174/32;
    }
    family iso {
      address 47.0005.80ff.f800.0000.0108.0001.0102.4501.4174.00;
    }
  }
}
}
}
routing-options {
  graceful-restart;
  router-id 10.245.14.174;
  autonomous-system 69;
}
protocols {
  mpls {
    interface all;
  }
  ospf {
    area 0.0.0.0 {
      interface t1-0/2/0.0;
      interface t3-0/1/3.0;
      interface fxp0.0 {
        disable;
      }
      interface lo0.0 {
        passive;
      }
    }
  }
  ldp {
    interface all;
  }
}
}

```

Router PE2 On Router PE2, configure BGP, OSPF, MPLS, LDP, and graceful restart in the master instance. Configure the following protocols in individual instances on the logical interfaces of **t1-0/1/3**: a static route on unit 200, OSPF on unit 201, RIP on unit 202, BGP on unit 203, and Frame Relay on unit 612 for the Layer 2 VPN instance. Also configure protocol-specific graceful restart in all routing instances, except the Layer 2 VPN instance.

[edit]

```
interfaces {
  t3-0/0/0 {
    unit 0 {
      family inet {
        address 10.96.0.6/30;
      }
      family mpls;
    }
  }
  t1-0/1/3 {
    dce;
    encapsulation frame-relay-ccc;
    unit 200 {
      dlci 200;
      family inet {
        address 10.96.200.1/30;
      }
      family mpls;
    }
    unit 201 {
      dlci 201;
      family inet {
        address 10.96.201.1/30;
      }
      family mpls;
    }
    unit 202 {
      dlci 202;
      family inet {
        address 10.96.202.1/30;
      }
      family mpls;
    }
    unit 203 {
      dlci 203;
      family inet {
        address 10.96.203.1/30;
      }
      family mpls;
    }
    unit 612 {
      encapsulation frame-relay-ccc;
      dlci 612;
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.245.14.182/32;
      }
      family iso {
        address 47.0005.80ff.f800.0000.0108.0001.0102.4501.4182.00;
      }
    }
  }
}
```

```
}
routing-options {
  graceful-restart;
  router-id 10.245.14.182;
  autonomous-system 69;
}
protocols {
  mpls {
    interface all;
  }
  bgp {
    group PEPE {
      type internal;
      neighbor 10.245.14.176 {
        local-address 10.245.14.182;
        family inet-vpn {
          unicast;
        }
        family l2vpn {
          unicast;
        }
      }
    }
  }
}
ospf {
  area 0.0.0.0 {
    interface t3-0/0/0.0;
    interface fxp0.0 {
      disable;
    }
    interface lo0.0 {
      passive;
    }
  }
}
ldp {
  interface all;
}
policy-options {
  policy-statement STATIC-import {
    from community STATIC;
    then accept;
  }
  policy-statement STATIC-export {
    then {
      community add STATIC;
      accept;
    }
  }
  policy-statement OSPF-import {
    from community OSPF;
    then accept;
  }
  policy-statement OSPF-export {
    then {
```

```
        community add OSPF;
        accept;
    }
}
policy-statement RIP-import {
    from community RIP;
    then accept;
}
policy-statement RIP-export {
    then {
        community add RIP;
        accept;
    }
}
policy-statement BGP-INET-import {
    from community BGP-INET;
    then accept;
}
policy-statement BGP-INET-export {
    then {
        community add BGP-INET;
        accept;
    }
}
policy-statement L2VPN-import {
    from community L2VPN;
    then accept;
}
policy-statement L2VPN-export {
    then {
        community add L2VPN;
        accept;
    }
}
community BGP-INET members target:69:103;
community L2VPN members target:69:512;
community OSPF members target:69:101;
community RIP members target:69:102;
community STATIC members target:69:100;
}
routing-instances {
    BGP-INET {
        instance-type vrf;
        interface t1-0/1/3.203;
        route-distinguisher 10.245.14.182:203;
        vrf-import BGP-INET-import;
        vrf-export BGP-INET-export;
        routing-options {
            graceful-restart;
            autonomous-system 65203;
        }
        protocols {
            bgp {
                group BGP-INET {
                    type external;
                }
            }
        }
    }
}
```

```

        export BGP-INET-import;
        neighbor 10.96.203.2 {
            local-address 10.96.203.1;
            family inet {
                unicast;
            }
            peer-as 65200;
        }
    }
}
}
}
}
L2VPN {
    instance-type l2vpn;
    interface t1-0/1/3.612;
    route-distinguisher 10.245.14.182:612;
    vrf-import L2VPN-import;
    vrf-export L2VPN-export;
    protocols {# There is no graceful-restart statement for Layer 2 VPN instances.
        l2vpn {
            encapsulation-type frame-relay;
            site CE2-ISIS {
                site-identifier 612;
                interface t1-0/1/3.612 {
                    remote-site-id 512;
                }
            }
        }
    }
}
OSPF {
    instance-type vrf;
    interface t1-0/1/3.201;
    route-distinguisher 10.245.14.182:201;
    vrf-import OSPF-import;
    vrf-export OSPF-export;
    routing-options {
        graceful-restart;
    }
    protocols {
        ospf {
            export OSPF-import;
            area 0.0.0.0 {
                interface all;
            }
        }
    }
}
RIP {
    instance-type vrf;
    interface t1-0/1/3.202;
    route-distinguisher 10.245.14.182:202;
    vrf-import RIP-import;
    vrf-export RIP-export;
    routing-options {

```

```

        graceful-restart;
    }
    protocols {
        rip {
            group RIP {
                export RIP-import;
                neighbor t1-0/1/3.202;
            }
        }
    }
}
STATIC {
    instance-type vrf;
    interface t1-0/1/3.200;
    route-distinguisher 10.245.14.182:200;
    vrf-import STATIC-import;
    vrf-export STATIC-export;
    routing-options {
        graceful-restart;
        static {
            route 10.96.210.1/32 next-hop t1-0/1/3.200;
        }
    }
}
}
}

```

Router CE2 On Router CE2, complete the Layer 2 and Layer 3 VPN configuration by mirroring the protocols already set on Routers PE2 and CE1. Specifically, configure the following on the logical interfaces of **t1-0/0/3**: OSPF on unit 201, RIP on unit 202, BGP on unit 203, and IS-IS on unit 612. Finally, configure graceful restart, BGP, IS-IS, OSPF, and RIP on the main instance to be able to connect to the routing instances on Router PE2.

```

[edit]
interfaces {
    t1-0/0/3 {
        encapsulation frame-relay;
        unit 200 {
            dlci 200;
            family inet {
                address 10.96.200.2/30;
            }
        }
        unit 201 {
            dlci 201;
            family inet {
                address 10.96.201.2/30;
            }
        }
        unit 202 {
            dlci 202;
            family inet {
                address 10.96.202.2/30;
            }
        }
    }
}

```



```

    }
  }
  unit 203 {
    dlci 203;
    family inet {
      address 10.96.203.2/30;
    }
  }
  unit 512 {
    dlci 512;
    family inet {
      address 10.96.252.2/30;
    }
  }
}
lo0 {
  unit 0 {
    family inet {
      address 10.245.14.180/32 {
        primary;
      }
      address 10.96.210.1/32;
      address 10.96.111.1/32;
      address 10.96.212.1/32;
      address 10.96.213.1/32;
      address 10.96.216.1/32;
    }
    family iso {
      address 47.0005.80ff.f800.0000.0108.0001.0102.4501.4180.00;
    }
  }
}
}
routing-options {
  graceful-restart;
  autonomous-system 65200;
}
protocols {
  bgp {
    group CE-PE-INET {
      type external;
      export BGP_INET_LB_DIRECT;
      neighbor 10.96.203.1 {
        local-address 10.96.203.2;
        family inet {
          unicast;
        }
      }
      peer-as 65203;
    }
  }
}
isis {
  export ISIS_L2VPN_LB_DIRECT;
  interface t1-0/0/3.612;
}

```

```
ospf {
  export OSPF_LB_DIRECT;
  area 0.0.0.0 {
    interface t1-0/0/3.201;
  }
}
rip {
  group RIP {
    export RIP_LB_DIRECT;
    neighbor t1-0/0/3.202;
  }
}
}
policy-options {
  policy-statement OSPF_LB_DIRECT {
    term direct {
      from {
        protocol direct;
        route-filter 10.96.201.0/30 exact;
        route-filter 10.96.211.1/32 exact;
      }
      then accept;
    }
    term final {
      then reject;
    }
  }
  policy-statement RIP_LB_DIRECT {
    term direct {
      from {
        protocol direct;
        route-filter 10.96.202.0/30 exact;
        route-filter 10.96.212.1/32 exact;
      }
      then accept;
    }
    term final {
      then reject;
    }
  }
  policy-statement BGP_INET_LB_DIRECT {
    term direct {
      from {
        protocol direct;
        route-filter 10.96.203.0/30 exact;
        route-filter 10.96.213.1/32 exact;
      }
      then accept;
    }
    term final {
      then reject;
    }
  }
  policy-statement ISIS_L2VPN_LB_DIRECT {
    term direct {
```

```

    from {
        protocol direct;
        route-filter 10.96.216.1/32 exact;
    }
    then accept;
}
term final {
    then reject;
}
}
}

```

Router PE1 Status Before a Restart

The following example displays neighbor relationships on Router PE1 before a restart happens:

```
user@PE1> show bgp neighbor
```

```

Peer: 10.96.103.2+3785 AS 65100 Local: 10.96.103.1+179 AS 65103
  Type: External   State: Established   Flags: <>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Export: [ BGP-INET-import ]
  Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily PeerAS
Refresh>
  Address families configured: inet-unicast
  Local Address: 10.96.103.1 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 10.96.110.1      Local ID: 10.96.103.1      Active Holdtime: 90
  Keepalive Interval: 30
  Local Interface: t3-0/0/0.103
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast
  NLRI for this session: inet-unicast
  Peer supports Refresh capability (2)
  Restart time configured on the peer: 120
  Stale routes from peer are kept for: 300
  Restart time requested by this peer: 120
  NLRI that peer supports restart for: inet-unicast
  NLRI peer can save forwarding state: inet-unicast
  NLRI that peer saved forwarding for: inet-unicast
  NLRI that restart is negotiated for: inet-unicast
  NLRI of all end-of-rib markers sent: inet-unicast
  Table BGP-INET.inet.0 Bit: 30001
    RIB State: BGP restart is complete
    RIB State: VPN restart is complete
    Send state: in sync
    Active prefixes:          0
    Received prefixes:        0
    Suppressed due to damping: 0
  Last traffic (seconds): Received 8    Sent 3    Checked 3
  Input messages: Total 15    Updates 0    Refreshes 0    Octets 321
  Output messages: Total 18    Updates 2    Refreshes 0    Octets 450
  Output Queue[2]: 0

Peer: 10.245.14.182+4701 AS 69    Local: 10.245.14.176+179 AS 69
  Type: Internal   State: Established   Flags: <>

```

```

Last State: OpenConfirm    Last Event: RecvKeepAlive
Last Error: None
Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily
Rib-group Refresh>
Address families configured: inet-vpn-unicast l2vpn
Local Address: 10.245.14.176 Holdtime: 90 Preference: 170
Number of flaps: 1
Peer ID: 10.245.14.182    Local ID: 10.245.14.176    Active Holdtime: 90
Keepalive Interval: 30
NLRI for restart configured on peer: inet-vpn-unicast l2vpn
NLRI advertised by peer: inet-vpn-unicast l2vpn
NLRI for this session: inet-vpn-unicast l2vpn
Peer supports Refresh capability (2)
Restart time configured on the peer: 120
Stale routes from peer are kept for: 300
Restart time requested by this peer: 120
NLRI that peer supports restart for: inet-vpn-unicast l2vpn
NLRI peer can save forwarding state: inet-vpn-unicast l2vpn
NLRI that peer saved forwarding for: inet-vpn-unicast l2vpn
NLRI that restart is negotiated for: inet-vpn-unicast l2vpn
NLRI of all end-of-rib markers sent: inet-vpn-unicast l2vpn
Table bgp.l3vpn.0 Bit: 10000
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: in sync
  Active prefixes:          0
  Received prefixes:        0
  Suppressed due to damping: 0
Table bgp.l2vpn.0 Bit: 20000
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: in sync
  Active prefixes:          1
  Received prefixes:        1
  Suppressed due to damping: 0
Table BGP-INET.inet.0 Bit: 30000
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: in sync
  Active prefixes:          0
  Received prefixes:        0
  Suppressed due to damping: 0
Table OSPF.inet.0 Bit: 60000
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: in sync
  Active prefixes:          0
  Received prefixes:        0
  Suppressed due to damping: 0
Table RIP.inet.0 Bit: 70000
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: in sync
  Active prefixes:          0
  Received prefixes:        0
  Suppressed due to damping: 0
Table STATIC.inet.0 Bit: 80000
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: in sync

```

```

Active prefixes:          0
Received prefixes:        0
Suppressed due to damping: 0
Table L2VPN.l2vpn.0 Bit: 90000
RIB State: BGP restart is complete
RIB State: VPN restart is complete
Send state: in sync
Active prefixes:          1
Received prefixes:        1
Suppressed due to damping: 0
Last traffic (seconds): Received 28   Sent 28   Checked 28
Input messages: Total 2     Updates 0     Refreshes 0     Octets 86
Output messages: Total 13   Updates 10    Refreshes 0     Octets 1073
Output Queue[0]: 0
Output Queue[1]: 0
Output Queue[2]: 0
Output Queue[3]: 0
Output Queue[4]: 0
Output Queue[5]: 0
Output Queue[6]: 0
Output Queue[7]: 0
Output Queue[8]: 0

user@PE1> show route instance detail
master:
  Router ID: 10.245.14.176
  Type: forwarding          State: Active
  Restart State: Complete Path selection timeout: 300
  Tables:
    inet.0                  : 17 routes (15 active, 0 holddown, 1 hidden)
    Restart Complete
    inet.3                  : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
    iso.0                   : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Complete
    mpls.0                  : 19 routes (19 active, 0 holddown, 0 hidden)
    Restart Complete
    bgp.l3vpn.0             : 10 routes (10 active, 0 holddown, 0 hidden)
    Restart Complete
    inet6.0                 : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
    bgp.l2vpn.0             : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Complete
BGP-INET:
  Router ID: 10.96.103.1
  Type: vrf                 State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.103
  Route-distinguisher: 10.245.14.176:103
  Vrf-import: [ BGP-INET-import ]
  Vrf-export: [ BGP-INET-export ]
  Tables:
    BGP-INET.inet.0         : 4 routes (4 active, 0 holddown, 0 hidden)
    Restart Complete
L2VPN:
  Router ID: 0.0.0.0
  Type: l2vpn               State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:

```

```

    t3-0/0/0.512
    Route-distinguisher: 10.245.14.176:512
    Vrf-import: [ L2VPN-import ]
    Vrf-export: [ L2VPN-export ]
    Tables:
      L2VPN.l2vpn.0      : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
OSPF:
  Router ID: 10.96.101.1
  Type: vrf      State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.101
  Route-distinguisher: 10.245.14.176:101
  Vrf-import: [ OSPF-import ]
  Vrf-export: [ OSPF-export ]
  Tables:
    OSPF.inet.0      : 8 routes (7 active, 0 holddown, 0 hidden)
  Restart Complete
RIP:
  Router ID: 10.96.102.1
  Type: vrf      State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.102
  Route-distinguisher: 10.245.14.176:102
  Vrf-import: [ RIP-import ]
  Vrf-export: [ RIP-export ]
  Tables:
    RIP.inet.0      : 6 routes (6 active, 0 holddown, 0 hidden)
  Restart Complete
STATIC:
  Router ID: 10.96.100.1
  Type: vrf      State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.100
  Route-distinguisher: 10.245.14.176:100
  Vrf-import: [ STATIC-import ]
  Vrf-export: [ STATIC-export ]
  Tables:
    STATIC.inet.0   : 4 routes (4 active, 0 holddown, 0 hidden)
  Restart Complete
__juniper_private1__:
  Router ID: 0.0.0.0
  Type: forwarding      State: Active

user@PE1> show route protocol l2vpn
inet.0: 16 destinations, 17 routes (15 active, 0 holddown, 1 hidden)
Restart Complete
inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
BGP-INET.inet.0: 5 destinations, 6 routes (5 active, 0 holddown, 0 hidden)
Restart Complete
OSPF.inet.0: 7 destinations, 8 routes (7 active, 0 holddown, 0 hidden)
Restart Complete
RIP.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
Restart Complete
STATIC.inet.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

```

```

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete
mpls.0: 20 destinations, 20 routes (20 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
800003          *[L2VPN/7] 00:06:00
                 > via t3-0/0/0.512, Pop          Offset: 4
t3-0/0/0.512     *[L2VPN/7] 00:06:00
                 > via t1-0/1/0.0, Push 800003, Push 100004(top) Offset: -4
bgp.l3vpn.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Restart Complete
inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
L2VPN.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
10.245.14.176:512:512:611/96
                 *[L2VPN/7] 00:06:01
                 Discard

bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

```

Router PE1 Status During a Restart

Before you can verify that graceful restart is working, you must simulate a router restart. To cause the routing process to refresh and simulate a restart, use the **restart routing** operational mode command:

```
user@PE1> restart routing
```

```
Routing protocol daemon started, pid 3558
```

The following sample output is captured during the router restart:

```
user@PE1> show bgp neighbor
```

```

Peer: 10.96.103.2      AS 65100 Local: 10.96.103.1      AS 65103
  Type: External      State: Active      Flags: <ImportEval>
  Last State: Idle      Last Event: Start
  Last Error: None
  Export: [ BGP-INET-import ]
  Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily PeerAS
Refresh>
  Address families configured: inet-unicast
  Local Address: 10.96.103.1 Holdtime: 90 Preference: 170
  Number of flaps: 0
Peer: 10.245.14.182+179 AS 69      Local: 10.245.14.176+2131 AS 69
  Type: Internal      State: Established      Flags: <ImportEval>
  Last State: OpenConfirm      Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily
Rib-group Refresh>
  Address families configured: inet-vpn-unicast l2vpn
  Local Address: 10.245.14.176 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 10.245.14.182      Local ID: 10.245.14.176      Active Holdtime: 90
  Keepalive Interval: 30

```

```

NLRI for restart configured on peer: inet-vpn-unicast 12vpn
NLRI advertised by peer: inet-vpn-unicast 12vpn
NLRI for this session: inet-vpn-unicast 12vpn
Peer supports Refresh capability (2)
Restart time configured on the peer: 120
Stale routes from peer are kept for: 300
Restart time requested by this peer: 120
NLRI that peer supports restart for: inet-vpn-unicast 12vpn
NLRI peer can save forwarding state: inet-vpn-unicast 12vpn
NLRI that peer saved forwarding for: inet-vpn-unicast 12vpn
NLRI that restart is negotiated for: inet-vpn-unicast 12vpn
NLRI of received end-of-rib markers: inet-vpn-unicast 12vpn
Table bgp.13vpn.0 Bit: 10000
  RIB State: BGP restart in progress
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          10
  Received prefixes:        10
  Suppressed due to damping: 0
Table bgp.12vpn.0 Bit: 20000
  RIB State: BGP restart in progress
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          1
  Received prefixes:        1
  Suppressed due to damping: 0
Table BGP-INET.inet.0 Bit: 30000
  RIB State: BGP restart in progress
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          2
  Received prefixes:        2
  Suppressed due to damping: 0
Table OSPF.inet.0 Bit: 60000
  RIB State: BGP restart is complete
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          2
  Received prefixes:        2
  Suppressed due to damping: 0
Table RIP.inet.0 Bit: 70000
  RIB State: BGP restart is complete
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          2
  Received prefixes:        2
  Suppressed due to damping: 0
Table STATIC.inet.0 Bit: 80000
  RIB State: BGP restart is complete
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          1
  Received prefixes:        1
  Suppressed due to damping: 0
Table L2VPN.12vpn.0 Bit: 90000
  RIB State: BGP restart is complete
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          1
  Received prefixes:        1

```



```

    Suppressed due to damping: 0
    Last traffic (seconds): Received 0    Sent 0    Checked 0
    Input messages: Total 14    Updates 13    Refreshes 0    Octets 1053
    Output messages: Total 3    Updates 0    Refreshes 0    Octets 105
    Output Queue[0]: 0
    Output Queue[1]: 0
    Output Queue[2]: 0
    Output Queue[3]: 0
    Output Queue[4]: 0
    Output Queue[5]: 0
    Output Queue[6]: 0
    Output Queue[7]: 0
    Output Queue[8]: 0

user@PE1> show route instance detail
master:
  Router ID: 10.245.14.176
  Type: forwarding      State: Active
  Restart State: Pending Path selection timeout: 300
  Tables:
    inet.0                : 17 routes (15 active, 1 holddown, 1 hidden)
    Restart Pending: OSPF LDP
    inet.3                : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Pending: OSPF LDP
    iso.0                 : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Complete
    mpls.0                : 23 routes (23 active, 0 holddown, 0 hidden)
    Restart Pending: LDP VPN
    bgp.l3vpn.0           : 10 routes (10 active, 0 holddown, 0 hidden)
    Restart Pending: BGP VPN
    inet6.0               : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
    bgp.l2vpn.0           : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Pending: BGP VPN
BGP-INET:
  Router ID: 10.96.103.1
  Type: vrf              State: Active
  Restart State: Pending Path selection timeout: 300
  Interfaces:
    t3-0/0/0.103
  Route-distinguisher: 10.245.14.176:103
  Vrf-import: [ BGP-INET-import ]
  Vrf-export: [ BGP-INET-export ]
  Tables:
    BGP-INET.inet.0       : 6 routes (5 active, 0 holddown, 0 hidden)
    Restart Pending: VPN
L2VPN:
  Router ID: 0.0.0.0
  Type: l2vpn            State: Active
  Restart State: Pending Path selection timeout: 300
  Interfaces:
    t3-0/0/0.512
  Route-distinguisher: 10.245.14.176:512
  Vrf-import: [ L2VPN-import ]
  Vrf-export: [ L2VPN-export ]
  Tables:
    L2VPN.l2vpn.0         : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Pending: VPN L2VPN
OSPF:
  Router ID: 10.96.101.1

```

```

Type: vrf                      State: Active
Restart State: Pending Path selection timeout: 300
Interfaces:
  t3-0/0/0.101
Route-distinguisher: 10.245.14.176:101
Vrf-import: [ OSPF-import ]
Vrf-export: [ OSPF-export ]
Tables:
  OSPF.inet.0                  : 8 routes (7 active, 1 holddown, 0 hidden)
Restart Pending: OSPF VPN
RIP:
Router ID: 10.96.102.1
Type: vrf                      State: Active
Restart State: Pending Path selection timeout: 300
Interfaces:
  t3-0/0/0.102
Route-distinguisher: 10.245.14.176:102
Vrf-import: [ RIP-import ]
Vrf-export: [ RIP-export ]
Tables:
  RIP.inet.0                   : 8 routes (6 active, 2 holddown, 0 hidden)
Restart Pending: RIP VPN
STATIC:
Router ID: 10.96.100.1
Type: vrf                      State: Active
Restart State: Pending Path selection timeout: 300
Interfaces:
  t3-0/0/0.100
Route-distinguisher: 10.245.14.176:100
Vrf-import: [ STATIC-import ]
Vrf-export: [ STATIC-export ]
Tables:
  STATIC.inet.0                : 4 routes (4 active, 0 holddown, 0 hidden)
Restart Pending: VPN
__juniper_private1__:
Router ID: 0.0.0.0
Type: forwarding              State: Active

```

```
user@PE1> show route instance summary
```

Instance	Type	Primary rib	Active/holddown/hidden
master	forwarding	inet.0	15/0/1
		iso.0	1/0/0
		mpls.0	35/0/0
		l3vpn.0	0/0/0
		inet6.0	2/0/0
		l2vpn.0	0/0/0
		l2circuit.0	0/0/0
BGP-INET	vrf	BGP-INET.inet.0	5/0/0
		BGP-INET.iso.0	0/0/0
		BGP-INET.inet6.0	0/0/0
L2VPN	l2vpn	L2VPN.inet.0	0/0/0
		L2VPN.iso.0	0/0/0
		L2VPN.inet6.0	0/0/0
		L2VPN.l2vpn.0	2/0/0
OSPF	vrf		

```

RIP          vrf          OSPF.inet.0          7/0/0
                OSPF.iso.0          0/0/0
                OSPF.inet6.0        0/0/0
RIP          vrf          RIP.inet.0          6/0/0
                RIP.iso.0          0/0/0
                RIP.inet6.0        0/0/0
STATIC       vrf          STATIC.inet.0        4/0/0
                STATIC.iso.0        0/0/0
                STATIC.inet6.0      0/0/0
__juniper_private1__ forwarding
                __juniper_priva.inet.0 0/0/0
                __juniper_privat.iso.0 0/0/0
                __juniper_priv.inet6.0 0/0/0

user@PE1> show route protocol l2vpn

inet.0: 16 destinations, 17 routes (15 active, 1 holddown, 1 hidden)
Restart Pending: OSPF LDP

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Pending: OSPF LDP

BGP-INET.inet.0: 5 destinations, 6 routes (5 active, 0 holddown, 0 hidden)
Restart Pending: VPN

OSPF.inet.0: 7 destinations, 8 routes (7 active, 1 holddown, 0 hidden)
Restart Pending: OSPF VPN

RIP.inet.0: 6 destinations, 8 routes (6 active, 2 holddown, 0 hidden)
Restart Pending: RIP VPN

STATIC.inet.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Pending: VPN

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 24 destinations, 24 routes (24 active, 0 holddown, 0 hidden)
Restart Pending: LDP VPN
+ = Active Route, - = Last Active, * = Both

800001          *[L2VPN/7] 00:00:13
                > via t3-0/0/0.512, Pop          Offset: 4
t3-0/0/0.512    *[L2VPN/7] 00:00:13
                > via t1-0/1/0.0, Push 800003, Push 100004(top) Offset: -4

bgp.l3vpn.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Restart Pending: BGP VPN

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

L2VPN.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Pending: VPN L2VPN
+ = Active Route, - = Last Active, * = Both

10.245.14.176:512:512:611/96
                *[L2VPN/7] 00:00:13

```

Discard

bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Pending: BGP VPN

Related Documentation

- [Enabling Graceful Restart on page 243](#)
- [Configuring Routing Protocols Graceful Restart on page 270](#)
- [Configuring Graceful Restart for MPLS-Related Protocols on page 277](#)
- [Configuring VPN Graceful Restart on page 280](#)
- [Configuring Logical System Graceful Restart on page 281](#)
- [Verifying Graceful Restart Operation on page 290](#)

Configuring Routing Protocols Graceful Restart

This topic includes the following sections:

- [Enabling Graceful Restart on page 270](#)
- [Configuring Graceful Restart Options for BGP on page 271](#)
- [Using Control Plane Dependent BFD along with Graceful Restart Helper Mode on page 272](#)
- [Configuring Graceful Restart Options for ES-IS on page 273](#)
- [Configuring Graceful Restart Options for IS-IS on page 273](#)
- [Configuring Graceful Restart Options for OSPF and OSPFv3 on page 274](#)
- [Configuring Graceful Restart Options for RIP and RIPng on page 275](#)
- [Configuring Graceful Restart Options for PIM Sparse Mode on page 276](#)
- [Tracking Graceful Restart Events on page 277](#)

Enabling Graceful Restart

By default, graceful restart is disabled. To enable graceful restart, include the **graceful-restart** statement at the **[edit routing-instance *instance-name* routing-options]** or **[edit routing-options]** hierarchy level.

For example:

```
routing-options {  
  graceful-restart;  
}
```

To configure the duration of the graceful restart period, include the **restart-duration** at the **[edit routing-options graceful-restart]** hierarchy level.



NOTE: Helper mode (the ability to assist a neighboring router attempting a graceful restart) is enabled by default when you start the routing platform, even if graceful restart is not enabled. You can disable helper mode on a per-protocol basis.

```
[edit]
routing-options {
  graceful-restart {
    disable;
    restart-duration seconds;
  }
}
```

To disable graceful restart globally, include the **disable** statement at the **[edit routing-options graceful-restart]** hierarchy level.

When graceful restart is enabled for all routing protocols at the **[edit routing-options graceful-restart]** hierarchy level, you can disable graceful restart on a per-protocol basis.



NOTE: If you configure graceful restart after a BGP or LDP session has been established, the BGP or LDP session restarts and the peers negotiate graceful restart capabilities. Also, the BGP peer routing statistics are reset to zero.

Configuring Graceful Restart Options for BGP

To configure the duration of the BGP graceful restart period, include the **restart-time** statement at the **[edit protocols bgp graceful-restart]** hierarchy level. To set the length of time the router waits to receive messages from restarting neighbors before declaring them down, include the **stale-routes-time** statement at the **[edit protocols bgp graceful-restart]** hierarchy level.

```
[edit]
protocols {
  bgp {
    graceful-restart {
      disable;
      restart-time seconds;
      stale-routes-time seconds;
    }
  }
}
routing-options {
  graceful-restart;
}
```

To disable BGP graceful restart capability for all BGP sessions, include the **disable** statement at the **[edit protocols bgp graceful-restart]** hierarchy level.



NOTE: To set BGP graceful restart properties or disable them for a group, include the desired statements at the `[edit protocols bgp group group-name graceful-restart]` hierarchy level.

To set BGP graceful restart properties or disable them for a specific neighbor in a group, include the desired statements at the `[edit protocols bgp group group-name neighbor ip-address graceful-restart]` hierarchy level.



NOTE: Configuring graceful restart for BGP resets the BGP peer routing statistics to zero. Also, existing BGP sessions restart, and the peers negotiate graceful restart capabilities.



NOTE: Do not configure both Bidirectional Forwarding Detection (BFD) for BGP and graceful restart for BGP. Routing performance may be sub-optimal if you do this.

Using Control Plane Dependent BFD along with Graceful Restart Helper Mode

When BFD is control plane dependent and the device detects a BFD down event and is not already entering the graceful restart helper mode, this is treated as a regular BFD down event and the device enters the graceful restart helper mode. This behavior makes the control plane dependent BFD unusable in conjunction with graceful restart.

Include the `dont-help-shared-fate-bfd-down` statement at the `[edit protocols bgp graceful-restart]` hierarchy to ensure that the device does not enter the graceful restart helper mode and data traffic continues to be forwarded to an alternate path even if there is an interface failure (without a control plane restart on the BGP neighbor).

```
[edit]
protocols {
  bgp {
    graceful-restart {
      disable;
      dont-help-shared-fate-bfd-down;
      restart-time seconds;
      stale-routes-time seconds;
    }
  }
}
routing-options {
  graceful-restart;
}
```

Starting in Junos OS Release 18.3R1, you can prevent SRX Series devices from entering the graceful restart helper mode when the device is configured with BFD with a single-hop external BGP (EBGP), by including the `dont-help-shared-fate-bfd-down` statement at the `[edit protocols bgp graceful-restart]` hierarchy.

See Also • [dont-help-shared-fate-bfd-down on page 520](#)

Configuring Graceful Restart Options for ES-IS

On J Series Services Routers, to configure the duration of the ES-IS graceful restart period, include the **restart-duration** statement at the **[edit protocols esis graceful-restart]** hierarchy level.

```
[edit]
protocols {
  esis {
    graceful-restart {
      disable;
      restart-duration seconds;
    }
  }
}
routing-options {
  graceful-restart;
}
```

To disable ES-IS graceful restart capability, include the **disable** statement at the **[edit protocols esis graceful-restart]** hierarchy level.

Configuring Graceful Restart Options for IS-IS

To configure the duration of the IS-IS graceful restart period, include the **restart-duration** statement at the **[edit protocols isis graceful-restart]** hierarchy level.

```
[edit]
protocols {
  isis {
    graceful-restart {
      disable;
      helper-disable;
      restart-duration seconds;
    }
  }
}
routing-options {
  graceful-restart;
}
```

To disable IS-IS graceful restart helper capability, include the **helper-disable** statement at the **[edit protocols isis graceful-restart]** hierarchy level. To disable IS-IS graceful restart capability, include the **disable** statement at the **[edit protocols isis graceful-restart]** hierarchy level.



NOTE: Starting with Junos OS Release 12.3, if adjacencies between the Routing Engine and the neighboring peer 'helper' routers time out, graceful restart protocol extensions are unable to notify the peer 'helper' routers about the impending restart. Graceful restart can then stop and cause interruptions in traffic.

To ensure that these adjacencies are kept, change the hold-time for IS-IS protocols from the default of 27 seconds to a value higher than 40 seconds.



NOTE: You can also track graceful restart events with the `traceoptions` statement at the `[edit protocols isis]` hierarchy level. For more information, see [“Tracking Graceful Restart Events” on page 277](#).

Configuring Graceful Restart Options for OSPF and OSPFv3

To configure the duration of the OSPF/OSPFv3 graceful restart period, include the `restart-duration` statement at the `[edit protocols (ospf | ospf3) graceful-restart]` hierarchy level. To specify the length of time for which the router notifies helper routers that it has completed graceful restart, include the `notify-duration` at the `[edit protocols (ospf | ospf3) graceful-restart]` hierarchy level. Strict OSPF link-state advertisement (LSA) checking results in the termination of graceful restart by a helping router. To disable strict LSA checking, include the `no-strict-lsa-checking` statement at the `[edit protocols (ospf | ospf3) graceful-restart]` hierarchy level.

```
[edit]
protocols {
  ospf | ospfv3 {
    graceful-restart {
      disable;
      helper-disable
      no-strict-lsa-checking;
      notify-duration seconds;
      restart-duration seconds;
    }
  }
}
routing-options {
  graceful-restart;
}
```

To disable OSPF/OSPFv3 graceful restart, include the `disable` statement at the `[edit protocols (ospf | ospf3) graceful-restart]` hierarchy level.

Starting with Release 11.3, the Junos OS supports both the standard (based on RFC 3623, *Graceful OSPF Restart*) and the restart signaling-based (as specified in RFC 4811, RFC 4812, and RFC 4813) helper modes for OSPF version 2 graceful restart configurations. Both the standard and restart signaling-based helper modes are enabled by default. To disable the helper mode for OSPF version 2 graceful restart configurations, include the `helper-disable <both | restart-signaling | standard>` statement at the `[edit protocols ospf`

graceful-restart] hierarchy level. Note that the last committed statement always takes precedence over the previous one.

```
[edit protocols ospf]
  graceful-restart {
    helper-disable <both | restart-signaling | standard>
  }
```

To reenabling the helper mode, delete the **helper-disable** statement from the configuration by using the **delete protocols ospf graceful-restart helper-disable <restart-signaling | standard | both>** command. In this case also, the last executed command takes precedence over the previous ones.



NOTE:

Restart signaling-based helper mode is not supported for OSPFv3 configurations. To disable helper mode for OSPFv3 configurations, include the **helper-disable** statement at the **[edit protocols ospfv3 graceful-restart]** hierarchy level.



TIP: You can also track graceful restart events with the **traceoptions** statement at the **[edit protocols (ospf | ospfv3)]** hierarchy level. For more information, see [“Tracking Graceful Restart Events” on page 277](#).



NOTE: You cannot enable OSPFv3 graceful restart between a routing platform running Junos OS Release 7.5 and earlier and a routing platform running Junos OS Release 7.6 or later. As a workaround, make sure both routing platforms use the same Junos OS version.

Configuring Graceful Restart Options for RIP and RIPng

To configure the duration of the RIP or RIPng graceful restart period, include the **restart-time** statement at the **[edit protocols (rip | ripng) graceful-restart]** hierarchy level.

```
[edit]
  protocols {
    (rip | ripng) {
      graceful-restart {
        disable;
        restart-time seconds;
      }
    }
  }
  routing-options {
    graceful-restart;
  }
```

To disable RIP or RIPng graceful restart capability, include the **disable** statement at the **[edit protocols (rip | ripng) graceful-restart]** hierarchy level.

Configuring Graceful Restart Options for PIM Sparse Mode

PIM sparse mode continues to forward existing multicast packet streams during a graceful restart, but does not forward new streams until after the restart is complete. After a restart, the routing platform updates the forwarding state with any updates that were received from neighbors and occurred during the restart period. For example, the routing platform relearns the join and prune states of neighbors during the restart, but does not apply the changes to the forwarding table until after the restart.

PIM sparse mode-enabled routing platforms generate a unique 32-bit random number called a generation identifier. Generation identifiers are included by default in PIM hello messages, as specified in the IETF Internet draft *Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)*. When a routing platform receives PIM hellos containing generation identifiers on a point-to-point interface, Junos OS activates an algorithm that optimizes graceful restart.

Before PIM sparse mode graceful restart occurs, each routing platform creates a generation identifier and sends it to its multicast neighbors. If a PIM sparse mode-enabled routing platform restarts, it creates a new generation identifier and sends it to its neighbors. When a neighbor receives the new identifier, it resends multicast updates to the restarting router to allow it to exit graceful restart efficiently. The restart phase completes when either the PIM state becomes stable or when the restart interval timer expires.

If a routing platform does not support generation identifiers or if PIM is enabled on multipoint interfaces, the PIM sparse mode graceful restart algorithm does not activate, and a default restart timer is used as the restart mechanism.

To configure the duration of the PIM graceful restart period, include the **restart-duration** statement at the **[edit protocols pim graceful-restart]** hierarchy level:

```
[edit]
protocols {
  pim {
    graceful-restart {
      disable;
      restart-duration seconds;
    }
  }
}
routing-options {
  graceful-restart;
}
```

To disable PIM sparse mode graceful restart capability, include the **disable** statement at the **[edit protocols pim graceful-restart]** hierarchy level.



NOTE: Multicast forwarding can be interrupted in two ways. First, if the underlying routing protocol is unstable, multicast reverse-path-forwarding (RPF) checks can fail and cause an interruption. Second, because the forwarding table is not updated during the graceful restart period, new multicast streams are not forwarded until graceful restart is complete.

Tracking Graceful Restart Events

To track the progress of a graceful restart event, you can configure graceful restart trace options flags for IS-IS and OSPF/OSPFv3. To configure graceful restart trace options, include the **graceful-restart** statement at the **[edit protocols *protocol* traceoptions flag]** hierarchy level:

```
[edit protocols]
isis {
  traceoptions {
    flag graceful-restart;
  }
}
(ospf | ospf3) {
  traceoptions {
    flag graceful-restart;
  }
}
```

Release History Table

Release	Description
12.3	Starting with Junos OS Release 12.3, if adjacencies between the Routing Engine and the neighboring peer 'helper' routers time out, graceful restart protocol extensions are unable to notify the peer 'helper' routers about the impending restart.

Related Documentation

- [Graceful Restart Concepts on page 233](#)
- [Graceful Restart System Requirements on page 241](#)
- [Graceful Restart and Routing Protocols on page 235](#)
- [Verifying Graceful Restart Operation on page 290](#)
- [Configuring Graceful Restart on page 244](#)

Configuring Graceful Restart for MPLS-Related Protocols

This section contains the following topics:

- [Configuring Graceful Restart Globally on page 278](#)
- [Configuring Graceful Restart Options for RSVP, CCC, and TCC on page 278](#)
- [Configuring Graceful Restart Options for LDP on page 279](#)

Configuring Graceful Restart Globally

To configure graceful restart globally for all MPLS-related protocols, include the **graceful-restart** statement at the **[edit routing-options]** hierarchy level. To configure the duration of the graceful restart period, include the **restart-duration** at the **[edit routing-options graceful-restart]** hierarchy level:

```
[edit]
routing-options {
  graceful-restart {
    disable;
    restart-duration seconds;
  }
}
```

To disable graceful restart globally, include the **disable** statement at the **[edit routing-options graceful-restart]** hierarchy level.

Configuring Graceful Restart Options for RSVP, CCC, and TCC

Because CCC and TCC rely on RSVP, you must modify these three protocols as a single group.

To configure how long the router retains the state of its RSVP neighbors while they undergo a graceful restart, include the **maximum-helper-recovery-time** statement at the **[edit protocols rsvp graceful-restart]** hierarchy level. This value is applied to all neighboring routers, so it should be based on the time required by the slowest RSVP neighbor to recover.

To configure the delay between when the router discovers that a neighboring router has gone down and when it declares the neighbor down, include the **maximum-helper-restart-time** statement at the **[edit protocols rsvp graceful-restart]** hierarchy level. This value is applied to all neighboring routers, so it should be based on the time required by the slowest RSVP neighbor to restart.

```
[edit]
protocols {
  rsvp {
    graceful-restart {
      disable;
      helper-disable;
      maximum-helper-recovery-time;
      maximum-helper-restart-time;
    }
  }
}
routing-options {
  graceful-restart;
}
```

To disable RSVP, CCC, and TCC graceful restart, include the **disable** statement at the **[edit protocols rsvp graceful-restart]** hierarchy level. To disable RSVP, CCC, and TCC

helper capability, include the **helper-disable** statement at the **[edit protocols rsvp graceful-restart]** hierarchy level.

Configuring Graceful Restart Options for LDP

When configuring graceful restart for LDP, you can include the following optional statements at the **[edit protocols ldp graceful-restart]** hierarchy level:

```
[edit protocols ldp graceful-restart]
disable;
helper-disable;
maximum-neighbor-reconnect-time seconds;
maximum-neighbor-recovery-time seconds;
reconnect-time seconds;
recovery-time seconds;
[edit routing-options]
graceful-restart;
```

The statements have the following effects on the graceful restart process:

- To configure the length of time required to reestablish a session after a graceful restart, include the **reconnect-time** statement; the range is 30 through 300 seconds. To limit the maximum reconnect time allowed from a restarting neighbor router, include the **maximum-neighbor-reconnect-time** statement; the range is 30 through 300 seconds.
- To configure the length of time that helper routers are required to maintain the old forwarding state during a graceful restart, include the **recovery-time** statement; the range is 120 through 1800 seconds. On the helper router, you can configure a statement that overrides the request from the restarting router and sets the maximum length of time the helper router will maintain the old forwarding state. To configure this feature, include the **maximum-neighbor-recovery-time** statement; the range is 140 through 1900 seconds.



NOTE: The value for the **recovery-time** and **maximum-neighbor-recovery-time** statements at the **[edit protocols ldp graceful-restart]** hierarchy level should be approximately 80 seconds longer than the value for the **restart-duration** statement at the **[edit routing-options graceful-restart]** hierarchy level. Otherwise, a warning message appears when you try to commit the configuration.

- To disable LDP graceful restart capability, include the **disable** statement. To disable LDP graceful restart helper capability, include the **helper-disable** statement.

- See Also**
- [Graceful Restart Concepts on page 233](#)
 - [Graceful Restart System Requirements on page 241](#)
 - [Graceful Restart and MPLS-Related Protocols on page 237](#)
 - [Verifying Graceful Restart Operation on page 290](#)
 - [Configuring Graceful Restart on page 244](#)

Configuring VPN Graceful Restart

Graceful restart allows a router whose VPN control plane is undergoing a restart to continue to forward traffic while recovering its state from neighboring routers. Without graceful restart, a control plane restart disrupts any VPN services provided by the router. Graceful restart is supported on Layer 2 VPNs, Layer 3 VPNs, virtual-router routing instances, and VPLS.

To implement graceful restart for a Layer 2 VPN or Layer 3 VPN, perform the configuration tasks described in the following sections:

- [Configuring Graceful Restart Globally on page 280](#)
- [Configuring Graceful Restart for the Routing Instance on page 280](#)

Configuring Graceful Restart Globally

To enable graceful restart, include the **graceful-restart** statement at the **[edit routing-options]** hierarchy level. To configure a global duration for the graceful restart period, include the **restart-duration** statement at the **[edit routing-options graceful-restart]** hierarchy level.

```
[edit]
routing-options {
  graceful-restart {
    disable;
    restart-duration seconds;
  }
}
```

To disable graceful restart globally, include the **disable** statement at the **[edit routing-options graceful-restart]** hierarchy level.

Configuring Graceful Restart for the Routing Instance

For Layer 3 VPNs only, you must also configure graceful restart for all routing and MPLS-related protocols within a routing instance by including the **graceful-restart** statement at the **[edit routing-instances instance-name routing-options]** hierarchy level. Because you can configure multi-instance BGP and multi-instance LDP, graceful restart for a carrier-of-carriers scenario is supported. To configure the duration of the graceful restart period for the routing instance, include the **restart-duration** statement at the **[edit routing-instances instance-name routing-options]**.

```
[edit]
routing-instances {
  instance-name {
    routing-options {
      graceful-restart {
        disable;
        restart-duration seconds;
      }
    }
  }
}
```

```
}
```

You can disable graceful restart for individual protocols with the **disable** statement at the **[edit routing-instances *instance-name* protocols *protocol-name* graceful-restart]** hierarchy level.

Related Documentation

- [Graceful Restart Concepts on page 233](#)
- [Graceful Restart System Requirements on page 241](#)
- [Graceful Restart and Layer 2 and Layer 3 VPNs on page 239](#)
- [Verifying Graceful Restart Operation on page 290](#)
- [Configuring Graceful Restart on page 244](#)

Configuring Logical System Graceful Restart

Graceful restart for a logical system functions much as graceful restart does in the main router. The only difference is the location of the **graceful-restart** statement.

The following topics describe what to configure to implement graceful restart in a logical system:

- [Enabling Graceful Restart Globally on page 281](#)
- [Configuring Graceful Restart for a Routing Instance on page 282](#)

Enabling Graceful Restart Globally

To enable graceful restart in a logical system, include the **graceful-restart** statement at the **[edit logical-systems *logical-system-name* routing-options]** hierarchy level. To configure a global duration of the graceful restart period, include the **restart-duration** statement at the **[edit logical-systems *logical-system-name* routing-options graceful-restart]** hierarchy level.

```
[edit]
logical-systems {
  logical-system-name {
    routing-options {
      graceful-restart {
        disable;
        restart-duration seconds;
      }
    }
  }
}
```

To disable graceful restart globally, include the **disable** statement at the **[edit logical-systems *logical-system-name* routing-options graceful-restart]** hierarchy level.

Configuring Graceful Restart for a Routing Instance

For Layer 3 VPNs only, you must also configure graceful restart globally for a routing instance inside a logical system. To configure, include the **graceful-restart** statement at the **[edit logical-systems *logical-system-name* routing-instances *instance-name* routing-options]** hierarchy level. Because you can configure multi-instance BGP and multi-instance LDP, graceful restart for a carrier-of-carriers scenario is supported. To configure the duration of the graceful restart period for the routing instance, include the **restart-duration** statement at the **[edit logical-systems *logical-system-name* routing-instances *instance-name* routing-options]**.

```
[edit]
logical-systems {
  logical-system-name {
    routing-instances {
      instance-name {
        routing-options {
          graceful-restart {
            disable;
            restart-duration seconds;
          }
        }
      }
    }
  }
}
```

To disable graceful restart for individual protocols with the **disable** statement at the **[edit logical-systems *logical-system-name* routing-instances *instance-name* protocols *protocol-name* graceful-restart]** hierarchy level.

Related Documentation

- [Graceful Restart Concepts on page 233](#)
- [Graceful Restart System Requirements on page 241](#)
- [Graceful Restart on Logical Systems on page 240](#)
- [Verifying Graceful Restart Operation on page 290](#)
- [Configuring Graceful Restart on page 244](#)

Configuring Graceful Restart for QFabric Systems

When you configure graceful restart in the QFabric CLI, the QFabric system applies the configuration to the network Node group to participate in graceful restart operations with devices external to the QFabric system. Such configuration preserves routing table state and helps neighboring routing devices to resume routing operations more quickly after a system restart. This also enables the network Node group to resume routing operations rapidly if there is a restart in the QFabric system (such as a software upgrade). As a result, we recommend enabling graceful restart for routing protocols in the QFabric CLI.



NOTE: The QFabric system also uses graceful restart internally within the fabric to facilitate interfabric resiliency and recovery. This internal feature is enabled by default with no configuration required.

- [Enabling Graceful Restart on page 283](#)
- [Configuring Graceful Restart Options for BGP on page 284](#)
- [Configuring Graceful Restart Options for OSPF and OSPFv3 on page 285](#)
- [Tracking Graceful Restart Events on page 286](#)

Enabling Graceful Restart

By default, graceful restart is disabled. To enable graceful restart, include the **graceful-restart** statement at the **[edit routing-instance *instance-name* routing-options]** or **[edit routing-options]** hierarchy level.

For example:

```
routing-options {
  graceful-restart;
}
```

To configure the duration of the graceful restart period, include the **restart-duration** at the **[edit routing-options graceful-restart]** hierarchy level.



NOTE: Helper mode (the ability to assist a neighboring router attempting a graceful restart) is enabled by default when you start the routing platform, even if graceful restart is not enabled. You can disable helper mode on a per-protocol basis.

```
[edit]
routing-options {
  graceful-restart {
    disable;
    restart-duration seconds;
  }
}
```

To disable graceful restart globally, include the **disable** statement at the **[edit routing-options graceful-restart]** hierarchy level.

When graceful restart is enabled for all routing protocols at the **[edit routing-options graceful-restart]** hierarchy level, you can disable graceful restart on a per-protocol basis.



NOTE: If you configure graceful restart after a BGP or LDP session has been established, the BGP or LDP session restarts and the peers negotiate graceful restart capabilities. Also, the BGP peer routing statistics are reset to zero.

Configuring Graceful Restart Options for BGP

To configure the duration of the BGP graceful restart period, include the **restart-time** statement at the **[edit protocols bgp graceful-restart]** hierarchy level. To set the length of time the router waits to receive messages from restarting neighbors before declaring them down, include the **stale-routes-time** statement at the **[edit protocols bgp graceful-restart]** hierarchy level.

```
[edit]
protocols {
  bgp {
    graceful-restart {
      disable;
      restart-time seconds;
      stale-routes-time seconds;
    }
  }
}
routing-options {
  graceful-restart;
}
```

To disable BGP graceful restart capability for all BGP sessions, include the **disable** statement at the **[edit protocols bgp graceful-restart]** hierarchy level.



NOTE: To set BGP graceful restart properties or disable them for a group, include the desired statements at the **[edit protocols bgp group *group-name* graceful-restart]** hierarchy level.

To set BGP graceful restart properties or disable them for a specific neighbor in a group, include the desired statements at the **[edit protocols bgp group *group-name* neighbor *ip-address* graceful-restart]** hierarchy level.



NOTE: Configuring graceful restart for BGP resets the BGP peer routing statistics to zero. Also, existing BGP sessions restart, and the peers negotiate graceful restart capabilities.

Configuring Graceful Restart Options for OSPF and OSPFv3

To configure the duration of the OSPF/OSPFv3 graceful restart period, include the **restart-duration** statement at the **[edit protocols (ospf | ospf3) graceful-restart]** hierarchy level. To specify the length of time for which the router notifies helper routers that it has completed graceful restart, include the **notify-duration** at the **[edit protocols (ospf | ospf3) graceful-restart]** hierarchy level. Strict OSPF link-state advertisement (LSA) checking results in the termination of graceful restart by a helping router. To disable strict LSA checking, include the **no-strict-lsa-checking** statement at the **[edit protocols (ospf | ospf3) graceful-restart]** hierarchy level.

```
[edit]
protocols {
  ospf | ospfv3 {
    graceful-restart {
      disable;
      helper-disable
      no-strict-lsa-checking;
      notify-duration seconds;
      restart-duration seconds;
    }
  }
}
routing-options {
  graceful-restart;
}
```

To disable OSPF/OSPFv3 graceful restart, include the **disable** statement at the **[edit protocols (ospf | ospf3) graceful-restart]** hierarchy level.

Starting with Release 11.3, the Junos OS supports both the standard (based on RFC 3623, *Graceful OSPF Restart*) and the restart signaling-based (as specified in RFC 4811, RFC 4812, and RFC 4813) helper modes for OSPF version 2 graceful restart configurations. Both the standard and restart signaling-based helper modes are enabled by default. To disable the helper mode for OSPF version 2 graceful restart configurations, include the **helper-disable <both | restart-signaling | standard>** statement at the **[edit protocols ospf graceful-restart]** hierarchy level. Note that the last committed statement always takes precedence over the previous one.

```
[edit protocols ospf]
graceful-restart {
  helper-disable <both | restart-signaling | standard>
}
```

To reenabling the helper mode, delete the **helper-disable** statement from the configuration by using the **delete protocols ospf graceful-restart helper-disable <restart-signaling | standard | both>** command. In this case also, the last executed command takes precedence over the previous ones.

**NOTE:**

Restart signaling-based helper mode is not supported for OSPFv3 configurations. To disable helper mode for OSPFv3 configurations, include the **helper-disable** statement at the [edit protocols ospfv3 graceful-restart] hierarchy level.



TIP: You can also track graceful restart events with the traceoptions statement at the [edit protocols (ospf | ospf3)] hierarchy level. For more information, see “Tracking Graceful Restart Events” on page 277.



NOTE: If you configure BFD and graceful restart for OSPF, graceful restart might not work as expected.

Tracking Graceful Restart Events

To track the progress of a graceful restart event, you can configure graceful restart trace options flags for IS-IS and OSPF/OSPFv3. To configure graceful restart trace options, include the **graceful-restart** statement at the [edit protocols *protocol* traceoptions flag] hierarchy level:

```
[edit protocols]
isis {
  traceoptions {
    flag graceful-restart;
  }
}
(ospf | ospf3) {
  traceoptions {
    flag graceful-restart;
  }
}
```

**Related
Documentation**

- [Graceful Restart Concepts on page 233](#)
- [Verifying Graceful Restart Operation on page 290](#)

Example: Managing Helper Modes for OSPF Graceful Restart

- [Requirements on page 287](#)
- [Overview on page 287](#)
- [Configuration on page 287](#)
- [Verification on page 288](#)

Requirements

M Series or T Series routers running Junos OS Release 11.4 or later and EX Series switches.

Overview

Junos OS Release 11.4 extends OSPF graceful restart support to include restart signaling-based helper mode. Both standard (RFC 3623-based) and restart signaling-based helper modes are enabled by default, irrespective of the graceful-restart configuration status on the routing device.

Junos OS, however, enables you to choose between the helper modes with the **helper-disable <standard | restart-signaling | both>** statement.

Configuration

Both standard and restart signaling-based helper modes are enabled by default, irrespective of the graceful-restart configuration status on the routing device. Junos OS allows you to disable or enable the helper modes based on your requirements.

To configure the helper mode options for graceful restart:

1. To enable graceful restart, add the **graceful-restart** statement at the **[edit routing-options]** hierarchy level.

```
[edit routing-options]
user@host# set graceful-restart
```

The helper modes, both standard and restart signaling-based, are enabled by default.

2. To disable one or both of the helper modes, add the **helper-disable <both | restart-signaling | standard>** statement at the **[edit protocols ospf graceful-restart]** hierarchy level.

- To disable both standard and restart signaling-based helper modes:

```
[edit protocols ospf graceful-restart]
user@host# set helper-disable both
```

- To disable only the restart signaling-based helper mode:

```
[edit protocols ospf graceful-restart]
user@host# set helper-disable restart-signaling
```

- To disable only the standard helper mode:

```
[edit protocols ospf graceful-restart]
user@host# set helper-disable standard
```



NOTE: You must commit the configuration before the change takes effect.

The last committed statement always takes precedence over the previous one.

3. To enable one or both of the helper modes when the helper modes are disabled, delete the **helper-disable <both | restart-signaling | standard>** statement from the **[edit protocols ospf graceful-restart]** hierarchy level.

- To enable both standard and restart signaling-based helper modes:

```
[edit protocols ospf graceful-restart]
user@host# delete helper-disable
```

- To enable the restart signaling-based helper mode:

```
[edit protocols ospf graceful-restart]
user@host# delete helper-disable restart-signaling
```

- To enable the standard helper mode:

```
[edit protocols ospf graceful-restart]
user@host# delete helper-disable standard
```



NOTE: You must commit the configuration before the change takes effect.

The last committed statement always takes precedence over the previous one.

Verification

Confirm that the configuration is working properly.

- [Verifying OSPF Graceful Restart and Helper Mode Configuration on page 288](#)

Verifying OSPF Graceful Restart and Helper Mode Configuration

Purpose Verify the OSPF graceful restart and helper mode configuration on a router.

- Action**
- Enter the `run show ospf overview` command from configuration mode.

```
user@host# run show ospf overview

~
~
~
Restart: Enabled
Restart duration: 180 sec
Restart grace period: 210 sec
Graceful restart helper mode: Enabled
Restart-signaling helper mode: Enabled
~
~
~
```

Meaning The output shows that graceful restart and both of the helper modes are enabled.

- Related Documentation**
- [Understanding Restart Signaling-Based Helper Mode Support for OSPF Graceful Restart on page 238](#)
 - [Tracing Restart Signaling-Based Helper Mode Events for OSPF Graceful Restart on page 289](#)
 - [helper-disable \(OSPF\) on page 528](#)

Tracing Restart Signaling-Based Helper Mode Events for OSPF Graceful Restart

Junos OS provides a tracing option to log restart signaling-based helper mode events for OSPF graceful restart. To enable tracing for restart signaling-based helper mode events, include the `traceoptions flag restart-signaling` statement at the `[edit protocols ospf]` hierarchy level.

To enable tracing for restart signaling-based events:

1. Create a log file for saving the log.

```
[edit protocols ospf]
user@host# set traceoptions file ospf-log
```

where *ospf-log* is the name of the log file.

2. Enable tracing for restart signaling-based helper mode events.

```
[edit protocols ospf]
user@host# set traceoptions flag restart-signaling
```

3. Commit the configuration.

```
[edit protocols ospf]
user@host# commit
```

The logs are saved to the **ospf-log** file in the **/var/log** folder.

Viewing the Log File

To view the restart signaling-based events from the log file, type:

```
user@host> file show /var/log/ospf-log | match "restart signaling"
Jun 25 14:44:08.890216 OSPF Restart Signaling: Start helper mode for nbr ip
14.19.3.2 id 10.10.10.1
Jun 25 14:44:11.358636 OSPF restart signaling: Received DBD with R bit set from
nbr ip=14.19.3.2 id=10.10.10.1. Start oob-resync.
Jun 25 14:44:11.380198 OSPF restart signaling: Received DBD with LR bit on from
nbr ip=14.19.3.2 id=10.10.10.1. Save its oob-resync capability 1
Jun 25 14:44:11.467200 OSPF restart signaling: nbr fsm for nbr ip=14.19.3.2
id=10.10.10.1 moving to state Full. Reset oob-resync parameters.
```

Related Documentation

- [Understanding Restart Signaling-Based Helper Mode Support for OSPF Graceful Restart on page 238](#)
- [Example: Managing Helper Modes for OSPF Graceful Restart on page 286](#)
- [helper-disable \(OSPF\) on page 528](#)

Verifying Graceful Restart Operation

This topic contains the following sections:

- [Graceful Restart Operational Mode Commands on page 290](#)
- [Verifying BGP Graceful Restart on page 291](#)
- [Verifying IS-IS and OSPF Graceful Restart on page 291](#)
- [Verifying CCC and TCC Graceful Restart on page 292](#)

Graceful Restart Operational Mode Commands

To verify proper operation of graceful restart, use the following commands:

- **show bgp neighbor** (for BGP graceful restart)
- **show log** (for IS-IS and OSPF/OSPFv3 graceful restart)
- **show (ospf | ospfv3) overview** (for OSPF/OSPFv3 graceful restart)
- **show rsvp neighbor detail** (for RSVP graceful restart—helper router)
- **show rsvp version** (for RSVP graceful restart—restarting router)
- **show ldp session detail** (for LDP graceful restart)
- **show connections** (for CCC and TCC graceful restart)

- **show route instance detail** (for Layer 3 VPN graceful restart and for any protocols using graceful restart in a routing instance)
- **show route protocol l2vpn** (for Layer 2 VPN graceful restart)

For more information about these commands and a description of their output fields, see the [CLI Explorer](#).

Verifying BGP Graceful Restart

To view graceful restart information for BGP sessions, use the **show bgp neighbor** command:

```
user@PE1> show bgp neighbor 192.0.2.10

Peer: 192.0.2.10+179 AS 64496 Local: 192.0.2.5+1106 AS 64496
  Type: Internal   State: Established   Flags: <>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Export: [ static ]
  Options:<Preference LocalAddress HoldTime GracefulRestart Damping PeerAS Refresh>

  Local Address: 192.0.2.5 Holdtime: 90 Preference: 170
  IPsec SA Name: hope
  Number of flaps: 0
  Peer ID: 192.0.2.10   Local ID: 192.0.2.5   Active Holdtime: 90
  Keepalive Interval: 30
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast
  NLRI for this session: inet-unicast
  Peer supports Refresh capability (2)
  Restart time configured on the peer: 180
  Stale routes from peer are kept for: 180
  Restart time requested by this peer: 300
  NLRI that peer supports restart for: inet-unicast
  NLRI that peer saved forwarding for: inet-unicast
  NLRI that restart is negotiated for: inet-unicast
  NLRI of received end-of-rib markers: inet-unicast
  NLRI of all end-of-rib markers sent: inet-unicast
  Table inet.0 Bit: 10000
  RIB State: restart is complete
  Send state: in sync
  Active prefixes: 0
  Received prefixes: 0
  Suppressed due to damping: 0
  Last traffic (seconds): Received 19   Sent 19   Checked 19
  Input messages: Total 2   Updates 1   Refreshes 0   Octets 42
  Output messages: Total 3   Updates 0   Refreshes 0   Octets 116
  Output Queue[0]: 0
```

Verifying IS-IS and OSPF Graceful Restart

To view graceful restart information for IS-IS and OSPF, configure traceoptions (see [“Tracking Graceful Restart Events” on page 277](#)).

Here is the output of a traceoptions log from an OSPF restarting router:

```
Oct  8 05:20:12 Restart mode - sending grace lsas
Oct  8 05:20:12 Restart mode - estimated restart duration timer triggered
Oct  8 05:20:13 Restart mode - Sending more grace lsas
```

Here is the output of a traceoptions log from an OSPF helper router:

```
Oct  8 05:20:14 Helper mode for neighbor 192.0.2.5
Oct  8 05:20:14 Received multiple grace lsa from 192.0.2.5
```

Verifying CCC and TCC Graceful Restart

To view graceful restart information for CCC and TCC connections, use the **show connections** command. The following example assumes four remote interface CCC connections between CE1 and CE2:

```
user@PE1> show connections
```

```
CCC and TCC connections [Link Monitoring On]
Legend for status (St)           Legend for connection types
UN -- uninitialized              if-sw: interface switching
NP -- not present               rmt-if: remote interface switching
WE -- wrong encapsulation       lsp-sw: LSP switching
DS -- disabled
Dn -- down
-> -- only outbound conn is up
<- -- only inbound conn is up
Up -- operational
RmtDn -- remote CCC down
Restart -- restarting

Legend for circuit types
intf -- interface
tlsp -- transmit LSP
rlsp -- receive LSP

CCC Graceful restart : Restarting
```

Connection/Circuit	Type	St	Time last up	# Up trans
CE1-CE2-0	rmt-if	Restart	-----	0
fe-1/1/0.0	intf	Up		
PE1-PE2-0	tlsp	Up		
PE2-PE1-0	rlsp	Up		
CE1-CE2-1	rmt-if	Restart	-----	0
fe-1/1/0.1	intf	Up		
PE1-PE2-1	tlsp	Up		
PE2-PE1-1	rlsp	Up		
CE1-CE2-2	rmt-if	Restart	-----	0
fe-1/1/0.2	intf	Up		
PE1-PE2-2	tlsp	Up		
PE2-PE1-2	rlsp	Up		
CE1-CE2-3	rmt-if	Restart	-----	0
fe-1/1/0.3	intf	Up		
PE1-PE2-3	tlsp	Up		
PE2-PE1-3	rlsp	Up		

- Related Documentation**
- [Graceful Restart Concepts on page 233](#)
 - [Configuring Graceful Restart for QFabric Systems on page 282](#)

PART 9

Power Management Overview

- [Understanding Power Management on page 295](#)
- [Redundant Power System Overview on page 305](#)

CHAPTER 23

Understanding Power Management

- [Understanding Power Management on EX Series Switches on page 295](#)
- [Configuring the Power Priority of Line Cards \(CLI Procedure\) on page 302](#)
- [Configuring Power Supply Redundancy \(CLI Procedure\) on page 303](#)

Understanding Power Management on EX Series Switches

The power management feature for Juniper Networks Ethernet Switches helps ensure that normal operation of the system is not disrupted because of insufficient power to the switch. For example:

- Power management ensures that operating line cards continue to receive power if a user installs a new line card in an operating switch when power is insufficient for both the new and existing line cards.
- Power management reserves a certain amount of power to power supply redundancy, so that if a power supply fails, the switch can continue to operate normally. If power management must use some of this reserved power to provide power to switch components, it raises an alarm to indicate that power supply redundancy no longer exists and that normal operations might be disrupted if a power supply fails.
- If power supply failure requires power management to power down some components, it does so gracefully by powering down line cards and PoE ports in the order specified by the user.

Power management manages power to switch components by employing a power budget policy. In its power budget policy, power management:

- Budgets power for each installed switch component that requires power. With the exception of PoE power for line cards that support PoE, the amount that power management budgets for each component is the maximum power that component might consume under worst case operating conditions. For example, for the fan tray, power management budgets the amount of power required to run the fans at their maximum speed setting, even if the current fan speed is much lower.
- Reserves a set amount of power for power supply redundancy. In its default configuration, power management manages the switch for $N+1$ power redundancy, which ensures uninterrupted system operation if one power supply fails. For example, if a switch has four online 3000 W power supplies, power management reserves 3000

W in its power budget policy for redundancy. It allocates the remaining 9000 W to normal operating power.

- Specifies the rules under which components receive power. These rules are designed to ensure the least disruption to switch operation under conditions of insufficient power. For example, power management provides power to core system components, such as the Routing Engines, before it provides power to line cards.

You can configure certain aspects of power management's budget policy, specifically:

- The power priority of individual line cards. By assigning different power priorities to the line cards, you can determine which line cards are more likely to receive power in the event of insufficient power.
- The power redundancy configuration. The default power redundancy configuration is $N+1$; you can optionally configure $N+N$. For example, if you have deployed two independent AC power feeds to the switch, configure $N+N$ redundancy. When you configure power management for $N+N$ redundancy, it reserves the appropriate amount of power in its power budget and reports insufficient power conditions accordingly.

These configurable items are discussed further in:

- [Power Priority of Line Cards on page 296](#)
- [Power Supply Redundancy on page 299](#)

Power Priority of Line Cards

The power priority of line cards determines:

- The order in which line cards are allocated power
- The order in which line cards that support PoE are allocated power for PoE
- How power is reallocated in cases of changes in power availability or demand in an operating switch



NOTE: On EX6200 switches, the four 10-Gigabit Ethernet SFP+ uplink ports on a Switch Fabric and Routing Engine (SRE) module are treated like a line card in the power budget.

This section covers:

- [How a Line Card's Power Priority Is Determined on page 297](#)
- [Line Card Priority and Line Card Power on page 297](#)
- [Line Card Priority and PoE Power on page 297](#)
- [Line Card Priority and Changes in the Power Budget on page 298](#)

How a Line Card's Power Priority Is Determined

Using the CLI, you can assign an explicit power priority to a line-card slot. If more than one slot has the same assigned priority, the power priority is determined by slot number, with the lowest-numbered slots receiving power first.

By default, all slots in an EX8200 switch are assigned the lowest priority. Thus if you do not explicitly assign priorities to slots, power priority is determined by slot number, with slot 0 having the highest priority.

In an EX6200 switch, all slots are assigned the lowest priority, except for the slots containing an SRE module. Slots containing an SRE module are automatically assigned the highest priority. This means that the line cards that represent the 10-Gigabit Ethernet SFP+ ports on SRE modules have the highest priority among the line cards.

Line Card Priority and Line Card Power

When an EX6200 or EX8200 switch is powered on, power management allocates power to components according to its power budget policy. After power management has allocated power to the base chassis components, it allocates the remaining available power to the line cards. It powers on the line cards in priority order until all line cards are powered on or the available power (including reserved power, if necessary) is exhausted. Thus if available power is exhausted before all line cards receive power, higher-priority cards are powered on while lower-priority cards remain powered off.

A lower-priority card might receive power while a higher-priority card does not if the remaining available power is sufficient to power on the lower-priority card but not the higher-priority card. For example, if a line card requiring 450 W is in a higher-priority slot than line card requiring 330 W, the line card requiring 330 W receives the power if there is less than 450 W but more than 330 W remaining in the power budget.

Line cards that have been administratively taken offline are not allocated power.



NOTE: Because power management does not allocate power to a line card that has been administratively taken offline, a line card that has been taken offline in an EX6200 or EX8200 switch is not automatically brought online when you commit a configuration. You must explicitly use the `request chassis fpc slot slot-number online` command to bring a line card online that was taken offline previously. This behavior differs from other platforms running Juniper Networks Junos operating system (Junos OS), which automatically bring an offline FPC online when you commit a configuration.

If power management cannot power on a line card because of insufficient power, it raises a major (red) alarm.

Line Card Priority and PoE Power

After all line cards have been powered on, power management allocates any remaining available power, including reserved power, to the PoE power budgets of line cards that have PoE ports. Power management allocates PoE power to line cards in the order of

power priority. If enough power is available, a line card receives its full PoE power budget before power management allocates PoE power to the next highest-priority line card. If not enough power is available, a line card receives partial PoE power and lower-priority line cards receive no PoE power.

If power management is unable to allocate enough power to meet the PoE power budget for a line card, it logs a message to the system log.

The default PoE power budget for a line card is the amount of power needed to supply the maximum supported power to all PoE ports. In cases where powered devices do not require the maximum power or in which some PoE ports are not used for powered devices, you can configure a smaller PoE power budget for a line card. By configuring a smaller PoE power budget, you make more power available for the PoE power budgets of lower-priority line cards.

You can also configure the power priority of the PoE ports on a line card. If power management is unable to allocate enough power to a line card to meet its PoE power budget, the line card PoE controller will turn off power to PoE ports in reverse priority order as required to meet the reduced power allocation.

See *Configuring PoE on EX Series Switches (CLI Procedure)* for more information on how to configure the PoE power budget for a line card and how to configure PoE port priorities.

Line Card Priority and Changes in the Power Budget

In an operating switch, power management dynamically reallocates power in response to changes in power availability or demand or changes in line card priority. Power management uses line card priority to determine how to reallocate power in response to the following events:

- A power supply fails, is removed, or is taken offline:
 - If power is insufficient to meet the PoE power allocations of all PoE line cards, power management deallocates PoE power from the line cards in reverse priority order until power is sufficient to meet the remaining PoE power allocations.
 - If power is insufficient to meet the base (non-PoE) power requirements of all the line cards, all PoE power is deallocated. If, after the deallocation of PoE power, power is still not sufficient, power management turns off line cards in reverse priority order until power is sufficient for the remaining line cards.
- A new line card is inserted or a line card is brought online:
 - If the line card supports PoE and there is insufficient power to meet its PoE power budget, PoE power is reallocated from lower-priority line cards. If not enough PoE power can be reallocated from lower-priority line cards, the new line card receives a partial PoE power allocation.
 - If there is insufficient power to power on the new line card, PoE power is removed from PoE line cards in reverse priority order until the new line card can be powered on.
 - If the removal of all PoE power is insufficient to free up enough power to power on the line card, the line card remains powered off and the PoE line cards continue to

receive their PoE power allocations. To minimize disruption on an operating switch, lower-priority line cards are not turned off to provide power to the new line card. However, if you restart the switch, power management reruns the current power budget policy and powers line cards on or off based on their priority. As a result, line cards receive power strictly by priority order and previously operating line cards might no longer receive power.

- A new power supply is brought online:
 - Any line cards that were powered off because of insufficient power are powered on in priority order.
 - After all line cards are powered on, remaining power is allocated to the PoE power budgets of line cards in priority order.
- A line card is removed or taken offline, freeing up power:
 - Any line cards that were powered down because of insufficient power are powered on in priority order.
 - After all line cards are powered on, any remaining power is allocated to the PoE power budgets of line cards in priority order.
- A user changes the assigned power priority of one or more line cards when power is insufficient to meet the power budget:
 - PoE power to the line cards is reallocated based on the new power priorities.
 - Base power allocation to the line cards is not changed—in other words, power management does not power down line cards that had been receiving power because they are now a lower priority. However, if you restart the switch, power management reruns the current power budget policy and powers line cards on or off based on their priority. As a result, line cards receive power strictly by priority order and previously operating line cards might no longer receive power.

If, because of insufficient power, power management reduces or eliminates the PoE power budget for a line card, it logs a message to the system log. If power management must power down a line card because of insufficient power, it raises a major (red) alarm.

Power Supply Redundancy

By default, power management in EX Series switches is configured to manage the power supplies for $N+1$ redundancy, in which one power supply is held in reserve for backup if one of the other power supplies is removed or fails.

You can configure power management to manage the power supplies for $N+N$ redundancy. In $N+N$ redundancy, power management holds N power supplies in reserve for backup. For example, if your switch has six power supplies and you configure $N+N$ redundancy, power management makes three power supplies available for normal operating power and reserves three power supplies for redundancy (3+3). If you have an odd number of power supplies, power management allocates one more power supply to normal operating power than to redundant power. For example, if you have five power supplies, the $N+N$ configuration is 3+2.

Given the same number of power supplies, an $N+N$ configuration usually provides less normal operating power than an $N+1$ configuration because the $N+N$ configuration holds more power in reserve for backup. [Table 10 on page 300](#) shows the effect on normal operating power in $N+1$ and $N+N$ configurations.

Table 10: Available Operating Power in $N+1$ and $N+N$ Redundancy Configurations

Number of Power Supplies at n W Each	Normal Operating Power in $N+1$ Configuration	Normal Operating Power in $N+N$ Configuration
2	1 x (n W)	1 x (n W)
3	2 x (n W)	2 x (n W)
4	3 x (n W)	2 x (n W)
5 (EX8200 switches only)	4 x (n W)	3 x (n W)
6 (EX8200 switches only)	5 x (n W)	3 x (n W)

To compensate for the reduced normal operating power, power management on EX8200 switches allocates less power to the chassis in an $N+N$ configuration than in an $N+1$ configuration. This reduction in allocated chassis power allows a switch in an $N+N$ configuration to power more line cards than it could without the reduction. For the EX8208 switch, the power allocated for the chassis is reduced to 1200 W from 1600 W; for the EX8216 switch, it is reduced to 1800 W from 2400 W.



NOTE: To achieve the reduction in allocated chassis power in an EX8200 switch, power management reduces the maximum fan speed to 60 percent in an $N+N$ configuration from 80 percent in an $N+1$ configuration. Because the maximum fan speed is reduced, it is possible that a line card that overheats would be shut down sooner in an $N+N$ configuration than in an $N+1$ configuration.

On EX6200 switches, the same amount of power is allocated for the chassis in $N+N$ configurations as in $N+1$ configurations.

Power management automatically recalculates the reserved power and normal operating power as power supplies go online or offline. For example, if you have an $N+N$ configuration with three online 2000 W power supplies, power management allocates 2000 W to reserved power. If you bring a fourth 2000 W power supply online, power management then allocates 4000 W to reserved power. If a power supply goes offline again, power management once again allocates 2000 W to reserved power.

When power is insufficient to meet the budgeted power requirements, power management raises alarms as follows:

- A minor (yellow) alarm is raised when insufficient power exists to maintain the configured $N+1$ or $N+N$ power reserves, but all line cards are still receiving their base

and PoE power allocations. If this condition persists for 5 minutes, the alarm becomes a major (red) alarm. Even though operation of the switch is unaffected in this condition, you should remedy it as quickly as possible because a power supply failure might cause a disruption in switch operation.

- A major (red) alarm is raised when insufficient power exists to provide all the line cards with their base and PoE power allocations. One or more PoE ports might be down or one or more line cards might be down.

Power management clears all alarms when sufficient power is available to meet normal operating and reserved power requirements.

**Related
Documentation**

- [*Understanding Alarm Types and Severity Levels on EX Series Switches*](#)
- [Configuring the Power Priority of Line Cards \(CLI Procedure\) on page 302](#)
- [Configuring Power Supply Redundancy \(CLI Procedure\) on page 303](#)
- [Verifying Power Configuration and Use on page 658](#)

Configuring the Power Priority of Line Cards (CLI Procedure)

The power management facility on EX6200 and EX8200 switches allows you to assign power priorities to the slots occupied by line cards. Power management provides power to the slots in priority order, which means that line cards in higher priority slots are more likely to receive power than line cards in lower priority slots if power to the switch is insufficient to power all the line cards.

The power priority you assign to a PoE line card affects both the order in which it receives base power and the order in which it receives PoE power. Base power is allocated first to all line cards in priority order. PoE power is then allocated to the PoE line cards in priority order.

When assigning power priority to slots, keep these points in mind:

- 0 is the highest priority. The number of priority levels depends on the number of slots in a switch—for example, for an EX8208 switch, which has eight slots, you can assign a priority of 0 through 7 to a slot.
- All slots are assigned the lowest priority by default.
- If a group of slots shares the same assigned priority, each slot's power priority within the group is based on its slot number, with the lowest-numbered slots receiving power first. For example, if slot 3 and slot 7 each have an assigned power priority of 2, slot 3 has the higher power priority.
- On EX6200 switches, slots containing a Switch Fabric and Routing Engine (SRE) module are automatically assigned the highest priority. If you assign a priority of 0 to a slot that has a lower number than a slot an SRE module is in, the slot with an SRE module still receives power first. You cannot change the power priority of slot containing an SRE module.

To assign or change the power priority for a slot:

```
[edit chassis]
user@switch# set fpc slot power-budget-priority priority
```

For example, to set slot 6 to priority 0, enter:

```
[edit chassis]
user@switch# set fpc 6 power-budget-priority 0
```

Related Documentation

- [Configuring Power Supply Redundancy \(CLI Procedure\) on page 303](#)
- [Verifying Power Configuration and Use on page 658](#)
- [Understanding Power Management on EX Series Switches on page 295](#)

Configuring Power Supply Redundancy (CLI Procedure)

By default, the power management feature in EX Series switches is configured to manage the power supplies for $N+1$ redundancy, in which one power supply is held in reserve for backup if any one of the other power supplies is removed or fails.

You can configure power management to manage the power supplies for $N+N$ redundancy. For example, to set up your AC power supplies for dual power feed, $N+N$ redundancy is required. In $N+N$ redundancy, power management allocates half of the online power supplies to normal operating power and half to redundant power. If you have an odd number of online power supplies, power management allocates one more power supply to normal operating power than to redundant power.

This topic describes how to configure power management for $N+N$ redundancy and how to revert back to $N+1$ redundancy if your deployment needs change.

Before you configure power management for $N+N$ redundancy, ensure that you have sufficient power supplies to meet the power requirements of an $N+N$ configuration. Use the `show chassis power-budget-statistics` command to display your current power budget.



NOTE: To allow more power to be available to line cards in an EX8200 switch, power management compensates for the reduced normal operating power in an $N+N$ configuration by allocating less power to the chassis than it does in an $N+1$ configuration. For the EX8208 switch, the power allocated to the chassis is reduced to 1200 W from 1600 W. For the EX8216 switch, it is reduced to 1800 W from 2400 W. In determining whether you have enough power for an $N+N$ configuration, take this reduction of allocated chassis power into account.

The reduction in allocated chassis power is achieved by reducing the maximum fan speed to 60 percent in an $N+N$ configuration from 80 percent in an $N+1$ configuration. Because the maximum fan speed is reduced, it is possible that a line card that overheats would be shut down sooner in an $N+N$ configuration than in an $N+1$ configuration.

On EX6200 switches, the same amount of power is allocated for the chassis in $N+N$ configurations as in $N+1$ configurations.

To configure *N+N* redundancy:

```
[edit chassis]
user@switch# set psu redundancy n-plus-n
```

To revert back to *N+1* redundancy:

```
[edit chassis]
user@switch# delete chassis psu redundancy n-plus-n
```

**Related
Documentation**

- [Verifying Power Configuration and Use on page 658](#)
- [Understanding Power Management on EX Series Switches on page 295](#)

Redundant Power System Overview

- [EX Series Redundant Power System Hardware Overview on page 305](#)
- [Understanding How Power Priority Is Determined and Set for Switches Connected to the EX Series Redundant Power System on page 308](#)
- [Determining and Setting Priority for Switches Connected to an EX Series RPS on page 310](#)

EX Series Redundant Power System Hardware Overview

You can use the EX Series Redundant Power System (RPS) to provide backup power for Juniper Networks EX2200 Ethernet Switches, (except Juniper Networks EX2200-C Ethernet Switches) and Juniper Networks EX3300 Ethernet Switches that are standalone switches or are members of a Virtual Chassis.

Most EX Series switches have a built-in capability for redundant power supplies—therefore, if one power supply fails on those switches, the other power supply takes over. However, EX2200 switches and EX3300 switches have only one internal fixed power supply. If an EX2200 switch or EX3300 switch is deployed in a critical situation, we recommend that you connect a an RPS to that switch to supply backup power during a loss of power.

RPS is not a primary power supply—it only provides backup power to switches when the single dedicated power supply fails. An RPS operates in parallel with the single dedicated power supplies of the switches connected to it and provides all connected switches enough power to support either Power over Ethernet (PoE) or non-PoE devices when the power supplies on the switches fail.

An RPS can hold up to three power supplies connected to as many as six switches—how that power is allocated is up to you. You determine whether or not to connect switches that provide PoE and you determine which switches have priority. Priority becomes an issue when you connect more than three switches that provide PoE to a fully loaded RPS because a switch providing PoE requires more power than a switch that does not provide PoE. Because a power supply can support only one switch providing PoE, the RPS can become oversubscribed when too many switches that must have enough power for PoE have a power failure.

- [Benefits of the EX Series Redundant Power System on page 306](#)
- [Switch Models and Configurations Supported by the RPS on page 306](#)

- [When a Switch's Power Supply Fails on page 307](#)
- [Components of the RPS on page 307](#)

Benefits of the EX Series Redundant Power System

Provides power backup—You connect up to six EX2200, EX3300, or a combination of these switches and supply power to any three of them.

Protection from high-voltage input and short circuits—RPS provides protection from high-voltage input and short circuits.

Switch Models and Configurations Supported by the RPS

The RPS supports all EX3300 switches and EX2200 switches except EX2200-C switches. You can simultaneously connect any supported switches to the same RPS, whether the switches are standalone switches or are configured in a Virtual Chassis.

All power provided by RPS is either PoE or non-PoE. By default, RPS supports switches that provide PoE. If even one switch provides PoE, then the RPS must be configured to provide enough power for PoE. When enough power for PoE is supplied, one switch can be powered by each power supply. If the switches are not providing PoE power, two switches can be powered by one RPS power supply—you can reconfigure an RPS to provide non-PoE power using a feature called multi-backup.

[Table 11 on page 306](#) lists some possible scenarios and RPS solutions. These examples assume that each RPS is fully loaded with three power supplies.

Table 11: Sample Requirements and RPS Solutions

Switches Requiring Backup	You need this RPS configuration:
Six switches that do not provide PoE to attached devices	One RPS can simultaneously provide power to all six switches if you change the power default to multi-backup—this indicates that no attached switch provides PoE to any devices.
One switch that provides PoE to other devices or two switches that do not provide PoE to any devices	One RPS will always back up all three switches, whether or not they provide PoE to connected devices. Leave the power at the default setting (no multi-backup) and let RPS determine that two switches need only minimum power and one switch provides PoE and therefore needs extra power. RPS automatically supplies the correct level of power.
One EX Series Virtual Chassis member that supplies PoE, one switch that supplies PoE, and one switch that does not supply PoE to any connected devices	One RPS will always back up all three switches. Leave the power default setting (no multi-backup) and let RPS determine that one switch needs only minimum power, one switch needs extra power because it supplies PoE, and the Virtual Chassis member also provides PoE to connected devices.

Table 11: Sample Requirements and RPS Solutions (continued)

Switches Requiring Backup	You need this RPS configuration:
One switch that supplies PoE and five switches that do not supply PoE	<p>You have two options.</p> <p>Option 1—Use one RPS: Up to three switches that do or do not supply PoE can be backed up simultaneously. You can prioritize the six switches to determine which three are most important if all six fail at once. You must leave the power default setting (no multi-backup) because you have one switch that supplies PoE to attached devices and therefore requires more power.</p> <p>Option 2—Use Two RPSs: In this case, you can connect three switches to each RPS and all switches will be backed up if they all fail at once. Alternatively, you can change the power default to multi-backup on one RPS and connect all five switches that do not supply PoE to that RPS, leaving the other RPS to back up the switch that supplies PoE.</p>
EX Series Virtual Chassis	Use as many RPSs as needed to back up all members of the Virtual Chassis.

When a Switch's Power Supply Fails

Because the power supplies for both EX3300 switches and EX2200 switches are internal, if the switch's power supply fails, you must replace the switch. You should remove or replace a switch with a failed power supply as soon as possible.

Do not try to use an RPS as a primary power supply because an RPS cannot boot or reboot a switch. Each switch connected to the RPS must have its own dedicated power supply and must have booted up using the internal power supply.

If a switch is deployed in a large network center where RPS has a separate source of electricity than the switches it supports, the RPS supplies power when only the switch's electricity fails. In this case, you would not have to replace the switch because the power supply is still functional. The switch will resume using its own internal power supply when electricity to the switch is restored.

Components of the RPS

Table 12 on page 307 lists and describes the components of an RPS:

Table 12: Redundant Power System Components

Component	Value
Power supplies that can be installed	Up to three EX-PWR3-930-AC power supplies. One is included and additional power supplies must be ordered separately.
Switch connector ports on RPS	6 (2 per power supply)
Power cords (for connecting power supplies to the AC power source outlet)	Up to three power cords, one per power supply.

Table 12: Redundant Power System Components (continued)

Component	Value
RPS cables (for connecting a switch to a power supply installed in the RPS)	6 (1 for each RPS-to-switch connection). One cable is supplied with the RPS. Additional cables must be ordered separately.


Understanding How Power Priority Is Determined and Set for Switches Connected to the EX Series Redundant Power System

The Redundant Power System (RPS) is designed to provide backup power to switches that lack built-in redundant power supplies. The RPS provides backup power to switches that either supply power over Ethernet (PoE), which require more power, or switches that do not supply PoE, which require less power. A power supply can either power one PoE device or two non-PoE devices. That means if an RPS is fully loaded with three power supplies, supports PoE switches, and more than three PoE switches have a power failure, some switches will not be powered. You can, however, determine which switches will be powered when an RPS is oversubscribed. When too many connected switches fail, the switches are given power based on their priority. Priority is also reconfigured when any power change takes place. For example, if three switches are already being backed up and another switch has a power failure, the RPS detects this, reconfigures the current top priorities, and allots power accordingly.

- [Default RPS Priority on page 308](#)
- [Changing the Priority of Switches on an EX Series RPS on page 309](#)

Default RPS Priority


While six non-PoE switches can all simultaneously be backed up with three power supplies, only three PoE switches can be backed up (because PoE uses more power). This means that an RPS with four or more PoE switches connected will have to select three of them for backup. You can determine priority by the connector positions you use to connect the switches. By default, an RPS assigns priority to switches based on their switch connector port location, with the leftmost port having the lowest priority and the rightmost port having the highest priority. If the PoE switches shown in [Figure 20 on page 309](#) all fail, the manufacturing, support, and finance switches will be backed up because they are connected to the rightmost connectors.

Figure 20: Default PoE Switch Priority Is Determined by Connector Port Location

Changing the Priority of Switches on an EX Series RPS

There is a way to alter the priority of PoE switches on an RPS without disconnecting the cables. You can optionally reconfigure any of the attached switches from their CLIs to establish a switch's RPS priority—this CLI configuration overcomes the priority determined by the switch connector port location. Priority ranges from zero (off) to 1 (lowest) through 6 (highest). By default, all switches are configured to 1, the lowest priority. Let's say that the sales switch is reconfigured from the switch's CLI for priority 5 (second highest).

Now in [Figure 21 on page 309](#), with the sales switch configured for RPS 5 from the CLI, the highest priority changes to sales (because 5 is higher than 1), then manufacturing, and then support.

Figure 21: Switch Priority After CLI Configuration

When assigning power priority to switches by using the CLI on the switch, keep these points in mind:

- By default, all switches are assigned priority 1 (lowest) and derive precedence from the location of their connector port on the RPS, with the rightmost port having highest priority.
- Priority 0 assigned from a switch CLI means that the RPS does not provide any backup power to the switch. Essentially, this turns off RPS support.

- Priority 6 assigned from a switch CLI is the highest priority and priority 1 is the lowest priority.
- The CLI command that assigns priority to EX2200 switches is slightly different from the CLI command that assigns priority to EX3300 switches because EX3300 switches can be configured as a Virtual Chassis.
- If two or more switches are assigned the same priority value from the switches' CLIs, then the power priority for those switches is determined by the RPS switch connector port location, with the ports to the right receiving priority.
- If a single power supply is installed, the RPS can provide backup power to one switch out of all the switches connected to the RPS. If you do not need any PoE power backup on any switch, you can increase the number of supported switches to two per power supply. Switches connected to an RPS must be either all PoE or all non-PoE.
- The RPS discontinues supplying backup power to a lower-priority switch if it detects a backup power need for a higher-priority switch at the same time.

**Related
Documentation**

- [Determining and Setting Priority for Switches Connected to an EX Series RPS on page 310](#)

Determining and Setting Priority for Switches Connected to an EX Series RPS

A Redundant Power System (RPS) provides backup power according to the RPS priority configured on the standalone EX Series switches or Virtual Chassis member switches connected to it. If all switches connected to the RPS are set to the default priority of 1, the priority is determined on the basis of the RPS port to which they are connected, with higher port numbers having the higher priorities.

The number of switches for which an RPS can provide backup power depends on whether the switches provide power over Ethernet (PoE).

- PoE: A fully loaded RPS provides backup power to a maximum of three switches that are enabled for PoE—the result in this case is one switch powered per power supply. If more than three PoE-enabled switches are connected to the RPS and the RPS is already providing backup power to three switches when another switch's power supply fails, the RPS detects this and re-allots backup power as required. It would then stop providing backup power to a low-priority switch to provide backup power to a higher-priority switch.
- Non-PoE: If you changed the RPS power setting to non-PoE with the command [request redundant-power-system multi-backup](#), your RPS is configured to provide back up power to as many as six non-PoE switches on a fully loaded RPS. Each power supply can support two switches when the switches do not need enough power for PoE.



NOTE: Before an RPS can back up a switch connected to it, the switch's RPS status must be ARMED. There are two ways to determine whether a switch's RPS status is ARMED—either check that the corresponding port LED on the RPS is lit and on steady or issue this command from the switch's CLI: [show chassis redundant-power-system](#).

This topic describes how to determine and set the power priority for a switch connected to an RPS.

- [Using RPS Default Configuration on page 311](#)
- [Setting the EX Series RPS Priority for a Switch \(CLI\) on page 311](#)

Using RPS Default Configuration

No configuration is required on an RPS if you:

- Plan to back up as many as six non-PoE switches
- Back up three PoE switches with three RPS power supplies
- Back up four or more PoE switches with RPS three power supplies and let the RPS port to which the switch is connected determine the priority

By default, an RPS assigns priority to switches on the basis of their switch connector port location, with the with higher port numbers having the higher priorities. By default, all switches are themselves configured with the same RPS priority (priority 1, the lowest), which is why priority is derived from the RPS connector port numbers.

Setting the EX Series RPS Priority for a Switch (CLI)

Each switch connected to RPS has an RPS priority value—that priority value determines which PoE switches receive power first from the RPS. By default, all switches are configured for priority 1 so priority is then determined by switch connector port location, left (lowest) to right (highest).

You can change the priority of a switch to 0 (off), or 1 (lowest) through 6 (highest) from the switch itself—this configuration takes precedence over switch connector port location.

To set or change the priority for a switch that does not support Virtual Chassis:

```
[edit]
user@switch# set redundant-power-system priority
```

To set or change the priority for a switch that supports Virtual Chassis:

```
[edit]
user@switch# set redundant-power-system membervc-member-id priority priority-number
```

Where member is 0 for a switch that has never been configured in a Virtual Chassis.

- Related Documentation**
- [Understanding How Power Priority Is Determined and Set for Switches Connected to the EX Series Redundant Power System on page 308](#)

PART 10

Configuring Virtual Router Redundancy Protocol (VRRP)

- [Understanding How the VRRP Router Failover Mechanism Prevents Network Failures on page 315](#)
- [Configuring VRRP on page 331](#)

CHAPTER 25

Understanding How the VRRP Router Failover Mechanism Prevents Network Failures

- [Understanding VRRP on page 315](#)
- [Understanding VRRP Between QFabric Systems on page 319](#)
- [Junos OS Support for VRRPv3 on page 322](#)
- [VRRP failover-delay Overview on page 328](#)

Understanding VRRP

For Ethernet, Fast Ethernet, Gigabit Ethernet, 10-Gigabit Ethernet, and logical interfaces, you can configure the Virtual Router Redundancy Protocol (VRRP) or VRRP for IPv6. VRRP enables hosts on a LAN to make use of redundant routing platforms on that LAN without requiring more than the static configuration of a single default route on the hosts. The VRRP routing platforms share the IP address corresponding to the default route configured on the hosts. At any time, one of the VRRP routing platforms is the master (active) and the others are backups. If the master routing platform fails, one of the backup routing platforms becomes the new master routing platform, providing a virtual default routing platform and enabling traffic on the LAN to be routed without relying on a single routing platform. Using VRRP, a backup device can take over a failed default device within a few seconds. This is done with minimum VRRP traffic and without any interaction with the hosts. Virtual Router Redundancy Protocol is not supported on management interfaces.

Devices running VRRP dynamically elect master and backup devices. You can also force assignment of master and backup devices using priorities from 1 through 255, with 255 being the highest priority. In VRRP operation, the default master device sends advertisements to backup devices at regular intervals. The default interval is 1 second. If a backup device does not receive an advertisement for a set period, the backup device with the next highest priority takes over as master and begins forwarding packets.



NOTE: Priority 255 cannot be set for routed VLAN interfaces (RVIs).



NOTE: To minimize network traffic, VRRP is designed in such a way that only the device that is acting as the master sends out VRRP advertisements at any given point in time. The backup devices do not send any advertisement until and unless they take over mastership.

VRRP for IPv6 provides a much faster switchover to an alternate default router than IPv6 neighbor discovery procedures. Typical deployments use only one backup router.



NOTE: Do not confuse the VRRP master and backup routing platforms with the master and backup member switches of a Virtual Chassis configuration. The master and backup members of a Virtual Chassis configuration compose a single host. In a VRRP topology, one host operates as the master routing platform and another operates as the backup routing platform, as shown in [Figure 24 on page 318](#).

VRRP is defined in RFC 3768, *Virtual Router Redundancy Protocol*. VRRP for IPv6 is defined in draft-ietf-vrrp-ipv6-spec-08.txt, *Virtual Router Redundancy Protocol for IPv6*. See also draft-ietf-vrrp-unified-mib-06.txt, *Definitions of Managed Objects for the VRRP over IPv4 and IPv6*.




NOTE: Even though VRRP, as defined in RFC 3768, does not support authentication, the Junos OS implementation of VRRP supports authentication as defined in RFC 2338. This support is achieved through the backward compatibility options in RFC 3768.



NOTE: On EX2300 and EX3400 switches, the VRRP protocol must be configured with a Hello interval of 2 seconds or more with dead interval not less than 6 seconds to prevent flaps during CPU intensive operations events such as routing engine switchover, interface flaps, and exhaustive data collection from the packet forwarding engine.

[Figure 22 on page 317](#) illustrates a basic VRRP topology. In this example, Routers A, B, and C are running VRRP and together make up a virtual router. The IP address of this virtual router is 10.10.0.1 (the same address as the physical interface of Router A).

Figure 22: Basic VRRP



Because the virtual router uses the IP address of the physical interface of Router A, Router A is the master VRRP router, while routers B and C function as backup VRRP routers. Clients 1 through 3 are configured with the default gateway IP address of 10.10.0.1. As the master router, Router A forwards packets sent to its IP address. If the master virtual router fails, the router configured with the higher priority becomes the master virtual router and provides uninterrupted service for the LAN hosts. When Router A recovers, it becomes the master virtual router again.



NOTE: In some cases, during an inherit session, there is a small time frame during which two routers are in Master-Master state. In such cases, the VRRP groups that inherit the state do send out VRRP advertisements every 120 seconds. So, it takes the routers up to 120 seconds to recover after moving to Master-Backup state from Master-Master state.

ACX series routers can support up to 64 VRRP group entries. These can be a combination of IPv4 or IPv6 families. If either of the family (IPv4 or IPv6) is solely configured for VRRP, then 64 unique VRRP group identifiers are supported. If both IPv4 and IPv6 families share the same VRRP group, then only 32 unique VRRP identifiers are supported.



NOTE: ACX Series routers support VRRP version 3 for IPv6 addresses.

ACX5448 router supports RFC 3798 VRRP version 2 and RFC 5798 VRRP version 3. ACX5448 router also supports configuring VRRP over aggregated Ethernet and integrated routing and bridging (IRB) interfaces.

The following limitations apply while configuring VRRP on ACX5448 router:

- Configure a maximum of 16 VRRP groups.
- Interworking of VRRP version 2 and VRRP version 3 is not supported.
- VRRP delegate processing is not supported.
- VRRP version 2 authentication is not supported.

Figure 22 on page 317 illustrates a basic VRRP topology with EX Series switches. In this example, Switches A, B, and C are running VRRP and together they make up a virtual routing platform. The IP address of this virtual routing platform is 10.10.0.1 (the same address as the physical interface of Switch A).

Figure 23: Basic VRRP on EX Series Switches


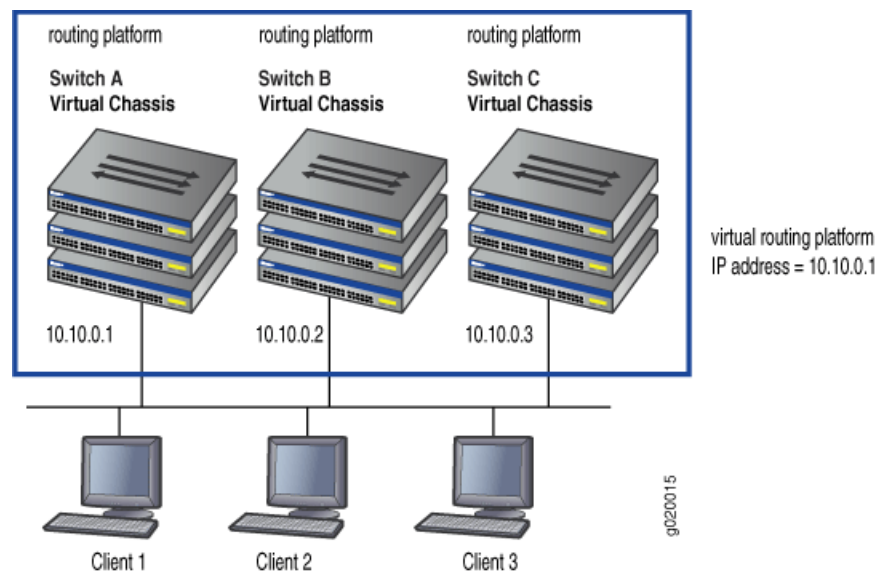


Figure 24 on page 318 illustrates a basic VRRP topology using Virtual Chassis configurations. Switch A, Switch B, and Switch C are each composed of multiple interconnected Juniper Networks EX4200 Ethernet Switches. Each Virtual Chassis configuration operates as a single switch, which is running VRRP, and together they make up a virtual routing platform. The IP address of this virtual routing platform is 10.10.0.1 (the same address as the physical interface of Switch A).

Figure 24: VRRP on Virtual Chassis Switches



Because the virtual routing platform uses the IP address of the physical interface of Switch A, Switch A is the master VRRP routing platform, while Switch B and Switch C

function as backup VRRP routing platforms. Clients 1 through 3 are configured with the default gateway IP address of 10.10.0.1 as the master router, Switch A, forwards packets sent to its IP address. If the master routing platform fails, the switch configured with the higher priority becomes the master virtual routing platform and provides uninterrupted service for the LAN hosts. When Switch A recovers, it becomes the master virtual routing platform again.

**Related
Documentation**

- [Understanding High Availability Features on Juniper Networks Routers on page 3](#)
- [High Availability Features for EX Series Switches Overview on page 9](#)
- [Junos OS Support for VRRPv3 on page 322](#)
- [Configuring Basic VRRP Support on page 332](#)
- [Configuring VRRP on page 337](#)
- [Configuring VRRP for IPv6 \(CLI Procedure\) on page 341](#)

Understanding VRRP Between QFabric Systems

Juniper Networks QFabric systems support the Virtual Router Redundancy Protocol (VRRP). This topic covers:

- [VRRP Differences on QFabric Systems on page 319](#)
- [Configuration Details on page 320](#)

VRRP Differences on QFabric Systems

Configuring servers on your network with static routes to a default gateway minimizes configuration effort and complexity and reduces processing overhead. However, a failure of the default gateway normally results in a catastrophic event, isolating the servers. Using Virtual Router Redundancy Protocol (VRRP) enables you to dynamically provide alternative gateways for servers if the primary gateway fails.

Switches configured with VRRP share a virtual IP (VIP) address, which is the address you configure as the default route on the servers. In normal VRRP operation, one of the switches is the VRRP master, meaning that it owns the VIP and is the active default gateway. The other devices are backups. The switches dynamically assign master and backup roles based on priorities that you configure. If the master fails, the backup switch with the highest priority becomes the master and takes ownership of the VIP within a few seconds. This is done without any interaction with the servers.

You can configure two QFabric systems to participate in a VRRP configuration as if they were two standalone switches. However, in normal VRRP operation, only one system can be the master for a given VRRP group at any one time, which means that only one system can act as a default gateway using the VIP configured for the group. When running VRRP over two QFabric systems, you might want both systems to simultaneously use the VIP to act as a gateway and forward traffic. To achieve this, you can configure a firewall filter to block the VRRP advertisement packets between the QFabric systems on the link between the two network Node groups. When you do this, both QFabric systems act as master and forward traffic received by the VIP (which is the default

gateway address that you configure on servers connected to both QFabric systems). If you use VMware's vMotion, this configuration allows virtual machines to transition between servers connected to the QFabric systems without updating their default gateway information. For example, a virtual machine running on a server connected to a QFabric system in data center A can transition to a server connected to a QFabric system in data center B without needing to resolve a new gateway IP address and MAC address because both QFabric systems use the same VIP.



NOTE: To use a firewall filter to block VRRP traffic, create a firewall term that matches traffic for protocol vrrp and discards that traffic.

Configuration Details

Configuring a VRRP group across two QFabric systems is similar to configuring VRRP on two switches. The main differences are listed here:

- All the interfaces in both QFabric systems that participate in VRRP must be members of the same VLAN.
- You must create routed VLAN interfaces (RVIs) in that VLAN on both QFabric systems.
- The IP addresses that you assign to both RVIs must be in the same subnet.
- You must configure VRRP on the RVIs.
- Both RVIs must be members of the same VRRP group. This is what allows the two QFabric systems to share a virtual IP address.

The following tables list the elements of an example VRRP configuration running on two QFabric systems—QFabric system A and QFabric system B. This example is configured so that both QFabric systems act as the VRRP master for VIP 10.1.1.50/24 and assumes that a firewall filter blocks the VRRP advertisements between the systems.

[Table 13 on page 320](#) lists the required characteristics of the RVIs in the example configuration.



NOTE: Most of the configuration settings in the following tables would also apply in a traditional VRRP configuration. However, the advertisement interval and priority settings would need to be different (as noted).

Table 13: RVIs on QFabric systems in example VRRP configuration

RVI on QFabric System A	RVI on QFabric System B
vlan.100	vlan.200
Member of VLAN 100. (Note that the VLAN is the same on both QFabric systems.)	Member of VLAN 100
IP address 10.1.1.100/24	IP address 10.1.1.200/24

Table 13: RVIs on QFabric systems in example VRRP configuration (continued)

Member of VRRP group 500	Member of VRRP group 500
Virtual IP address 10.1.1.50/24	Virtual IP address 10.1.1.50/24

You must configure VRRP on the RVIs on both QFabric systems. [Table 14 on page 321](#) lists the elements of a sample VRRP configuration on each RVI. Note that with the exception of the priority, the parameters *must* be the same on both systems.

Table 14: Sample VRRP configuration each RVI

VRRP on RVI on QFabric System A	VRRP on RVI on QFabric System B
VRRP group 500	VRRP group 500
Virtual IP address 10.1.1.50/24	Virtual IP address 10.1.1.50/24
Advertisement interval 60 seconds. (In a normal VRRP configuration, you would set this interval to be much smaller, such as 1 second. However, in this configuration these packets are blocked by the firewall filter on the interface that connects to QFabric system B, so there is no need to send them frequently.)	Advertisement interval 60 seconds
Authentication type md5	Authentication type md5
Authentication key \$9\$1.4ElMVb2aGi4aZjkqzFRhSeWx7-wY2aM8	Authentication key \$9\$1.4ElMVb2aGi4aZjkqzFRhSeWx7-wY2aM8
Priority 254. (In a normal VRRP configuration, this value would be different on the two systems and the system with the higher value would be the master. However, in this configuration both systems are acting as master, so you do not have to configure different values.)	Priority 254



NOTE: Priority 255 is not supported for RVIs.

[Table 15 on page 321](#) lists all the interfaces on QFabric system A in the example configuration and identifies what they connect to.

Table 15: Interfaces on QFabric system A. All interfaces are members of VLAN 100.

VLAN 100 Interfaces on QFabric System A	Connects To
vlan.100	vlan.200
Network Node group interface QFA-NNG:xe-0/0/0	QFB-NNG:xe-0/0/0 on QFabric system B
Network Node group interface QFA-NNG:xe-0/0/1	Redundant server Node group interface QFA-RSNG:xe-0/0/0
Redundant server Node group interface QFA-RSNG:xe-0/0/0	Connects to a network Node group interface QFA-NNG:xe-0/0/1

Table 15: Interfaces on QFabric system A. All interfaces are members of VLAN 100. (continued)

Redundant server Node group interface QFA-RSNG:xe-0/0/1	LAN with servers running virtual machines
---	---

Table 16 on page 322 lists all the interfaces on QFabric system B in the example configuration and identifies what they connect to.

Table 16: Interfaces on QFabric system B. All interfaces are members of VLAN 100 (same as on QFabric system A).

VLAN 100 Interfaces on QFabric System B	Connects To
vlan.200	vlan.100
Network Node group interface QFB-NNG:xe-0/0/0	QFA-NNG:xe-0/0/0 on QFabric system A
Network Node group interface QFB-NNG:xe-0/0/1	Redundant server Node group interface QFB-RSNG:xe-0/0/0
Redundant server Node group interface QFB-RSNG:xe-0/0/0	Connects to a network Node group interface QFB-NNG:xe-0/0/1
Redundant server Node group interface QFB-RSNG:xe-0/0/1	LAN with servers running virtual machines

- Related Documentation**
- [Understanding VRRP on page 315](#)
 - [Configuring Basic VRRP Support for QFX](#)
 - [Example: Configuring VRRP for Load Sharing on page 378](#)

Junos OS Support for VRRPv3

The advantage of using VRRPv3 is that VRRPv3 supports both IPv4 and IPv6 address families, whereas VRRPv2 supports only IPv4 addresses.

The following topics describe the Junos OS support for and interoperability of VRRPv3, as well as some differences between VRRPv3 and its precursors:

- [Junos OS VRRP Support on page 322](#)
- [IPv6 VRRP Checksum Behavioral Differences on page 323](#)
- [VRRP Interoperability on page 324](#)
- [Upgrading from VRRPv2 to VRRPv3 on page 324](#)
- [Functionality of VRRPv3 Features on page 326](#)

Junos OS VRRP Support

In releases earlier than Release 12.2, Junos OS supported RFC 3768, *Virtual Router Redundancy Protocol (VRRP)* (for IPv4) and Internet draft draft-ietf-vrrp-ipv6-spec-08, *Virtual Router Redundancy Protocol for IPv6*.

VRRPv3 is not supported on routers that use releases earlier than Junos OS Release 12.2 and is also not supported for IPv6 on QFX10000 switches.

Starting with Release 12.2, Junos OS supports:

- RFC 3768, *Virtual Router Redundancy Protocol (VRRP)*
- RFC 5798, *Virtual Router Redundancy Protocol (VRRP) Version 3 for IPv4 and IPv6*
- RFC 6527, *Definitions of Managed Objects for Virtual Router Redundancy Protocol Version 3 (VRRPv3)*



NOTE: VRRP (for IPv6) on routers that use Junos OS Release 12.2 and later releases does not interoperate with VRRP (for IPv6) on routers with earlier Junos OS releases because of the differences in VRRP checksum calculations. See [“IPv6 VRRP Checksum Behavioral Differences” on page 323](#).

IPv6 VRRP Checksum Behavioral Differences

You must consider the following checksum differences when enabling IPv6 VRRP networks:

- In releases earlier than Junos OS Release 12.2, when VRRP for IPv6 is configured, the VRRP checksum is calculated according to section 5.3.8 of RFC 3768, *Virtual Router Redundancy Protocol (VRRP)*.
- Starting with Junos OS Release 12.2, when VRRP for IPv6 is configured, irrespective of VRRPv3 being enabled or not, the VRRP checksum is calculated according to section 5.2.8 of RFC 5798, *Virtual Router Redundancy Protocol (VRRP) Version 3 for IPv4 and IPv6*.

Moreover, the pseudoheader is included only when calculating the IPv6 VRRP checksum. The pseudoheader is not included when calculating the IPv4 VRRP checksum.

To make the router with Junos OS Release 12.2 (or later Junos OS releases) IPv6 VRRP interoperate with the router running a Junos OS release earlier than Release 12.2, include the **checksum-without-pseudoheader** configuration statement at the **[edit protocols vrrp]** hierarchy level in the router running Junos OS Release 12.2 or later.

- The **tcpdump** utility in Junos OS Release 12.2 and later calculates the VRRP checksum according to RFC 5798, *Virtual Router Redundancy Protocol (VRRP) Version 3 for IPv4 and IPv6*. Therefore, when **tcpdump** parses IPv6 VRRP packets that are received from older Junos OS releases (earlier than Junos OS Release 12.2), the **bad vrrp cksum** message is displayed:

```
23:20:32.657328 Out
...
-----original packet-----
00:00:5e:00:02:03 > 33:33:00:00:00:12, ethertype IPv6 (0x86dd), length
94: (class 0xc0, hlim 255, next-header: VRRP (112), length: 40)
fe80::224:dcff:fe47:57f > ff02::12: VRRPv3-advertisement 40: vrid=3 prio=100
intvl=100(centisec) (bad vrrp cksum b4e2!) addrs(2):
```

```

fe80::200:5eff:fe00:3,2001:4818:f000:14::1
      3333 0000 0012 0000 5e00 0203 86dd 6c00
      0000 0028 70ff fe80 0000 0000 0000 0224
      dcff fe47 057f ff02 0000 0000 0000 0000
      0000 0000 0012 3103 6402 0064 b4e2 fe80
      0000 0000 0000 0200 5eff fe00 0003 2001
      4818 f000 0014 0000 0000 0000 0001

```

You can ignore this message because it does not indicate VRRP failure.

VRRP Interoperability

In releases earlier than Junos OS Release 12.2, VRRP (IPv6) followed Internet draft draft-ietf-vrrp-ipv6-spec-08, but checksum was calculated based on RFC 3768 section 5.3.8. Starting with Release 12.2, VRRP (IPv6) follows RFC 5798 and checksum is calculated based on RFC 5798 section 5.2.8. Because of the differences in VRRP checksum calculations, IPv6 VRRP configured on routers that use Junos OS Release 12.2 and later releases does not interoperate with IPv6 VRRP configured in releases before Junos OS Release 12.2.

To make the router with Junos OS Release 12.2 (or later Junos OS releases) IPv6 VRRP interoperate with the router running Junos OS releases earlier than Release 12.2, include the **checksum-without-pseudoheader** configuration statement at the **[edit protocols vrrp]** hierarchy level in the router with Junos OS Release 12.2 or later.

Here are some general points to know about VRRP interoperability:

- If you have configured VRRPv3 (IPv4 or IPv6) on routers that use Junos OS Release 12.2 or later releases, it will not operate with routers that use Junos OS Release 12.1 or earlier releases. This is because only Junos OS Release 12.2 and later releases support VRRPv3.
- VRRP (IPv4 or IPv6) configured on routers that use Junos OS Release 12.2 and later releases interoperate with VRRP (IPv4 or IPv6) configured on routers that use releases earlier than Junos OS Release 12.2.
- VRRPv3 for IPv4 does not interoperate with the previous versions of VRRP. If VRRPv2 IPv4 advertisement packets are received by a router on which VRRPv3 is enabled, the router transitions itself to the backup state to avoid creating multiple masters in the network. Due to this behavior, you must be cautious when enabling VRRPv3 on your existing VRRPv2 networks. See [“Upgrading from VRRPv2 to VRRPv3” on page 324](#) for more information.



NOTE: VRRPv3 advertisement packets are ignored by the routers on which previous versions of VRRP are configured.

Upgrading from VRRPv2 to VRRPv3

Enable VRRPv3 in your network only if VRRPv3 can be enabled on all the VRRP routers in your network.

Enable VRRPv3 on your VRRPv2 network only when upgrading from VRRPv2 to VRRPv3. Mixing the two versions of VRRP is not a permanent solution.



CAUTION: VRRP version change is considered catastrophic and disruptive and may not be hitless. The packet loss duration depends on many factors, such as number of VRRP groups, the interfaces and FPCs involved, and the load of other services and protocols running on the router.

Upgrading from VRRPv2 to VRRPv3 must be done very carefully to avoid traffic loss, due to these considerations:

- It is not possible to configure VRRPv3 on all routers simultaneously.
- During the transition period, both VRRPv2 and VRRPv3 operate in the network.
- Changing VRRP versions restarts the state machine for all VRRP groups.
- VRRPv3 (for IPv4) routers default to the backup state when they get VRRPv2 (for IPv4) advertisement packets.
- VRRPv2 (for IPv4) packets are always given the highest priority.
- Checksum differences between VRRPv2 and VRRPv3 (for IPv6) can create multiple master routers.

Disable VRRPv3 (for IPv6) on the backup routers while upgrading to avoid creating multiple master routers.

Table 17 on page 325 illustrates the steps and events that take place during a VRRPv2 to VRRPv3 transition. In Table 17 on page 325, two VRRPv2 routers, R1 and R2, are configured in two groups, G1 and G2. Router R1 acts as the master for G1, and Router R2 acts as the master for G2.

Table 17: VRRPv2 to VRRPv3 Transition Steps and Events

1. Upgrade Router R1 with Junos OS Release 12.2 or later.
 - Router R2 becomes the master for both G1 and G2.
 - After the upgrade of Router R1 is completed, Router R1 becomes the master for G1.
 - Router R2 remains the master for G2.
2. Upgrade Router R2 with Junos OS Release 12.2 or later.
 - Router R1 becomes the master for both G1 and G2.
 - After the upgrade of Router R2 is completed, Router R2 becomes the master for G2.
 - Router R1 remains the master for G1.

For IPv4	For IPv6
3. Enable VRRPv3 on Router R1. <ul style="list-style-type: none"> • Router R1 becomes the backup for both G1 and G2 because VRRPv2 IPv4 advertisement packets are given higher priority. 	3. Deactivate G1 and G2 on Router R2. <ul style="list-style-type: none"> • G1 and G2 on Router R1 become master.
4. Enable VRRPv3 on Router R2.	4. Enable VRRPv3 on Router R1. <ul style="list-style-type: none"> • Router R1 becomes the master for both G1 and G2.

Table 17: VRRPv2 to VRRPv3 Transition Steps and Events (continued)

- | | |
|--|--|
| • Router R1 becomes the master for G1. | 5. Enable VRRPv3 on Router R2. |
| • Router R2 becomes the master for G2. | 6. Activate G1 and G2 on Router R2. |
| | • Router R2 becomes the master for G2. |
| | • Router R1 remains the master for G1. |

When enabling VRRPv3, make sure that VRRPv3 is enabled on all the VRRP routers in the network because VRRPv3 (IPv4) does not interoperate with the previous versions of VRRP. For example, if VRRPv2 IPv4 advertisement packets are received by a router on which VRRPv3 is enabled, the router transitions itself to the backup state to avoid creating multiple masters in the network.

You can enable VRRPv3 by configuring the **version-3** statement at the **[edit protocols vrrp]** hierarchy level (for IPv4 or IPv6 networks). Configure the same protocol version on all VRRP routers on the LAN.

Functionality of VRRPv3 Features

Some Junos OS features differ in VRRPv3 from previous VRRP versions.

- [VRRPv3 Authentication on page 326](#)
- [VRRPv3 Advertisement Intervals on page 326](#)
- [Unified ISSU for VRRPv3 on page 326](#)

VRRPv3 Authentication

When VRRPv3 (for IPv4) is enabled, it does not allow authentication.

- The **authentication-type** and **authentication-key** statements cannot be configured for any VRRP groups.
- You must use non-VRRP authentication.

VRRPv3 Advertisement Intervals

VRRPv3 (for IPv4 and IPv6) advertisement intervals must be set with the **fast-interval** statement at the **[edit interfaces interface-name unit 0 family inet address ip-address vrrp-group group-name]** hierarchy level.

- Do not use the **advertise-interval** statement (for IPv4).
- Do not use the **inet6-advertise-interval** statement (for IPv6).

Unified ISSU for VRRPv3

Design changes for VRRP unified in-service software upgrade (ISSU) are made in Junos OS Release 15.1 to achieve the following functionality:

- Maintain protocol adjacency with peer routers during unified ISSU. Protocol adjacency created on peer routers for the router undergoing unified ISSU should not flap, which means that VRRP on the remote peer router should not flap.

- Maintain interoperability with competitive or complementary equipment.
- Maintain interoperability with other Junos OS releases and other Juniper Network products.

The values of the following configurations (found at the **[edit interfaces *interface-name* unit 0 family inet address *ip-address* vrrp-group *group-name*]** hierarchy level) need to be kept at maximum values to support unified ISSU:

- On the master router, the advertisement interval (the **fast-interval** statement) needs to be kept at 40950 milliseconds.
- On the backup router, the master-down interval (the **advertisements-threshold** statement) needs to be kept at the largest threshold value.

This VRRP unified ISSU design only works for VRRPv3. It is not supported on VRRPv1 or VRRPv2. Other limitations include the following:

- The VRRP unified ISSU takes care of VRRP only. Packet forwarding is the responsibility of the Packet Forwarding Engine. The Packet Forwarding Engine unified ISSU should ensure uninterrupted traffic flow.
- VRRP is not affected by any change event during unified ISSU, for example, the switchover of the master Routing Engine to backup or the backup Routing Engine to master.
- VRRP might stop and discard any running timer before entering into unified ISSU. This means the expected action upon the expiry of the timer never takes place. However, you can defer unified ISSU until the expiration of all running timers.
- Unified ISSU at both local and remote routers cannot be done simultaneously.

Release History Table

Release	Description
12.2	Junos OS Release 12.2 and later releases support VRRPv3.

**Related
Documentation**

- [Understanding VRRP on page 315](#)
- [Configuring Basic VRRP Support on page 332](#)

VRRP failover-delay Overview

Failover is a backup operational mode in which the functions of a network device are assumed by a secondary device when the primary device becomes unavailable because of a failure or a scheduled down time. Failover is typically an integral part of mission-critical systems that must be constantly available on the network.



NOTE: VRRP does not support session synchronization between members. If the master device fails, the backup device with the highest priority takes over as master and will begin forwarding packets. Any existing sessions will be dropped on the backup device as out-of-state.

A fast failover requires a short delay. Thus, failover-delay configures the failover delay time, in milliseconds, for VRRP and VRRP for IPv6 operations. Junos OS supports a range of 50 through 100000 milliseconds for delay in failover time.

The VRRP process (vrripd) running on the Routing Engine communicates a VRRP mastership change to the Packet Forwarding Engine for every VRRP session. Each VRRP group can trigger such communication to update the Packet Forwarding Engine with its own state or the state inherited from an active VRRP group. To avoid overloading the Packet Forwarding Engine with such messages, you can configure a failover-delay to specify the delay between subsequent Routing Engine to Packet Forwarding Engine communications.

The Routing Engine communicates a VRRP mastership change to the Packet Forwarding Engine to facilitate necessary state change on the Packet Forwarding Engine, such as reprogramming of Packet Forwarding Engine hardware filters, VRRP sessions and so on. The following sections elaborate the Routing Engine to Packet Forwarding Engine communication in two scenarios:

- [When failover-delay Is Not Configured on page 328](#)
- [When failover-delay Is Configured on page 329](#)

When failover-delay Is Not Configured

Without failover-delay configured, the sequence of events for VRRP sessions operated from the Routing Engine is as follows:

1. When the first VRRP group detected by the Routing Engine changes state, and the new state is master, the Routing Engine generates appropriate VRRP announcement messages. The Packet Forwarding Engine is informed about the state change, so that hardware filters for that group are reprogrammed without delay. The new master then sends gratuitous ARP message to the VRRP groups.
2. The delay in failover timer starts. By default, failover-delay timer is:
 - 500 milliseconds—when the configured VRRP announcement interval is less than 1 second.

- 2 seconds—when the configured VRRP announcement interval is 1 second or more, and the total number of VRRP groups on the router is 255.
 - 10 seconds—when the configured VRRP announcement interval is 1 second or more, and the number of VRRP groups on the router is more than 255.
3. The Routing Engine performs one-by-one state change for subsequent VRRP groups. Every time there is a state change, and the new state for a particular VRRP group is master, the Routing Engine generates appropriate VRRP announcement messages. However, communication toward the Packet Forwarding Engine is suppressed until the failover-delay timer expires.
 4. After failover-delay timer expires, the Routing Engine sends message to the Packet Forwarding Engine about all VRRP groups that managed to change the state. As a consequence, hardware filters for those groups are reprogrammed, and for those groups whose new state is master, gratuitous ARP messages are sent.

This process repeats until state transition for all VRRP groups is complete.

Thus, without configuring failover-delay, the full state transition (including states on the Routing Engine and the Packet Forwarding Engine) for the first VRRP group is performed immediately, while state transition on the Packet Forwarding Engine for remaining VRRP groups is delayed by at least 0.5-10 seconds, depending on the configured VRRP announcement timers and the number of VRRP groups. During this intermediate state, receiving traffic for VRRP groups for state changes that were not yet completed on the Packet Forwarding Engine might be dropped at the Packet Forwarding Engine level due to deferred reconfiguration of hardware filters.

When failover-delay Is Configured

When failover-delay is configured, the sequence of events for VRRP sessions operated from the Routing Engine is modified as follows:

1. The Routing Engine detects that some VRRP groups require a state change.
2. The failover-delay starts for the period configured. The allowed failover-delay timer range is 50 through 100000 milliseconds.
3. The Routing Engine performs one-by-one state change for the VRRP groups. Every time there is a state change, and the new state for a particular VRRP group is master, the Routing Engine generates appropriate VRRP announcement messages. However, communication toward the Packet Forwarding Engine is suppressed until the failover-delay timer expires.
4. After failover-delay timer expires, the Routing Engine sends message to the Packet Forwarding Engine about all VRRP groups that managed to change the state. As a consequence, hardware filters for those groups are reprogrammed, and for those groups whose new state is master, gratuitous ARP messages are sent.

This process repeats until state transition for all VRRP groups is complete.

Thus, when failover-delay is configured even the Packet Forwarding Engine state for the first VRRP group is deferred. However, the network operator has the advantage of configuring a failover-delay value that best suits the need of the network deployment to ensure minimal outage during VRRP state change.



NOTE: failover-delay influences only VRRP sessions operated by the VRRP process (vrrpd) running on the Routing Engine. For VRRP sessions distributed to the Packet Forwarding Engine, failover-delay configuration has no effect.

Related Documentation

- *failover-delay*

CHAPTER 26

Configuring VRRP

- [Configuring Basic VRRP Support on page 332](#)
- [Configuring VRRP on page 337](#)
- [VRRP and VRRP for IPv6 Overview on page 339](#)
- [Configuring VRRP and VRRP for IPv6 on page 339](#)
- [Configuring VRRP for IPv6 \(CLI Procedure\) on page 341](#)
- [Example: Configuring VRRP for IPv6 on page 342](#)
- [Configuring VRRP Authentication \(IPv4 Only\) on page 348](#)
- [Configuring VRRP Preemption and Hold Time on page 349](#)
- [Configuring the Advertisement Interval for the VRRP Master Router on page 350](#)
- [Configuring the Startup Period for VRRP Operations on page 353](#)
- [Configuring a Backup Router to Preempt the VRRP Master Router on page 353](#)
- [Configuring a Backup to Accept Packets Destined for the Virtual IP Address on page 354](#)
- [Modifying the Preemption Hold-Time Value for the VRRP Master Router on page 355](#)
- [Configuring the Asymmetric Hold Time for VRRP Routers on page 355](#)
- [Configuring Passive ARP Learning for Backup VRRP Routers on page 356](#)
- [Configuring VRRP Route Tracking on page 357](#)
- [Configuring a Logical Interface to Be Tracked for a VRRP Group on page 358](#)
- [Configuring a Route to Be Tracked for a VRRP Group on page 361](#)
- [Example: Configuring Multiple VRRP Owner Groups on page 363](#)
- [Configuring Inheritance for a VRRP Group on page 369](#)
- [Configuring an Interface to Accept All Packets Destined for the Virtual IP Address of a VRRP Group on page 371](#)
- [Configuring the Silent Period to Avoid Alarms Due to Delay in Receiving VRRP Advertisement Packets on page 372](#)
- [Enabling the Distributed Periodic Packet Management Process for VRRP on page 373](#)
- [Improving the Convergence Time for VRRP on page 374](#)
- [Configuring VRRP to Improve Convergence Time on page 376](#)
- [Tracing VRRP Operations on page 377](#)

- [Example: Configuring VRRP for Load Sharing on page 378](#)
- [Troubleshooting VRRP on page 383](#)

Configuring Basic VRRP Support



NOTE: Starting in Junos OS Release 13.2, VRRP nonstop active routing (NSR) is enabled only when you configure the `nonstop-routing` statement at the `[edit routing-options]` or `[edit logical system logical-system-name routing-options]` hierarchy level.

The Virtual Router Redundancy Protocol (VRRP) groups multiple routing devices into a virtual router. At any time, one of the VRRP routing platforms is the master (active) and the others are backups. If the master fails, one of the backup routing platforms becomes the new master router.

To configure basic VRRP support, configure VRRP groups on interfaces by including the `vrrp-group` statement:

```
vrrp-group group-id {  
  priority number;  
  virtual-address [ addresses ];  
}
```

An interface can be a member of multiple VRRP groups. Within a VRRP group, the master virtual router and the backup virtual router must be configured on different routing platforms.

You can include this statement at the following hierarchy level:

- `[edit interfaces interface-name unit logical-unit-number family inet address address]`

Mandatory parameters to configure a VRRP group are as follows (examples will follow):

1. Configure the group identifier (mandatory).
2. Configure the group:
 - Configure the virtual IP address of one or more virtual routers that are members of the VRRP group (mandatory).
 - Configure the virtual link-local address (VRRP for IPv6 only). The virtual link-local address is autogenerated when you enable VRRPv3 on the interface. You may explicitly define a virtual link-local address for each VRRP for the IPv6 group. The virtual link-local address must be on the same subnet as the physical interface address.
 - Configure the priority for the routing platform to become the master virtual router (mandatory).

When choosing a VRRP group identifier, consider the following:

- In Junos OS releases prior to 17.3R1, you should not use the same VRRP group identifier on more than one subinterface on a given physical interface. This causes the VRRP virtual MAC address to be deleted from the packet forwarding engine, resulting in packet drops due to unknown MAC address. If your VRRP configuration needs to scale beyond 255 groups, consider configuring VRRP over an integrated routing and bridging (IRB) interface, since this restriction does not apply to IRB interfaces.
- Starting in Junos OS release 17.3R1, if network-services is configured in IP mode, don't configure the same VRRP group ID for multiple VRRP sessions on the same physical interface unless VRRP delegation is disabled. If multiple VRRP sessions are configured on the same physical interface with the same VRRP group ID while VRRP delegation is enabled, the other VRRP virtual IP addresses become unreachable when one of the logical interfaces is deleted.
- Starting in Junos OS release 17.3R1, if network-services is configured in enhanced-ip mode, you can use the same VRRP group ID for multiple VRRP sessions.

When configuring a virtual IP address, consider the following:

- The virtual IP address must be the same for all routing platforms in the VRRP group.
- If you configure a virtual IP address to be the same as the physical interface's address, the interface becomes the master virtual router for the group. In this case, you must configure the priority to be 255, and you must configure preemption by including the **preempt** statement.
- If the virtual IP address you choose is not the same as the physical interface's address, you must ensure that the virtual IP address does not appear anywhere else in the routing platform's configuration. Verify that you do not use this address for other interfaces, for the IP address of a tunnel, or for the IP address of static ARP entries.
- You cannot configure a virtual IP address to be the same as the interface's address for an aggregated Ethernet interface. This configuration is not supported.
- For VRRP for IPv6, the **EUI-64** option cannot be used. In addition, the Duplicate Address Detection (DAD) process will not run for virtual IPv6 addresses.
- You cannot configure the same virtual IP address on interfaces that belong to the same logical system and routing instance combination. However, you can configure the same virtual IP address on interfaces that belong to different logical systems and routing instance combinations.

In determining what priority will make a given routing platform in a VRRP group a master or backup, consider the following:

- You can force assignment of master and backup routers using priorities from 1 through 255, where 255 is the highest priority.
- The priority value for the VRRP router that owns the IP address(es) associated with the virtual router must be 255.
- VRRP routers backing up a virtual router must use priority values from 1 through 254.
- The default priority value for VRRP routers backing up a virtual router is 100.
- Are there tracked interfaces or routes with priority costs?

The priority cost is the value associated with a tracked logical interface or route that is to be subtracted from the configured VRRP priority when the tracked logical interface or route goes down, forcing a new master router election. The value of a priority cost can be from 1 through 254. The sum of the priority costs for all tracked logical interfaces or routes must be less than or equal to the configured priority of the VRRP group.



NOTE: Mixed tagging (configuring two logical interfaces on the same Ethernet port, one with single-tag framing and one with dual-tag framing) is supported only for interfaces on Gigabit Ethernet IQ2 and IQ PICs. If you include the `flexible-vlan-tagging` statement at the `[edit interfaces interface-name]` hierarchy level for a VRRP-enabled interface on a PIC that does not support mixed tagging, VRRP on that interface is disabled. In the output of the `show vrrp summary operational` command, the interface status is listed as `Down`.



NOTE: If you enable MAC source address filtering on an interface, you must include the virtual MAC address in the list of source MAC addresses that you specify in the `source-address-filter` statement at the `[edit interfaces interface-name]` hierarchy level. (For more information, see the [Junos OS Network Interfaces Library for Routing Devices](#).) MAC addresses ranging from `00:00:5e:00:01:00` through `00:00:5e:00:01:ff` are reserved for VRRP, as defined in RFC 2378. The VRRP group number must be the decimal equivalent of the last hexadecimal byte of the virtual MAC address.

Here are specific examples of configuring a VRRP group.

Configuring for VRRP IPv4 Groups

To configure basic VRRP (IPv4) groups on interfaces:



NOTE: You can also configure a VRRP IPv4 group at the `[edit logical-systems logical-system-name]` hierarchy level.

1. Configure the group identifier.

```
[edit interfaces interface-name unit logical-unit-number family inet address address]
user@device# set vrrp-group group-id
```

Assign a value from 0 through 255.

2. Configure the VRRP for IPv4 group:

- Configure the virtual IP address of one or more virtual routers that are members of the VRRP group.

```
[edit interfaces interface-name unit logical-unit-number family inet address address]
```

```
user@device# set vrrp-group group-id virtual-address [ addresses ]
```

Normally, you configure only one virtual IP address per group. However, you can configure up to eight addresses. Do not include a prefix length in a virtual IP address.

- Configure the priority for this routing platform to become the master virtual router.

```
[edit interfaces interface-name unit logical-unit-number family inet address address]
user@device# set vrrp-group group-id priority number
```

Configure the value used to elect the master virtual router in the VRRP group. It can be a number from 1 through 255. The default value for backup routers is 100. A larger value indicates a higher priority. The routing platform with the highest priority within the group becomes the master router. Master router sends periodic VRRP advertisement messages to each virtual routers. The backup routers do not attempt to preempt the master router unless it has higher priority. This eliminates service disruption unless a more preferred path becomes available. It is possible to administratively prohibit all preemption attempts, with the exception of a VRRP router becoming master router of any virtual router associated with addresses it owns.

Configuring VRRP for IPv6 Groups

To configure basic VRRP for IPv6 groups on interfaces:



NOTE: You can also configure a VRRP IPv6 group at the [edit logical-systems *logical-system-name*] hierarchy level.

1. Configure the group identifier.

```
[edit interfaces interface-name unit logical-unit-number family inet6 address
ipv6-address]
user@device# set vrrp-inet6-group group-id
```

Assign a value from 0 through 255.

2. Configure the VRRP for IPv6 group:

- Configure the virtual IP address of one or more virtual routers that are members of the VRRP group.

```
[edit interfaces interface-name unit logical-unit-number family inet6 address
ipv6-address]
user@device# set vrrp-inet6-group group-id virtual-inet6-address [ ipv6-addresses
]
```

Normally, you configure only one virtual IP address per group. However, you can configure up to eight addresses. Do not include a prefix length in a virtual IP address.

- Configure the virtual link-local address.

```
[edit interfaces interface-name unit logical-unit-number family inet6 address
  ipv6-address]
user@device# set vrrp-inet6-group group-id virtual-link-local-address ipv6-address
```

You must explicitly define a virtual link-local address for each VRRP for IPv6 group. Otherwise, when you attempt to commit the configuration, the commit request fails. The virtual link-local address must be on the same subnet as the physical interface address.

- Configure the priority for this routing platform to become the master virtual router.

```
[edit interfaces interface-name unit logical-unit-number family inet6 address
  ipv6-address]
user@device# set vrrp-inet6-group group-id priority number
```

Configure the value used to elect the master virtual router in the VRRP group. It can be a number from 1 through 255. The default value for backup routers is 100. A larger value indicates a higher priority. The routing platform with the highest priority within the group becomes the master router. If there are two or more backup routers with the same priority, the router that has the highest primary address becomes the master.

Release History Table

Release	Description
18.1R1	Master router sends periodic VRRP advertisement messages to each virtual routers. The backup routers do not attempt to preempt the master router unless it has higher priority. This eliminates service disruption unless a more preferred path becomes available. It is possible to administratively prohibit all preemption attempts, with the exception of a VRRP router becoming master router of any virtual router associated with addresses it owns.
17.3R1	Starting in Junos OS release 17.3R1, if network-services is configured in IP mode, don't configure the same VRRP group ID for multiple VRRP sessions on the same physical interface unless VRRP delegation is disabled.
17.3R1	Starting in Junos OS release 17.3R1, if network-services is configured in enhanced-ip mode, you can use the same VRRP group ID for multiple VRRP sessions.
13.2	Starting in Junos OS Release 13.2, VRRP nonstop active routing (NSR) is enabled only when you configure the nonstop-routing statement at the [edit routing-options] or [edit logical system <i>logical-system-name</i> routing-options] hierarchy level.

Related Documentation

- [Configuring a Logical Interface to Be Tracked for a VRRP Group on page 358](#)
- [Configuring a Route to Be Tracked for a VRRP Group on page 361](#)
- [Junos OS Support for VRRPv3 on page 322](#)
- [Understanding VRRP on page 315](#)
- [Configuring the Startup Period for VRRP Operations on page 353](#)

- [Configuring VRRP Authentication \(IPv4 Only\) on page 348](#)
- [Configuring the Advertisement Interval for the VRRP Master Router on page 350](#)
- [Configuring VRRP on page 337](#)

Configuring VRRP

Configure one master (Router A) and one backup (Router B) routing platform. The address configured in the **virtual-address** statements differs from the addresses configured in the **address** statements. When you configure multiple VRRP groups on an interface, you configure one to be the master virtual router for that group.

On Router A

```
[edit interfaces]
ge-0/0/0 {
  unit 0 {
    family inet {
      address 192.168.1.20/24 {
        vrrp-group 27 {
          virtual-address 192.168.1.15;
          priority 254;
          authentication-type simple;
          authentication-key booJUM;
        }
      }
    }
  }
}
```

On Router B

```
[edit interfaces]
ge-4/2/0 {
  unit 0 {
    family inet {
      address 192.168.1.24/24 {
        vrrp-group 27 {
          virtual-address 192.168.1.15;
          priority 200;
          authentication-type simple;
          authentication-key booJUM;
        }
      }
    }
  }
}
```

Configuring One Router to Be the Master Virtual Router for the Group

```
[edit interfaces]
ge-0/0/0 {
  unit 0 {
    family inet {
      address 192.168.1.20/24 {
```

```
        vrrp-group 2 {  
            virtual-address 192.168.1.20;  
            priority 255;  
            advertise-interval 3;  
            preempt;  
        }  
        vrrp-group 10 {  
            virtual-address 192.168.1.55;  
            priority 201;  
            advertise-interval 3;  
        }  
        vrrp-group 1 {  
            virtual-address 192.168.1.54;  
            priority 22;  
            advertise-interval 4;  
        }  
    }  
}
```

Configuring VRRP and MAC Source Address Filtering

The VRRP group number is the decimal equivalent of the last byte of the virtual MAC address.

```
[edit interfaces]  
ge-5/2/0 {  
    gigeether-options {  
        source-filtering;  
        source-address-filter {  
            00:00:5e:00:01:0a; # Virtual MAC address  
        }  
    }  
    unit 0 {  
        family inet {  
            address 192.168.1.10/24 {  
                vrrp-group 10; # VRRP group number  
                virtual-address 192.168.1.10;  
                priority 255;  
                preempt;  
            }  
        }  
    }  
}
```

Related Documentation

- [Understanding VRRP on page 315](#)
- [Example: Configuring VRRP for IPv6 on page 342](#)
- [Configuring VRRP Route Tracking on page 357](#)

VRRP and VRRP for IPv6 Overview

You can configure the Virtual Router Redundancy Protocol (VRRP) and VRRP for IPv6 for the following interfaces:

- Ethernet
- Fast Ethernet
- Tri-Rate Ethernet copper
- Gigabit Ethernet
- 10-Gigabit Ethernet LAN/WAN PIC
- Ethernet logical interfaces

VRRP and VRRP for IPv6 allow hosts on a LAN to make use of redundant routers on that LAN without requiring more than the static configuration of a single default route on the hosts. The VRRP routers share the IP address corresponding to the default route configured on the hosts. At any time, one of the VRRP routers is the master (active) and the others are backups. If the master fails, one of the backup routers becomes the new master router, thus always providing a virtual default router and allowing traffic on the LAN to be routed without relying on a single router.

VRRP is defined in RFC 3768, *Virtual Router Redundancy Protocol*.

For VRRP and VRRP for IPv6 overview information, configuration guidelines, and statement summaries, see the *High Availability Feature Guide*.

Related Documentation

- [Configuring VRRP and VRRP for IPv6 on page 339](#)
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring VRRP and VRRP for IPv6

To configure VRRP or VRRP for IPv6, include the **vrrp-group** or **vrrp-inet6-group** statement, respectively. These statements are available at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet address *address*]

The VRRP and VRRP IPv6 configuration statements are as follows:

```
(vrrp-group | vrrp-inet6-group) group-number {
  (accept-data | no-accept-data);
  advertise-interval seconds;
  authentication-key key;
  authentication-type authentication;
  fast-interval milliseconds;
```

```
(preempt | no-preempt) {  
    hold-time seconds;  
}  
priority-number number;  
track {  
    priority-hold-time;  
    interface interface-name {  
        priority-cost priority;  
        bandwidth-threshold bits-per-second {  
            priority-cost;  
        }  
    }  
}  
virtual-address [ addresses ];  
}
```

You can configure VRRP IPv6 with a global unicast address.

To trace VRRP and VRRP for IPv6 operations, include the **traceoptions** statement at the **[edit protocols vrrp]** hierarchy level:

```
[edit protocols vrrp]  
traceoptions {  
    file <filename> <files number <match regular-expression <microsecond-stamp>  
        <size size> <world-readable | no-world-readable>;  
    flag flag;  
    no-remote-trace;  
}
```

When there are multiple VRRP groups, there is a few seconds delay between the time the first gratuitous ARP is sent out and the rest of the gratuitous ARP are sent. Configuring failover-delay compensates for this delay. To configure the failover delay from 500 to 2000 milliseconds for VRRP and VRRP for IPv6 operations, include the **failover-delay milliseconds** statement at the **[edit protocols vrrp]** hierarchy level:

```
[edit protocols vrrp]  
failover-delay milliseconds;
```

To configure the startup period for VRRP and VRRP for IPv6 operations, include the **startup-silent-period** statement at the **[edit protocols vrrp]** hierarchy level:

```
[edit protocols vrrp]  
startup-silent-period seconds;
```

To enable VRRPv3, set the **version-3** statement at the **[edit protocols vrrp]** hierarchy level:

```
[edit protocols vrrp]  
version-3;
```

- Related Documentation**
- *failover-delay*
 - *traceoptions*
 - *failover-delay*
 - [vrrp-group on page 635](#)
 - [VRRP and VRRP for IPv6 Overview on page 339](#)
 - *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring VRRP for IPv6 (CLI Procedure)

By configuring the Virtual Router Redundancy Protocol (VRRP) on EX Series switches, you can enable hosts on a LAN to make use of redundant routing platforms on that LAN without requiring more than the static configuration of a single default route on the hosts. You can configure VRRP for IPv6 on Gigabit Ethernet, 10-Gigabit Ethernet, and logical interfaces.

To configure VRRP for IPv6:

1. Configure VRRP group support on interfaces:

```
[edit interfaces interface-name unit logical-unit-number family inet6 address
address]
user@switch# set vrrp-inet6-group group-id priority number virtual-inet6-address address
virtual-link-local-address ipv6-address
```

You must explicitly define a virtual link local address for each VRRP for IPv6 group. Otherwise, when you attempt to commit the configuration, the commit request fails. The virtual link local address must be on the same subnet as the physical interface address.

2. If you want to configure the priority order in which this switch functioning as a backup router becomes the master router if the master router becomes nonoperational, configure a priority for this switch:

```
[edit interfaces interface-name unit logical-unit-number family inet6 address
address vrrp-inet6-group group-id]
user@switch# set priority number
```

3. Specify the interval in milliseconds in which the master router sends advertisement packets to the members of the VRRP group:

```
[edit interfaces interface-name unit logical-unit-number family inet6 address
address vrrp-inet6-group group-id]
user@switch# set inet6-advertise-interval milliseconds
```

4. By default, a higher-priority backup router preempts a lower-priority master router.
 - To explicitly enable the master router to be preempted:

```
[edit interfaces interface-name unit logical-unit-number family inet6 address
address vrrp-inet6-group group-id]
user@switch# set preempt
```

- To prohibit a higher-priority backup router from preempting a lower priority master router:

```
[edit interfaces interface-name unit logical-unit-number family inet6 address
address vrrp-inet6-group group-id]
user@switch# set no-preempt
```

**Related
Documentation**

- [show vrrp](#)
- [Understanding VRRP on page 315](#)

Example: Configuring VRRP for IPv6

This example shows how to configure VRRP properties for IPv6 in one master (Router A) and one backup (Router B).

- [Requirements on page 342](#)
- [Overview on page 342](#)
- [Configuring VRRP on page 342](#)
- [Verification on page 347](#)

Requirements

This example uses the following hardware and software components:

- Two routers
- Junos OS Release 11.3 or later
- Junos OS Release 18.1 R1 or later for SRX Series Services Gateways.
- Static routing or a dynamic routing protocol enabled on both routers.

Overview

This example uses a VRRP group, which has its own virtual IPv6 address. Devices on the LAN use this virtual IPv6 address as their default gateway. If the master router fails, the backup router takes over for it.

Configuring VRRP

- [Configuring Router A on page 343](#)
- [Configuring Router B on page 345](#)

Configuring Router A

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set interfaces ge-1/0/0 unit 0 family inet6 address 2001:db8::5:0:0:6/64
set interfaces ge-1/0/0 unit 0 family inet6 address 2001:db8::6:0:0:6/64 vrrp-inet6-group
  3 virtual-inet6-address 2001:db8::6:0:0:7
set interfaces ge-1/0/0 unit 0 family inet6 address 2001:db8::6:0:0:6/64 vrrp-inet6-group
  3 virtual-link-local-address 2001:db8::5:0:0:7
set interfaces ge-1/0/0 unit 0 family inet6 address 2001:db8::6:0:0:6/64 vrrp-inet6-group
  3 priority 200
set interfaces ge-1/0/0 unit 0 family inet6 address 2001:db8::6:0:0:6/64 vrrp-inet6-group
  3 preempt
set protocols router-advertisement interface ge-1/0/0.0 prefix 2001:db8::/64
set protocols router-advertisement interface ge-1/0/0.0 max-advertisement-interval 4
set protocols router-advertisement interface ge-1/0/0.0 virtual-router-only
```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure this example:

1. Configure the interfaces.

```
[edit interfaces]
user@hostA# set ge-1/0/0 unit 0 family inet6 address 2001:db8::5:0:0:6/64
user@hostA# set ge-1/0/0 unit 0 family inet6 address 2001:db8::6:0:0:6/64
```

2. Configure the IPv6 VRRP group identifier.

```
[edit interfaces ge-1/0/0 unit 0 family inet6 address 2001:db8::6:0:0:6/64]
user@hostA# set vrrp-inet6-group 3
```

3. Configure the virtual IP address of a virtual router that is a member of the VRRP group.

```
[edit interfaces ge-1/0/0 unit 0 family inet6 address 2001:db8::6:0:0:6/64]
user@hostA# set vrrp-inet6-group 3 virtual-inet6-address 2001:db8::6:0:0:7
```

4. Configure the virtual link-local address.

```
[edit interfaces ge-1/0/0 unit 0 family inet6 address 2001:db8::6:0:0:6/64]
user@hostA# set vrrp-inet6-group 3 virtual-link-local-address 2001:db8::5:0:0:7
```

5. Configure the priority for this routing platform to become the master virtual router.

```
[edit interfaces ge-1/0/0 unit 0 family inet6 address 2001:db8::6:0:0:6/64]
user@hostA# set vrrp-inet6-group 3 priority 200
```

6. By default, a higher-priority backup router preempts a lower-priority master router. To explicitly enable the master router to be preempted, include the **preempt** statement.

```
[edit interfaces ge-1/0/0 unit 0 family inet6 address 2001:db8::6:0:0:6/64]
user@hostA# set vrrp-inet6-group 3 preempt
```

7. For VRRP for IPv6, you must configure the interface on which VRRP is configured to send IPv6 router advertisements for the VRRP group. When an interface receives an IPv6 router solicitation message, it sends an IPv6 router advertisement to all VRRP groups configured on it.

```
[edit protocols router-advertisement interface ge-1/0/0.0]
user@hostA# set prefix 2001:db8::/64
user@hostA# set max-advertisement-interval 4
```

8. Configure router advertisements to be sent only for VRRP IPv6 groups configured on the interface if the groups are in the master state.

```
[edit protocols router-advertisement interface ge-1/0/0.0]
user@hostA# set virtual-router-only
```

Results From configuration mode, confirm your configuration by entering the **show interfaces** and **show protocols router-advertisement** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@hostA# show interfaces
ge-1/0/0 {
  unit 0 {
    family inet6 {
      address 2001:db8::5:0:0:6/64;
      address 2001:db8::6:0:0:6/64 {
        vrrp-inet6-group 3;
        vrrp-inet6-group 3 virtual-inet6-address 2001:db8::6:0:0:7;
        vrrp-inet6-group 3 virtual-link-local-address 2001:db8::5:0:0:7;
        vrrp-inet6-group 3 priority 200;
        vrrp-inet6-group 3 preempt;
      }
    }
  }
}
```

```
[edit]
user@hostA# show protocols router-advertisement
interface ge-1/0/0.0 {
  prefix 2001:db8::/64;
  max-advertisement-interval 4;
  virtual-router-only;
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Router B

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set interfaces ge-1/0/0 unit 0 family inet6 address 2001:db8::5:0:0:8/64
set interfaces ge-1/0/0 unit 0 family inet6 address 2001:db8::6:0:0:8/64 vrrp-inet6-group
  3 virtual-inet6-address 2001:db8::6:0:0:7
set interfaces ge-1/0/0 unit 0 family inet6 address 2001:db8::6:0:0:8/64 vrrp-inet6-group
  3 virtual-link-local-address 2001:db8::5:0:0:7
set interfaces ge-1/0/0 unit 0 family inet6 address 2001:db8::6:0:0:8/64 vrrp-inet6-group
  3 priority 100
set interfaces ge-1/0/0 unit 0 family inet6 address 2001:db8::6:0:0:8/64 vrrp-inet6-group
  3 preempt
set protocols router-advertisement interface ge-1/0/0.0 prefix 2001:db8::/64
set protocols router-advertisement interface ge-1/0/0.0 max-advertisement-interval 4
set protocols router-advertisement interface ge-1/0/0.0 virtual-router-only
```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure this example:

1. Configure the interfaces.

```
[edit interfaces]
user@hostB# set ge-1/0/0 unit 0 family inet6 address 2001:db8::5:0:0:8/64
user@hostB# set ge-1/0/0 unit 0 family inet6 address 2001:db8::6:0:0:8/64
```

2. Configure the IPv6 VRRP group identifier.

```
[edit interfaces ge-1/0/0 unit 0 family inet6 address 2001:db8::6:0:0:8/64]
user@hostB# set vrrp-inet6-group 3
```

3. Configure the virtual IP address of a virtual router that is a member of the VRRP group.

```
[edit interfaces ge-1/0/0 unit 0 family inet6 address 2001:db8::6:0:0:8/64]
user@hostB# set vrrp-inet6-group 3 virtual-inet6-address 2001:db8::6:0:0:7
```

4. Configure the virtual link-local address.

```
[edit interfaces ge-1/0/0 unit 0 family inet6 address 2001:db8::6:0:0:8/64]
user@hostB# set vrrp-inet6-group 3 virtual-link-local-address 2001:db8::5:0:0:7
```

5. Configure the priority for this routing platform to become the master virtual router.

```
[edit interfaces ge-1/0/0 unit 0 family inet6 address 2001:db8::6:0:0:8/64]
user@hostB# set vrrp-inet6-group 3 priority 100
```

6. By default, a higher-priority backup router preempts a lower-priority master router. To explicitly enable the master router to be preempted, include the **preempt** statement.

```
[edit interfaces ge-1/0/0 unit 0 family inet6 address 2001:db8::6:0:0:8/64]
user@hostB# set vrrp-inet6-group 3 preempt
```

7. Configure the interface on which VRRP is configured to send IPv6 router advertisements for the VRRP group. When an interface receives an IPv6 router solicitation message, it sends an IPv6 router advertisement to all VRRP groups configured on it.

```
[edit protocols router-advertisement interface ge-1/0/0.0]
user@hostB# set prefix 2001:db8::/64
user@hostB# set max-advertisement-interval 4
```

8. Configure router advertisements to be sent only for VRRP IPv6 groups configured on the interface if the groups are in the master state.

```
[edit protocols router-advertisement interface ge-1/0/0.0]
user@hostB# set virtual-router-only
```

Results From configuration mode, confirm your configuration by entering the **show interfaces** and **show protocols router-advertisement** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@hostB# show interfaces
ge-1/0/0 {
  unit 0 {
    family inet6 {
      address 2001:db8::5:0:0:8/64;
```



```

address 2001:db8::6:0:0:8/64 {
    vrrp-inet6-group 3;
    vrrp-inet6-group 3 virtual-inet6-address 2001:db8::6:0:0:7;
    vrrp-inet6-group 3 virtual-link-local-address 2001:db8::5:0:0:7;
    vrrp-inet6-group 3 priority 100;
    vrrp-inet6-group 3 preempt;
}
}
}
}

```

```

[edit]
user@hostB# show protocols router-advertisement
interface ge-1/0/0.0 {
    prefix 2001:db8::/64;
    max-advertisement-interval 4;
    virtual-router-only;
}

```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

- [Verifying that VRRP Is Working on Router A on page 347](#)
- [Verifying that VRRP Is Working on Router B on page 348](#)

Verifying that VRRP Is Working on Router A

Purpose Verify that VRRP is active on Router A and that its role in the VRRP group is correct.

Action Use the following command to verify that VRRP is active on Router A and that the router is master for group 3.

```
user@hostA> show vrrp
```

Interface	State	Group	VR state	Timer	Type	Address
ge-1/0/0.0	up					3
master	A .0327	1c1	2001:db8::6:0:0:6/64	2001:db8::6:0:0:7		
				vip		

Meaning The **show vrrp** command displays fundamental information about the VRRP configuration. This output shows that the VRRP group is active and that this router has assumed the master role. The **lcl** address is the physical address of the interface and the **vip** address is the virtual address shared by both routers. The **Timer** value (**A .0327**) indicates the remaining time (in seconds) in which this router expects to receive a VRRP advertisement from the other router.

Verifying that VRRP Is Working on Router B

Purpose Verify that VRRP is active on Router B and that its role in the VRRP group is correct.

Action Use the following command to verify that VRRP is active on Router B and that the router is backup for group 3.

```
user@hostB> show vrrp
```

Interface	State	Group	VR state	Timer	Type	Address
backup	A .0327	1c1	2001:db8::6:0:0:8/64	ge-1/0/0.0 vip 2001:db8::6:0:0:7	up	3

Meaning The `show vrrp` command displays fundamental information about the VRRP configuration. This output shows that the VRRP group is active and that this router has assumed the backup role. The `lcl` address is the physical address of the interface and the `vip` address is the virtual address shared by both routers. The **Timer** value (**A.0327**) indicates the remaining time (in seconds) in which this router expects to receive a VRRP advertisement from the other router.

- Related Documentation**
- [Understanding VRRP on page 315](#)
 - [Configuring VRRP on page 337](#)
 - [Configuring VRRP Route Tracking on page 357](#)

Configuring VRRP Authentication (IPv4 Only)

VRRP (IPv4 only) protocol exchanges can be authenticated to guarantee that only trusted routing platforms participate in routing in an autonomous system (AS). By default, VRRP authentication is disabled. You can configure one of the following authentication methods. Each VRRP group must use the same method.

- Simple authentication—Uses a text password included in the transmitted packet. The receiving routing platform uses an authentication key (password) to verify the packet.
- Message Digest 5 (MD5) algorithm—Creates the authentication data field in the IP authentication header. This header is used to encapsulate the VRRP PDU. The receiving routing platform uses an authentication key (password) to verify the authenticity of the IP authentication header and VRRP PDU.

To enable authentication and specify an authentication method, include the **authentication-type** statement:

```
authentication-type authentication;
```

authentication can be **simple** or **md5**. The authentication type must be the same for all routing platforms in the VRRP group.

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group *group-id*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group *group-id*]

If you include the **authentication-type** statement, you can configure a key (password) on each interface by including the **authentication-key** statement:

authentication-key *key*;

key (the password) is an ASCII string. For simple authentication, it can be from 1 through 8 characters long. For MD5 authentication, it can be from 1 through 16 characters long. If you include spaces, enclose all characters in quotation marks (" "). The key must be the same for all routing platforms in the VRRP group.

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group *group-id*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group *group-id*]



NOTE: When VRRPv3 is enabled, the **authentication-type** and **authentication-key** statements cannot be configured for any VRRP groups. Therefore, if authentication is required, you need to configure alternative non-VRRP authentication mechanisms.

Related Documentation

- [Understanding VRRP on page 315](#)
- [Junos OS Support for VRRPv3 on page 322](#)
- [Configuring Basic VRRP Support on page 332](#)
- [Configuring VRRP on page 337](#)

Configuring VRRP Preemption and Hold Time

- [Configuring VRRP Preemption on page 350](#)
- [Configuring the Preemption Hold Time on page 350](#)

Configuring VRRP Preemption

By default, a higher-priority VRRP backup switch preempts a lower-priority master switch. To explicitly enable this behavior, include the following statement:

```
preempt;
```

To prohibit a higher-priority VRRP backup switch from preempting a lower-priority master switch, include the following statement on the lower-priority switch:

```
no-preempt;
```

You can include these statements at the following hierarchy level:

- [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group *group-id*]

Configuring the Preemption Hold Time

You can also configure a preemption hold time, which is the number of seconds a higher-priority backup router that has just started up waits before preempting the master router. You might want to configure a hold time so that routing protocols or other Junos OS components converge before preemption occurs.

The hold time is applied only on startup. By default, the hold-time value is 0 seconds, meaning that preemption can occur immediately after the backup router starts up.

To modify the preemption hold-time value, configure the following statement:

```
hold-time seconds;
```

The hold time can be from 0 through 3600 seconds.

You can include this statement at the following hierarchy level:

- [edit interfaces *interface-name* unit *logical-unit-number* family inet address vrrp-group *group-id*] preempt

Related Documentation

- [Understanding VRRP on page 315](#)
- [Configuring Basic VRRP Support on page 332](#)
- [Example: Configuring VRRP for Load Sharing on page 378](#)
- [asymmetric-hold-time on page 602](#)

Configuring the Advertisement Interval for the VRRP Master Router

By default, the master router sends VRRP advertisement packets every second to all members of the VRRP group. These packets indicate that the master router is still

operational. If the master router fails or becomes unreachable, the backup router with the highest priority value becomes the new master router.

You can modify the advertisement interval in seconds or in milliseconds. The interval must be the same for all routing platforms in the VRRP group.

For VRRP for IPv6, you must configure IPv6 router advertisements for the interface on which VRRP is configured to send IPv6 router advertisements for the VRRP group. To do so, include the **interface *interface-name*** statement at the **[edit protocols router-advertisement]** hierarchy level. (For information about this statement and guidelines, see the *Junos OS Routing Protocols Library*.) When an interface receives an IPv6 router solicitation message, it sends an IPv6 router advertisement to all VRRP groups configured on it. In the case of logical systems, IPv6 router advertisements are not sent to VRRP groups.



NOTE: The master VRRP for an IPv6 router must respond to a router solicitation message with the virtual IP address of the router. However, when the **interface *interface-name*** statement is included at the **[edit protocols router-advertisement]** hierarchy level, the backup VRRP for an IPv6 router might send a response before the VRRP master responds, so that the default route of the client is not set to the master VRRP router's virtual IP address. To avoid this situation, include the **virtual-router-only** statement at the **[edit protocols router-advertisement interface *interface-name*]** hierarchy level. When this statement is included, router advertisements are sent only for VRRP IPv6 groups configured on the interface (if the groups are in the master state). You must include this statement on both the master and backup VRRP for IPv6 routers.

This topic contains the following sections:

- [Modifying the Advertisement Interval in Seconds on page 351](#)
- [Modifying the Advertisement Interval in Milliseconds on page 352](#)

Modifying the Advertisement Interval in Seconds

To modify the time, in seconds, between the sending of VRRP advertisement packets, include the **advertise-interval** statement:

```
advertise-interval seconds;
```

The interval can be from 1 through 255 seconds.

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group *group-id*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group *group-id*]**



NOTE: When VRRPv3 is enabled, the `advertise-interval` statement cannot be used to configure advertisement intervals. Instead, use the `fast-interval` statement to configure advertisement intervals.

Modifying the Advertisement Interval in Milliseconds

To modify the time, in milliseconds, between the sending of VRRP advertisement packets, include the **fast-interval** statement:

```
fast-interval milliseconds;
```

The interval can be from 10 through 40,950 milliseconds.

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family (inet | inet6) address *address* (vrrp-group | vrrp-inet6-group) *group-id*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family (inet | inet6) address *address* (vrrp-group | vrrp-inet6-group) *group-id*]



NOTE: In the VRRP PDU, Junos OS sets the advertisement interval to 0. When you configure VRRP with other vendors' routers, the `fast-interval` statement works correctly only when the other routers also have an advertisement interval set to 0 in the VRRP PDUs. Otherwise, Junos OS interprets other routers' settings as advertisement timer errors.

To modify the time, in milliseconds, between the sending of VRRP for IPv6 advertisement packets, include the **inet6-advertise-interval** statement:

```
inet6-advertise-interval ms;
```

The range of values is from 100 through 40,950 milliseconds (ms).

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family inet6 address *address* vrrp-inet6-group *group-id*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet6 address *address* vrrp-inet6-group *group-id*]



NOTE: When VRRPv3 is enabled, the `inet6-advertise-interval` statement cannot be used to configure advertisement intervals. Instead, use the `fast-interval` statement to configure advertisement intervals.

- Related Documentation**
- [Understanding VRRP on page 315](#)
 - [Junos OS Support for VRRPv3 on page 322](#)
 - [Configuring Basic VRRP Support on page 332](#)
 - [Configuring a Backup Router to Preempt the VRRP Master Router on page 353](#)
 - [Modifying the Preemption Hold-Time Value for the VRRP Master Router on page 355](#)
 - [Configuring the Asymmetric Hold Time for VRRP Routers on page 355](#)
 - [Configuring the Silent Period to Avoid Alarms Due to Delay in Receiving VRRP Advertisement Packets on page 372](#)
 - [Configuring VRRP on page 337](#)

Configuring the Startup Period for VRRP Operations

To configure the startup period for VRRP operations, include the **startup-silent-period** statement at the **[edit protocols vrrp]** hierarchy level:

```
[edit protocols vrrp]
  startup-silent-period seconds;
```



NOTE: During the silent startup period, the **show vrrp detail** command output shows a value of 0 for Master priority, and your own IP address for Master router. These values indicate that the Master selection is not completed yet, and these values can be ignored.

- Related Documentation**
- [Understanding VRRP on page 315](#)
 - [Configuring Basic VRRP Support on page 332](#)
 - [Configuring VRRP Authentication \(IPv4 Only\) on page 348](#)
 - [Configuring VRRP on page 337](#)

Configuring a Backup Router to Preempt the VRRP Master Router

By default, a higher-priority backup router preempts a lower-priority master router. To explicitly enable the master router to be preempted, include the **preempt** statement:

```
preempt;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number* family (inet | inet6) address *address* (vrrp-group | vrrp-inet6-group) *group-id*]**

- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family (Inet | inet6) address *address* (vrrp-group | vrrp-inet6-group) *group-id*]

To prohibit a higher-priority backup router from preempting a lower-priority master router, include the **no-preempt** statement:

```
no-preempt;
```

**Related
Documentation**

- [Understanding VRRP on page 315](#)
- [Configuring the Advertisement Interval for the VRRP Master Router on page 350](#)
- [Modifying the Preemption Hold-Time Value for the VRRP Master Router on page 355](#)
- [Configuring the Asymmetric Hold Time for VRRP Routers on page 355](#)
- [Configuring VRRP on page 337](#)

Configuring a Backup to Accept Packets Destined for the Virtual IP Address

By default, a switch configured to be a VRRP backup but acting as the master does not process packets sent to the virtual IP address—that is, packets in which the destination address is the virtual IP address. To configure a backup switch to process packets sent to the virtual IP address while it is acting as the master, include the **accept-data** statement on the backup:

```
accept-data;
```

You can include this statement at the following hierarchy level:

- [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group] *group-id*

To explicitly prohibit the backup from accepting packets destined for the virtual IP address while acting as master, include the **no-accept-data** statement:

```
no-accept-data;
```

If you include the **accept-data** statement, configure the connected hosts so that they:

- Process gratuitous ARP requests.
- Do not use packets other than ARP replies to update their ARP cache.

This statement is disabled by default. If you enable it, your configuration does not comply with RFC 3768.

To restrict incoming IP packets to ICMP only, you must configure firewall filters to accept only ICMP packets.

- Related Documentation**
- [Understanding VRRP on page 315](#)
 - [Configuring Basic VRRP Support on page 332](#)
 - [Example: Configuring VRRP for Load Sharing on page 378](#)

Modifying the Preemption Hold-Time Value for the VRRP Master Router

The hold time is the maximum number of seconds that can elapse before a higher-priority backup router preempts the master router. You might want to configure a hold time so that all Junos OS components converge before preemption.

By default, the hold-time value is 0 seconds. A value of 0 means that preemption can occur immediately after the backup router comes online. Note that the hold time is counted from the time the backup router comes online. The hold time is only valid when the VRRP router is just coming online.

To modify the preemption hold-time value, include the **hold-time** statement:

```
hold-time seconds;
```

The hold time can be from 0 through 3600 seconds.

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family (inet | inet6) address *address* (vrrp-group | vrrp-inet6-group) *group-id* preempt]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family (inet | inet6) address *address* (vrrp-group | vrrp-inet6-group) *group-id* preempt]

- Related Documentation**
- [Configuring the Advertisement Interval for the VRRP Master Router on page 350](#)
 - [Configuring a Backup Router to Preempt the VRRP Master Router on page 353](#)
 - [Configuring the Asymmetric Hold Time for VRRP Routers on page 355](#)
 - [Configuring VRRP on page 337](#)

Configuring the Asymmetric Hold Time for VRRP Routers

In Junos OS Release 9.5 and later, the **asymmetric-hold-time** statement at the [edit protocols vrrp] hierarchy level enables you to configure a VRRP master router to switch over to the backup router immediately—that is, without waiting for the priority hold time to expire—when a tracked interface or route goes down or when the bandwidth of a tracked interface decreases. Such events can cause an immediate reduction in the priority based on the configured priority cost for the event, and trigger a mastership election.

However, when the tracked route or interface comes up again, or when the bandwidth for a tracked interface increases, the backup (original master) router waits for the hold

time to expire before it updates the priority and initiates the switchover if the priority is higher than the priority for the VRRP master (original backup) router.

If the **asymmetric-hold-time** statement is not configured, the VRRP master waits for the hold time to expire before it initiates a switchover when a tracked route goes down or when the bandwidth of a tracked interface decreases.

**Example: Configuring
Asymmetric Hold Time**

```
[edit]
user@host# set protocols vrrp asymmetric-hold-time
[edit]
user@host# show protocols vrrp
asymmetric-hold-time;
```

**Related
Documentation**

- [Configuring the Advertisement Interval for the VRRP Master Router on page 350](#)
- [Configuring a Backup Router to Preempt the VRRP Master Router on page 353](#)
- [Modifying the Preemption Hold-Time Value for the VRRP Master Router on page 355](#)
- [Configuring VRRP on page 337](#)

Configuring Passive ARP Learning for Backup VRRP Routers

By default, the backup VRRP router drops ARP requests for the VRRP-IP to VRRP-MAC address translation. This means that the backup router does not learn the ARP (IP-to-MAC address) mappings for the hosts sending the requests. When it detects a failure of the master router and transitions to become the new master router, the backup router must re-learn all the entries that were present in the ARP cache of the master router. In environments with many directly attached hosts, such as metro Ethernet environments, the number of ARP entries to learn can be high. This can cause a significant transition delay, during which the traffic transmitted to some of the hosts might be dropped.

Passive ARP learning enables the ARP cache in the backup router to hold approximately the same contents as the ARP cache in the master router, thus preventing the problem of learning ARP entries in a burst. To enable passive ARP learning, include the **passive-learning** statement at the **[edit system arp]** hierarchy level:

```
[edit system arp]
passive-learning;
```

We recommend setting passive learning on both the backup and master VRRP routers. Doing so prevents the need to manually intervene when the master router becomes the backup router. While a router is operating as the master router, the passive learning configuration has no operational impact. The configuration takes effect only when the router is operating as a backup router.

For information about configuring gratuitous ARP and the ARP aging timer, see the *Junos OS Administration Library*.

Related Documentation • [Understanding VRRP on page 315](#)

Configuring VRRP Route Tracking

Configure Routers R1 and R2 to run VRRP. Configure static routes and a policy for exporting the static routes on Router R3. The VRRP routing instances on R2 track the routes that are advertised by R3.

On Router R1

```
[edit interfaces]
ge-1/0/3 {
  unit 0 {
    vlan-id 1;
    family inet {
      address 200.100.50.2/24 {
        vrrp-group 0 {
          virtual-address 200.100.50.101;
          priority 195;
        }
      }
    }
  }
}
```

On Router R2

```
[edit interfaces]
ge-1/0/1 {
  unit 0 {
    vlan-id 1;
    family inet {
      address 200.100.50.1/24 {
        vrrp-group 0 {
          virtual-address 200.100.50.101;
          priority 200;
          track {
            route 59.0.58.153/32 routing-instance default priority-cost 5;
            route 59.0.58.154/32 routing-instance default priority-cost 5;
            route 59.0.58.155/32 routing-instance default priority-cost 5;
          }
        }
      }
    }
  }
}
```

On Router R3

```
[edit]
policy-options {
  policy-statement static-policy {
    term term1 {
      then accept;
    }
  }
}
```

```
    }  
  }  
  protocols {  
    ospf {  
      export static-policy;  
      reference-bandwidth 4g;  
      area 0.0.0.0 {  
        interface all;  
        interface fxp0.0 {  
          disable;  
        }  
      }  
    }  
  }  
  }  
  routing-options {  
    static {  
      route 59.0.0.153/32 next-hop 45.45.45.46;  
      route 59.0.0.154/32 next-hop 45.45.45.46;  
      route 59.0.0.155/32 next-hop 45.45.45.46;  
    }  
  }  
}
```

**Related
Documentation**

- [Understanding VRRP on page 315](#)
- [Configuring a Route to Be Tracked for a VRRP Group on page 361](#)
- [Configuring VRRP on page 337](#)
- [Example: Configuring VRRP for IPv6 on page 342](#)

Configuring a Logical Interface to Be Tracked for a VRRP Group

VRRP can track whether a logical interface is up, down, or not present, and can also dynamically change the priority of the VRRP group based on the state of the tracked logical interface, triggering a new master router election. VRRP can also track the operational speed of a logical interface and dynamically update the priority of the VRRP group when the speed crosses a configured threshold.

When interface tracking is enabled, you cannot configure a priority of 255 (a priority of 255 designates the master router). For each VRRP group, you can track up to 10 logical interfaces.

To configure a logical interface to be tracked, include the following statements:

```
track {  
  interface interface-name {  
    bandwidth-threshold bits-per-second priority-cost priority;  
    priority-cost priority;  
  }  
  priority-hold-time seconds;  
}
```

```
interface et-0/0/0 {
  priority-cost 30;
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* *vrrp-group group-id*]
- [edit interfaces *interface-name* unit *logical-unit-number* family inet6 address *address* *vrrp-inet6-group group-id*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet address *address* *vrrp-group group-id*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet6 address *address* *vrrp-inet6-group group-id*]

The interface specified is the interface to be tracked for the VRRP group. The priority hold time is the minimum length of time that must elapse between dynamic priority changes. A tracking event, such as an interface state change (up or down) or a change in bandwidth, triggers one of the following responses:

- The first tracking event initiates the priority hold timer, and also initializes the pending priority based on the current priority and the priority cost. However, the current priority remains unchanged.
- A tracking event or a manual configuration change that occurs while the priority hold timer is on triggers a pending priority update. However, the current priority remains unchanged.

This ensures that Junos OS does not initiate mastership elections every time a tracked interface flaps.

When the priority hold time expires, the current priority inherits the value from the pending priority, and the pending priority ceases.



NOTE: If you have configured `asymmetric-hold-time`, VRRP does not wait for the priority hold time to expire before initiating mastership elections if a tracked interface fails (state changes from up to down), or if the available bandwidth for a tracked interface decreases. For more information about `asymmetric-hold-time`, see [“Configuring the Asymmetric Hold Time for VRRP Routers” on page 355](#).

There are two `priority-cost` statements that show at this hierarchy level. The `bandwidth-threshold` statement specifies a threshold for the tracked interface. When the bandwidth of the tracked interface drops below the configured bandwidth threshold value, the VRRP group uses the bandwidth threshold priority cost. You can track up to five bandwidth threshold statements for each tracked interface. Just under the `interface` statement there is a `priority-cost` statement that gives the value to subtract from priority when the interface is down.

The sum of the priority costs for all tracked logical interfaces must be less than or equal to the configured priority of the VRRP group. If you are tracking more than one interface, the router applies the sum of the priority costs for the tracked interfaces (at most, only one priority cost for each tracked interface) to the VRRP group priority.

Prior to Junos OS Release 15.1, an adjusted priority could not be zero. If the difference between the priority costs and the configured priority of the VRRP group was zero, the adjusted priority would become 1.



NOTE: In Junos OS Release 15.1 and later, an adjusted priority can be zero.

The priority value zero (0) indicates that the current master router has stopped participating in VRRP. Such a priority value is used to trigger one of the backup routers to quickly transition to the master router without having to wait for the current master to time out.

If you are tracking more than one interface, the router applies the sum of the priority costs for the tracked interfaces (at most, only one priority cost for each tracked interface) to the VRRP group priority. However, the interface priority cost and bandwidth threshold priority cost values for each VRRP group are not cumulative. The router uses only one priority cost to a tracked interface as indicated in [Table 18 on page 360](#).

Table 18: Interface State and Priority Cost Usage

Tracked Interface State	Priority Cost Usage
Down	priority-cost priority
Not down; media speed below one or more bandwidth thresholds	Priority cost of the lowest applicable bandwidth threshold

You must configure an interface priority cost only if you have configured no bandwidth thresholds. If you have not configured an interface priority cost value, and the interface is down, the interface uses the bandwidth threshold priority cost value of the lowest bandwidth threshold.

Release History Table

Release	Description
15.1	In Junos OS Release 15.1 and later, an adjusted priority can be zero.

Related Documentation

- [Understanding VRRP on page 315](#)
- [Configuring a Route to Be Tracked for a VRRP Group on page 361](#)
- [Configuring VRRP on page 337](#)

Configuring a Route to Be Tracked for a VRRP Group

VRRP can track whether a route is reachable (that is, the route exists in the routing table of the routing instance included in the configuration) and dynamically change the priority of the VRRP group based on the reachability of the tracked route, triggering a new master router election.

To configure a route to be tracked, include the following statements:

```
track {
  priority-hold-time seconds;
  route prefix/prefix-length routing-instance instance-name priority-cost priority;
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* *vrrp-group group-id*]
- [edit interfaces *interface-name* unit *logical-unit-number* family inet6 address *address* *vrrp-inet6-group group-id*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet address *address* *vrrp-group group-id*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet6 address *address* *vrrp-inet6-group group-id*]

The route prefix specified is the route to be tracked for the VRRP group. The priority hold time is the minimum length of time that must elapse between dynamic priority changes. A route tracking event, such as adding a route to or removing a route from the routing table, might trigger one or more of the following:

- The first tracking event initiates the priority hold timer, and also initializes the pending priority based on the current priority and the priority cost. However, the current priority remains unchanged.
- A tracking event or a manual configuration change that occurs while the priority hold timer is on triggers a pending priority update. However, the current priority remains unchanged.

When the priority hold time expires, the current priority inherits the value from the pending priority, and the pending priority ceases.

This ensures that Junos OS does not initiate mastership elections every time a tracked route flaps.



NOTE: If you have configured **asymmetric-hold-time**, VRRP does not wait for the priority hold time to expire before initiating mastership elections if a tracked route is removed from the routing table. For more information about **asymmetric-hold-time**, see [“Configuring the Asymmetric Hold Time for VRRP Routers” on page 355](#).

The routing instance is the routing instance in which the route is to be tracked. If the route is in the default, or global, routing instance, specify the instance name as **default**.



NOTE: Tracking a route that belongs to a routing instance from a different logical system is not supported.

The priority cost is the value to be subtracted from the configured VRRP priority when the tracked route goes down, forcing a new master router election. The value can be from 1 through 254.

The sum of the priority costs for all tracked routes must be less than or equal to the configured priority of the VRRP group. If you are tracking more than one route, the router applies the sum of the priority costs for the tracked routes (at most, only one priority cost for each tracked route) to the VRRP group priority.

Prior to Junos OS Release 15.1, an adjusted priority could not be zero. If the difference between the priority costs and the configured priority of the VRRP group was zero, the adjusted priority would become 1.



NOTE: In Junos OS Release 15.1 and later, an adjusted priority can be zero.

The priority value zero (0) indicates that the current master router has stopped participating in VRRP. Such a priority value is used to trigger one of the backup routers to quickly transition to the master router without having to wait for the current master to time out.

Release History Table

Release	Description
15.1	Prior to Junos OS Release 15.1, an adjusted priority could not be zero.

Related Documentation

- [Understanding VRRP on page 315](#)
- [Configuring a Logical Interface to Be Tracked for a VRRP Group on page 358](#)
- [Configuring VRRP Route Tracking on page 357](#)

Example: Configuring Multiple VRRP Owner Groups

These examples show how to configure multiple virtual router redundancy protocol (VRRP) IPv4 and IPv6 owner groups.

- [Requirements on page 363](#)
- [Overview on page 363](#)
- [Configuration on page 363](#)
- [Verification on page 369](#)

Requirements

This example uses the following hardware and software components:

- A EX-Series, M-Series, MX-Series, or T-Series router.
- Junos OS release 12.3 or later

Overview

Multiple VRRP owner groups allows users to reuse interface address identifiers (IFAs) as virtual IP addresses (VIPs). You can configure multiple IPv4 owner groups, multiple IPv6 owner groups, or a mix of IPv4 and IPv6 owner groups.

Configuration

CLI Quick Configuration To quickly configure this section of the example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Multiple IPv4 owner groups	<pre>edit interfaces ge-1/0/0 unit 0 family inet set address 10.0.0.2/24 vrrp-group 2 virtual-address 10.0.0.4 accept-data set address 20.0.0.2/24 vrrp-group 3 virtual-address 20.0.0.2 priority 255 set address 30.0.0.2/24 vrrp-group 4 virtual-address 30.0.0.2 priority 255</pre>
Multiple IPv6 owner groups	<pre>edit interfaces ge-1/0/0 unit 0 family inet6 set address 2001:4818:f000:20::1/64 vrrp-inet6-group 1 virtual-inet6-address 2001:4818:f000:20::1 set address 2001:4818:f000:20::1/64 vrrp-inet6-group 1 virtual-link-local-address fe80:4818:f000:20::1 set address 2001:4818:f000:20::1/64 vrrp-inet6-group 1 priority 255 set address fe80:4818:f000:13::2/64 set address 2001:1000:f000:20::1/64 vrrp-inet6-group 2 virtual-inet6-address 2001:1000:f000:20::1 set address 2001:1000:f000:20::1/64 vrrp-inet6-group 2 virtual-link-local-address fe80:1000:f000:20::1 set address 2001:1000:f000:20::1/64 vrrp-inet6-group 2 priority 255</pre>

```
set address 2001:2000:f000:20::1/64 vrrp-inet6-group 3 virtual-inet6-address
2001:2000:f000:20::2
set address 2001:2000:f000:20::1/64 vrrp-inet6-group 3 virtual-link-local-address
fe80:2000:f000:20::2
set address 2001:2000:f000:20::1/64 vrrp-inet6-group 3 priority 250
```

Multiple IPv4 and IPv6 owner groups

```
edit interfaces ge-1/0/0 unit 0
set family inet address 10.0.0.1/24 vrrp-group 5 virtual-address 10.0.0.1
set family inet address 10.0.0.1/24 vrrp-group 5 priority 255
set family inet6 address 2001:4818:f000:20::1/64 vrrp-inet6-group 1 virtual-inet6-address
2001:4818:f000:20::1
set family inet6 address 2001:4818:f000:20::1/64 vrrp-inet6-group 1
virtual-link-local-address fe80:4818:f000:20::1
set family inet6 address 2001:4818:f000:20::1/64 vrrp-inet6-group 1 priority 255
set family inet6 address 2001:1000:f000:20::1/64 vrrp-inet6-group 2 virtual-inet6-address
2001:1000:f000:20::1
set family inet6 address 2001:1000:f000:20::1/64 vrrp-inet6-group 2
virtual-link-local-address fe80:1000:f000:20::1
set family inet6 address 2001:1000:f000:20::1/64 vrrp-inet6-group 2 priority 255
set family inet6 address 2001:2000:f000:20::1/64 vrrp-inet6-group 3 virtual-inet6-address
2001:2000:f000:20::2
set family inet6 address 2001:2000:f000:20::1/64 vrrp-inet6-group 3
virtual-link-local-address fe80:2000:f000:20::2
set family inet6 address 2001:2000:f000:20::1/64 vrrp-inet6-group 3 priority 250
```

Configuring multiple IPv4 owner groups

Step-by-Step Procedure

To configure multiple IPv4 owner groups:

1. Create an IPv4 interface on the device

```
[edit]
user@host# edit interfaces ge-1/0/0 unit 0 family inet
```

2. Configure the first IPv4 owner group

```
[edit interfaces ge-1/0/0 unit 0 family inet]
user@host# set address 10.0.0.2/24 vrrp-group 2 virtual-address 10.0.0.4
accept-data
```

3. Configure the second IPv4 owner group

```
[edit interfaces ge-1/0/0 unit 0 family inet]
user@host# set address 20.0.0.2/24 vrrp-group 3 virtual-address 20.0.0.2 priority
255
```

4. Configure the third IPv4 owner group

```
[edit interfaces ge-1/0/0 unit 0 family inet]
```

```
user@host# set address 30.0.0.2/24 vrrp-group 4 virtual-address 30.0.0.2 priority
255
```

Configuring multiple IPv6 owner groups

Step-by-Step Procedure

To configure multiple IPv6 owner groups:

1. Create an IPv6 interface on the device

```
[edit]
user@host# edit interfaces ge-1/0/0 unit 0 family inet6
```

2. Configure the inet6 address for the first IPv6 owner group

```
[edit interfaces ge-1/0/0 unit 0 family inet6]
user@host# set address 2001:4818:f000:20::1/64 vrrp-inet6-group 1
virtual-inet6-address 2001:4818:f000:20::1
```

3. [edit interfaces ge-1/0/0 unit 0 family inet6]
user@host# set address 2001:4818:f000:20::1/64 vrrp-inet6-group 1
virtual-link-local-address fe80:4818:f000:20::1

4. [edit interfaces ge-1/0/0 unit 0 family inet6]
user@host# set address 2001:4818:f000:20::1/64 vrrp-inet6-group 1 priority 255

5. [edit interfaces ge-1/0/0 unit 0 family inet6]
user@host# set family inet6 address 2001:1000:f000:20::1/64 vrrp-inet6-group 2
virtual-inet6-address 2001:1000:f000:20::1

6. [edit interfaces ge-1/0/0 unit 0 family inet6]
user@host# set family inet6 address 2001:1000:f000:20::1/64 vrrp-inet6-group 2
virtual-link-local-address fe80:1000:f000:20::1

7. [edit interfaces ge-1/0/0 unit 0 family inet6]
user@host# set family inet6 address 2001:1000:f000:20::1/64 vrrp-inet6-group 2
priority 255

8. [edit interfaces ge-1/0/0 unit 0 family inet6]

```
user@host# set family inet6 address 2001:2000:f000:20::1/64 vrrp-inet6-group 3
virtual-inet6-address 2001:2000:f000:20::2
```

9. [edit interfaces ge-1/0/0 unit 0 family inet6]
user@host# set family inet6 address 2001:2000:f000:20::1/64 vrrp-inet6-group 3
virtual-link-local-address fe80:2000:f000:20::2
10. [edit interfaces ge-1/0/0 unit 0 family inet6]
user@host# set address 2001:2000:f000:20::1/64 vrrp-inet6-group 3 priority 250

Configuring multiple IPv4 and IPv6 owner groups

Step-by-Step Procedure

To configure multiple IPv4 and IPv6 owner groups:

1. Create an interface on the device

[edit]
user@host# edit interfaces ge-1/0/0 unit 0
2. Configure the family inet address and virtual address for the IPv4 owner group

[edit interfaces ge-1/0/0 unit 0]
user@host# set family inet address 10.0.0.1/24 vrrp-group 5 virtual-address 10.0.0.1
3. Set the priority of the IPv4 owner group to 255

[edit interfaces ge-1/0/0 unit 0]
set family inet address 10.0.0.1/24 vrrp-group 5 priority 255
4. Configure the inet6 address for the first IPv6 owner group

[edit interfaces ge-1/0/0 unit 0]
set family inet6 address 2001:4818:f000:20::1/64 vrrp-inet6-group 1
virtual-inet6-address 2001:4818:f000:20::1
5. Set the virtual link local address for the first IPv6 owner group

[edit interfaces ge-1/0/0 unit 0]
set family inet6 address 2001:4818:f000:20::1/64 vrrp-inet6-group 1
virtual-link-local-address fe80:4818:f000:20::1
6. Set the first IPv6 owner group's priority to 255

```
[edit interfaces ge-1/0/0 unit 0]
set family inet6 address 2001:4818:f000:20::1/64 vrrp-inet6-group 1 priority 255
```

7. Configure the inet6 address for the second IPv6 owner group

```
[edit interfaces ge-1/0/0 unit 0]
set family inet6 address 2001:1000:f000:20::1/64 vrrp-inet6-group 2
virtual-inet6-address 2001:1000:f000:20::1
```

8. Set the virtual link local address for the second IPv6 owner group

```
[edit interfaces ge-1/0/0 unit 0]
set family inet6 address 2001:1000:f000:20::1/64 vrrp-inet6-group 2
virtual-link-local-address fe80:1000:f000:20::1
```

9. Set the second IPv6 owner group's priority to 255

```
[edit interfaces ge-1/0/0 unit 0]
set family inet6 address 2001:1000:f000:20::1/64 vrrp-inet6-group 2 priority 255
```

10. Configure the inet6 address for the third IPv6 owner group

```
[edit interfaces ge-1/0/0 unit 0]
set family inet6 address 2001:2000:f000:20::1/64 vrrp-inet6-group 3
virtual-inet6-address 2001:2000:f000:20::2
```

11. Set the virtual link local address for the third IPv6 owner group

```
[edit interfaces ge-1/0/0 unit 0]
set family inet6 address 2001:2000:f000:20::1/64 vrrp-inet6-group 3
virtual-link-local-address fe80:2000:f000:20::2
```

12. Set the third IPv6 owner group's priority to 250

```
[edit interfaces ge-1/0/0 unit 0]
set family inet6 address 2001:2000:f000:20::1/64 vrrp-inet6-group 3 priority 250
```

Results

Multiple IPv4 owner groups

```
[edit interfaces]
ge-1/0/0
unit 0 {
  family inet {
    address 10.0.0.2/24 {
```

```
        vrrp-group 2 {
            virtual-address 10.0.0.4;
            accept-data;
        }
    }
    address 20.0.0.2/24 {
        vrrp-group 3 {
            virtual-address 20.0.0.2;
            priority 255;
        }
    }
    address 30.0.0.2/24 {
        vrrp-group 4 {
            virtual-address 30.0.0.2;
            priority 255;
        }
    }
}
```

Multiple IPv6 owner groups

```
[edit interfaces]
ge-1/0/0
  unit 0 {
    family inet6 {
      address 2001:4818:f000:20::1/64 {
        vrrp-inet6-group 1 {
          virtual-inet6-address 2001:4818:f000:20::1;
          virtual-link-local-address fe80:4818:f000:20::1;
          priority 255;
        }
      }
      address fe80:4818:f000:13::2/64;
      address 2001:1000:f000:20::1/64 {
        vrrp-inet6-group 2 {
          virtual-inet6-address 2001:1000:f000:20::1;
          virtual-link-local-address fe80:1000:f000:20::1;
          priority 255;
        }
      }
      address 2001:2000:f000:20::1/64 {
        vrrp-inet6-group 3 {
          virtual-inet6-address 2001:2000:f000:20::2;
          virtual-link-local-address fe80:2000:f000:20::2;
          priority 250;
        }
      }
    }
  }
}
```

Multiple IPv4 and IPv6 owner groups

```
[edit interfaces]
ge-1/0/0
  unit 0 {
```

```
family inet {
  address 10.0.0.1/24 {
    vrrp-group 5 {
      virtual-address 10.0.0.1;
      priority 255;
    }
  }
}
family inet6 {
  address 2001:4818:f000:20::1/64 {
    vrrp-inet6-group 1 {
      virtual-inet6-address 2001:4818:f000:20::1;
      virtual-link-local-address fe80:4818:f000:20::1;
      priority 255;
    }
  }
  address 2001:1000:f000:20::1/64 {
    vrrp-inet6-group 2 {
      virtual-inet6-address 2001:1000:f000:20::1;
      virtual-link-local-address fe80:1000:f000:20::1;
      priority 255;
    }
  }
  address 2001:2000:f000:20::1/64 {
    vrrp-inet6-group 3 {
      virtual-inet6-address 2001:2000:f000:20::2;
      virtual-link-local-address fe80:2000:f000:20::2;
      priority 250;
    }
  }
}
```

Verification

To verify the configuration, run the **show interfaces ge-1/0/0** command, or use whichever name you assigned to the interface.

- Related Documentation**
- [Tracing VRRP Operations on page 377](#)
 - [Configuring Inheritance for a VRRP Group on page 369](#)

Configuring Inheritance for a VRRP Group

Junos OS enables you to configure VRRP groups on the various subnets of a VLAN to inherit the state and configuration of one of the groups, which is known as the *active VRRP group*. When the **vrrp-inherit-from** configuration statement is included in the configuration, only the active VRRP group, from which the other VRRP groups are inheriting the state, sends out frequent VRRP advertisements, and processes incoming VRRP advertisements. The groups that are inheriting the state do not process any incoming VRRP advertisement because the state is always inherited from the active VRRP group. However, the groups that are inheriting the state do send out VRRP advertisements once

every 2 to 3 minutes to facilitate MAC address learning on the switches placed between the VRRP routers.

If the **vrrp-inherit-from** statement is not configured, each of the VRRP master groups in the various subnets on the VLAN sends out separate VRRP advertisements and adds to the traffic on the VLAN.

To configure inheritance for a VRRP group, include the **vrrp-inherit-from** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* *vrrp-group group-id*]** hierarchy level.

```
[edit interfaces interface-name unit logical-unit-number family inet address address
  vrrp-group group-id]
vrrp-inherit-from vrrp-group;
```

When you configure a group to inherit a state from another group, the inheriting groups and the active group must be on the same physical interface and logical system. However, the groups do not need to necessarily be on the same routing instance (as was in Junos OS releases earlier than 9.6), VLAN, or logical interface.

When you include the **vrrp-inherit-from** statement for a VRRP group, the VRRP group inherits the following parameters from the active group:

- **advertise-interval**
- **authentication-key**
- **authentication-type**
- **fast-interval**
- **preempt | no-preempt**
- **priority**
- **track interfaces**
- **track route**

However, you can configure the **accept-data | no-accept-data** statement for the group to specify whether the interface should accept packets destined for the virtual IP address.

Related Documentation

- [Understanding VRRP on page 315](#)

Configuring an Interface to Accept All Packets Destined for the Virtual IP Address of a VRRP Group

In VRRP implementations where the router acting as the master router is not the IP address owner—the IP address owner is the router that has the interface whose actual IP address is used as the virtual router's IP address (virtual IP address)—the master router accepts only the ARP packets from the packets that are sent to the virtual IP address. Junos OS enables you to override this limitation with the help of the **accept-data** configuration. When the **accept-data** statement is included in the configuration, the master router accepts all packets sent to the virtual IP address even when the master router is not the IP address owner.



NOTE: If the master router is the IP address owner or has its priority set to 255, the master router, by default, accepts all packets addressed to the virtual IP address. In such cases, the **accept-data** configuration is not required.

To configure an interface to accept all packets sent to the virtual IP address, include the **accept-data** statement:

```
accept-data;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family (inet | inet6) address *address* (vrrp-group | vrrp-inet6-group) *group-id*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family (inet | inet6) address *address* (vrrp-group | vrrp-inet6-group) *group-id*]

To prevent a master router that is the IP address owner or has its priority set to 255 from accepting packets other than the ARP packets addressed to the virtual IP address, include the **no-accept-data** statement:

```
no-accept-data;
```



NOTE:

- If you want to restrict the incoming IP packets to ICMP packets only, you must configure firewall filters to accept only ICMP packets.
- If you include the **accept-data** statement, your routing platform configuration does not comply with RFC 3768 (see section 6.4.3 of RFC 3768, *Virtual Router Redundancy Protocol (VRRP)*).

- Related Documentation**
- [Understanding VRRP on page 315](#)
 - [Configuring VRRP on page 337](#)

Configuring the Silent Period to Avoid Alarms Due to Delay in Receiving VRRP Advertisement Packets

The silent period starts when the interface state is changed from down to up. During this period, the Master Down Event is ignored. Configure the silent period interval to avoid alarms caused by the delay or interruption of the incoming VRRP advertisement packets during the interface startup phase.

To configure the silent period interval that the Master Down Event timer ignores, include the **startup-silent-period** statement at the **[edit protocols vrrp]** hierarchy level:

```
[edit protocols vrrp]
  startup-silent-period seconds;
```



NOTE: During the silent startup period, the **show vrrp detail** command output shows a value of 0 for Master priority and your IP address for Master router. These values indicate that the Master selection is not completed yet, and these values can be ignored.

When you have configured **startup-silent-period**, the Master Down Event is ignored until the **startup-silent-period** expires.

For example, configure a VRRP group, *vrrp-group1*, with an advertise interval of 1 second, startup silent period of 10 seconds, and an interface *interface1* with a priority less than 255.

When *interface1* transitions from down to up:

- The *vrrp-group1* group moves to the backup state, and starts the Master Down Event timer (3 seconds; three times the value of the advertise interval, which is 1 second in this case).
- If no VRRP PDU is received during the 3-second period, the **startup-silent-period** (10 seconds in this case) is checked, and if the startup silent period has not expired, the Master Down Event timer is restarted. This is repeated until the **startup-silent-period** expires. In this example, the Master Down Event timer runs four times (12 seconds) by the time the 10-second startup silent period expires.
- If no VRRP PDU is received by the end of the fourth 3-second cycle, *vrrp-group1* takes over mastership.

Related Documentation

- [Understanding VRRP on page 315](#)
- [startup-silent-period on page 623](#)

Enabling the Distributed Periodic Packet Management Process for VRRP

Typically, VRRP advertisements are sent by the VRRP process (vrrpd) on the master VRRP router at regular intervals to let other members of the group know that the VRRP master router is operational.

When the vrrpd process is busy and does not send VRRP advertisements, the backup VRRP routers might assume that the master router is down and take over as the master router, causing unnecessary flaps. This takeover might occur even though the original master router is still active and available and might resume sending advertisements after the traffic has decreased. To address this problem and to reduce the load on the vrrpd process, Junos OS uses the periodic packet management process (ppmd) to send VRRP advertisements on behalf of the vrrpd process. However, you can further delegate the job of sending VRRP advertisements to the distributed ppm process that resides on the Packet Forwarding Engine.

The ability to delegate the sending of VRRP advertisements to the distributed ppm process ensures that the VRRP advertisements are sent even when the ppm process—which is now responsible for sending VRRP advertisements—is busy. Such delegation prevents the possibility of false alarms when the ppm process is busy. The ability to delegate the sending of VRRP advertisements to distributed ppm also adds to scalability because the load is shared across multiple ppm instances and is not concentrated on any single unit.



NOTE: CPU-intensive VRRP advertisements, such as advertisements with MD5 authentication, continue to be processed by the VRRP process on the Routing Engine even when distributed ppm is enabled.



NOTE: VRRP is supported by graceful Routing Engine switchover only in the case that PPM delegation is enabled (the default).



NOTE: Aggregated Ethernet and integrated routing and bridging (IRB) delegation is supported only for MPC line cards. Routing devices with inbuilt MPCs such as the MX104 and below do not support this feature.

To configure the distributed ppm process to send VRRP advertisements, include the **delegate-processing** statement at the **[edit protocols vrrp]** hierarchy level:

```
[edit protocols vrrp]
  delegate-processing;
```

To configure the distributed `ppmd` process to send VRRP advertisements over aggregated Ethernet and IRB interfaces, include the `delegate-processing ae-irb` statement at the `[edit protocols vrrp]` hierarchy level:

```
[edit protocols vrrp]
  delegate-processing ae-irb;
```

- Related Documentation
- [Understanding VRRP on page 315](#)
 - [delegate-processing \(VRRP\) on page 606](#)

Improving the Convergence Time for VRRP

You can enable faster convergence time for the configured Virtual Router Redundancy Protocol (VRRP), thereby reducing the traffic restoration time to less than 1 second. To improve the convergence time for the VRRP, perform the following tasks:

- **Configure the distributed periodic packet management process**—When the VRRP process is busy and does not send VRRP advertisements, the backup VRRP routers might assume that the master router is down and take over as the master router, causing unnecessary flaps. To address this problem and to reduce the load on the VRRP process, Junos OS uses the distributed periodic packet management (PPM) process to send VRRP advertisements on behalf of the VRRP process.

To configure the distributed PPM process, include the `delegate-processing` statement at the `[edit protocols vrrp]` hierarchy level.

- **Disable the skew timer**—The skew timer in VRRP is used to ensure that two backup routers do not switch to the master state at the same time in case of a failover situation. When there is only one master router and one backup router in the network deployment, you can disable the skew timer, thereby reducing the time required to transition to the master state.

To disable the skew timer, include the `skew-timer-disable` statement at the `[edit protocols vrrp]` hierarchy level.

- **Configure the number of fast advertisements that can be missed by a backup router before it starts transitioning to the master state**—The backup router waits until a certain number of advertisement packets are lost after which it transitions to the master state. This waiting time can be fatal in scenarios such as router failure or link failure. To avoid such a situation and to enable faster convergence time, in Junos OS Release 12.2 and later, you can configure a fast advertisement interval value that specifies the number of fast advertisements that can be missed by a backup router before it starts transitioning to the master state.

To configure the fast advertisement interval, include the `global-advertisements-threshold` statement at the `[edit protocols vrrp]` hierarchy level.

- **Configure inheritance of VRRP groups**—Junos OS enables you to configure VRRP groups on the various subnets of a virtual LAN (VLAN) to inherit the state and configuration of one of the groups, which is known as the active VRRP group. When the `vrrp-inherit-from` statement is included in the configuration, only the active VRRP

group, from which the other VRRP groups inherit the state, sends out frequent VRRP advertisements and processes incoming VRRP advertisements. Use inherit groups for scaled configurations. For example, if you have 1000 VRRP groups with an advertisement interval of 100 ms, then use inherit groups.

To configure inheritance for a VRRP group, include the **vrrp-inherit-from** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group *group-id*]** hierarchy level.

- **Disable duplicate address detection for IPv6 interfaces**—Starting with Junos OS Release 15.1, duplicate address detection is a feature of the Neighbor Discovery Protocol for IPv6. Duplicate address detection is enabled by default and determines whether an address is already in use by another node. When detection address detection is enabled, convergence time is high after an IPv6 interface that has been configured for VRRP tracking comes up. To disable duplicate address detection, include the **ipv6-duplicate-addr-transmits 0** statement at the **[edit system internet-options]** hierarchy level. To disable duplicate address detection only for a specific interface, include the **dad-disable** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet6]** hierarchy level.



NOTE:

- Inheritance of VRRP groups is supported with all types of interfaces. Other measures to reduce convergence time, such as VRRP distribution, disabling skew timer, and reducing advertisement threshold.
- Compared to other routers, the convergence time and the traffic restoration time are less for MX Series routers with MPCs.
- Reduction in convergence time is applicable for all types of configurations at the physical interface but the convergence time might not be less than 1 second for all the configurations. The convergence time depends on the number of groups that are transitioning from the backup to the master state and the interval at which these groups are transitioning.

Related Documentation

- [Configuring Inheritance for a VRRP Group on page 369](#)
- [Configuring VRRP to Improve Convergence Time on page 376](#)
- [delegate-processing on page 606](#)
- [global-advertisements-threshold on page 609](#)
- [skew-timer-disable on page 622](#)

Configuring VRRP to Improve Convergence Time

You can enable faster convergence time for the configured Virtual Router Redundancy Protocol (VRRP), thereby reducing the traffic restoration time to less than 1 second. To improve the convergence time for VRRP, perform the following tasks.

Before you begin, configure VRRP. See [“Configuring VRRP” on page 337](#).

1. Configure the distributed periodic packet management (PPM) process to send VRRP advertisements when the VRRP process is busy.

```
[edit]
user@host# set protocols vrrp delegate-processing
```

2. Disable the skew timer to reduce the time required to transition to the master state.

```
[edit]
user@host# set protocols vrrp skew-timer-disable
```



NOTE: When there is only one master router and one backup router in the network deployment, you can disable the skew timer, thereby reducing the time required to transition to the master state.

3. Configure the number of fast advertisements that can be missed by a backup router before it starts transitioning to the master state.

```
[edit]
user@host# set protocols vrrp global-advertisement-threshold advertisement-value
```

4. Configure VRRP groups on the various subnets of a VLAN to inherit the state and to configure one of the groups.

```
[edit]
user@host# set interfaces interface-name unit logical-unit-number family inet address
address vrrp-group group-id
```

5. Verify the configuration.

```
[edit]
user@host# show protocols vrrp
```

**NOTE:**

- Inheritance of VRRP groups is supported with all types of interfaces. Other measures to reduce convergence time, such as VRRP distribution, disabling skew timer, and reducing advertisement threshold, are not applicable when VRRP is configured over integrated routing and bridging (IRB) interfaces, aggregated Ethernet interfaces, and multichassis link aggregation group (MC-LAG) interfaces.
- Compared to other routers, the convergence time and the traffic restoration time are less for MX Series routers with MPCs.
- Reduction in convergence time is applicable for all types of configurations at the physical interface, but the convergence time might not be less than 1 second for all the configurations. The convergence time depends on the number of groups that are transitioning from the backup to the master state and the interval at which these groups are transitioning.

Related Documentation

- [Improving the Convergence Time for VRRP on page 374](#)
- [Configuring Inheritance for a VRRP Group on page 369](#)
- [delegate-processing on page 606](#)
- [global-advertisements-threshold on page 609](#)
- [skew-timer-disable on page 622](#)

Tracing VRRP Operations

To trace VRRP operations, include the **traceoptions** statement at the **[edit protocols vrrp]** hierarchy level.

By default, VRRP logs the error, data carrier detect (DCD) configuration, and routing socket events in a file in the **/var/log** directory. By default, this file is named **/var/log/vrrpd**. The default file size is 1 megabyte (MB), and three files are created before the first one gets overwritten.

To change the configuration of the logging file, include the **traceoptions** statement at the **[edit protocols vrrp]** hierarchy level:

```
[edit protocols vrrp]
traceoptions {
  file filename <files number> <match regular-expression> <microsecond-stamp>
    <size size> <world-readable | no-world-readable>;
  flag flag;
  no-remote-trace;
}
flag flag;
```

You can specify the following VRRP tracing flags:

- **all**—Trace all VRRP operations.
- **database**—Trace all database changes.
- **general**—Trace all general events.
- **interfaces**—Trace all interface changes.
- **normal**—Trace all normal events.
- **packets**—Trace all packets sent and received.
- **state**—Trace all state transitions.
- **timer**—Trace all timer events.

Related Documentation

- [Understanding VRRP on page 315](#)

Example: Configuring VRRP for Load Sharing

If you do not want to dedicate a switch to be a VRRP backup (and therefore leave it idle unless the master fails), you can create a load-sharing configuration in which each participating switch simultaneously acts as a master and a backup.

One reason to use a load-sharing (active-active) configuration is that you are more likely to actively monitor and maintain both switches and notice if a problem occurs on either of them. If you use a configuration in which one switch is only a backup (an active-backup configuration), you might be less likely to pay attention to the backup switch while it is idle. In the worst case, this could lead to the backup switch developing an undetected problem and not being able to perform adequately when a failover occurs.

- [Requirements on page 378](#)
- [Overview and Topology on page 378](#)
- [Configuring VRRP on Both Switches on page 380](#)
- [Verification on page 382](#)

Requirements

This example uses the following hardware and software components:

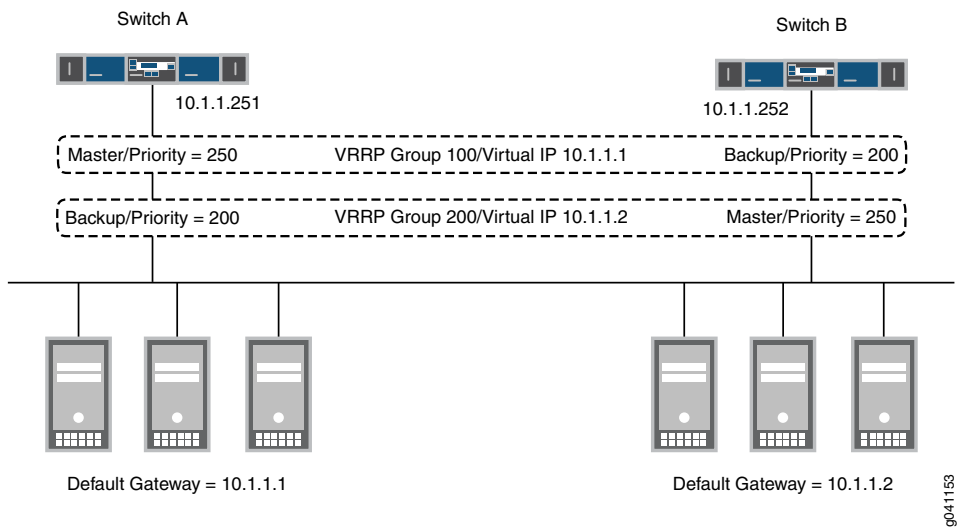
- Two switches
- Junos OS Release 11.3 or later
- Static routing or a dynamic routing protocol enabled on both switches.

Overview and Topology

This example uses two VRRP groups, each of which has its own virtual IP address. Devices on the LAN use one of these virtual IP addresses as their default gateway. If one of the

switches fails, the other switch takes over for it. In the topology shown in [Figure 25 on page 379](#), for example, Switch A is the master for VRRP group 100. If Switch A fails, Switch B takes over and forwards traffic that the end devices send to the default gateway address 10.1.1.1.

Figure 25: VRRP Load-Sharing Configuration



This example shows a simple configuration to illustrate the basic steps for configuring two switches running VRRP to back each other up. [Table 19 on page 379](#) lists VRRP settings for each switch.

Table 19: Settings for VRRP Load-Sharing Example

Switch A	Switch B
VRRP Group 100: <ul style="list-style-type: none">Interface address: 10.1.1.251VIP: 10.1.1.1Priority: 250	VRRP Group 100: <ul style="list-style-type: none">Interface address: 10.1.1.252VIP: 10.1.1.1Priority: 200
VRRP Group 200: <ul style="list-style-type: none">Interface address: 10.1.1.251VIP: 10.1.1.2Priority: 200	VRRP Group 200: <ul style="list-style-type: none">Interface address: 10.1.1.252VIP: 10.1.1.2Priority: 250

In addition to configuring the two switches as shown, you must configure your end devices so that some of them use one of the virtual IP addresses as their default gateway and the remaining end devices use the other virtual IP address as their default gateway.

Note that if a failover occurs, the remaining switch might be unable to handle all of the traffic, depending on the demand.

Configuring VRRP on Both Switches

CLI Quick Configuration

Enter the following on Switch A:

```
[edit]
set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.251/24 vrrp-group 100 virtual-address 10.1.1.1
set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.251/24 vrrp-group 100 priority 250
set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.251/24 vrrp-group 200 virtual-address 10.1.1.2
set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.251/24 vrrp-group 200 priority 200
```

Enter the following on Switch B:

```
[edit]
set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.252/24 vrrp-group 100 virtual-address 10.1.1.1
set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.252/24 vrrp-group 100 priority 200
set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.252/24 vrrp-group 200 virtual-address 10.1.1.2
set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.252/24 vrrp-group 200 priority 250
```

Step-by-Step Procedure

Configure the VRRP groups and priorities on Switch A:

1. Create VRRP group 100 on Switch A and configure the virtual IP address for the group:

```
[edit]
user@switch# set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.252/24 vrrp-group 100 virtual-address 10.1.1.1
```

2. Assign the VRRP priority for this interface in this group:

```
[edit]
user@switch# set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.251/24 vrrp-group 100 priority 250
```

3. Create VRRP group 200 on Switch A and configure the virtual IP address for the group:

```
[edit]
user@switch# set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.251/24 vrrp-group 200 virtual-address 10.1.1.2
```

4. Assign the VRRP priority for this interface in this group:

```
[edit]
user@switch# set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.251/24 vrrp-group 100 priority 200
```

Step-by-Step Procedure Configure the VRRP groups and priorities on Switch B:

1. Create VRRP group 100 on Switch B and configure the virtual IP address for the group:


```
[edit]
user@switch# set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.252/24 vrrp-group 100 virtual-address 10.1.1.1
```
2. Assign the VRRP priority for this interface in this group:


```
[edit]
user@switch# set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.252/24 vrrp-group 100 priority 200
```

Switch A remains the master for group 100 because it has the highest priority for this group.
3. Create VRRP group 200 on Switch A and configure the virtual IP address for the group:


```
[edit]
user@switch# set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.252/24 vrrp-group 200 virtual-address 10.1.1.2
```
4. Assign the VRRP priority for this interface in this group:


```
[edit]
user@switch# set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.251/24 vrrp-group 100 priority 250
```

Switch B becomes the master for group 200 because it has the highest priority for this group.

Results Display the results of the configuration on Switch A:

```
user@switch> show configuration
interfaces {
  xe-0/0/0 {
    unit 0 {
      family inet {
        address 10.1.1.251 {
          vrrp-group 100 {
            virtual address 10.1.1.1
            priority 250
          }
          vrrp-group 200 {
            virtual address 10.1.1.2
            priority 200
          }
        }
      }
    }
  }
}
```

Display the results of the configuration on Switch B:

```
user@switch> show configuration
interfaces {
  xe-0/0/0 {
    unit 0 {
      family inet {
        address 10.1.1.252 {
          vrrp-group 100 {
            virtual address 10.1.1.1
            priority 200
          }
          vrrp-group 200 {
            virtual address 10.1.1.2
            priority 250
          }
        }
      }
    }
  }
}
```

Verification

- [Verifying that VRRP Is Working on Switch A on page 382](#)
- [Verifying that VRRP Is Working on Switch B on page 383](#)

Verifying that VRRP Is Working on Switch A

Purpose Verify that VRRP is active on Switch A and that the master and backup roles are correct.

Action Use the following command to verify that VRRP is active on Switch A and that the switch is master for group 100 and backup for group 200.

```
user@switch> show vrrp
```

Interface Address	State	Group	VR state	Timer	Type
xe-0/0/0.0	up	100	master	A .0327 1c1 10.1.1.251 vip 10.1.1.1	
xe-0/0/0.0	up	200	backup	A .0327 1c1 10.1.1.251 vip 10.1.1.2	

Meaning The `show vrrp` command displays fundamental information about the VRRP configuration. This output shows that both VRRP groups are active and that this switch has assumed the correct master and backup roles. The `lcl` address is the physical address of the interface and the `vip` address is the virtual address shared by both switches. The **Timer** value (**A .0327**) indicates the remaining time (in seconds) in which this switch expects to receive a VRRP advertisement from the other switch. If an advertisement for group

200 does not arrive before the timer expires, Switch A asserts itself as the master for this group.

Verifying that VRRP Is Working on Switch B

Purpose Verify that VRRP is active on Switch B and that the master and backup roles are correct.

Action Use the following command to verify that VRRP is active on Switch B and that the switch is backup for group 100 and master for group 200.

```
user@switch> show vrrp
```

Interface Address	State	Group	VR state	Timer	Type
xe-0/0/0.0	up	100	backup	A .0327 1c1 10.1.1.252 vip 10.1.1.1	
xe-0/0/0.0	up	200	master	A .0327 1c1 10.1.1.252 vip 10.1.1.2	

Meaning The `show vrrp` command displays fundamental information about the VRRP configuration. This output shows that both VRRP groups are active and that this switch has assumed the correct master and backup roles. The `lcl` address is the physical address of the interface and the `vip` address is the virtual address shared by both switches. The **Timer** value (**A .0327**) indicates the remaining time (in seconds) in which this switch expects to receive a VRRP advertisement from the other switch. If an advertisement for group 100 does not arrive before the timer expires, Switch B asserts itself as the master for this group.

Related Documentation

- [Understanding VRRP on page 315](#)
- [Configuring Basic VRRP Support for QFX](#)

Troubleshooting VRRP

Problem **Description:** If you configure multiple VRRP groups on an interface (using multiple VLANs), traffic for some of the groups might be briefly dropped if a failover occurs. This can happen because the new master must send gratuitous ARP replies for each VRRP group to update the ARP tables in the connected devices, and there is a short delay between each gratuitous ARP reply. Traffic sent by devices that have not yet received the gratuitous ARP reply is dropped (until the device receives the reply and learns the MAC address of the new master).

Solution Configure a failover delay so that the new master delays sending gratuitous ARP replies for the period that you set. This allows the new master to send the ARP replies for all of the VRRP groups simultaneously.

Related Documentation • [failover-delay on page 607](#)

PART 11

Performing Unified In-Service Software Upgrade (ISSU)

- [Getting Started with Unified ISSU and Understanding How Unified ISSU Works on page 387](#)
- [Unified ISSU System Requirements on page 401](#)
- [Performing a Unified ISSU on page 423](#)
- [Performing an ISSR on page 465](#)

CHAPTER 27

Getting Started with Unified ISSU and Understanding How Unified ISSU Works

- [Getting Started with Unified In-Service Software Upgrade on page 387](#)
- [Understanding the Unified ISSU Process on page 388](#)
- [Understanding In-Service Software Upgrade \(ISSU\) on page 399](#)
- [Understanding In-Service Software Upgrade \(ISSU\) in ACX5000 Series Routers on page 400](#)

Getting Started with Unified In-Service Software Upgrade

The unified in-service software upgrade (ISSU) feature enables you to upgrade between two different Junos OS releases with no disruption on the control plane and with minimal disruption of traffic.

To quickly access the information you need, click on the link in [Table 20 on page 387](#).

Table 20: Locating the Information You Need to Work With ISSU

Task You Need to Perform	Where The Information Is Located
Verify unified ISSU support for your device	"Unified ISSU System Requirements" on page 401
Perform a unified ISSU	"Example: Performing a Unified ISSU" on page 424
Verify that the unified ISSU is successful	"Verifying a Unified ISSU" on page 461
Understand how the unified ISSU process works	"Understanding the Unified ISSU Process" on page 388

Unified ISSU takes advantage of the redundancy provided by dual Routing Engines and works in conjunction with the graceful Routing Engine switchover feature and the nonstop active routing feature.

Unified ISSU provides the following benefits:

- Eliminates network downtime during software image upgrades
- Reduces operating costs, while delivering higher service levels

- Allows fast implementation of new features

**Related
Documentation**

- [Understanding High Availability Features on Juniper Networks Routers on page 3](#)

Understanding the Unified ISSU Process

This topic explains the unified ISSU processes that take place on a router, on a TX Matrix router, on a TX Matrix Plus router and its connected line-card chassis (LCCs), as well as on a TX Matrix Plus router with 3D SIBs and its connected LCCs.

- [Understanding the Unified ISSU Process on a Router on page 388](#)
- [Understanding the Unified ISSU Process on the TX Matrix Router on page 393](#)
- [Understanding the Unified ISSU Process on the TX Matrix Plus Router and on the TX Matrix Plus Router with 3D SIBs on page 396](#)

Understanding the Unified ISSU Process on a Router

This topic describes the processes that take place on a router with dual Routing Engines when you initiate a unified in-service software upgrade (ISSU).

Unified ISSU Process on a Router

After you use the `request system software in-service-upgrade` command, the following process occurs.

In Figure 26 on page 390 through Figure 31 on page 393 that follow:

- A solid line indicates the high-speed internal link between a Routing Engine and a Packet Forwarding Engine.
- A dotted line indicates the messages exchanged between the Packet Forwarding Engine and the chassis process (chassisd) on the Routing Engine.
- RE0m and RE1b indicate master and backup Routing Engines, respectively.
- The check mark indicates that the device is running the new version of software.



NOTE: Unified ISSU can only upgrade up to three major releases ahead of the current release on a device. To upgrade to a release more than three releases ahead of the current release on a device, use the unified ISSU process to upgrade the device to one or more intermediate releases until the device is within three major releases of the target release.



NOTE: The following process pertains to all supported routing platforms except the TX Matrix router and TX Matrix Plus router. On most routers, the Packet Forwarding Engine resides on a Flexible PIC Concentrator (FPC). However, on an M120 router, the Forwarding Engine Board (FEB) replaces the functions of a Packet Forwarding Engine. In the illustrations and steps, when considering an M120 router, you can regard the Packet Forwarding Engine as an FPC. As an additional step on an M120 router, after the FPCs and PICs have been upgraded, the FEBs are upgraded.


1. The master Routing Engine validates the router configuration to ensure that it can be committed when you use the new software version.

Checks are made for the following:


- Disk space is available for the `/var` file system on both Routing Engines.
- The configuration is supported by a unified ISSU.
- The PICs are supported by a unified ISSU.
- Graceful Routing Engine switchover is enabled.
- Nonstop active routing is enabled.

These checks are the same as the checks made when you enter the **request system software validate in-service-upgrade** command. If there is insufficient disk space available on either of the Routing Engines, the unified ISSU process fails and returns an error message. However, unsupported PICs do not prevent a unified ISSU. If there are unsupported PICs, the system issues a warning to indicate that these PICs will restart during the upgrade. Similarly, if there is an unsupported protocol configured, the system issues a warning that packet loss might occur for the unsupported protocol during the upgrade.


Figure 26: Device Status Before Starting a Unified ISSU



2. After the validation succeeds, the management process installs (copies) the new software image to the backup Routing Engine.
3. The backup Routing Engine is rebooted.
4. After the backup Routing Engine is rebooted and is running the new software, the kernel state synchronization process (ksyncd) synchronizes (copies) the configuration file and the kernel state from the master Routing Engine.

Figure 27: Device Status After the Backup Routing Engine Is Upgraded

5. After the configuration file and the kernel state are synchronized to the backup Routing Engine, the chassis process (chassisd) on the master Routing Engine prepares other software processes for the unified ISSU. The chassis process informs the various software processes (such as rpd, apsd, bfdd, and so on) about the unified ISSU and waits for responses from them. When all the processes are ready, the chassis process sends an ISSU_PREPARE message to the FPCs installed in the router. You can display the unified ISSU process messages by using the **show log messages** command.
6. The Packet Forwarding Engine on each FPC saves its state and downloads the new software image from the backup Routing Engine. Next, each Packet Forwarding Engine sends an ISSU_READY message to the chassis process.

Figure 28: Device Status After One Packet Forwarding Engine Downloads the New Software


7. After receiving an ISSU_READY message from a Packet Forwarding Engine, the chassis process sends an ISSU_REBOOT message to the FPC on which the Packet Forwarding Engine resides. The FPC reboots with the new software image. After the FPC is rebooted, the Packet Forwarding Engine restores the FPC state, and a high-speed internal link is established with the backup Routing Engine running the new software. The chassis process link is also reestablished with the master Routing Engine.



NOTE: The Packet Forwarding Engine reboots that occur during an unified ISSU are designed to have a very short window of down time.

8. After all Packet Forwarding Engines have sent a READY message using the chassis process on the master Routing Engine, other software processes are prepared for a Routing Engine switchover. The system is ready for a switchover at this point.


Figure 29: Device Status Before the Routing Engine Switchover




NOTE: For M120 routers, the FEBs are upgraded at this point. When all FEBs have been upgraded, the system is ready for a switchover.

9. The Routing Engine switchover occurs, and the Routing Engine (re1) that was the backup now becomes the master Routing Engine.

Figure 30: Device Status After the Routing Engine Switchover



10. The new backup Routing Engine is now upgraded to the new software image. (This step is skipped if you have specified the **no-old-master-upgrade** option in the **request system software in-service-upgrade** command.)

Figure 31: Device Status After the Unified ISSU Is Complete

11. When the backup Routing Engine has been successfully upgraded, the unified ISSU is complete.

Understanding the Unified ISSU Process on the TX Matrix Router

This topic describes the processes that take place on a TX Matrix router when you initiate a unified in-service software upgrade (ISSU).

- [Unified ISSU Process on the TX Matrix Router on page 394](#)

Unified ISSU Process on the TX Matrix Router

This section describes the processes that take place on a TX Matrix router and the routers acting as connected line-card chassis (LCCs).



NOTE: A routing matrix is a multichassis architecture that consists of a TX Matrix router and from one to four T640 routers. From the perspective of the user interface, the routing matrix appears as a single router. The TX Matrix router controls all the T640 routers in the routing matrix.

Each router has dual Routing Engines.

After you use the **request system software in-service-upgrade** command on a TX Matrix router, the following process occurs:

1. The management process (mgd) on the master Routing Engine of the TX Matrix router (global master) checks the current configuration.
Checks are made for the following:
 - Disk space is available for the **/var** file system on all Routing Engines.
 - The configuration is supported by a unified ISSU.
 - The PICs are supported by a unified ISSU.
 - Graceful Routing Engine switchover is enabled.
 - Nonstop active routing is enabled.
2. After successful validation of the configuration, the management process copies the new image to the backup Routing Engines on the TX Matrix router and the T640 routers.
3. The kernel synchronization process (ksyncd) on the backup Routing Engines synchronizes the kernels on the backup Routing Engines with the kernels on the master Routing Engines.
4. The global backup Routing Engine is upgraded with the new software. Next the global backup Routing Engine is rebooted. Then the global backup Routing Engine synchronizes the configuration and kernel state from the global master Routing Engine.
5. The LCC backup Routing Engines are upgraded and rebooted. Then the LCC backup Routing Engines connect with the upgraded global backup Routing Engine and synchronize the configuration and kernel state.
6. The unified ISSU control moves from the management process to the chassis process (chassisd). The chassis process informs the various software processes (such as rpd, apsd, bfdd, and so on) about the unified ISSU and waits for responses from them.

7. After receiving messages from the software processes indicating that the processes are ready for unified ISSU, the chassis process on the global master Routing Engine sends messages to the chassis process on the routing nodes to start the unified ISSU.
8. The chassis process on the routing nodes sends ISSU_PREPARE messages to the field-replaceable units (FRUs), such as FPCs and intelligent PICs.
9. After receiving an ISSU_PREPARE message, the Packet Forwarding Engines save the current state information and download the new software image from the backup Routing Engines. Next, each Packet Forwarding Engine sends ISSU_READY messages to the chassis process. You can display the unified ISSU process messages by using the **show log messages** command.
10. After receiving an ISSU_READY message from the Packet Forwarding Engines, the chassis process sends an ISSU_REBOOT message to the FRUs. While the upgrade is in progress, the FRUs keep sending ISSU_IN_PROGRESS messages to the chassis process on the routing nodes. The chassis process on each routing node, in turn, sends an ISSU_IN_PROGRESS message to the chassis process on the global master Routing Engine.



NOTE: The Packet Forwarding Engine reboots that occur during a unified ISSU are designed to have a very short window of down time.

11. After the unified ISSU reboot, the Packet Forwarding Engines restore the saved state information and connect back to the routing nodes. The chassis process on each routing node sends an ISSU_READY message to the chassis process on the global master Routing Engine. The CM_MSG_READY message from the chassis process on the routing nodes indicate that the unified ISSU is complete on the FRUs.
12. The unified ISSU control moves back to the management process on the global master Routing Engine.
13. The management process initiates Routing Engine switchover on the master Routing Engines.
14. Routing Engine switchover occurs on the TX Matrix router and the T640 routers.
15. After the switchover, the FRUs connect to the new master Routing Engines. Then the chassis manager and Packet Forwarding Engine manager on the T640 router FRUs connect to the new master Routing Engines on the T640 routers.
16. The management process on the global master Routing Engine initiates the upgrade process on the old master Routing Engines on the T640 routers. (This step is skipped

if you have specified the **no-old-master-upgrade** option in the **request system software in-service-upgrade** command.)

17. After the Routing Engines that were previously the masters on the T640 routers are upgraded, the management process initiates the upgrade of the Routing Engine that was previously the global master on the TX Matrix router.
18. After a successful unified ISSU, the TX Matrix router and the T640 routers are rebooted if you specified the **reboot** option in the **request system software in-service-upgrade** command.

Understanding the Unified ISSU Process on the TX Matrix Plus Router and on the TX Matrix Plus Router with 3D SIBs

This topic describes the processes that take place on a TX Matrix Plus router and the routers acting as connected line-card chassis (LCCs) as well as on a TX Matrix Plus router with 3D SIBs and its connected routers acting as LCCs.



NOTE: A routing matrix is a multichassis architecture. In this topic, the term TX Matrix Plus router denotes a routing matrix based on a Juniper Networks TX Matrix Plus router and its connected T1600 LCCs. The term TX Matrix Plus router with 3D SIBs denotes a routing matrix based on a Juniper Networks TX Matrix Plus router and its connected T1600 and T4000 LCCs.

Each router has dual Routing Engines.

-
- [Unified ISSU Process on the TX Matrix Plus Router and on the TX Matrix Plus Router with 3D SIBs on page 397](#)

Unified ISSU Process on the TX Matrix Plus Router and on the TX Matrix Plus Router with 3D SIBs

After you use the `request system software in-service-upgrade` command, the following process occurs:

1. The management process (mgd) on the master Routing Engine of the TX Matrix Plus router (global master) checks the current configuration.

Checks are made for the following:

- Disk space is available for the `/var` file system on both Routing Engines.
 - The configuration is supported by a unified ISSU.
 - The PICs are supported by a unified ISSU.
 - Graceful Routing Engine switchover is enabled.
 - Nonstop active routing is enabled.
2. After successful validation of the configuration, the management process copies the new image to the backup Routing Engines on the TX Matrix Plus router and the connected T1600 router LCCs.
 3. The kernel synchronization process (ksyncd) on the backup Routing Engines synchronizes the kernels on the backup Routing Engines with the kernels on the master Routing Engines.
 4. The global backup Routing Engine is upgraded with the new software. Next the global backup Routing Engine is rebooted. Then the global backup Routing Engine synchronizes the configuration and kernel state from the global master Routing Engine.
 5. The unified ISSU control moves from the management process to the chassis process (chassisd). The chassis process informs the various software processes (such as rpd, apsd, bfdd, and so on) about the unified ISSU and waits for responses from them.
 6. After receiving messages from the software processes indicating that the processes are ready for unified ISSU, the chassis process on the global master Routing Engine sends messages to the chassis process on the routers to start the unified ISSU.
 7. The chassis process on the routers sends ISSU_PREPARE messages to the field-replaceable units (FRUs), such as FPCs and intelligent PICs.
 8. After receiving an ISSU_PREPARE message, the Packet Forwarding Engines save the current state information and download the new software image from the backup Routing Engines. Next, each Packet Forwarding Engine sends ISSU_READY messages to the chassis process. You can display the unified ISSU process messages by using the `show log messages` command.

9. After receiving an ISSU_READY message from the Packet Forwarding Engines, the chassis process sends an ISSU_REBOOT message to the FRUs. While the upgrade is in progress, the FRUs keep sending ISSU_IN_PROGRESS messages to the chassis process. The chassis process on each router, in turn, sends an ISSU_IN_PROGRESS message to the chassis process on the global master Routing Engine.
10. After the unified ISSU reboot, the Packet Forwarding Engines restore the saved state information and connect back to the router. Then the chassis process on each router sends an ISSU_READY message to the chassis process on the global master Routing Engine. The CM_MSG_READY message (this message is sent from the LCC chassisd to the global master's chassisd) indicates that the unified ISSU is complete on the FRUs.
11. The unified ISSU control moves back to the management process on the global master Routing Engine.
12. The management process initiates a Routing Engine switchover on the master Routing Engines.
13. Routing Engine switchover occurs on the TX Matrix Plus router and all the connected LCCs.
14. After the switchover, the FRUs connect to the new master Routing Engines, and the chassis manager and Packet Forwarding Engine manager on the connected LCC FRUs connect to the new master Routing Engines on the connected LCCs.
15. The management process on the global master Routing Engine initiates the upgrade process on the Routing Engines that were previously the masters on the connected T1600 router LCCs. (This step is skipped if you have specified the **no-old-master-upgrade** option in the **request system software in-service-upgrade** command.)
16. After the Routing Engines that were previously the masters on the connected T1600 router LCCs are upgraded, the management process initiates the upgrade of the Routing Engine that was previously the global master on the TX Matrix Plus router and all its connected LCCs.
17. After a successful unified ISSU, the TX Matrix Plus global Routing Engine (re1) that was previously the master and is now the backup and the LCC Routing Engines that were previously the masters and are now the backups are rebooted if you specified the **reboot** option in the **request system software in-service-upgrade** command.

**Related
Documentation**

- [Getting Started with Unified In-Service Software Upgrade on page 387](#)
- [Best Practices for Performing a Unified ISSU on page 423](#)

- [Unified ISSU System Requirements on page 401](#)
- [Example: Performing a Unified ISSU on page 424](#)
- [request system software validate in-service-upgrade on page 740](#)

Understanding In-Service Software Upgrade (ISSU)

An in-service software upgrade (ISSU) enables you to upgrade between two different Junos OS releases with minimal disruption on the control plane and with minimal disruption of traffic. During an ISSU, the Junos OS runs in two separate virtual machines (VMs)—one VM is in the master role acting as the master Routing Engine, and the other VM is in the backup role acting as the backup Routing Engine. The Junos OS is upgraded on the backup VM. After a successful software upgrade, the backup VM then becomes the master VM, and the original master VM is no longer needed and is shut down.



Video: [How Does ISSU Work on the QFX5100?](#)

ISSU provides the following benefits:

- Eliminates network downtime during software image upgrades
- Reduces operating costs, while delivering higher service levels
- Allows fast implementation of new features
- [In-Service Software Upgrade Process on page 399](#)

In-Service Software Upgrade Process

When you request an ISSU on a standalone device:

1. The management process (mgd) verifies that non-stop routing (NSR), graceful Routing Engine switchover (GRES), and non-stop bridging (NSB) are enabled.
2. The switch downloads and validates the software package.
3. The ISSU state machine spawns the backup Routing Engine (RE) with the newer software.
4. The ISSU state machine checks to see if the backup RE has synchronized all of the data with the master RE.
5. The ISSU state machine moves the devices (for example, forwarding ASIC, FPGA, management port and serial console) from the master RE to the backup RE.
6. The mastership is switched between the REs, so the backup RE becomes the master RE.
7. The old master RE is shut down.

- Related Documentation**
- [In-Service Software Upgrade \(ISSU\) System Requirements](#)
 - [Performing an In-Service Software Upgrade \(ISSU\) with Non-Stop Routing on page 452](#)

Understanding In-Service Software Upgrade (ISSU) in ACX5000 Series Routers

An in-service software upgrade (ISSU) enables you to upgrade between two different Junos OS releases with minimal disruption on the control plane and with minimal disruption of traffic. During an ISSU, the Junos OS runs in two separate virtual machines (VMs)—one VM is in the master role acting as the master Routing Engine, and the other VM is in the backup role acting as the backup Routing Engine. The Junos OS is upgraded on the backup VM. After a successful software upgrade, the backup VM then becomes the master VM, and the original master VM is no longer needed and is shut down.



NOTE: ISSU is supported in Junos OS Release 15.1X54–D60 or later for ACX5000 Series routers.

ISSU provides the following benefits:

- Eliminates network downtime during software image upgrades
- Reduces operating costs, while delivering higher service levels
- Allows fast implementation of new features
- [In-Service Software Upgrade Process on page 400](#)

In-Service Software Upgrade Process

When you request an ISSU on a standalone device:

1. The management process (mgd) verifies that non-stop routing (NSR), graceful Routing Engine switchover (GRES), and non-stop bridging (NSB) are enabled.
2. The router downloads and validates the software package.
3. The ISSU state machine spawns the backup Routing Engine (RE) with the newer software.
4. The ISSU state machine checks to see if the backup RE has synchronized all of the data with the master RE.
5. The ISSU state machine moves the devices (for example, forwarding ASIC, FPGA, management port and serial console) from the master RE to the backup RE.
6. The mastership is switched between the REs, so the backup RE becomes the master RE.
7. The old master RE is shut down.

Related Documentation

Unified ISSU System Requirements

- [Unified ISSU System Requirements on page 401](#)

Unified ISSU System Requirements

The unified in-service software upgrade (ISSU) feature enables you to upgrade your device between two different Junos OS releases with no disruption on the control plane and with minimal disruption of traffic. Unified ISSU is supported only on dual Routing Engine platforms. In addition, the graceful Routing Engine switchover (GRES) and nonstop active routing (NSR) features must be enabled.

To access an interactive tool for verifying hardware support for unified ISSU, see the [Juniper Networks Feature Explorer](#).

This section contains the following topics:

- [General Unified ISSU Considerations for All Platforms on page 402](#)
- [Unified ISSU Considerations for MX Series Routers on page 403](#)
- [Unified ISSU Considerations for PTX Series Routers on page 404](#)
- [Unified ISSU Considerations for T Series Routers on page 404](#)
- [Unified ISSU Considerations for EX Series Switches on page 405](#)
- [Unified ISSU Platform Support on page 405](#)
- [Unified ISSU Protocol Support for M Series, MX Series, and T Series Routers and EX9200 Switches on page 406](#)
- [Unified ISSU Feature Support on page 407](#)
- [Unified ISSU PIC Support Considerations on page 407](#)

General Unified ISSU Considerations for All Platforms

Unified ISSU has the following caveats:

- We recommend that you not use unified ISSU to upgrade from an earlier Junos OS release to Junos OS Release 14.2.R1 or 15.1.R1. ISSU is not supported in Junos OS Release 14.2. For more information about Junos OS Release 14.2, see the [Release Notes for Junos OS Release 14.2](#). For more information about Junos OS Release 15.1, see the [Release Notes for Junos OS Release 15.1](#).
- Using unified ISSU to upgrade from an earlier Junos OS release to Junos OS Release 17.1R1 or later does not work if VPLS dynamic profiles are configured and enhanced subscriber management is not configured.
- The master Routing Engine and backup Routing Engine must be running the same software version before you can perform a unified ISSU.
- The unified ISSU process is aborted and a message is displayed if the Junos OS version specified for installation is a version earlier than the one currently running on the device.
- The unified ISSU process is aborted if the specified upgrade has conflicts with the current configuration, components supported, and so forth.
- You cannot take PICs offline or bring them online during a unified ISSU.
- User-initiated GRES is blocked when the device is undergoing a unified ISSU.
- Unified ISSU does not support extension application packages developed with the Junos SDK.
- To downgrade from a unified ISSU-capable release to a previous software release (unified ISSU-capable or not), use the **request system software add package-name** command. Unlike an upgrade using the unified ISSU process, a downgrade using the **request system software add package-name** command can cause network disruptions and loss of data. For more information about the use of the **request system software add package-name** command, see the *Software Installation and Upgrade Guide*.
- Unicast reverse-path-forwarding (RPF)-related statistics are not saved across a unified ISSU, and the unicast RPF counters are reset to zero during a unified ISSU.
- BGP session uptime and downtime statistics are not synchronized between the master and backup Routing Engines during a unified ISSU. The backup Routing Engine maintains its own session uptime based on the time when the backup first becomes aware of the established sessions. For example, if the backup Routing Engine is rebooted (or if you run **restart routing** on the backup Routing Engine), the backup Routing Engine uptime is a short duration, because the backup has just learned about the established sessions. If the backup is operating when the BGP sessions first come up on the master, the uptime on the master and the uptime on the backup are almost the same duration. After a Routing Engine switchover, the new master continues from the time left on the backup Routing Engine.
- If proxy ARP is enabled on your device, you must delete the **unconditional-src-learn** statement from the **[edit interfaces interface-name unit 0 family inet]** hierarchy level

before the unified ISSU process begins and include it after the unified ISSU process is complete. Note that the **unconditional-src-learn** statement is not included by default.

Unified ISSU Considerations for MX Series Routers

Unified ISSU has the following caveats for MX Series routers:

- On MX Series 3D Universal Edge Routers (with Modular Port Concentrator/Modular Interface Card (MPC/MIC) interfaces), unified ISSU is supported starting with Junos OS Release 11.2.
- On MX Series 3D Universal Edge Routers with MPC3E and MPC4E interfaces, unified ISSU is supported starting with Junos OS Release 13.3.
- Unified ISSU is supported with Junos OS Release 17.4R1 for MX Series routers with MPC-3D-16XGE-SFPP, MPC-3D-16XGE-SFPP-R-B, MPC-SEPTUM-S, MPC2E-3D-NG, MPC2E-3D-NG-IR-B, MPC2E-3D-NG-Q, MPC2E-3D-NG-Q-IR-B, MPC2E-3D-NG-Q-R-B, MPC2E-3D-NG-R-B, MPC3E-3D-NG, MPC3E-3D-NG-IR-B, MPC3E-3D-NG-Q, MPC3E-3D-NG-Q-IR-B, MPC3E-3D-NG-Q-R-B, MPC3E-3D-NG-R-B, MPC4E-3D-2CGE-8XGE, MPC4E-3D-2CGE8XGE-IR-B, MPC4E-3D-2CGE8XGE-R-B, MPC4E-3D-32XGE-IR-B, MPC4E-3D-32XGE-R-B, MPC4E-3D-32XGE-SFPP, MPC5E-100G10G, MPC5E-100G10G-IRB, MPC5E-100G10G-RB, MPC5E-40G10G, MPC5E-40G10G-IRB, MPC5E-40G10G-RB, MPC5EQ-100G10G, MPC5EQ-100G10G-IRB, MPC5EQ-100G10G-RB, MPC5EQ-40G10G, MPC5EQ-40G10G-IRB, MPC5EQ-40G10G-RB, MPC7E-10G, MPC7E-10G-IRB, MPC7E-10G-RB, MPC7E-MRATE, MPC7E-MRATE-IRB, MPC7E, MRATE-RB, MPC7EQ-10G-B, MPC7EQ-10G-IRB, MPC7EQ-10G-RB, MPC7EQ-MRATE-B, MPC7EQ-MRATE-IRB, MPC7EQ-MRATE-RB Flexible Port Concentrators (FPCs). If you perform a unified ISSU on a MX Series router with these FPCs installed, the FPCs need to be rebooted in order to complete the unified ISSU process.
- Unified ISSU for MX Series routers does not support the IEEE 802.1ag OAM and IEEE 802.3ah protocols.
- Unified ISSU is not supported when clock synchronization is configured for Synchronous Ethernet, Precision Time Protocol (PTP), and hybrid mode on the MICs and MPCEs on MX240, MX480, and MX960 routers. If clock synchronization is configured, the unified ISSU process aborts.
- On MX Series routers with MPC/MIC interfaces, the policers for transit traffic and statistics are disabled temporarily during the unified ISSU process.
- On MX Series MPCs, interface-specific and firewall filter statistics are preserved across a unified ISSU. During the unified ISSU, counter and policer operations are disabled.
- To preserve statistics across a unified ISSU on MX Series routers with MPC/MIC interfaces, the router stores the statistics data as binary large objects. The router collects the statistics before the unified ISSU is initialized, and restores the statistics after the unified ISSU completes. No statistics are collected during the unified ISSU process.
- After a unified ISSU operation is completed, an MPC reboot is required for MACsec to work. If you upgrade a router from an earlier Junos OS release to Release 14.2R2 or

Release 15.1R1 using unified ISSU and MACsec is configured on that router, you must reboot the MPC for MACsec to function properly.

- When there is a large number of subscribers configured, the Layer 2 scheduler can become oversubscribed. The unified ISSU process might abort when the system runs out of schedulers. The system generates log messages with ISSU failures and CRC errors on the control plane. If you encounter this issue, please contact JTAC for assistance in eliminating the Layer 2 scheduler oversubscription in your configuration.
- MX Series routers support Link Aggregation Control Protocol (LACP) with fast hellos during unified ISSU. This support is disabled by default. You must enable the `fast-hello-issu` option on the main router and on the peer routers before starting unified ISSU. Note that the peer router must also be an MX Series router for this functionality to work.

Unified ISSU Considerations for PTX Series Routers

Unified ISSU has the following caveats for PTX Series routers:

- Starting with Junos OS Release 13.2, unified ISSU is supported on the PTX5000 and PTX3000 with the FPC-PTX-P1-A FPC. However, you can perform unified ISSU only from Junos OS Release 13.2 to 13.3 and from Junos OS Release 14.1 to a later release. You must *not* perform unified ISSU from Junos OS Release 13.2 or 13.3 to 14.1 and later releases.
- Link Aggregation Control Protocol (LACP) is not supported during unified ISSU on PTX Series routers. You must disable the `lacp` statement at the `[edit interfaces interface-name aggregated-ether-options]` hierarchy level before the unified ISSU process begins and enable it after the unified ISSU process is complete.

Unified ISSU Considerations for T Series Routers

Unified ISSU has the following caveats for T Series devices:

- During the unified ISSU process on a routing matrix with TX Matrix Plus routers with 3D SIBs, only 75 percent of the traffic remains uninterrupted.
- The scale supported on T640-FPC2-E, T640-FPC2-E2, T640-FPC3-E, and T640-FPC3-E2 Flexible Port Concentrators (FPCs) is less than that supported on T640-FPC1-ES, T640-FPC2-ES, T640-FPC3-ES, T1600-FPC4-ES, and T640-FPC4-1P-ES FPCs because of differences in hardware configuration. Therefore, when a unified ISSU is performed, if the configured scale on any of the FPCs is more than what is supported on that FPC, field-replaceable unit (FRU) upgrade of that FPC fails. To check the current hardware configuration of an FPC, use the `show chassis fpc operational` command.
- The PD-4XGE-XFP PIC goes offline during a unified ISSU if the PIC is installed in a T-1600-FPC4-ES with part number 710-013037 revision 12 or earlier.
- In the FPCs on T4000 routers, interface-specific and firewall filter statistics are preserved across a unified ISSU. During the unified ISSU, counter and policer operations are disabled.

- To preserve statistics across a unified ISSU on T4000 routers with FPC/PIC interfaces, the router stores the statistics data as binary large objects. The router collects the statistics before the unified ISSU is initialized, and restores the statistics after the unified ISSU completes. No statistics are collected during the unified ISSU process.
- To verify that statistics are preserved across the unified ISSU, you can issue CLI operational commands such as **show interfaces statistics** after the unified ISSU completes.
- When you configure the unified ISSU feature on the T4000 Core Router, you can also configure LACP. However, LACP periodic fast mode is not supported. If you configure LACP periodic transmission, set it to slow mode at both sides before initiating a unified ISSU. If fast mode is configured, the configuration can be committed without any commit or system log error messages, but you might notice that a larger than expected amount of traffic drops because of the LACP links going down during a unified ISSU.

Unified ISSU Considerations for EX Series Switches

Unified ISSU has the following caveats for EX Series devices:

- EX9204, EX9208, EX9214, and EX9251, and EX9253 switches do not support LACP fast timer configuration starting with Junos OS Release 17.4. If the LACP fast timer is configured, there will be LAG interface flaps traffic loss during ISSU. We recommend moving to LACP slow before beginning ISSU on these devices.

Unified ISSU Platform Support

Table 21 on page 405 lists the platforms that support unified ISSU when dual Routing Engines are installed and the first Junos OS release that supports unified ISSU on those platforms. In addition to verifying that your platform supports unified ISSU, you need to verify that the field-replaceable unit, such as PICs, that are installed also support unified ISSU.

To access an interactive tool for verifying hardware support for unified ISSU, see the Juniper Networks Feature Explorer (<https://pathfinder.juniper.net/feature-explorer/>).

Table 21: Unified ISSU Support for Dual Routing Engine Platforms

Platform	Junos OS Release
EX9200 switch	<ul style="list-style-type: none"> • 12.3R3 or later • 14.2R1 or later on EX9200-32XS, EX9200-4QS, and EX9200-2C-8XS • 17.1R1 or later on EX9200-6QS
M10i router	9.5R1
M120 router	9.2R1
M320 router	9.0R1
MX240 router	9.3R1

Table 21: Unified ISSU Support for Dual Routing Engine Platforms (continued)

Platform	Junos OS Release
MX480 router	9.3R1
MX960 router	9.3R1
MX2010 router	13.2R1
MX2020 router	13.2R1
MX104 router	14.1R1
MX Series Virtual Chassis	14.1R1
MX10003 router	18.2R1
PTX5000 router	13.2R1
PTX3000 router	13.2R1
T320 router	9.0R1
T640 router	9.0R1
T1600 router	9.1R1
T4000 router	12.3R1
TX Matrix router	9.3R1
TX Matrix Plus router	12.3R2
TX Matrix Plus routers with 3D SIBs	14.1R1

Unified ISSU Protocol Support for M Series, MX Series, and T Series Routers and EX9200 Switches

To find out which releases support ISSU, please use the [ISSU Feature Explorer](#) tool on the Juniper Networks website. The ISSU Feature Explorer tool contains information about the Juniper Networks devices that support ISSU, the releases that support ISSU for each device, and the SKUs that support ISSU for each release.



NOTE: To gain access to the ISSU Feature Explorer tool, you need to log in with a customer or partner account on the Juniper Networks website. For more information on setting up a Juniper Networks account, please see the [Juniper Networks Guide to Creating a User Account](#).

Unified ISSU Feature Support

Unified ISSU supports most Junos OS features starting in Junos OS Release 9.0. However, the following constraints apply:

- Link Aggregation Control Protocol (LACP)—Link changes are not processed until after the unified ISSU is complete.
- Automatic Protection Switching (APS)—Network changes are not processed until after the unified ISSU is complete.
- Ethernet Operation, Administration, and Management (OAM) as defined by IEEE 802.3ah and by IEEE 802.1ag—When a Routing Engine switchover occurs, the OAM hello message times out, triggering protocol convergence.
- Ethernet circuit cross-connect (CCC) encapsulation—Circuit changes are not processed until after the unified ISSU is complete.
- Logical systems—On devices that have logical systems configured on them, only the master logical system supports unified ISSU.



NOTE: Starting with Junos OS Release 16.1R1, while performing a unified ISSU from a FreeBSD 6.1-based Junos OS to an upgraded FreeBSD 10.x-based Junos OS, the configuration must be validated on a remote host or on a Routing Engine. The remote host or the Routing Engine must be running a Junos OS with an upgraded FreeBSD. In addition, only a few selected directories and files are preserved while upgrading from FreeBSD 6.1-based Junos OS to FreeBSD 10.x-based Junos OS. See [Upgrading Junos OS with Upgraded FreeBSD](#).

Unified ISSU PIC Support Considerations

The following sections list the PICs that are supported by unified ISSU.

- [PIC Considerations on page 408](#)
- [SONET/SDH PICs on page 409](#)
- [Fast Ethernet and Gigabit Ethernet PICs on page 410](#)
- [Channelized PICs on page 413](#)
- [Tunnel Services PICs on page 414](#)
- [ATM PICs on page 414](#)
- [Serial PICs on page 415](#)
- [DS3, E1, E3, and T1 PICs on page 415](#)
- [Enhanced IQ PICs on page 416](#)
- [Enhanced IQ2 Ethernet Services Engine \(ESE\) PIC on page 416](#)
- [Unified ISSU FPC Support on T4000 Routers on page 417](#)
- [Unified ISSU Support on MX Series 3D Universal Edge Routers on page 417](#)



NOTE: For information about ISSU support on individual PICs based on device and release, use the [ISSU Feature Explorer](#) tool.



NOTE: For information about Flexible PIC Concentrator (FPC) types, FPC/PIC compatibility, and the initial Junos OS release in which a particular PIC is supported on an FPC, see the PIC guide for your platform.

PIC Considerations

Take the following PIC restrictions into consideration before performing a unified ISSU:

- **Unsupported PICs**—If a PIC is not supported by unified ISSU, at the beginning of the upgrade, the software issues a warning that the PIC will be taken offline. After the PIC is brought offline and the unified ISSU is complete, the PIC is brought back online with the new firmware.
- **PIC combinations**—For some PICs, newer Junos OS services can require significant Internet Processor ASIC memory, and some configuration rules might limit certain combinations of PICs on particular platforms. With a unified ISSU:
 - If a PIC combination is not supported by the software version that the device is being upgraded from, the validation check displays a message and aborts the upgrade.
 - If a PIC combination is not supported by the software version to which the device is being upgraded, the validation check displays a message and aborts the upgrade, even if the PIC combination is supported by the software version from which the device is being upgraded.
- **Interface statistics**—Interface statistics might be incorrect because:
 - During bootup of the new microkernel on the Packet Forwarding Engine, host-bound traffic is not handled and might be dropped, causing packet loss.
 - During the hardware update of the Packet Forwarding Engine and its interfaces, traffic is halted and discarded. (The duration of the hardware update depends on the number and type of interfaces and on the device configuration.)
 - During a unified ISSU, periodic statistics collection is halted. If hardware counters saturate or wrap around, the software does not display accurate interface statistics.
- **CIR oversubscription**—If oversubscription of the committed information rate (CIR) is configured on logical interfaces:
 - And the sum of the CIR exceeds the physical interface's bandwidth, after a unified ISSU is performed, each logical interface might not be given its original CIR.
 - And the sum of the delay buffer rate configured on logical interfaces exceeds the physical interface's bandwidth, after a unified ISSU is performed, each logical interface might not receive its original delay-buffer-rate calculation.

SONET/SDH PICs

Table 22 on page 409 lists the SONET/SDH PICs that are supported during a unified ISSU.

Table 22: Unified ISSU PIC Support: SONET/SDH

PIC Type	Number of Ports	Model Number	Device
OC3c/STM1	4	PB-4OC3-SON-MM—(EOL)	M120 M320, T320, T640, T1600
		PB-4OC3-SON-SMIR—(EOL)	
		PE-4OC3-SON-MM—(EOL)	M10i
		PE-4OC3-SON-SMIR—(EOL)	
	2	PE-2OC3-SON-MM—(EOL)	
		PE-2OC3-SON-SMIR—(EOL)	
OC3c/STM1 with SFP	2	PE-2OC3-SON-SFP	M10i
OC3c/STM1, SFP (Multi-Rate)	4 OC3 ports, 4 OC12 ports	PB-4OC3-4OC12-SON-SFP	M120 M320, MX Series, T320, T640, T1600, T4000, TX Matrix Plus, TX Matrix Plus with 3D SIBs
	4 OC3 ports, 1 OC12 port	PB-4OC3-1OC12-SON-SFP	
		PB-4OC3-1OC12-SON2-SFP	M10i
		PE-4OC3-1OC12-SON-SFP	
OC12c/STM4	1	PE-1OC12-SON-SFP	M10i
		PE-1OC12-SON-MM—(EOL)	
		PE-1OC12-SON-SMIR—(EOL)	
		PB-1OC12-SON-MM—(EOL)	M120, M320, T320, T640, T1600, T4000, TX Matrix, TX Matrix Plus with 3D SIBs
		PB-1OC12-SON-SMIR—(EOL)	
	4	PB-4OC12-SON-MM PB-4OC12-SON-SMIR	
OC12c/STM4, SFP	1	PB-1OC12-SON-SFP	M120, M320, T320, T640, T1600, TX Matrix, TX Matrix Plus
OC48c/STM16, SFP	1	PB-1OC48-SON-SFP PB-1OC48-SON-B-SFP	M120, M320, MX Series, T320, T640, T1600, TX Matrix, T4000, TX Matrix Plus, TX Matrix Plus with 3D SIBs
	4	PC-4OC48-SON-SFP	

Table 22: Unified ISSU PIC Support: SONET/SDH (continued)

PIC Type	Number of Ports	Model Number	Device
OC192/STM64	1	PC-1OC192-SON-VSR	MX Series routers
OC192/STM64, XFP	1	PC-1OC192-SON-LR PC-1OC192-SON-SR2 PC-1OC192-VSR	M320, T320, T640, T1600, T4000, TX Matrix Plus with 3D SIBs
OC192/STM64, XFP	4	PD-4OC192-SON-XFP	M120, T640, T1600, T4000, TX Matrix Plus with 3D SIBs
	1	PC-1OC192-SON-XFP	T4000, MX Series routers, TX Matrix Plus with 3D SIBs
OC768/STM256	1	PD-1OC768-SON-SR	T640, T1600, T4000, TX Matrix Plus, TX Matrix Plus with 3D SIBs

Fast Ethernet and Gigabit Ethernet PICs

Table 23 on page 411 lists the Fast Ethernet and Gigabit Ethernet PICs that are supported during a unified ISSU.



NOTE: Starting with Junos OS Release 9.2, new Ethernet IQ2 PIC features might cause the software to reboot the PIC when a unified ISSU is performed. For information about applicable new Ethernet IQ2 PIC features, refer to the release notes for the specific Junos OS release.

Table 23: Unified ISSU PIC Support: Fast Ethernet and Gigabit Ethernet

PIC Type	Number of Ports	Model Number	Device
Fast Ethernet	4	PB-4FE-TX	M120, M320, T320, T640, T1600, TX Matrix
		PE-4FE-TX	M10i
	8	PB-8FE-FX	M120, M320
		PE-8FE-FX	M10i
	12	PB-12FE-TX-MDI	M120, M320, T320
		PB-12FE-TX-MDIX	
		PE-12FE-TX-MDI	M10i
		PE-12FE-TX-MDIX	
	48	PB-48FE-TX-MDI	M120, M320, T320
		PB-48FE-TX-MDIX	
Gigabit Ethernet, RJ-45	40	EX9200-40T	EX9200
Gigabit Ethernet, SFP	1	PE-1GE-SFP	M10i
		PB-1GE-SFP	M120, M320, T320, T640, T1600, T4000, TX Matrix, TX Matrix Plus, TX Matrix Plus with 3D SIBs
	2	PB-2GE-SFP	
	4	PB-4GE-SFP	
	10	PC-10GE-SFP	EX9200
	40	EX9200-40F	
Gigabit Ethernet IQ, SFP	1	PE-1GE-SFP-QPP	M10i
		PB-1GE-SFP-QPP	M120, M320, T320, T640, T1600, T4000, TX Matrix, TX Matrix Plus, TX Matrix Plus with 3D SIBs
	2	PB-2GE-SFP-QPP	
Gigabit Ethernet IQ2, SFP	4	PB-4GE-TYPE1-SFP-IQ2	M120, M320, T320, T640, T1600, T4000, TX Matrix, TX Matrix Plus, TX Matrix Plus with 3D SIBs
	8	PB-8GE-TYPE2-SFP-IQ2	
		PC-8GE-TYPE3-SFP-IQ2	

Table 23: Unified ISSU PIC Support: Fast Ethernet and Gigabit Ethernet (continued)

PIC Type	Number of Ports	Model Number	Device
Gigabit Ethernet IQ2, XFP	1	PC-1XGE-TYPE3-XFP-IQ2	M120, M320, T320, T640, T1600, TX Matrix, TX Matrix Plus, TX Matrix Plus with 3D SIBs
10-Gigabit Ethernet XFP	4	PD-4XGE-XFP <i>NOTE:</i> This PIC goes offline during a unified ISSU if the PIC is inserted on T-1600-FPC4-ES with part number 710-013037 revision 12 or below.	T640, T1600, T4000, TX Matrix, TX Matrix Plus, TX Matrix Plus with 3D SIBs
10-Gigabit Ethernet SFP+	10	PD-5-10XGE-SFPP	T640, T1600, T4000, TX Matrix Plus, TX Matrix Plus with 3D SIBs
	24	P1-PTX-24-10GE-SFPP	PTX5000
		EX9200-6QS	EX9200
	32	EX9200-32XS	EX9200
10-Gigabit Ethernet, DWDM	1	PC-1XGE-DWDM-CBAND	M120, M320, T320, T640, T1600, TX Matrix, TX Matrix Plus, TX Matrix Plus with 3D SIBs
10-Gigabit Ethernet, DWDM OTN	1	PC-1XGE-DWDM-OTN	T4000, TX Matrix Plus with 3D SIBs
10-Gigabit Ethernet LAN/WAN PIC with SFP+	12	PF-12XGE-SFPP	T4000, TX Matrix Plus with 3D SIBs
	24	PF-24XGE-SFPP	T4000, TX Matrix Plus with 3D SIBs
10-Gigabit Ethernet, SFP+	32	14.2R1 or later EX9200-32XS	EX9200
10-Gigabit Ethernet, XENPAK	1	PC-1XGE-XENPAK	M120, M320, T320, T640, T1600, T4000, TX Matrix, TX Matrix Plus, TX Matrix Plus with 3D SIBs
40-Gigabit Ethernet, CFP	2	P1-PTX-2-40GE-CFP	PTX5000
10-Gigabit Ethernet, 40-Gigabit Ethernet, QFSP+	16/4	14.2R1 or later EX9200-4QS	EX9200
	24/6	17.1R1 or later EX9200-6QS	
	48/12	P2-10G-40G-QSFPP	PTX5000

Table 23: Unified ISSU PIC Support: Fast Ethernet and Gigabit Ethernet (continued)

PIC Type	Number of Ports	Model Number	Device
100-Gigabit Ethernet, CFP	1	PF-1CGE-CFP	T4000, TX Matrix Plus with 3D SIBs
	2	P1-PTX-2-100GE-CFP	PTX5000
	4	P2-100GE-CFP2	PTX5000
100-Gigabit Ethernet CFP/10-Gigabit Ethernet SFP+	2/8	EX9200-2C-8XS	EX9200
100-Gbps DWDM OTN	2	P1-PTX-2-100G-WDM	PTX5000
100-Gbps OTN, CFP2	4	P2-100GE-OTN	PTX5000

Channelized PICs

Table 24 on page 413 lists the channelized PICs that are supported during a unified ISSU.

Table 24: Unified ISSU PIC Support: Channelized

PIC Type	Number of Ports	Model Number	Platform
Channelized E1 IQ	10	PB-10CHE1-RJ48-QPP	M120, M320, T320, T640, T1600, TX Matrix
		PB-10CHE1-RJ48-QPP-N	M120
		PE-10CHE1-RJ48-QPP	M10i
		PE-10CHE1-RJ48-QPP-N	
Channelized T1 IQ	10	PB-10CHT1-RJ48-QPP	M320, T320, T640, T1600
		PE-10CHT1-RJ48-QPP	M10i
Channelized OC IQ	1	PB-1CHOC12SMIR-QPP	M120, M320, T320, T640, T1600, TX Matrix, TX Matrix Plus
		PB-1CHSTM1-SMIR-QPP	
		PB-1CHOC3-SMIR-QPP	
		PE-1CHOC12SMIR-QPP	M10i
		PE-1CHOC3-SMIR-QPP	

Table 24: Unified ISSU PIC Support: Channelized (continued)

PIC Type	Number of Ports	Model Number	Platform
Channelized DS3 to DS0 IQ	4	PB-4CHDS3-QPP	M120, M320, T320, T640, T1600, TX Matrix, TX Matrix Plus
		PE-4CHDS3-QPP	M10i
Channelized STM 1	1	PE-1CHSTM1-SMIR-QPP	M10i

Tunnel Services PICs

[Table 25 on page 414](#) lists the Tunnel Services PICs that are supported during a unified ISSU.

Table 25: Unified ISSU PIC Support: Tunnel Services

PIC Type	Model Number	Platform
1-Gbps Tunnel	PE-TUNNEL	M10i
	PB-TUNNEL-1	M120, M320, T320, T640, T1600, TX Matrix, TX Matrix Plus
4-Gbps Tunnel	PB-TUNNEL	
10-Gbps Tunnel	PC-TUNNEL	

ATM PICs

[Table 26 on page 414](#) lists the ATM PICs that are supported during a unified ISSU. The table includes support on Enhanced III FPCs.

Table 26: Unified ISSU PIC Support: ATM

PIC Type	Number of Ports	Model Number	Platform
DS3	4	PB-4DS3-ATM2	M120, M320, T320, T640, T1600, TX Matrix
		PE-4DS3-ATM2	M10i
E3	4	PB-4E3-ATM2	M120, M320, T320, T640, T1600, TX Matrix, TX Matrix Plus
	2	PE-2E3-ATM2	M10i

Table 26: Unified ISSU PIC Support: ATM (continued)

PIC Type	Number of Ports	Model Number	Platform
OC3/STM1	2	PB-2OC3-ATM2-MM	M120, M320, T320, T640, T1600, TX Matrix, TX Matrix Plus
		PB-2OC3-ATM2-SMIR	
		PE-2OC3-ATM2-MM	M10i
		PE-2OC3-ATM2-SMIR	
OC12/STM4	1	PB-1OC12-ATM2-MM	M120, M320, T320, T640, T1600, TX Matrix, TX Matrix Plus
		PB-1OC12-ATM2-SMIR	
	2	PB-2OC12-ATM2-MM	M120, M320, T320, T640, T1600, T4000, TX Matrix, TX Matrix Plus, TX Matrix Plus with 3D SIBs
		PB-2OC12-ATM2-SMIR	
	1	PE-1OC12-ATM2-MM	M10i
		PE-1OC12-ATM2-SMIR	
OC48/STM16	1	PB-1OC48-ATM2-SFP	M120, M320, T320, T640, T1600, TX Matrix, TX Matrix Plus

Serial PICs

Unified ISSU supports the following 2-port EIA-530 serial PICs:

- PB-2EIA530 on M320 routers with Enhanced III FPCs, and on M120 routers
- PE-2EIA530 on M10i routers

DS3, E1, E3, and T1 PICs

Unified ISSU supports the following PICs on M120, M320, and T320 routers; T640 and T1600 routers; and the TX Matrix router:

- 4-Port DS3 PIC (PB-4DS3)
- 4-Port E1 Coaxial PIC (PB-4E1-COAX)
- 4-Port E1 RJ48 PIC (PB-4E1-RJ48)
- 4-Port E3 IQ PIC (PB-4E3-QPP)
- 4-Port T1 PIC (PB-4T1-RJ48)



NOTE: Unified ISSU is also supported on the 4-Port DS3 PIC (PB-4DS3) and the 4-Port E3 IQ PIC (PB-4E3-QPP) on the TX Matrix Plus router.

Unified ISSU supports the following PICs on M10i routers:

- 2-Port DS3 PIC (PE-2DS3)
- 4-Port DS3 PIC (PE-4DS3)
- 4-Port E1 PICs (PE-4E1-COAX and PE-4E1-RJ48)
- 2-Port E3 PIC (PE-2E3)
- 4-Port T1 PIC (PE-4T1-RJ48)
- 4-Port E3 IQ PIC (PE-4E3-QPP)

Enhanced IQ PICs

Unified ISSU supports the following PICs on M120 router, M320 router, and on T320 routers; T640 routers, T1600 routers, TX Matrix router, and the TX Matrix Plus router:

- 1-Port Channelized OC12/STM4 Enhanced IQ PIC (PB-1CHOC12-STM4-IQE-SFP)
- 1-Port nonchannelized OC12/STM4 Enhanced IQ PIC (PB-1OC12-STM4-IQE-SFP)
- 4-Port Channelized DS3/E3 Enhanced IQ PIC (PB-4CHDS3-E3-IQE-BNC)
- 4-Port nonchannelized DS3/E3 Enhanced IQ PIC (PB-4DS3-E3-IQE-BNC)
- 4-Port nonchannelized SONET/SDH OC48/STM16 Enhanced IQ (IQE) PIC with SFP (PC-4OC48-STM16-IQE-SFP)

Unified ISSU supports 1-port Channelized OC48/STM16 Enhanced IQ (IQE) PIC with SFP (PB-1CHOC48-STM16-IQE-SFP) on MX Series routers.

Enhanced IQ2 Ethernet Services Engine (ESE) PIC

Unified ISSU supports the enhanced IQ2 ESE PICs listed in [Table 27 on page 416](#).

Table 27: Unified ISSU Support: Enhanced IQ2 Ethernet Services Engine (ESE) PIC

Model Number	Number of Ports	Platform
PC-8GE-TYPE3-SFP-IQ2E	8	M120, M320, T320, T640, T4000 TX Matrix, TX Matrix Plus, TX Matrix Plus with 3D SIBs
PB-8GE-TYPE2-SFP-IQ2E	8	M120, M320, T320, T640, TX Matrix, TX Matrix Plus, TX Matrix Plus with 3D SIBs
PB-4GE-TYPE1-SFP-IQ2E	4	M120, M320, T320, T640
PC-1XGE-TYPE3-XFP-IQ2E	1	M120, M320, T320, T640, T4000, TX Matrix, TX Matrix Plus, TX Matrix Plus with 3D SIBs
PB-1CHOC48-STM16-IQE	1	M120, M320, T320, T640, T4000, TX Matrix, TX Matrix Plus, TX Matrix Plus with 3D SIBs
PE-4GE-TYPE1-SFP-IQ2E	4	M10i

Table 27: Unified ISSU Support: Enhanced IQ2 Ethernet Services Engine (ESE) PIC (continued)

Model Number	Number of Ports	Platform
PE-4GE-TYPE1-SFP-IQ2	4	M10i

Unified ISSU FPC Support on T4000 Routers

In the FPCs on T4000 routers, interface-specific and firewall filter statistics are preserved across a unified ISSU. During the unified ISSU, counter and policer operations are disabled.

To preserve statistics across a unified ISSU on T4000 routers with FPC/PIC interfaces, the router stores the statistics data as binary large objects. The router collects the statistics before the unified ISSU is initialized, and restores the statistics after the unified ISSU completes. No statistics are collected during the unified ISSU process.

To verify that statistics are preserved across the unified ISSU, you can issue CLI operational commands such as **show interfaces statistics** after the unified ISSU completes.

Unified ISSU is supported on the following FPCs:

- T4000 FPC5 (model numbers—T4000-FPC5-3D and T4000-FPC5-LSR)
- Enhanced Scaling FPC4-1P (model number—T640-FPC4-1P-ES)
- Enhanced Scaling FPC4 (T1600-FPC4-ES)
- Enhanced Scaling FPC3 (T640-FPC3-ES)
- Enhanced Scaling FPC2 (T640-FPC2-ES)



NOTE: The aforementioned FPCs are also supported on TX Matrix Plus routers with 3D SIBs.

Unified ISSU Support on MX Series 3D Universal Edge Routers

The following sections list the Dense Port Concentrators (DPCs), Flexible PIC Concentrators (FPCs), Modular Port Concentrators (MPCs), and Modular Interface Cards (MICs) that are supported during a unified ISSU on MX Series routers.

- [Unified ISSU DPC and FPC Support on MX Series Routers on page 417](#)
- [Unified ISSU MIC and MPC Support on MX Series Routers on page 418](#)
- [Unified ISSU Limitations on MX Series Routers on page 421](#)

Unified ISSU DPC and FPC Support on MX Series Routers

Unified ISSU supports all DPCs except the Multiservices DPC on MX Series routers. Unified ISSU also supports Type 2 FPC (**MX-FPC2**) and Type 3 FPC (**MX-FPC3**) on MX Series routers. For more information about DPCs and FPCs on MX Series routers, go to https://www.juniper.net/documentation/en_US/release-independent/junos/information-products/pathway-pages/mx-series/.

Unified ISSU MIC and MPC Support on MX Series Routers

Unified ISSU supports all the Modular Port Concentrators (MPCs) and Modular Interface Cards (MICs) listed in [Table 28 on page 418](#) and [Table 29 on page 419](#). Unified ISSU is not supported on MX80 routers.

In the MPCs on MX Series routers, interface-specific and firewall filter statistics are preserved across a unified ISSU. During the unified ISSU, counter and policer operations are disabled.

To preserve statistics across a unified ISSU on MX Series routers with MPC/MIC interfaces, the router stores the statistics data as binary large objects. The router collects the statistics before the unified ISSU is initialized, and restores the statistics after the unified ISSU completes. No statistics are collected during the unified ISSU process.

To verify that statistics are preserved across the unified ISSU, you can issue CLI operational commands such as **show interfaces statistics** after the unified ISSU completes.

Table 28: Unified ISSU Support: MX Series Router MPCs

MPC Type	Number of Ports	Model Number	Platform
MPC1	—	MX-MPC1-3D	MX Series routers
MPC1E	—	MX-MPC1E-3D	MX Series routers
MPC1 Q	—	MX-MPC1-3D-Q	MX Series routers
MPC1E Q	—	MX-MPC1E-3D-Q	MX Series routers
MPC2	—	MX-MPC2-3D	MX Series routers
MPC2E	—	MX-MPC2E-3D	MX Series routers
MPC2 Q	—	MX-MPC2-3D-Q	MX Series routers
MPC2E Q	—	MX-MPC2E-3D-Q	MX Series routers
MPC2 EQ	—	MX-MPC2-3D-EQ	MX Series routers
MPC2E EQ	—	MX-MPC2E-3D-EQ	MX Series routers
16x10GE MPC	16	MPC-3D-16XGE-SFPP	MX Series routers
MPC3E	—	MX-MPC3E-3D	MX Series routers
32x10GE MPC4E	32	MPC4E-3D-32XGE-SFPP	MX Series routers
2x100GE + 8x10GE MPC4E	10	MPC4E-3D-2CGE-8XGE	MX Series routers
6x40GE + 24x10GE MPC5E	30	MPC5E-40G10G	MX Series routers

Table 28: Unified ISSU Support: MX Series Router MPCs (continued)

MPC Type	Number of Ports	Model Number	Platform
6x40GE + 24x10GE MPC5EQ	30	MPC5EQ-40G10G	MX Series routers
2x100GE + 4x10GE MPC5E	6	MPC5E-100G10G	MX Series routers
2x100GE + 4x10GE MPC5EQ	6	MPC5EQ-100G10G	MX Series routers
MPC6E	2	MX2K-MPC6E	MX Series routers
MPC7E (multi-rate)	12	MPC7E-MRATE	MX Series routers
MPC7E 10G	40	MPC7E-10G	MX Series routers
MPC8E	—	MX2K-MPC8E	MX Series routers
MPC9E	—	MX2K-MPC9E	MX Series routers

Table 29: Unified ISSU Support: MX Series Router MICs

MIC Type	Number of Ports	Model Number	Platform
ATM MIC with SFP	8	MIC-3D-8OC3-2OC12-ATM	MX Series routers
Channelized SONET/SDH OC192/STM64 MIC with XFP	4	MIC-3D-1OC192-XFP	MX Series routers
Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP	4	MIC-3D-4COC3-1COC12-CE	MX Series routers
Channelized E1/T1 Circuit Emulation MIC	16	MIC-3D-16CHE1-T1-CE	MX Series routers
Channelized SONET/SDH OC3/STM1 (Multi-Rate) MICs with SFP	4	MIC-3D-4CHOC3-2CHOC12	MX Series routers
Channelized SONET/SDH OC3/STM1 (Multi-Rate) MICs with SFP	8	MIC-3D-4CHOC3-2CHOC12	MX Series routers
Channelized DS3/E3 MIC	8	MIC-3D-8CHDS3-E3-B	MX Series routers
DS3/E3	8	MIC-3D-8DS3-E3	MX Series routers
40-Gigabit Ethernet MIC with QSFP	2	MIC3-3D-2X40GE-QSFP	MX Series routers
10-Gigabit Ethernet MIC with SFPP	10	MIC3-3D-10XGE-SFPP	MX Series routers
100-Gigabit Ethernet MIC with CXP	1	MIC3-3D-1X100GE-CXP	MX Series routers

Table 29: Unified ISSU Support: MX Series Router MICs (continued)

MIC Type	Number of Ports	Model Number	Platform
100-Gigabit Ethernet MIC with CFP	1	MIC3-3D-1X100GE-CFP	MX Series routers
Gigabit Ethernet MIC with SFP	20	MIC-3D-20GE-SFP	MX Series routers
10-Gigabit Ethernet MIC with SFP+ (24 Ports)	24	MIC6-10G	MX Series routers
10-Gigabit Ethernet DWDM OTN MIC (non-OTN mode only)	24	MIC6-10G-OTN	MX Series routers
100-Gigabit Ethernet MIC with CFP2 (non-OTN mode only)	2	MIC6-100G-CFP2	MX Series routers
100-Gigabit Ethernet MIC with CXP (4 Ports)	4	MIC6-100G-CXP	MX Series routers
10-Gigabit Ethernet MICs with XFP	2	MIC-3D-2XGE-XFP	MX Series routers
10-Gigabit Ethernet MICs with XFP	4	MIC-3D-4XGE-XFP	MX Series routers
SONET/SDH OC3/STM1 (Multi-Rate) MICs with SFP	4	MIC-3D-4OC3OC12-1OC48	MX Series routers
SONET/SDH OC3/STM1 (Multi-Rate) MICs with SFP	8	MIC-3D-8OC3OC12-4OC48	MX Series routers
Tri-Rate Copper Ethernet MIC	40	MIC-3D-40GE-TX	MX Series routers
100-Gigabit DWDM OTn MIC with CFP2-ACO	1	MIC3-100G-DWDM	MX960 routers



NOTE: Note that unified ISSU is supported only by the MICs listed in [Table 29 on page 419](#).



NOTE: Consider the following guidelines before performing a unified ISSU on an MX Series router with ATM interfaces at scale:

- The PPP keepalive interval must be 10 seconds or greater. PPP requires three keepalives to fail before it brings down the session. Thirty seconds (ten seconds multiplied by three) provides a safe margin to maintain PPP sessions across the unified ISSU in case of any traffic loss during the operation. Configure the interval with the `keepalives` statement at the `[edit interfaces at-interface-name] or [edit interfaces at-interface-name unit logical-unit-number]` hierarchy level.
- The OAM F5 loopback cell period must be 20 seconds or greater to maintain ATM connectivity across the unified ISSU. Configure the interval with the `oam-period` statement at the `[edit interfaces at-interface-name unit logical-unit-number]` hierarchy level.

Unified ISSU Limitations on MX Series Routers

Unified ISSU is currently not supported when clock synchronization is configured for Synchronous Ethernet, Precision Time Protocol (PTP), and hybrid mode on MX80 routers and on the MICs and MPCs on MX240, MX480, and MX960 routers.



NOTE: Before enabling ISSU on MX routers, when upgrading from a Junos OS Release 14.1 or earlier to Junos OS Release 14.2 or later, you must disable IGMP snooping, and PIM snooping, in all protocol hierarchies. This includes the bridge-domain and routing-instances hierarchies.



NOTE: On MX Series routers with MPC/MIC interfaces, the policers for transit traffic and statistics are disabled temporarily during the unified ISSU process.

Release History Table

Release	Description
17.4	Unified ISSU is supported with Junos OS Release 17.4R1 for MX Series routers
17.4	EX9204, EX9208, EX9214, and EX9251, and EX9253 switches do not support LACP fast timer configuration starting with Junos OS Release 17.4.
16.1R1	Starting with Junos OS Release 16.1R1, while performing a unified ISSU from a FreeBSD 6.1-based Junos OS to an upgraded FreeBSD 10.x-based Junos OS, the configuration must be validated on a remote host or on a Routing Engine.
13.3	On MX Series 3D Universal Edge Routers with MPC3E and MPC4E interfaces, unified ISSU is supported starting with Junos OS Release 13.3.
13.2	Starting with Junos OS Release 13.2, unified ISSU is supported on the PTX5000 and PTX3000 with the FPC-PTX-P1-A FPC.
11.2	On MX Series 3D Universal Edge Routers (with Modular Port Concentrator/Modular Interface Card (MPC/MIC) interfaces), unified ISSU is supported starting with Junos OS Release 11.2.

Related Documentation

- [Getting Started with Unified In-Service Software Upgrade on page 387](#)
- [Best Practices for Performing a Unified ISSU on page 423](#)
- [Understanding the Unified ISSU Process on page 388](#)
- [Example: Performing a Unified ISSU on page 424](#)
- [Configuring LACP for Aggregated Ethernet Interfaces](#)
- [request system software validate on \(Junos OS with Upgraded FreeBSD\)](#)

CHAPTER 29

Performing a Unified ISSU

- [Best Practices for Performing a Unified ISSU on page 423](#)
- [Example: Performing a Unified ISSU on page 424](#)
- [Performing an In-Service Software Upgrade \(ISSU\) with Non-Stop Routing on page 452](#)
- [Performing an In-Service Software Upgrade \(ISSU\) in ACX5000 Series Routers on page 457](#)
- [Verifying a Unified ISSU on page 461](#)
- [Troubleshooting Unified ISSU Problems on page 462](#)
- [Managing and Tracing BFD Sessions During Unified ISSU Procedures on page 462](#)

Best Practices for Performing a Unified ISSU

When you are planning to perform a unified in-service software upgrade (ISSU), choose a time when your network is as stable as possible. As with a normal upgrade, Telnet sessions, SNMP, and CLI access are briefly interrupted. In addition, the following restrictions apply:

- The master Routing Engine and backup Routing Engine must be running the same software version before you can perform a unified ISSU.
- Verify that your platform supports the unified ISSU feature.
- Read the “Unified ISSU Considerations” topic in the chapter [“Unified ISSU System Requirements” on page 401](#) to anticipate any special circumstances that might affect your upgrade.

Related Documentation

- [Example: Performing a Unified ISSU on page 424](#)
- [Verifying a Unified ISSU on page 461](#)
- [Troubleshooting Unified ISSU Problems on page 462](#)

Example: Performing a Unified ISSU

This example shows how to perform a unified in-service software upgrade (ISSU).

- [Requirements on page 424](#)
- [Overview on page 425](#)
- [Configuration on page 426](#)
- [Verifying Dual Routing Engines and Enabling GRES and NSR on page 426](#)
- [Verifying the Software Versions and Backing Up the Device Software on page 428](#)
- [Adjusting Timers and Changing Feature-Specific Configuration on page 429](#)
- [Upgrading and Rebooting Both Routing Engines Automatically on page 431](#)
- [Restoring Feature-Specific Configuration on page 436](#)
- [Upgrading Both Routing Engines and Rebooting the New Backup Routing Engine Manually on page 438](#)
- [Upgrading and Rebooting Only One Routing Engine on page 445](#)

Requirements

This example uses the following hardware and software components:

- MX480 router with dual Routing Engines
- Junos OS Release 13.3R6 as the starting release
- Junos OS Release 14.1R4 as the ending release

Before You Begin

Before you perform a unified ISSU, be sure you:

- Perform a compatibility check to ensure that the software and hardware components and the configuration on the device support unified ISSU by using the [request system software validate in-service-upgrade](#) command
- Read the chapter “[Unified ISSU System Requirements](#)” on [page 401](#) to anticipate any special circumstances that might affect your upgrade.
 - Verify that your platform supports the unified ISSU feature.
 - Verify that the field-replaceable units (FRUs) installed in your platform support the unified ISSU feature or that you can accept the results of performing the upgrade with some FRUs that do not support unified ISSU.
 - Verify that the protocols and features configured on your platform support the unified ISSU feature or that you can accept the results of performing the upgrade with some protocols and features that do not support unified ISSU.
- Download the software package from the Juniper Networks Support website at <https://www.juniper.net/support/> and place the package on your local server.



BEST PRACTICE: When you access the Download Software web page for your device, record the md5 checksum. After downloading the software package to your device, confirm that it is not modified in any way by using the file checksum md5 command. For more information about verifying the md5 checksum, see <https://kb.juniper.net/InfoCenter/index?page=content&id=KB17665>.



NOTE: Starting with Junos OS Release 16.1R1, while performing a unified ISSU from a FreeBSD 6.1 based Junos OS to an upgraded FreeBSD 10.x based Junos OS, the configuration must be validated on a remote host or on a routing engine. The remote host or the routing engine must be running a Junos OS with an upgraded FreeBSD. In addition, only a few selected directories and files will be preserved while upgrading from FreeBSD 6.1 based Junos OS to FreeBSD 10.x based Junos OS. See [Upgrading Junos OS with Upgraded FreeBSD](#) and [request system software validate on \(Junos OS with Upgraded FreeBSD\)](#)

Overview


This procedure can be used to upgrade M Series, T Series, MX Series, EX Series, and PTX Series devices that have dual Routing Engines installed and support unified ISSU.

In the example, the hostnames, filenames, and FRUs are representational. When you perform the procedure on your device, the hostnames, filenames, and FRUs are different. The command output is truncated to only show the text of interest in this procedure.

Topology

Figure 32 on page 425 shows the topology used in this example.

Figure 32: Unified ISSU Example Topology



Configuration

There are variations of the procedure depending on if you want to install the new software on one or both Routing Engines and if you want to automatically reboot both Routing Engines or manually reboot one of the Routing Engines.

In all cases, you must verify that dual Routing Engines are installed and that graceful Routing Engine switchover (GRES) and nonstop active routing (NSR) are enabled. We recommend that you back up the device software before the upgrade.

To perform a unified ISSU, select the appropriate tasks from the following list:

- [Verifying Dual Routing Engines and Enabling GRES and NSR on page 426](#)
- [Verifying the Software Versions and Backing Up the Device Software on page 428](#)
- [Adjusting Timers and Changing Feature-Specific Configuration on page 429](#)
- [Upgrading and Rebooting Both Routing Engines Automatically on page 431](#)
- [Restoring Feature-Specific Configuration on page 436](#)
- [Upgrading Both Routing Engines and Rebooting the New Backup Routing Engine Manually on page 438](#)
- [Upgrading and Rebooting Only One Routing Engine on page 445](#)

Verifying Dual Routing Engines and Enabling GRES and NSR

Step-by-Step Procedure Enabling GRES and NSR is required regardless of which variation of the unified ISSU procedure you use.

To verify that your device has dual Routing Engines and to enable GRES and NSR:

1. Log in to your device.
2. Verify that dual Routing Engines are installed in your device by using the **show chassis hardware** command.

```
user@host> show chassis hardware
```

Routing Engine 0	REV 01	740-051822	9013086837	RE-S-1800x4
Routing Engine 1	REV 01	740-051822	9013086740	RE-S-1800x4

The command output contains lines listing Routing Engine 0 and Routing Engine 1.

3. By default, GRES is disabled; if you have not already done so, enable GRES by including the **graceful-switchover** statement at the **[edit chassis redundancy]** hierarchy level on the master Routing Engine.

```
[edit ]
```

```
user@host# set chassis redundancy graceful-switchover
```


4. By default, NSR is disabled; if you have not already done so, enable NSR by including the **nonstop-routing** statement at the **[edit routing-options]** hierarchy level.

```
[edit]
user@host# set routing-options nonstop-routing
```

5. When you configure NSR, you must also include the **commit synchronize** statement at the **[edit system]** hierarchy level so that configuration changes are synchronized on both Routing Engines.

```
[edit]
user@host# set system commit synchronize
```

6. After you have verified your configuration and are satisfied with it, commit the changes by using the **commit** command.

```
[edit]
user@host# commit
commit complete
```

When you enable GRES and commit the configuration, the CLI prompt changes to indicate which Routing Engine you are using. For example:

```
{master} [edit]
user@host#
```

7. Exit configuration mode by using the **exit** command.

```
{master} [edit]
user@host# exit
Exiting configuration mode
```

8. Verify that NSR is configured on the master Routing Engine (**re0**) by using the **show task replication** command.

```
{master}
user@host> show task replication
```

```
Stateful Replication: Enabled
RE mode: Master
```

Protocol	Synchronization Status
OSPF	Complete
IS-IS	Complete

In the output, verify that the **Synchronization Status** field displays **Complete**.

9. Verify that GRES is enabled on the backup Routing Engine (**re1**) by using the **show system switchover** command.

```
user@host> request routing-engine login re1
{backup}
user@host> show system switchover

Graceful switchover: On
Configuration database: Ready
Kernel database: Ready
Peer state: Steady State
```

In the output, verify that the **Graceful switchover** field state displays **On**. For more information about the **show system switchover** command, see [show system switchover](#).

Verifying the Software Versions and Backing Up the Device Software

Step-by-Step Procedure Unified ISSU requires that both Routing Engines are running the same version of Junos OS before the upgrade. As a preventive measure in case any problems occur during an upgrade, it is a best practice to back up the system software to the device hard disk.

To verify the software versions and back up the device software:

1. Verify that the same version of Junos OS is installed and running on both Routing Engines by using the **show version** command.

```
{backup}
user@host> show version invoke-on all-routing-engines

re0:
-----
Hostname: host
Model: mx480
Junos: 13.3R6.5
JUNOS Base OS boot [13.3R6.5]
JUNOS Base OS Software Suite [13.3R6.5]
JUNOS 64-bit Kernel Software Suite [13.3R6.5]
JUNOS Crypto Software Suite [13.3R6.5]
JUNOS Packet Forwarding Engine Support (M/T/EX Common) [13.3R6.5]
JUNOS Packet Forwarding Engine Support (MX Common) [13.3R6.5]
JUNOS Online Documentation [13.3R6.5]

re1:
-----
Hostname: host
Model: mx480
Junos: 13.3R6.5
JUNOS Base OS boot [13.3R6.5]
JUNOS Base OS Software Suite [13.3R6.5]
JUNOS 64-bit Kernel Software Suite [13.3R6.5]
JUNOS Crypto Software Suite [13.3R6.5]
JUNOS Packet Forwarding Engine Support (M/T/EX Common) [13.3R6.5]
JUNOS Packet Forwarding Engine Support (MX Common) [13.3R6.5]
JUNOS Online Documentation [13.3R6.5]
```

2. Back up the system software to the device hard disk by using the **request system snapshot** command on *each* Routing Engine.



NOTE: The root file system is backed up to /altroot, and /config is backed up to /altconfig. After you issue the `request system snapshot` command, the device flash and hard disks are identical. You can return to the previous version of the software only by booting the device from removable media.

```
{backup}
user@host> request system snapshot
user@host> request routing-engine login re0
{master}
user@host> request system snapshot
```

Adjusting Timers and Changing Feature-Specific Configuration

Step-by-Step Procedure If you have any of the following feature-specific configuration on your device, perform the appropriate steps.

To adjust timers and change feature-specific configuration:

1. Bidirectional Forwarding Detection (BFD) sessions temporarily increase their detection and transmission timers during unified ISSU procedures. After the upgrade, these timers revert to the values in use before the unified ISSU started.

If BFD is enabled on your device and you want to disable the BFD timer negotiation during the unified ISSU, include the `no-issu-timer-negotiation` statement at the `[edit protocols bfd]` hierarchy level.

```
{master} [edit]
user@host# set protocols bfd no-issu-timer-negotiation
```



NOTE: If you include this statement, the BFD timers maintain their original values during the unified ISSU, and the BFD sessions might flap during the unified ISSU or Routing Engine switchover, depending on the detection intervals.

2. If proxy ARP is enabled on your M Series, MX Series, or EX 9200 Series device, remove the `unconditional-src-learn` statement from the `[edit interfaces interface-name unit 0 family inet]` hierarchy level.

By default the statement is not included. This example shows the ge-0/0/1 interface only.

```
{master} [edit]
user@host# delete interfaces ge-0/0/1 unit 0 family inet unconditional-src-learn
```

3. If LACP is enabled on your PTX Series device, remove the **lACP** statement from the **[edit interfaces *interface-name* aggregated-ether-options]** hierarchy level.

```
{master} [edit]
user@host# delete interfaces aex aggregated-ether-options lACP
```

4. If ATM Point-to-Point Protocol (PPP) is enabled on your M Series or T Series device, set the keepalive interval to 10 seconds or greater.

PPP requires three keepalives to fail before it brings down the session. Thirty seconds (10 seconds x three) provides a safe margin to maintain PPP sessions in case of any traffic loss during the unified ISSU operation.

This example shows the at-0/0/1 interface only.

```
{master} [edit]
user@host# set interfaces at-0/0/1 unit 0 keepalives interval 10
```

5. If ATM OAM is enabled on your M Series or T Series device, set the OAM F5 loopback cell period to 20 seconds or greater to maintain ATM connectivity across the unified ISSU.

Include the **oam-period** statement at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level and specify 20 seconds. This example shows the at-0/0/1 interface only.

```
{master} [edit]
user@host# set interfaces at-0/0/1 unit 0 oam-period 20
```

6. After you have verified your configuration and are satisfied with it, commit the changes by using the **commit** command.

```
{master} [edit]
user@host# commit
commit complete
```

7. Exit configuration mode by using the **exit** command.

```
{master} [edit]
user@host# exit
{master}
user@host>
```

Upgrading and Rebooting Both Routing Engines Automatically

Step-by-Step Procedure In this procedure, both Routing Engines automatically reboot. Rebooting both Routing Engines automatically is the most common scenario. Variations to this procedure are described in other sections.

Table 30 on page 431 shows the Routing Engine status prior to starting the unified ISSU.

Table 30: Routing Engine Status Before Upgrading

RE0	RE1
Master	Backup
Old software version installed	Old software version installed
Old software version running	Old software version running

To upgrade and reboot both Routing Engines automatically:

1. Copy the Junos OS software package to the device by using the **file copy** `ftp://username@hostname.net/filename /var/tmp/filename` command.

We recommend that you copy the package to the `/var/tmp` directory, which is a large file system on the hard disk.

```
{master}
user@host> file copy
ftp://myid@myhost.mydomain.net/jinstall64-14.1R4.10-domestic-signed.tgz
/var/tmp/jinstall64-14.1R4.10-domestic-signed.tgz
```



BEST PRACTICE: When you access the Download Software web page for your device, record the md5 checksum. After downloading the software package to your device, confirm that it is not modified in any way by using the `file checksum md5` command. For more information about verifying the md5 checksum, see <https://kb.juniper.net/InfoCenter/index?page=content&id=KB17665>.

2. On the master Routing Engine, start the upgrade by using the **request system software in-service-upgrade package-name reboot** command.



NOTE: Do not try running any additional commands until after the Connection closed message is displayed and your session is disconnected.

```
{master}
```

```
user@host> request system software in-service-upgrade
/var/tmp/jinstall64-14.1R4.10-domestic-signed.tgz reboot

Chassis ISSU Check Done
ISSU: Validating Image
FPC 0 will be offlined (In-Service-Upgrade not supported)
PIC 0/0 will be offlined (In-Service-Upgrade not supported)
PIC 0/1 will be offlined (In-Service-Upgrade not supported)
Do you want to continue with these actions being taken ? [yes,no] (no) yes

Checking compatibility with configuration
Initializing...
Using jbase-13.3R6.5
Verified manifest signed by PackageProductionEc_2015
Using /var/tmp/jinstall64-14.1R4.10-domestic-signed.tgz
Verified jinstall64-14.1R4.10-domestic.tgz signed by PackageProductionEc_2015
Using jinstall64-14.1R4.10-domestic.tgz
Using jbundle64-14.1R4.10-domestic.tgz
Checking jbundle requirements on /
Using jbase-14.1R4.10.tgz
Verified manifest signed by PackageProductionEc_2015
Verified jbase-14.1R4.10 signed by PackageProductionEc_2015
Using /var/v/c/tmp/jbundle/jboot-14.1R4.10.tgz
Using jcrypto64-14.1R4.10.tgz
Verified manifest signed by PackageProductionEc_2015
Verified jcrypto64-14.1R4.10 signed by PackageProductionEc_2015
Using jdocs-14.1R4.10.tgz
Verified manifest signed by PackageProductionEc_2015
Verified jdocs-14.1R4.10 signed by PackageProductionEc_2015
Using jkernel64-14.1R4.10.tgz
Using jpfe-14.1R4.10.tgz
Verified SHA1 checksum of jpfe-M10-14.1R4.10.tgz
Verified SHA1 checksum of jpfe-M120-14.1R4.10.tgz
Verified SHA1 checksum of jpfe-M160-14.1R4.10.tgz
Verified SHA1 checksum of jpfe-M320-14.1R4.10.tgz
Verified SHA1 checksum of jpfe-M40-14.1R4.10.tgz
Verified SHA1 checksum of jpfe-M7i-14.1R4.10.tgz
Verified SHA1 checksum of jpfe-T-14.1R4.10.tgz
Verified SHA1 checksum of jpfe-X2000-14.1R4.10.tgz
Verified SHA1 checksum of jpfe-X960-14.1R4.10.tgz
Verified SHA1 checksum of jpfe-common-14.1R4.10.tgz
Using jplatform-14.1R4.10.tgz
Verified manifest signed by PackageProductionEc_2015
Verified jplatform-14.1R4.10 signed by PackageProductionEc_2015
Using jroute-14.1R4.10.tgz
Verified manifest signed by PackageProductionEc_2015
Verified jroute-14.1R4.10 signed by PackageProductionEc_2015
Using jruntime-14.1R4.10.tgz
Verified manifest signed by PackageProductionEc_2015
Verified jruntime-14.1R4.10 signed by PackageProductionEc_2015
Using jruntime64-14.1R4.10.tgz
Verified manifest signed by PackageProductionEc_2015
Verified jruntime64-14.1R4.10 signed by PackageProductionEc_2015
Using jservices-14.1R4.10.tgz
Using jservices-crypto-14.1R4.10.tgz
Hardware Database regeneration succeeded
Validating against /config/juniper.conf.gz
mgd: commit complete
Validation succeeded
ISSU: Preparing Backup RE
Pushing /var/tmp/jinstall64-14.1R4.10-domestic-signed.tgz to
```

```

re1:/var/tmp/jinstall64-14.1R4.10-domestic-signed.tgz
Installing package '/var/tmp/jinstall64-14.1R4.10-domestic-signed.tgz' ...
Verified jinstall64-14.1R4.10-domestic.tgz signed by PackageProductionEc_2015
Verified jinstall64-14.1R4.10-domestic.tgz signed by PackageProductionRSA_2015
Adding jinstall64...
Verified manifest signed by PackageProductionEc_2015

WARNING:      This package will load JUNOS 14.1R4.10 software.
WARNING:      It will save JUNOS configuration files, and SSH keys
WARNING:      (if configured), but erase all other files and information
WARNING:      stored on this machine. It will attempt to preserve dumps
WARNING:      and log files, but this can not be guaranteed. This is the
WARNING:      pre-installation stage and all the software is loaded when
WARNING:      you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in
/var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

WARNING:      A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING:      'request system reboot' command when software installation is
WARNING:      complete. To abort the installation, do not reboot your system,
WARNING:      instead use the 'request system software delete jinstall'
WARNING:      command as soon as this operation completes.

Saving state for rollback ...
Backup upgrade done
Rebooting Backup RE

Rebooting re1
ISSU: Backup RE Prepare Done
Waiting for Backup RE reboot
GRES operational
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item           Status           Reason
  FPC 0          Offline          Offlined by cli command
Resolving mastership...
Complete. The other routing engine becomes the master.
ISSU: RE switchover Done
ISSU: Upgrading Old Master RE
Installing package '/var/tmp/jinstall64-14.1R4.10-domestic-signed.tgz' ...
Verified jinstall64-14.1R4.10-domestic.tgz signed by PackageProductionEc_2015
Verified jinstall64-14.1R4.10-domestic.tgz signed by PackageProductionRSA_2015
Adding jinstall64...
Verified manifest signed by PackageProductionEc_2015

WARNING:      This package will load JUNOS 14.1R4.10 software.
WARNING:      It will save JUNOS configuration files, and SSH keys
WARNING:      (if configured), but erase all other files and information
WARNING:      stored on this machine. It will attempt to preserve dumps
WARNING:      and log files, but this can not be guaranteed. This is the
WARNING:      pre-installation stage and all the software is loaded when

```

```

WARNING:      you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in
/var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

WARNING:      A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING:      'request system reboot' command when software installation is
WARNING:      complete. To abort the installation, do not reboot your system,
WARNING:      instead use the 'request system software delete jinstall'
WARNING:      command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall64-14.1R4.10-domestic-signed.tgz
...
Saving state for rollback ...
ISSU: Old Master Upgrade Done
ISSU: IDLE
Shutdown NOW!
[pid 10149]

{backup}
user@host>

{backup}
user@host>
*** FINAL System shutdown message from user@host ***

System going down IMMEDIATELY

Connection closed by foreign host.

```

When the Routing Engine that was previously the master is rebooted, you are logged out of the device.

- Wait a few minutes and then log in to the device again.

[Table 31 on page 434](#) shows the Routing Engine status after the unified ISSU.

Table 31: Routing Engine Status After Upgrading and Rebooting Both Routing Engines

RE0	RE1
Backup	Master
New software version installed	New software version installed
New software version running	New software version running

You are logged in to the new backup Routing Engine (**re0**).

- Verify that both Routing Engines have been upgraded by using the **show version** command.


```
{backup}
user@host> show version invoke-on all-routing-engines

re0:
-----
Hostname: host
Model: mx480
Junos: 14.1R4.10
JUNOS Base OS boot [14.1R4.10]
JUNOS Base OS Software Suite [14.1R4.10]
JUNOS Packet Forwarding Engine Support (M/T/EX Common) [14.1R4.10]
JUNOS Packet Forwarding Engine Support (MX Common) [14.1R4.10]
JUNOS platform Software Suite [14.1R4.10]
JUNOS Runtime Software Suite [14.1R4.10]
JUNOS Online Documentation [14.1R4.10]

re1:
-----
Hostname: host
Model: mx480
Junos: 14.1R4.10
JUNOS Base OS boot [14.1R4.10]
JUNOS Base OS Software Suite [14.1R4.10]
JUNOS Packet Forwarding Engine Support (M/T/EX Common) [14.1R4.10]
JUNOS Packet Forwarding Engine Support (MX Common) [14.1R4.10]
JUNOS platform Software Suite [14.1R4.10]
JUNOS Runtime Software Suite [14.1R4.10]
JUNOS Online Documentation [14.1R4.10]
```

5. If you want to, you can optionally display the unified ISSU log messages by using the **show log messages** command.
6. If you want to, you can optionally make **re0** the master Routing Engine by using the **request chassis routing-engine master acquire** command.

```
{backup}
user@host> request chassis routing-engine master acquire

Attempt to become the master routing engine ? [yes,no] (no) yes

Resolving mastership...
Complete. The local routing engine becomes the master.

{master}
user@host>
```

[Table 32 on page 435](#) shows the Routing Engine status after Step 5 is completed.

Table 32: Routing Engine Status After Upgrading, Rebooting, and Switching Mastership

RE0	RE1
Master	Backup
New software version installed	New software version installed
New software version running	New software version running

7. Perform the applicable steps in “[Restoring Feature-Specific Configuration](#)” on [page 436](#).
8. If you are satisfied with the results of your testing, you can optionally back up the system software to the device's hard disk by using the **request system snapshot** command on *each* Routing Engine.



NOTE: The root file system is backed up to /altroot, and /config is backed up to /altconfig. After you issue the **request system snapshot** command, you cannot easily return to the previous version of the software, because the device flash and hard disks are identical. To return to the previous version of the software, you must boot the device from removable media.

```
{master}
user@host> request system snapshot
user@host> request routing-engine login re1
{backup}
user@host> request system snapshot
```

Restoring Feature-Specific Configuration

Step-by-Step Procedure If you have any of the following feature-specific configuration on your device, perform the appropriate steps.

To restore feature-specific configuration:

1. If BFD is enabled on your device and you previously disabled the BFD timer negotiation, delete the **no-issu-timer-negotiation** statement at the **[edit protocols bfd]** hierarchy level.

```
{master} [edit]
user@host# delete protocols bfd no-issu-timer-negotiation
```

2. If proxy ARP is enabled on your M Series, MX Series, or EX9200 device and you previously removed the **unconditional-src-learn** statement, include the statement again.

This example shows the ge-0/0/1 interface only.

```
{master} [edit]
user@host# set interfaces ge-0/0/1 unit 0 family inet unconditional-src-learn
```

3. If LACP is enabled on your PTX Series device and you previously removed the **lACP** statement, include the statement again.

```
{master} [edit]
```

```
user@host# set interfaces aex aggregated-ether-options lacp
```

4. If ATM PPP is enabled on your M Series or T Series device and you previously set the keepalive interval to 10 seconds or greater, restore the original value.

This example shows the at-0/0/1 interface only and shows the interval being set to the default 3 seconds.

```
{master} [edit]
user@host# set interfaces at-0/0/1 unit 0 keepalives interval 3
```

5. If ATM OAM is enabled on your M Series or T Series device and you previously set the OAM F5 loopback cell period to 20 seconds or greater, change the configuration back to the original value.

This example shows the at-0/0/1 interface only and shows the period being set to 10 seconds.

```
{master} [edit]
user@host# set interfaces at-0/0/1 unit 0 oam-period 10
```

6. After you have verified your configuration and are satisfied with it, commit the changes by using the **commit** command.

```
{master} [edit]
user@host# commit
commit complete
```

7. Exit configuration mode by using the **exit** command.

```
{master} [edit]
user@host# exit
{master}
user@host>
```

Upgrading Both Routing Engines and Rebooting the New Backup Routing Engine Manually

Step-by-Step Procedure In certain circumstances, you might want to install the new software on only one Routing Engine and reboot only the master until after you can test the new software. A Routing Engine does not start running the new software until after it is rebooted.

The advantage is if the results of your testing requires you to downgrade the software, you can switch Routing Engines to run the old software on one Routing Engine and then install the old software on the other Routing Engine. This is not the typical scenario.

To upgrade both Routing Engines and to reboot the new backup Routing Engine manually:

1. Perform the steps in “Verifying Dual Routing Engines and Enabling GRES and NSR” on page 426.
2. Perform the steps in “Verifying the Software Versions and Backing Up the Device Software” on page 428.
3. Perform the steps in “Adjusting Timers and Changing Feature-Specific Configuration” on page 429.
4. Copy the Junos OS software package to the device using the **file copy** `ftp://username@hostname.net/filename /var/tmp/filename` command.

We recommend that you copy the package to the `/var/tmp` directory, which is a large file system on the hard disk.

```
{master}
user@host> file copy
ftp://myid@myhost.mydomain.net/jinstall64-14.1R4.10-domestic-signed.tgz
/var/tmp/jinstall64-14.1R4.10-domestic-signed.tgz
```



BEST PRACTICE: When you access the Download Software web page for your device, record the md5 checksum. After downloading the software package to your device, confirm that it is not modified in any way by using the file checksum md5 command. For more information about verifying the md5 checksum, see <https://kb.juniper.net/InfoCenter/index?page=content&id=KB17665>.

Table 33 on page 438 shows the Routing Engine status prior to starting the unified ISSU.

Table 33: Routing Engine Status Before Upgrading and Manually Rebooting the Backup Routing Engine

RE0	RE1
Master	Backup

Table 33: Routing Engine Status Before Upgrading and Manually Rebooting the Backup Routing Engine (continued)

RE0	RE1
Old software version installed	Old software version installed
Old software version running	Old software version running

- On the master Routing Engine, start the upgrade by using the **request system software in-service-upgrade package-name** command without the reboot option.

```
{master}
user@host> request system software in-service-upgrade
/var/tmp/jinstall64-14.1R4.10-domestic-signed.tgz

Chassis ISSU Check Done
ISSU: Validating Image
FPC 0 will be offlined (In-Service-Upgrade not supported)
PIC 0/0 will be offlined (In-Service-Upgrade not supported)
PIC 0/1 will be offlined (In-Service-Upgrade not supported)
Do you want to continue with these actions being taken ? [yes,no] (no) yes

Checking compatibility with configuration
Initializing...
Using jbase-13.3R6.5
Verified manifest signed by PackageProductionEc_2015
Using /var/tmp/jinstall64-14.1R4.10-domestic-signed.tgz
Verified jinstall64-14.1R4.10-domestic.tgz signed by PackageProductionEc_2015
Using jinstall64-14.1R4.10-domestic.tgz
Using jbundle64-14.1R4.10-domestic.tgz
Checking jbundle requirements on /
Using jbase-14.1R4.10.tgz
Verified manifest signed by PackageProductionEc_2015
Verified jbase-14.1R4.10 signed by PackageProductionEc_2015
Using /var/v/c/tmp/jbundle/jboot-14.1R4.10.tgz
Using jcrypto64-14.1R4.10.tgz
Verified manifest signed by PackageProductionEc_2015
Verified jcrypto64-14.1R4.10 signed by PackageProductionEc_2015
Using jdocs-14.1R4.10.tgz
Verified manifest signed by PackageProductionEc_2015
Verified jdocs-14.1R4.10 signed by PackageProductionEc_2015
Using jkernel64-14.1R4.10.tgz
Using jpfe-14.1R4.10.tgz
Verified SHA1 checksum of jpfe-M10-14.1R4.10.tgz
Verified SHA1 checksum of jpfe-M120-14.1R4.10.tgz

Verified SHA1 checksum of jpfe-M160-14.1R4.10.tgz
Verified SHA1 checksum of jpfe-M320-14.1R4.10.tgz
Verified SHA1 checksum of jpfe-M40-14.1R4.10.tgz
Verified SHA1 checksum of jpfe-M7i-14.1R4.10.tgz
Verified SHA1 checksum of jpfe-T-14.1R4.10.tgz
Verified SHA1 checksum of jpfe-X2000-14.1R4.10.tgz
Verified SHA1 checksum of jpfe-X960-14.1R4.10.tgz
Verified SHA1 checksum of jpfe-common-14.1R4.10.tgz
Using jplatform-14.1R4.10.tgz
Verified manifest signed by PackageProductionEc_2015
Verified jplatform-14.1R4.10 signed by PackageProductionEc_2015
```

```

Using jroute-14.1R4.10.tgz
Verified manifest signed by PackageProductionEc_2015
Verified jroute-14.1R4.10 signed by PackageProductionEc_2015
Using jruntime-14.1R4.10.tgz
Verified manifest signed by PackageProductionEc_2015
Verified jruntime-14.1R4.10 signed by PackageProductionEc_2015
Using jruntime64-14.1R4.10.tgz
Verified manifest signed by PackageProductionEc_2015
Verified jruntime64-14.1R4.10 signed by PackageProductionEc_2015
Using jservices-14.1R4.10.tgz
Using jservices-crypto-14.1R4.10.tgz
Hardware Database regeneration succeeded
Validating against /config/juniper.conf.gz
mgd: commit complete
Validation succeeded
ISSU: Preparing Backup RE
Pushing /var/tmp/jinstall64-14.1R4.10-domestic-signed.tgz to
re1:/var/tmp/jinstall64-14.1R4.10-domestic-signed.tgz
Installing package '/var/tmp/jinstall64-14.1R4.10-domestic-signed.tgz' ...
Verified jinstall64-14.1R4.10-domestic.tgz signed by PackageProductionEc_2015
Verified jinstall64-14.1R4.10-domestic.tgz signed by PackageProductionRSA_2015
Adding jinstall64...
Verified manifest signed by PackageProductionEc_2015

WARNING:      This package will load JUNOS 14.1R4.10 software.
WARNING:      It will save JUNOS configuration files, and SSH keys
WARNING:      (if configured), but erase all other files and information
WARNING:      stored on this machine. It will attempt to preserve dumps
WARNING:      and log files, but this can not be guaranteed. This is the
WARNING:      pre-installation stage and all the software is loaded when
WARNING:      you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in
/var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

WARNING:      A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING:      'request system reboot' command when software installation is
WARNING:      complete. To abort the installation, do not reboot your system,
WARNING:      instead use the 'request system software delete jinstall'
WARNING:      command as soon as this operation completes.

Saving state for rollback ...
Backup upgrade done
Rebooting Backup RE

Rebooting re1
ISSU: Backup RE Prepare Done
Waiting for Backup RE reboot
GRES operational
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item                Status                Reason

```

```

FPC 0          Offline          Offlined by cli command
Resolving mastership...
Complete. The other routing engine becomes the master.
ISSU: RE switchover Done
ISSU: Upgrading Old Master RE
Installing package '/var/tmp/jinstall64-14.1R4.10-domestic-signed.tgz' ...
Verified jinstall64-14.1R4.10-domestic.tgz signed by PackageProductionEc_2015
Verified jinstall64-14.1R4.10-domestic.tgz signed by PackageProductionRSA_2015
Adding jinstall64...
Verified manifest signed by PackageProductionEc_2015

WARNING:      This package will load JUNOS 14.1R4.10 software.
WARNING:      It will save JUNOS configuration files, and SSH keys
WARNING:      (if configured), but erase all other files and information
WARNING:      stored on this machine. It will attempt to preserve dumps
WARNING:      and log files, but this can not be guaranteed. This is the
WARNING:      pre-installation stage and all the software is loaded when
WARNING:      you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in
/var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

WARNING:      A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING:      'request system reboot' command when software installation is
WARNING:      complete. To abort the installation, do not reboot your system,
WARNING:      instead use the 'request system software delete jinstall'
WARNING:      command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall64-14.1R4.10-domestic-signed.tgz
...
Saving state for rollback ...
ISSU: Old Master Upgrade Done
ISSU: IDLE

```

Table 34 on page 441 shows the Routing Engine status after the unified ISSU and before manually rebooting the backup Routing Engine.

Table 34: Routing Engine Status After Upgrading and Before Manually Rebooting the Backup Routing Engine

RE0	RE1
Backup	Master
New software version installed	New software version installed
Old software version running	New software version running

- Verify that the new backup, (old master) Routing Engine (**re0**), is still running the previous software image and that the new master Routing Engine (**re1**) is running the new software image, by using the **show version** command.

```
{backup}
```

```

user@host> show version invoke-on all-routing-engines

re0:
-----
Hostname: host
Model: mx480
Junos: 13.3R6.5
JUNOS Base OS boot [13.3R6.5]
JUNOS Base OS Software Suite [13.3R6.5]
JUNOS 64-bit Kernel Software Suite [13.3R6.5]
JUNOS Crypto Software Suite [13.3R6.5]
JUNOS Packet Forwarding Engine Support (M/T/EX Common) [13.3R6.5]
JUNOS Packet Forwarding Engine Support (MX Common) [13.3R6.5]
JUNOS Online Documentation [13.3R6.5]

re1:
-----
Hostname: host
Model: mx480
Junos: 14.1R4.10
JUNOS Base OS boot [14.1R4.10]
JUNOS Base OS Software Suite [14.1R4.10]
JUNOS Packet Forwarding Engine Support (M/T/EX Common) [14.1R4.10]
JUNOS Packet Forwarding Engine Support (MX Common) [14.1R4.10]
JUNOS platform Software Suite [14.1R4.10]
JUNOS Runtime Software Suite [14.1R4.10]
JUNOS Online Documentation [14.1R4.10]

```

7. At this point, if you do not want to install the newer software version on the new backup Routing Engine (re0), issue the **request system software delete *package-name*** command on it.

Otherwise, to complete the upgrade, go to the next step.

8. Reboot the new backup Routing Engine (re0) by issuing the **request system reboot** command.

```

{backup}
user@host> request system reboot

Reboot the system ? [yes,no] (no) yes

*** FINAL System shutdown message from remote@host ***

System going down IMMEDIATELY

Shutdown NOW!
[pid 38432]

{backup}
user@home> Connection closed by foreign host.

```

If you are not on the console port, you are disconnected from the device session.

[Table 35 on page 443](#) shows the Routing Engine status after the unified ISSU, after rebooting the backup Routing Engine, but before switching mastership.

Table 35: Routing Engine Status After Upgrading, Manually Rebooting, and Before Switching Mastership

RE0	RE1
Backup	Master
New software version installed	New software version installed
New software version running	New software version running

9. Wait a few minutes, then log in to the device again.
You are logged in to the new backup Routing Engine (**re0**).
10. Verify that both Routing Engines have been upgraded by using the **show version** command.

```
{backup}
user@host> show version invoke-on all-routing-engines

re0:
-----
Hostname: host
Model: mx480
Junos: 14.1R4.10
JUNOS Base OS boot [14.1R4.10]
JUNOS Base OS Software Suite [14.1R4.10]
JUNOS Packet Forwarding Engine Support (M/T/EX Common) [14.1R4.10]
JUNOS Packet Forwarding Engine Support (MX Common) [14.1R4.10]
JUNOS platform Software Suite [14.1R4.10]
JUNOS Runtime Software Suite [14.1R4.10]
JUNOS Online Documentation [14.1R4.10]

re1:
-----
Hostname: host
Model: mx480
Junos: 14.1R4.10
JUNOS Base OS boot [14.1R4.10]
JUNOS Base OS Software Suite [14.1R4.10]
JUNOS Packet Forwarding Engine Support (M/T/EX Common) [14.1R4.10]
JUNOS Packet Forwarding Engine Support (MX Common) [14.1R4.10]
JUNOS platform Software Suite [14.1R4.10]
JUNOS Runtime Software Suite [14.1R4.10]
JUNOS Online Documentation [14.1R4.10]
```

11. If you want to, you can optionally display the unified ISSU log messages by using the **show log messages** command.
12. If you want to, you can optionally make **re0** the master Routing Engine by using the **request chassis routing-engine master acquire** command:

```
{backup}
```

```

user@host> request chassis routing-engine master acquire
Attempt to become the master routing engine ? [yes,no] (no) yes

Resolving mastership...
Complete. The local routing engine becomes the master.

{master}
user@host>

```

Table 36 on page 444 shows the Routing Engine status after the unified ISSU, after rebooting the backup Routing Engine, and after switching mastership.

Table 36: Routing Engine Status After Upgrading, Manually Rebooting, and Switching Mastership

RE0	RE1
Master	Backup
New software version installed	New software version installed
New software version running	New software version running

13. Perform the applicable steps in “[Restoring Feature-Specific Configuration](#)” on [page 436](#).
14. If you are satisfied with the results of your testing, you can optionally back up the system software to the device’s hard disk by using the **request system snapshot** command on *each* Routing Engine.



NOTE: The root file system is backed up to /altroot, and /config is backed up to /altconfig. After you issue the **request system snapshot** command, you cannot easily return to the previous version of the software, because the device flash and hard disks are identical. To return to the previous version of the software, you must boot the device from removable media.

```

{master}
user@host> request system snapshot
user@host> request routing-engine login re1
{backup}
user@host> request system snapshot

```

Upgrading and Rebooting Only One Routing Engine

Step-by-Step Procedure In certain circumstances you might want to install the new software on only one Routing Engine.

The advantage is if the results of your testing requires you to downgrade the software, you can switch Routing Engines to run the old software on one Routing Engine and then install the old software on the other Routing Engine. This is not the typical scenario.

Table 37 on page 445 shows the Routing Engine status prior to starting the unified ISSU.

Table 37: Routing Engine Status Before Upgrading and Rebooting One Routing Engine

RE0	RE1
Master	Backup
Old software version installed	Old software version installed
Old software version running	Old software version running

To upgrade and rebooting only one Routing Engine:

1. Perform the steps in “Verifying Dual Routing Engines and Enabling GRES and NSR” on page 426.
2. Perform the steps in “Verifying the Software Versions and Backing Up the Device Software” on page 428.
3. Perform the applicable steps in “Adjusting Timers and Changing Feature-Specific Configuration” on page 429.
4. Copy the Junos OS software package to the device by using the **file copy** `ftp://username@hostname.net/filename /var/tmp/filename` command.

We recommend that you copy the package to the `/var/tmp` directory, which is a large file system on the hard disk.

```
{master}
user@host> file copy
ftp://myid@myhost.mydomain.net/jinstall64-14.1R4.10-domestic-signed.tgz
/var/tmp/jinstall64-14.1R4.10-domestic-signed.tgz
```



BEST PRACTICE: When you access the Download Software web page for your device, record the md5 checksum. After downloading the software package to your device, confirm that it is not modified in any way by using the file checksum md5 command. For more information about verifying the md5 checksum, see <https://kb.juniper.net/InfoCenter/index?page=content&id=KB17665>.

5. On the master Routing Engine, start the upgrade by using the **request system software in-service-upgrade package-name no-old-master-upgrade** command.

```
{master}
user@host> request system software in-service-upgrade
/var/tmp/jinstall64-14.1R4.10-domestic-signed.tgz no-old-master-upgrade

Chassis ISSU Check Done
ISSU: Validating Image
FPC 0 will be offlined (In-Service-Upgrade not supported)
PIC 0/0 will be offlined (In-Service-Upgrade not supported)
PIC 0/1 will be offlined (In-Service-Upgrade not supported)
Do you want to continue with these actions being taken ? [yes,no] (no) yes

Checking compatibility with configuration
Initializing...
Using jbase-13.3R6.5
Verified manifest signed by PackageProductionEc_2015
Using /var/tmp/jinstall64-14.1R4.10-domestic-signed.tgz
Verified jinstall64-14.1R4.10-domestic.tgz signed by PackageProductionEc_2015
Using jinstall64-14.1R4.10-domestic.tgz
Using jbundle64-14.1R4.10-domestic.tgz
Checking jbundle requirements on /
Using jbase-14.1R4.10.tgz
Verified manifest signed by PackageProductionEc_2015
Verified jbase-14.1R4.10 signed by PackageProductionEc_2015
Using /var/v/c/tmp/jbundle/jboot-14.1R4.10.tgz
Using jcrypto64-14.1R4.10.tgz
Verified manifest signed by PackageProductionEc_2015
Verified jcrypto64-14.1R4.10 signed by PackageProductionEc_2015
Using jdocs-14.1R4.10.tgz
Verified manifest signed by PackageProductionEc_2015
Verified jdocs-14.1R4.10 signed by PackageProductionEc_2015
Using jkernel64-14.1R4.10.tgz
Using jpfe-14.1R4.10.tgz
Verified SHA1 checksum of jpfe-M10-14.1R4.10.tgz
Verified SHA1 checksum of jpfe-M120-14.1R4.10.tgz
Verified SHA1 checksum of jpfe-M160-14.1R4.10.tgz
Verified SHA1 checksum of jpfe-M320-14.1R4.10.tgz
Verified SHA1 checksum of jpfe-M40-14.1R4.10.tgz
Verified SHA1 checksum of jpfe-M7i-14.1R4.10.tgz
Verified SHA1 checksum of jpfe-T-14.1R4.10.tgz
Verified SHA1 checksum of jpfe-X2000-14.1R4.10.tgz
Verified SHA1 checksum of jpfe-X960-14.1R4.10.tgz
Verified SHA1 checksum of jpfe-common-14.1R4.10.tgz
Using jplatform-14.1R4.10.tgz
Verified manifest signed by PackageProductionEc_2015
Verified jplatform-14.1R4.10 signed by PackageProductionEc_2015
Using jroute-14.1R4.10.tgz
Verified manifest signed by PackageProductionEc_2015
Verified jroute-14.1R4.10 signed by PackageProductionEc_2015
Using jruntime-14.1R4.10.tgz
Verified manifest signed by PackageProductionEc_2015
Verified jruntime-14.1R4.10 signed by PackageProductionEc_2015
Using jruntime64-14.1R4.10.tgz
Verified manifest signed by PackageProductionEc_2015
Verified jruntime64-14.1R4.10 signed by PackageProductionEc_2015
Using jservices-14.1R4.10.tgz
Using jservices-crypto-14.1R4.10.tgz
Hardware Database regeneration succeeded
Validating against /config/juniper.conf.gz
```

```

mgd: commit complete
Validation succeeded
ISSU: Preparing Backup RE
Pushing /var/tmp/jinstall64-14.1R4.10-domestic-signed.tgz to
re1:/var/tmp/jinstall64-14.1R4.10-domestic-signed.tgz
Installing package '/var/tmp/jinstall64-14.1R4.10-domestic-signed.tgz' ...
Verified jinstall64-14.1R4.10-domestic.tgz signed by PackageProductionEc_2015
Verified jinstall64-14.1R4.10-domestic.tgz signed by PackageProductionRSA_2015
Adding jinstall64...
Verified manifest signed by PackageProductionEc_2015

WARNING:      This package will load JUNOS 14.1R4.10 software.
WARNING:      It will save JUNOS configuration files, and SSH keys
WARNING:      (if configured), but erase all other files and information
WARNING:      stored on this machine. It will attempt to preserve dumps
WARNING:      and log files, but this can not be guaranteed. This is the
WARNING:      pre-installation stage and all the software is loaded when
WARNING:      you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in
/var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

WARNING:      A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING:      'request system reboot' command when software installation is
WARNING:      complete. To abort the installation, do not reboot your system,
WARNING:      instead use the 'request system software delete jinstall'
WARNING:      command as soon as this operation completes.

Saving state for rollback ...
Backup upgrade done
Rebooting Backup RE

Rebooting re1
ISSU: Backup RE Prepare Done
Waiting for Backup RE reboot
GRES operational
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item          Status          Reason
  FPC 0         Offline         Offlined by cli command
Resolving mastership...
Complete. The other routing engine becomes the master.
ISSU: RE switchover Done
Skipping Old Master Upgrade
ISSU: IDLE

```

Table 38 on page 448 shows the Routing Engine status after the unified ISSU upgrades the master Routing Engine but before the backup Routing Engine is upgraded.

Table 38: Routing Engine Status After Upgrading One Routing Engine and Before Upgrading the Other Routing Engine

RE0	RE1
Backup	Master
Old software version installed	New software version installed
Old software version running	New software version running

- Verify that the new backup, (old master) Routing Engine (**re0**), is still running the previous software image and that the new master Routing Engine (**re1**) is running the new software image, by using the **show version** command.

```
{backup}
user@host> show version invoke-on all-routing-engines

re0:
-----
Hostname: host
Model: mx480
Junos: 13.3R6.5
JUNOS Base OS boot [13.3R6.5]
JUNOS Base OS Software Suite [13.3R6.5]
JUNOS 64-bit Kernel Software Suite [13.3R6.5]
JUNOS Crypto Software Suite [13.3R6.5]
JUNOS Packet Forwarding Engine Support (M/T/EX Common) [13.3R6.5]
JUNOS Packet Forwarding Engine Support (MX Common) [13.3R6.5]
JUNOS Online Documentation [13.3R6.5]

re1:
-----
Hostname: host
Model: mx480
Junos: 14.1R4.10
JUNOS Base OS boot [14.1R4.10]
JUNOS Base OS Software Suite [14.1R4.10]
JUNOS Packet Forwarding Engine Support (M/T/EX Common) [14.1R4.10]
JUNOS Packet Forwarding Engine Support (MX Common) [14.1R4.10]
JUNOS platform Software Suite [14.1R4.10]
JUNOS Runtime Software Suite [14.1R4.10]
JUNOS Online Documentation [14.1R4.10]
```

- If your testing is complete and you want to install the new software on the backup Routing Engine, you must first disable GRES and NSR on both Routing Engines and commit the configuration.

```
{backup} [edit ]
user@host# delete chassis redundancy graceful-switchover
user@host# delete routing-options nonstop-routing
user@host# commit
warning: Graceful-switchover is enabled, commit on backup is not recommended
Continue commit on backup RE? [yes,no] (no) yes
re0:
```

```

configuration check succeeds
re1:
commit complete
re0:
commit complete
[edit ]
user@host#

```

8. Install the new software on the backup Routing Engine (re0) by using the **request system software add /var/tmp/jinstall64-14.1R4.10-domestic-signed.tgz** command.

```

user@host> request system software add
/var/tmp/jinstall64-14.1R4.10-domestic-signed.tgz

NOTICE: Validating configuration against
jinstall64-14.1R4.10-domestic-signed.tgz.
NOTICE: Use the 'no-validate' option to skip this if desired.
Checking compatibility with configuration
Initializing...
Using jbase-13.3R6.5
Verified manifest signed by PackageProductionEc_2015
Using /var/tmp/jinstall64-14.1R4.10-domestic-signed.tgz
Verified jinstall64-14.1R4.10-domestic.tgz signed by PackageProductionEc_2015
Using jinstall64-14.1R4.10-domestic.tgz
Using jbundle64-14.1R4.10-domestic.tgz
Checking jbundle requirements on /
Using jbase-14.1R4.10.tgz
Verified manifest signed by PackageProductionEc_2015
Verified jbase-14.1R4.10 signed by PackageProductionEc_2015
Using /var/vc/tmp/jbundle/jboot-14.1R4.10.tgz
Using jcrypto64-14.1R4.10.tgz
Verified manifest signed by PackageProductionEc_2015
Verified jcrypto64-14.1R4.10 signed by PackageProductionEc_2015
Using jdocs-14.1R4.10.tgz
Verified manifest signed by PackageProductionEc_2015
Verified jdocs-14.1R4.10 signed by PackageProductionEc_2015
Using jkernel64-14.1R4.10.tgz
Using jpfe-14.1R4.10.tgz
Verified SHA1 checksum of jpfe-M10-14.1R4.10.tgz
Verified SHA1 checksum of jpfe-M120-14.1R4.10.tgz
Verified SHA1 checksum of jpfe-M160-14.1R4.10.tgz
Verified SHA1 checksum of jpfe-M320-14.1R4.10.tgz
Verified SHA1 checksum of jpfe-M40-14.1R4.10.tgz
Verified SHA1 checksum of jpfe-M7i-14.1R4.10.tgz
Verified SHA1 checksum of jpfe-T-14.1R4.10.tgz
Verified SHA1 checksum of jpfe-X2000-14.1R4.10.tgz
Verified SHA1 checksum of jpfe-X960-14.1R4.10.tgz
Verified SHA1 checksum of jpfe-common-14.1R4.10.tgz
Using jplatform-14.1R4.10.tgz
Verified manifest signed by PackageProductionEc_2015
Verified jplatform-14.1R4.10 signed by PackageProductionEc_2015
Using jroute-14.1R4.10.tgz
Verified manifest signed by PackageProductionEc_2015
Verified jroute-14.1R4.10 signed by PackageProductionEc_2015
Using jruntime-14.1R4.10.tgz
Verified manifest signed by PackageProductionEc_2015
Verified jruntime-14.1R4.10 signed by PackageProductionEc_2015
Using jruntime64-14.1R4.10.tgz

```

```

Verified manifest signed by PackageProductionEc_2015
Verified jruntime64-14.1R4.10 signed by PackageProductionEc_2015
Using jservices-14.1R4.10.tgz
Using jservices-crypto-14.1R4.10.tgz
Hardware Database regeneration succeeded
Validating against /config/juniper.conf.gz
mgd: commit complete
Validation succeeded
Installing package '/var/tmp/jinstall64-14.1R4.10-domestic-signed.tgz' ...
Verified jinstall64-14.1R4.10-domestic.tgz signed by PackageProductionEc_2015
Verified jinstall64-14.1R4.10-domestic.tgz signed by PackageProductionRSA_2015
Adding jinstall64...
Verified manifest signed by PackageProductionEc_2015

WARNING:      This package will load JUNOS 14.1R4.10 software.
WARNING:      It will save JUNOS configuration files, and SSH keys
WARNING:      (if configured), but erase all other files and information
WARNING:      stored on this machine. It will attempt to preserve dumps
WARNING:      and log files, but this can not be guaranteed. This is the
WARNING:      pre-installation stage and all the software is loaded when
WARNING:      you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in
/var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

WARNING:      A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING:      'request system reboot' command when software installation is
WARNING:      complete. To abort the installation, do not reboot your system,
WARNING:      instead use the 'request system software delete jinstall'
WARNING:      command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall64-14.1R4.10-domestic-signed.tgz
...
Saving state for rollback ...

```

9. Reboot **re0** by using the **request system reboot** command.

```

user@host> request system reboot
Reboot the system ? [yes,no] (no) yes

*** FINAL System shutdown message from user@host ***

System going down IMMEDIATELY

Shutdown NOW!
[pid 22857]

user@host> Connection closed by foreign host.

```

If you are not on the console port, you are disconnected from the router session.

10. After waiting a few minutes, log in to the device again.
You are logged in to the backup Routing Engine (**re0**).

11. Verify that both Routing Engines are running the new software image by using the **show version** command.

```
{backup}
user@host> show version invoke-on all-routing-engines

Hostname: host
Model: mx480
Junos: 14.1R4.10
JUNOS Base OS boot [14.1R4.10]
JUNOS Base OS Software Suite [14.1R4.10]
JUNOS Packet Forwarding Engine Support (M/T/EX Common) [14.1R4.10]
JUNOS Packet Forwarding Engine Support (MX Common) [14.1R4.10]
JUNOS platform Software Suite [14.1R4.10]
JUNOS Runtime Software Suite [14.1R4.10]
JUNOS Online Documentation [14.1R4.10]

re1:
-----
Hostname: host
Model: mx480
Junos: 14.1R4.10
JUNOS Base OS boot [14.1R4.10]
JUNOS Base OS Software Suite [14.1R4.10]
JUNOS Packet Forwarding Engine Support (M/T/EX Common) [14.1R4.10]
JUNOS Packet Forwarding Engine Support (MX Common) [14.1R4.10]
JUNOS platform Software Suite [14.1R4.10]
JUNOS Runtime Software Suite [14.1R4.10]
JUNOS Online Documentation [14.1R4.10]
```

12. If you want to, you can optionally display the unified ISSU log messages by using the **show log messages** command.
13. If you want to, make **re0** the master Routing Engine by using the **request chassis routing-engine master acquire** command.

```
{backup}
user@host> request chassis routing-engine master acquire

Attempt to become the master routing engine ? [yes,no] (no) yes

Resolving mastership...
Complete. The local routing engine becomes the master.

user@host>
```

[Table 39 on page 451](#) shows the Routing Engine status after the unified ISSU, after rebooting the backup Routing Engine, and after switching mastership.

Table 39: Routing Engine Status After Upgrading, Manually Rebooting, and Switching Mastership

RE0	RE1
Master	Backup
New software version installed	New software version installed

Table 39: Routing Engine Status After Upgrading, Manually Rebooting, and Switching Mastership (continued)

RE0	RE1
New software version running	New software version running

14. Enable GRES and NSR again by performing the steps in [“Verifying Dual Routing Engines and Enabling GRES and NSR”](#) on page 426.
15. Perform the applicable steps in [“Restoring Feature-Specific Configuration”](#) on page 436.
16. If you are satisfied with the results of your testing, you can optionally back up the system software to the device's hard disk by using the **request system snapshot** command on *each* Routing Engine.



NOTE: The root file system is backed up to /altroot, and /config is backed up to /altconfig. After you issue the **request system snapshot** command, you cannot easily return to the previous version of the software, because the device flash and hard disks are identical. To return to the previous version of the software, you must boot the device from removable media.

```
{master}
user@host> request system snapshot
user@host> request routing-engine login re1
{backup}
user@host> request system snapshot
```

Related Documentation

- [Getting Started with Unified In-Service Software Upgrade on page 387](#)
- [Understanding the Unified ISSU Process on page 388](#)
- [Unified ISSU System Requirements on page 401](#)
- [Best Practices for Performing a Unified ISSU on page 423](#)
- [Verifying a Unified ISSU on page 461](#)
- [Troubleshooting Unified ISSU Problems on page 462](#)
- [Managing and Tracing BFD Sessions During Unified ISSU Procedures on page 462](#)

Performing an In-Service Software Upgrade (ISSU) with Non-Stop Routing

You can use an in-service software upgrade with non-stop routing to upgrade the software running on the switch with minimal traffic disruption during the upgrade.



NOTE: Starting with Junos OS Release 18.2R1 on the QFX5200 switch, we recommend that you wait at least five minutes between in-service software upgrades.



NOTE: Starting with Junos OS Release 17.1R1, on QFX5100 and EX4600 switches, you cannot perform an ISSU from a Junos OS Release earlier than 17.1R1 to Junos OS Release 17.1R1.

This topic covers:

1. [Preparing the Switch for Software Installation on page 453](#)
2. [Upgrading the Software Using ISSU on page 454](#)

Preparing the Switch for Software Installation

Before you begin software installation using ISSU:



NOTE: Before you perform an in-service software upgrade, if applicable, remove the `set system internet-options no-tcp-reset drop-all-tcp` command from the configuration, otherwise the upgrade will fail and an error message will be displayed.

NSB and non-stop routing enable NSB-supported Layer 2 protocols to synchronize protocol information between the master and backup Routing Engines.

- Enable non-stop routing. See [“Configuring Nonstop Active Routing on Switches” on page 220](#) for information on how to enable it.
- Enable nonstop bridging (NSB). See [“Configuring Nonstop Bridging on Switches \(CLI Procedure\)” on page 193](#) for information on how to enable it.
- Configure the Bidirectional Forwarding Detection Protocol (BFD) timeout to be more than one second, otherwise you will receive an error.

Upgrading the Software Using ISSU

This procedure describes how to upgrade the software running on a standalone switch:



NOTE: If the Host OS software needs to be updated, you cannot perform an ISSU. Instead, perform a standard software upgrade.

To upgrade the switch using ISSU:

1. Download the software package by following the procedure in the Downloading Software Files with a Browser section in *Installing Software Packages on QFX Series Devices*.
2. Copy the software package or packages to the switch. We recommend that you copy the file to the `/var/tmp` directory.
3. Log in to the console connection. Using a console connection allows you to monitor the progress of the upgrade.
4. Start the ISSU:
 - On the switch, enter:

```
user@switch> request system software in-service-upgrade
/var/tmp/package-name.tgz
```

where `package-name.tgz` is, for example,
`jinstall-host-qfx-5e-18.1R1-secured-signed.tgz`.



NOTE: During the upgrade, you will not be able to access the Junos OS CLI.

The switch displays status messages similar to the following messages as the upgrade executes:

ISSU: Validating Image

PRE ISSU CHECK:

```
-----
PFE Status                : Online
Member Id zero            : Valid
VC not in mixed or fabric mode : Valid
Member is single node vc   : Valid
BFD minimum-interval check done : Valid
GRES enabled              : Valid
GR enabled                : Valid
drop-all-tcp not configured : Valid
Ready for ISSU            : Valid
```

warning: Do NOT use /user during ISSU. Changes to /user during ISSU may get lost!

```

Pushing Junos image package to the host...
Installing /var/tmp/install-media-qfx-5e-junos-2018-secure.tgz
Extracting the package ...
total 1110328
-rw-r--r-- 1 18735 758 237044439 Oct 26 05:11
jinstall-qfx-5e-junos-2018-secure-linux.tgz
-rw-r--r-- 1 18735 758 899918118 Oct 26 05:11
jinstall-qfx-5e-junos-2018-secure-app.tgz

=====
Current Host kernel version : 3.14.52-rt50-WR7.0.0.9_ovp
Package Host kernel version : 3.14.52-rt50-WR7.0.0.9_ovp
Current Host version       : 3.0.7
Package Host version       : 3.0.7
Min host version required for applications: 3.0.7
Min host version required for in-service-upgrade: 3.0.7
=====

Setting up Junos host applications for in-service-upgrade ...
-----
Running Junos application installer for in-service-upgrade
-----

-----
Installing /var/sw/applications/qfx-5e-flex-2018.tgz
-----
pkg_install_rpms: qfx-5e-base-1.0-0-2018.x86_64.rpm
Installing qfx-5e-control-plane-flex-1.0-0-2018.x86_64.rpm ...
=====
Loading cache...
Updating cache... ##### [100%]

Committing transaction...
Preparing... ##### [
0%]
  1:Installing qfx-5e-contro.. ##### [100%]

Output from qfx-5e-control-plane-flex-1.0-0@x86_64:
-----
Installing JUNOS image: jinstall-jcp-i386-flex-18.12018.img.gz
-----
Extracting jinstall-jcp-i386-flex-18.12018.img.gz to
/recovery/junos/jinstall-jcp-i386-flex-18.12018-2018.img
Prepare host for virtfs...
Integrity check passed for hash-control-plane.md5.

Installing packages (1):
  qfx-5e-control-plane-flex-1.0-0@x86_64

812.9MB of package files are needed. 821.5MB will be used.

Saving cache...

=====
Application installed.
Waiting to sync newly setup VM disk

```

```

VM ready after 200 seconds
[Oct 26 05:19:22]:ISSU: Preparing Backup RE
Prepare for ISSU
[Oct 26 05:19:27]:ISSU: Backup RE Prepare Done
Spawning the backup RE
Spawn backup RE, index 0 successful
Starting secondary dataplane
Second dataplane container started
GRES in progress
Waiting for backup RE switchover ready
GRES operational
Copying home directories
Copying home directories successful
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
[Oct 26 05:28:33]:ISSU: Preparing Daemons
[Oct 26 05:28:39]:ISSU: Daemons Ready for ISSU
[Oct 26 05:28:43]:ISSU: Starting Upgrade for FRUs
[Oct 26 05:28:54]:ISSU: FPC Warm Booting
[Oct 26 05:29:59]:ISSU: FPC Warm Booted
[Oct 26 05:30:10]:ISSU: Preparing for Switchover
[Oct 26 05:30:14]:ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item          Status          Reason
  FPC 0         Online (ISSU)
Send ISSU done to chassisd on backup RE
Chassis ISSU Completed
Removing dcpfe1 eth1 128.0.0.16 IP
Bringing down bme01
Post Chassis ISSU processing done
[Oct 26 05:30:17]:ISSU: IDLE
Stopping primary dataplane
Clearing ISSU states
Console and management sessions will be disconnected. Please login again.

```



NOTE: If the ISSU process stops, you can look at the CLI output when you issue the `request system software in-service-upgrade` command to diagnose the problem. You can also look at syslog files for more information.

5. Log in after the reboot of the switch completes. To verify that the software has been upgraded, enter the following command:

```
user@switch> show version
```

Release History Table

Release	Description
18.1R1	Starting with Junos OS Release 18.2R1 on the QFX5200 switch, we recommend that you wait at least five minutes between in-service software upgrades.
17.1R1	Starting with Junos OS Release 17.1R1, on QFX5100 and EX4600 switches, you cannot perform an ISSU from a Junos OS Release earlier than 17.1R1 to Junos OS Release 17.1R1.

**Related
Documentation**

- [Understanding In-Service Software Upgrade \(ISSU\) on page 399](#)
- [request system software in-service-upgrade on page 699](#)

Performing an In-Service Software Upgrade (ISSU) in ACX5000 Series Routers

You can use an in-service software upgrade to upgrade the software running on the router with minimal traffic disruption during the upgrade.



NOTE: ISSU is supported in Junos OS Release 15.1X54–D60 and later on ACX5000 Series routers.

This topic covers:

1. [Preparing the Router for Software Installation on page 457](#)
2. [Upgrading the Software Using ISSU on page 459](#)
3. [Verifying a Unified ISSU on page 461](#)

Preparing the Router for Software Installation

Before you begin software installation using ISSU:



NOTE: Before you perform an in-service software upgrade, if applicable, remove the `set system internet-options no-tcp-reset drop-all-tcp` command from the configuration, otherwise the upgrade will fail and an error message will be displayed.

- Ensure that nonstop active routing (NSR) and nonstop bridging (NSB) are enabled. If enabled, disable graceful restart (GR), because NSR and GR cannot be enabled simultaneously. NSB and GR enable NSB-supported Layer 2 protocols to synchronize protocol information between the master and backup Routing Engines.
- If NSR is not enabled (**Stateful Replication is Disabled**), then enable NSR. NSR requires you to configure graceful Routing Engine switchover (GRES). By default, NSR is disabled.

- To enable graceful Routing Engine switchover, include the **graceful-switchover** statement at the **[edit chassis redundancy]** hierarchy level as **user@host#set chassis redundancy graceful-switchover**.
- To enable NSR, include the **nonstop-routing** statement at the **[edit routing-options]** hierarchy level as **user@host#set routing-options nonstop-routing**.
- Enable nonstop bridging (NSB). Nonstop bridging requires you to configure graceful Routing Engine switchover (GRES). By default, NSB is disabled.
 - To enable graceful Routing Engine switchover, include the **graceful-switchover** statement at the **[edit chassis redundancy]** hierarchy level as **user@host#set chassis redundancy graceful-switchover**.
 - To enable NSB, include the **nonstop-bridging** statement at the **[edit protocols layer2-control]** hierarchy level as **user@host#set protocols layer2-control nonstop-bridging**.
- (Optional) Back up the system software—Junos OS, the active configuration, and log files—on the router to an external storage device with the **request system snapshot** command.

On ACX5000 line of routers, you need to consider the following feature before performing ISSU:

- ISSU supports link fault management (LFM) timeout sessions of 1 second interval. During ISSU, you may notice LFM flaps for sessions having timeout interval of less than 1 second.
- Bidirectional Forwarding Detection (BFD) sessions having timeout interval of less than 1 second need to be reconfigured to 1 second before starting the ISSU process. You can restore the timeout interval to its original value after completing the ISSU process.
- ISSU supports interval slow (every 30 seconds) for periodic transmission of Link Aggregation Control Protocol (LACP) packets.
- ISSU supports Virtual Router Redundancy Protocol (VRRP) version 3.

ISSU do not support the following ACX5000 features:

- Downgrade to an earlier version of Junos OS software. If you want to install an earlier version of Junos OS software, use the **request system software add** CLI command.
- Upgrade of Host OS software.
- Connectivity fault management (CFM).
- TWAMP, RPF, RFC2544, and clocksyncd daemon (timing functionality).
- Mirroring and pseudowire cross connect.
- IPv6 firewall, IPv6 COS (classification and rewrite), IPv6 VPN, and VPLS mesh group.
- Virtual Router Redundancy Protocol (VRRP) version 1 and 2.
- Interval fast (every second) for periodic transmission of Link Aggregation Control Protocol (LACP) packets. If the periodic interval fast is configured, then you may notice

traffic drops because of LACP links going down during ISSU. ACX5000 line of routers can support LACP with fast hello by configuring the **fast-hello-issu** option (**user@host# set protocols lacp fast-hello-issu**) on the main router and peer routers before starting ISSU.



NOTE: The peer router must have Junos OS software to support this functionality.

Upgrading the Software Using ISSU

This procedure describes how to upgrade the software running on a standalone router:



NOTE: If the Host OS software needs to be updated, you cannot perform an ISSU. Instead, perform a standard software upgrade.

It is recommended to cleanup any unwanted data from the `/var` directory (`/var/log`, `/var/tmp`) before initiating the ISSU process.

To upgrade the router using ISSU:

1. Download the software package from the Juniper Networks Support website <https://www.juniper.net/support/downloads/junos.html>.



NOTE: To access the download site, you must have a service contract with Juniper Networks and an access account. If you need help obtaining an account, complete the registration form at the Juniper Networks website <https://www.juniper.net/registration/Register.jsp>.

2. Go to ACX Series section and select the ACX5000 Series platform software you want to download.
3. Copy the software package or packages to the router. We recommend that you copy the file to the `/var/tmp` directory.
4. Log in to the console connection. Using a console connection allows you to monitor the progress of the upgrade.
5. Start the ISSU:
 - On the router, enter:

```
user@host> request system software in-service-upgrade /var/tmp/package-name.tgz
```

where **package-name.tgz** is, for example,
jinstall-acx5k-15.1X54-D60.9-domestic-signed.tgz.



NOTE: During the upgrade, you will not be able to access the Junos OS CLI.

The router displays status messages similar to the following messages as the upgrade executes:

```
PRE ISSU CHECK:
-----
PFE Status                               : Online
BFD minimum-interval check done         : Valid
GRES enabled                             : Valid
NSR enabled                              : Valid
drop-all-tcp not configured             : Valid
OVSDDB not configured                    : Valid

warning: Do NOT use /user during ISSU. Changes to /user during ISSU may get
lost!
[Oct 24 00:25:37]:ISSU: Validating Image
[Oct 24 00:25:44]:ISSU: Preparing Backup RE
Prepare for ISSU
[Oct 24 00:25:49]:ISSU: Backup RE Prepare Done
Extracting jinstall-acx5k-15.1X54-D60.3-domestic ...
Install jinstall-acx5k-15.1X54-D60.3-domestic completed
Spawning the backup RE
Spawn backup RE, index 0 successful
GRES in progress
GRES done in 0 seconds
Waiting for backup RE switchover ready
GRES operational
Copying home directories
Copying home directories successful
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
[Oct 24 00:31:56]:ISSU: Preparing Daemons
[Oct 24 00:32:57]:ISSU: Daemons Ready for ISSU
[Oct 24 00:33:02]:ISSU: Starting Upgrade for FRUs
[Oct 24 00:33:23]:ISSU: FPC Warm Booting
[Oct 24 00:34:41]:ISSU: FPC Warm Booted
[Oct 24 00:34:51]:ISSU: Preparing for Switchover
[Oct 24 00:34:57]:ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item          Status          Reason
  FPC 0          Online (ISSU)
Send ISSU done to chassisd on backup RE
Chassis ISSU Completed
[Oct 24 00:35:18]:ISSU: IDLE
Console and management sessions will be disconnected. Please login again.
```



NOTE: An ISSU might stop instead of abort if the FPC is at the warm boot stage. Also, any links that go down and up will not be detected during a warm boot of the Packet Forwarding Engine (PFE).



NOTE: If the ISSU process stops, you can look at the log files to diagnose the problem. The log files are located at `/var/log/vjunos-log.tgz`.

6. Log in after the router reboots. To verify that the software has been upgraded, enter the following command:

```
user@host> show version
```

7. Disable or delete the configuration done to enable the ISSU. This includes disabling nonstop active routing (NSR), nonstop bridging (NBR) and graceful Routing Engine (GRES).

Verifying a Unified ISSU

Verify the status of FPCs and their corresponding PICs after the most recent unified ISSU.

Issue the **show chassis in-service-upgrade** command on the master Routing Engine.

```
user@host> show chassis in-service-upgrade
```

Item	Status	Reason
FPC 0	Online	

Display the unified ISSU process messages by using the **show log messages** command.

Related Documentation

Verifying a Unified ISSU

Purpose Verify the status of FPCs and their corresponding PICs after the most recent unified ISSU.

Action Issue the **show chassis in-service-upgrade** command on the master Routing Engine.

```
user@host> show chassis in-service-upgrade
```

Item	Status	Reason
FPC 0	Online	
FPC 1	Online	
FPC 2	Online	
PIC 0	Online	
PIC 1	Online	
FPC 3	Offline	Offlined by CLI command
FPC 4	Online	
PIC 1	Online	
FPC 5	Online	
PIC 0	Online	
FPC 6	Online	
PIC 3	Online	
FPC 7	Online	

Display the unified ISSU process messages by using the **show log messages** command.

Meaning See *show chassis in-service-upgrade* for more information.

- Related Documentation**
- [Example: Performing a Unified ISSU on page 424](#)
 - [Troubleshooting Unified ISSU Problems on page 462](#)
 - [Understanding the Unified ISSU Process on page 388](#)
 - [Unified ISSU System Requirements on page 401](#)
 - [Managing and Tracing BFD Sessions During Unified ISSU Procedures on page 462](#)

Troubleshooting Unified ISSU Problems

If the unified ISSU procedure stops progressing:

1. Open a new session on the master Routing Engine and issue the **request system software abort in-service-upgrade** command.
2. Check the existing router session to verify that the upgrade has been aborted.

An “ISSU: aborted!” message is provided. Additional system messages provide you with information about where the upgrade stopped and recommendations for the next step to take.

See *request chassis cluster in-service-upgrade abort (ISSU)* for more information.

- Related Documentation**
- [Understanding the Unified ISSU Process on page 388](#)
 - [Unified ISSU System Requirements on page 401](#)
 - [Best Practices for Performing a Unified ISSU on page 423](#)
 - [Example: Performing a Unified ISSU on page 424](#)
 - [Verifying a Unified ISSU on page 461](#)
 - [Managing and Tracing BFD Sessions During Unified ISSU Procedures on page 462](#)

Managing and Tracing BFD Sessions During Unified ISSU Procedures

Bidirectional Forwarding Detection (BFD) sessions temporarily increase their detection and transmission timers during unified ISSU procedures. After the upgrade, these timers revert to the values in use before the unified ISSU started. The BFD process replicates the unified ISSU state and timer values to the backup Routing Engine for each session.

No additional configuration is necessary to enable unified ISSU for BFD. However, you can disable the BFD timer negotiation during the unified ISSU by including the **no-issu-timer-negotiation** statement at the **[edit protocols bfd]** hierarchy level.

```
[edit protocols bfd]
no-issu-timer-negotiation;
```

If you include this statement, the BFD timers maintain their original values during unified ISSU.



CAUTION: The BFD sessions might flap during unified ISSU or Routing Engine switchover, depending on the detection intervals.

For more information about BFD, see the *Junos OS Routing Protocols Library*.

To configure unified ISSU trace options for BFD sessions, include the **issu** statement at the **[edit protocols bfd traceoptions flag]** hierarchy level.

```
[edit protocols]
bfd {
  traceoptions {
    flag issu;
  }
}
```

Related Documentation

- [Getting Started with Unified In-Service Software Upgrade on page 387](#)
- [Understanding the Unified ISSU Process on page 388](#)
- [Unified ISSU System Requirements on page 401](#)
- [Best Practices for Performing a Unified ISSU on page 423](#)
- [Example: Performing a Unified ISSU on page 424](#)
- [Verifying a Unified ISSU on page 461](#)
- [Troubleshooting Unified ISSU Problems on page 462](#)

Performing an ISSR

- [Performing an In-Service Software Reboot on page 465](#)

Performing an In-Service Software Reboot



NOTE: We recommend that you wait at least five minutes between in-service software reboots.

When you request an in-service software reboot (ISSR) on a standalone device:

1. The management process (MGD) verifies that graceful restart (GR) or non-stop routing and graceful Routing Engine switchover (GRES) are enabled.
2. The ISSU state machine spawns the backup Routing Engine (RE) with the existing software version.
3. The ISSU state machine checks to see if the backup RE has synchronized all of the data with the master RE.
4. The ISSU state machine requests the routing protocol process (RPD) to notify its readiness for switchover.
5. RPD initiates the GR or non-stop routing procedures by notifying all of the registered protocols.
6. RPD notifies the ISSU state machine that its ready for switchover.
7. The mastership is switched between the REs, so the backup RE becomes the master RE.
8. The old master RE is shut down.
9. RPD is spawned on the new master and continues the GR or non-stop routing procedure and exits either GR or non-stop routing after the protocol state synchronizes.

To perform an ISSR:

1. Issue the **request system reboot in-service** command.

For example:

```

user@switch> request system reboot in-service

Reboot the system ? [yes,no]
[Feb 22 02:37:04]:ISSU: Validating Image

PRE ISSR CHECK:
-----
PFE Status                               : Online
Member Id zero                           : Valid
VC not in mixed or fabric mode           : Valid
Member is single node vc                 : Valid
BFD minimum-interval check done          : Valid
GRES enabled                             : Valid
NSR enabled                              : Valid
drop-all-tcp not configured             : Valid
Ready for ISSR                           : Valid

warning: Do NOT use /user during ISSR. Changes to /user during ISSR may get
lost!
Current image is jinstall-jcp-i386-flex-18.1.img
[Feb 22 02:37:14]:ISSU: Preparing Backup RE
Prepare for ISSR
[Feb 22 02:37:19]:ISSU: Backup RE Prepare Done
Spawning the backup RE
Spawn backup RE, index 1 successful
Starting secondary dataplane
Second dataplane container started
GRES in progress
Waiting for backup RE switchover ready
GRES operational
Copying home directories
Copying home directories successful
Initiating Chassis In-Service-Upgrade for ISSR
Chassis ISSU Started
[Feb 22 02:42:55]:ISSU: Preparing Daemons
[Feb 22 02:43:00]:ISSU: Daemons Ready for ISSU
[Feb 22 02:43:05]:ISSU: Starting Upgrade for FRUs
[Feb 22 02:43:15]:ISSU: FPC Warm Booting
[Feb 22 02:44:16]:ISSU: FPC Warm Booted
[Feb 22 02:44:27]:ISSU: Preparing for Switchover
[Feb 22 02:44:31]:ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item      Status      Reason
  FPC 0      Online (ISSU)
Send ISSR done to chassisd on backup RE
Chassis ISSU Completed
Removing dcpfe0 eth1 128.168.0.16 IP
Bringing down bme00
Post Chassis ISSU processing done
[Feb 22 02:44:33]:ISSU: IDLE
Stopping primary dataplane
Clearing ISSU states
Console and management sessions will be disconnected. Please login again.
device_handoff successful ret: 0
Shutdown NOW!
[pid 14305]

```

```
*** FINAL System shutdown message from root@sw-duckhorn-01 ***  
System going down IMMEDIATELY
```

Related Documentation • *request system reboot*

PART 12

Performing Nonstop Software Upgrade (NSSU)

- [Getting Started with NSSU and Understanding How NSSU Works on page 471](#)
- [Performing a NSSU on page 481](#)

Getting Started with NSSU and Understanding How NSSU Works

- [Understanding Nonstop Software Upgrade on EX Series Switches on page 471](#)

Understanding Nonstop Software Upgrade on EX Series Switches

Nonstop software upgrade (NSSU) enables you to upgrade the software running on Juniper Networks EX Series Ethernet Switches with redundant Routing Engines and all member switches in EX Series Virtual Chassis using a single command. During the upgrade there might be minimal network traffic disruption during mastership switchover, and the extent of disruption could be dependent on the network topology, configuration, network traffic, and other environment factors.



NOTE: When an EX Series switch in a mixed Virtual Chassis is upgraded to Junos OS Release 15.1 or later from a release earlier than Release 15.1, there might be a drop in traffic for up to 60 seconds.

The following EX Series Virtual Chassis support NSSU:

- EX3300 Virtual Chassis
- EX3400 Virtual Chassis
- EX4200 Virtual Chassis
- EX4300 Virtual Chassis
- EX4500 Virtual Chassis
- EX4550 Virtual Chassis
- All mixed Virtual Chassis composed of EX4200, EX4500, and EX4550 switches
- EX4600 Virtual Chassis
- EX4650 Virtual Chassis



NOTE: An EX4650 Virtual Chassis operates the same as a QFX5120 Virtual Chassis, so for details on upgrading an EX4650 Virtual Chassis using NSSU, see *Understanding Nonstop Software Upgrade on a Virtual Chassis and Mixed Virtual Chassis* and *Upgrading Software on a Virtual Chassis and Mixed Virtual Chassis Using Nonstop Software Upgrade* instead of this topic.

- EX6200 switches
- EX8200 switches
- EX8200 Virtual Chassis

Performing an NSSU provides these benefits:

- No disruption to the control plane—An NSSU takes advantage of graceful Routing Engine switchover (GRES) and nonstop active routing (NSR) to ensure no disruption to the control plane. During the upgrade process, interface, kernel, and routing protocol information is preserved.
- Minimal disruption to network traffic—An NSSU minimizes network traffic disruption by:
 - Upgrading line cards one at a time in an EX6200 switch, EX8200 switch, or EX8200 Virtual Chassis while permitting traffic to continue to flow through the line cards that are not being upgraded.
 - Upgrading member switches one at a time in other EX Series Virtual Chassis while permitting traffic to continue to flow through the members that are not being upgraded.

To achieve minimal disruption to traffic, you must configure link aggregation groups (LAGs) such that the member links of each LAG reside on different line cards or Virtual Chassis members. When one member link of a LAG is down, the remaining links are up, and traffic continues to flow through the LAG.



NOTE: Because NSSU upgrades the software on each line card or on each Virtual Chassis member one at a time, an upgrade using NSSU can take longer than an upgrade using the `request system software add` command.

In releases prior to Junos OS Release 16.1, for EX6200 switches, EX8200 switches, and EX8200 Virtual Chassis, you can reduce the amount of time an upgrade takes by configuring line-card upgrade groups. The line cards in an upgrade group are upgraded simultaneously, reducing the amount of time it takes to complete an upgrade. See [“Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade” on page 482](#).

This topic covers:

- [Requirements for Performing an NSSU on page 473](#)
- [How an NSSU Works on page 474](#)
- [NSSU Limitations on page 477](#)
- [NSSU and Junos OS Release Support on page 477](#)
- [Overview of NSSU Configuration and Operation on page 478](#)

Requirements for Performing an NSSU

The following requirements apply to all switches and Virtual Chassis:

- All Virtual Chassis members and all Routing Engines must be running the same Junos OS release.
- Graceful Routing Engine switchover (GRES) must be enabled.
- Nonstop active routing (NSR) must be enabled.



NOTE: Although nonstop bridging (NSB) does not have to be enabled to perform an NSSU, we recommend enabling NSB before performing an NSSU. Enabling NSB ensures that all NSB-supported Layer 2 protocols operate seamlessly during the Routing Engine switchover that is part of the NSSU. In releases prior to Junos OS Release 16.1, see [“Configuring Nonstop Bridging on Switches \(CLI Procedure\)” on page 193](#).

- For minimal traffic disruption, you must define link aggregation groups (LAGs) such that the member links reside on different Virtual Chassis members or on different line cards.

The following are requirements for performing NSSU on an EX Series Virtual Chassis (excluding EX6200 or EX8200 Virtual Chassis):

- The Virtual Chassis members must be connected in a ring topology so that no member is isolated as a result of another member being rebooted. This topology prevents the Virtual Chassis from splitting during an NSSU.
- The Virtual Chassis master and backup must be adjacent to each other in the ring topology. Adjacency permits the master and backup to always be in sync, even when the switches in linecard roles are rebooting.
- The Virtual Chassis must be preprovisioned so that the linecard role has been explicitly assigned to member switches acting in a linecard role. During an NSSU, the Virtual Chassis members must maintain their roles—the master and backup must maintain their master and backup roles (although mastership will change), and the remaining switches must maintain their linecard roles.
- A two-member Virtual Chassis must have **no-split-detection** configured so that the Virtual Chassis does not split when an NSSU upgrades a member.



.....

NOTE: For the EX4300 Virtual Chassis, you should enable the `vcp-no-hold-time` statement at the `[edit virtual-chassis]` hierarchy level before performing a software upgrade using NSSU. If you do not enable the `vcp-no-hold-time` statement, the Virtual Chassis might split during the upgrade. A split Virtual Chassis can cause disruptions to your network, and you might have to manually reconfigure your Virtual Chassis after the NSSU if the split and merge feature was disabled. For more information about a split Virtual Chassis, see *Understanding Split and Merge in a Virtual Chassis*

.....

How an NSSU Works

This section describes what happens when you request an NSSU on these switches and Virtual Chassis:

- [EX3300, EX3400, EX4200, EX4300, EX4500, EX4600, and Mixed Virtual Chassis on page 474](#)
- [EX6200 and EX8200 Switches on page 475](#)
- [EX8200 Virtual Chassis on page 476](#)

[EX3300, EX3400, EX4200, EX4300, EX4500, EX4600, and Mixed Virtual Chassis](#)

When you request an NSSU on an EX3300, EX3400, EX4200, EX4300, EX4500, or mixed Virtual Chassis:

1. The Virtual Chassis master verifies that:
 - The backup is online and running the same software version.
 - Graceful Routing Engine switchover (GRES) and nonstop active routing (NSR) are enabled.
 - The Virtual Chassis has a preprovisioned configuration.
2. The master installs the new software image on the backup and reboots it.
3. The master resynchronizes the backup.
4. The master installs the new software image on member switches that are in the linecard role and reboots them, one at a time. The master waits for each member to become online and active before starting the software upgrade on the next member.
5. When all members that are in the linecard role have been upgraded, the master performs a graceful Routing Engine switchover, and the upgraded backup becomes the master.
6. The software on the original master is upgraded and the original master is automatically rebooted. After the original master has rejoined the Virtual Chassis, you can optionally return control to it by requesting a graceful Routing Engine switchover.

EX6200 and EX8200 Switches

When you request an NSSU on a standalone switch with redundant Routing Engines:

1. The switch verifies that:
 - Both Routing Engines are online and running the same software version.
 - Both Routing Engines have sufficient storage space for the new software image.
 - Graceful Routing Engine switchover and nonstop active routing are enabled.
2. The switch installs the new software image on the backup Routing Engine and reboots it.
3. The switch resynchronizes the backup Routing Engine to the master Routing Engine.
4. The line cards in the first upgrade group (or the line card in slot 0, if no upgrade groups are defined) download the new image and then restart. Traffic continues to flow through the line cards in the other upgrade groups during this process.
5. When line cards restarted in Step 4 are online again, the line cards in the next upgrade group download the new image and restart. This process continues until all online line cards have restarted with the new software.



NOTE: If you have taken a line card offline with the CLI before you start the NSSU, the line card is not restarted and remains offline.

6. The switch performs a graceful Routing Engine switchover, so that the upgraded backup Routing Engine becomes the master.
7. The switch installs the new software on the original master Routing Engine.

To complete the upgrade process, the original master Routing Engine must be rebooted. You can do so manually or have the switch perform an automatic reboot by including the **reboot** option when you request the NSSU. After the original master has been rebooted, you can optionally return control to it by requesting a graceful Routing Engine switchover.
8. (EX6200 switch only) The original master Routing Engine reboots to complete the software upgrade.



NOTE: To complete the upgrade process on an EX8200 switch, you must intervene to reboot the original master Routing Engine. You can reboot the original master Routing Engine manually or have the switch perform an automatic reboot by including the `reboot` option when you request the NSSU.

9. (Optional) After the original master has been rebooted, you can return control to it by requesting a graceful Routing Engine switchover.

The switch can maintain normal operations with either Routing Engine acting as the master Routing Engine after the software upgrade, so you only have to perform this switchover if you want to return Routing Engine control to the original master Routing Engine.

EX8200 Virtual Chassis

When you request an NSSU on an EX8200 Virtual Chassis:

1. The master external Routing Engine verifies that:
 - It has a backup external Routing Engine that is online.
 - All Virtual Chassis members have redundant Routing Engines and the Routing Engines are online.
 - All Routing Engines are running the same software version.
 - All Routing Engines have sufficient storage space for the new software image.
 - Graceful Routing Engine switchover and nonstop active routing (NSR) are enabled.
2. The master external Routing Engine installs the new software image on the backup external Routing Engine and reboots it.
3. The backup external Routing Engine resynchronizes with the master external Routing Engine.
4. The master external Routing Engine installs the new software on the backup Routing Engines in the member switches and reboots the backup Routing Engines.
5. When the reboot of the backup Routing Engines complete, the line cards in the first upgrade group download the new image and then restart. (If no upgrade groups are defined, the line card in slot 0 of member 0 downloads the new image and restarts.) Traffic continues to flow through the line cards in the other upgrade groups during this process.
6. When line cards restarted in Step 5 are online again, the line cards in the next upgrade group (or the next sequential line card) download the new image and restart. This process continues until all online line cards have restarted with the new software.



NOTE: If you have taken a line card offline with the CLI before you start the NSSU, the line card is not restarted and remains offline.

7. The new software image is installed on the master Routing Engines, both external and internal.
8. The member switches perform a graceful Routing Engine switchover, so that the upgraded backup Routing Engines become masters.
9. The master external Routing Engine performs a graceful Routing Engine switchover so that the backup external Routing Engine is now the master.

To complete the upgrade process, the original master Routing Engines, both external and internal, must be rebooted. You can do so manually by establishing a console connection to each Routing Engine or have the reboot performed automatically by including the **reboot** option when you request the NSSU. After the original master external Routing Engine has been rebooted, you can optionally return control to it by requesting a graceful Routing Engine switchover.

NSSU Limitations

You cannot use an NSSU to downgrade the software—that is, to install an earlier version of the software than is currently running on the switch. To install an earlier software version, use the **request system software add** command.

You cannot roll back to the previous software version after you perform an upgrade using NSSU. If you need to roll back to the previous software version, you can do so by rebooting from the alternate root partition if you have not already copied the new software version into the alternate root partition.

NSSU and Junos OS Release Support

A Virtual Chassis must be running a Junos OS release that supports NSSU before you can perform an NSSU. If a Virtual Chassis is running a software version that does not support NSSU, use the **request system software add** command.

Table 40 on page 477 lists the EX Series switches and Virtual Chassis that support NSSU and the Junos OS release at which they began supporting it.

Table 40: Platform and Release Support for NSSU

Platform	Junos OS Release
EX3300 Virtual Chassis	12.2 or later
EX3400 Virtual Chassis	15.1X53-D55
EX4200 Virtual Chassis	12.1 or later

Table 40: Platform and Release Support for NSSU (continued)

Platform	Junos OS Release
EX4300 Virtual Chassis	13.2X51-D20 or later
EX4500 Virtual Chassis	12.1 or later
EX4550 Virtual Chassis	12.2 or later
Mixed EX4200 and EX4500 Virtual Chassis	12.1 or later
Mixed EX4200 and EX4550 Virtual Chassis	12.2 or later
Mixed EX4200, EX4500, and EX4550 Virtual Chassis	12.2 or later
Mixed EX4500 and EX4550 Virtual Chassis	12.2 or later
EX6200 switch	12.2 or later
EX8200 switch	10.4 or later
EX8200 Virtual Chassis	11.1 or later

Overview of NSSU Configuration and Operation

You must ensure that the configuration of the switch or Virtual Chassis meets the requirements described in [“Requirements for Performing an NSSU” on page 473](#). NSSU requires no additional configuration.

In releases prior to Junos OS Release 16.1, for EX6200 switches, EX8200 switches, and EX8200 Virtual Chassis, you can optionally configure line-card upgrade groups using the CLI. See [“Example: Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade on EX Series Switches” on page 494](#).

You perform an NSSU by executing the `request system software nonstop-upgrade` command. For detailed instructions on how to perform an NSSU, see the topics in Related Documentation.

Release History Table

Release	Description
16.1	In releases prior to Junos OS Release 16.1, for EX6200 switches, EX8200 switches, and EX8200 Virtual Chassis, you can reduce the amount of time an upgrade takes by configuring line-card upgrade groups.

Related Documentation

- [Upgrading Software Using Nonstop Software Upgrade on EX Series Virtual Chassis and Mixed Virtual Chassis \(CLI Procedure\) on page 654](#)
- [Upgrading Software on an EX6200 or EX8200 Standalone Switch Using Nonstop Software Upgrade \(CLI Procedure\) on page 485](#)

- [Upgrading Software on an EX8200 Virtual Chassis Using Nonstop Software Upgrade \(CLI Procedure\) on page 650](#)
- [Configuring Nonstop Active Routing on Switches on page 220](#)
- [Configuring Graceful Routing Engine Switchover in a Virtual Chassis on page 151](#)
- [Example: Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade on EX Series Switches on page 494](#)

CHAPTER 32

Performing a NSSU

- [Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade on page 482](#)
- [Upgrading Software on an EX6200 or EX8200 Standalone Switch Using Nonstop Software Upgrade \(CLI Procedure\) on page 485](#)
- [Example: Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade on EX Series Switches on page 494](#)

Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade

Nonstop software upgrade (NSSU) enables you to upgrade software using a single command and with minimal disruption to network traffic on supporting switches. To reduce the total time required to complete an NSSU operation on all switches being upgraded, you can configure line-card upgrade groups on the following supported platforms:

- an EX6200 or EX8200 switch with redundant Routing Engines
- an EX8200 Virtual Chassis
- a QFX3500, QFX3600, and QFX5100 Virtual Chassis
- a Virtual Chassis Fabric

In its default configuration, NSSU upgrades each line card in a switch or member in a Virtual Chassis or VCF one at a time. Traffic continues to flow through the other line cards or members while each one is being restarted as part of the upgrade. This behavior allows you to minimize disruption to traffic if you have configured link aggregation groups (LAGs) such that the member links of each LAG reside on different line cards or members; as a result, when one member link of a LAG is down, the remaining links are up, and traffic continues to flow through the LAG.

When you define an upgrade group for NSSU, NSSU upgrades the line cards or members in the upgrade group at the same time instead of sequentially, reducing the total time needed to complete the upgrade on all line cards or members. To achieve minimal traffic disruption, you must define the line-card upgrade groups such that the member links of the LAGs reside on line cards or members that are in different upgrade groups. For information on how to configure LAGs, see *Configuring Aggregated Ethernet Links (CLI Procedure)*.

To configure line-card upgrade groups on a standalone EX6200 or EX8200 switch:

- To create an upgrade group and add a line card to it:

```
[edit chassis]
user@switch# set nssu upgrade-group group-name fpcs slot-number
```

For example, to create an upgrade group called **group3** and add the line card in slot 5 to it:

```
[edit chassis]
user@switch# set nssu upgrade-group group3 fpcs 5
```

If **group3** already exists, this command adds line card 5 to **group3**.

- To create an upgrade group and add multiple line cards to it:

```
[edit chassis]
user@switch# set nssu upgrade-group group-name fpcs [list-of-slot-numbers]
```


For example, to create an upgrade group called **primary** and add line cards in slots 1, 4, and 7 to it:

```
[edit chassis]
user@switch# set nssu upgrade-group primary fpcs [1 4 7]
```

If **primary** already exists, this command adds line cards in slots 1, 4, and 7 to **primary**.

To configure line-card upgrade groups on an EX8200 Virtual Chassis:

- To create an upgrade group and add a line card on a Virtual Chassis member to it:

```
[edit chassis]
user@switch# set nssu upgrade-group group-name member member-id fpcs slot-number
```

For example, to create an upgrade group called **primary-ny** and add the line card on member 1 in slot 5 to it:

```
[edit chassis]
user@switch# set nssu upgrade-group primary-ny member 1 fpcs 5
```

If **primary-ny** already exists, this command adds line card 5 on member 1 to **primary-ny**.

- To create an upgrade group that contains multiple line cards on a Virtual Chassis member:

```
[edit chassis]
user@switch# set nssu upgrade-group group-name member member-id fpcs
[list-of-slot-numbers]
```

For example, to create an upgrade group called **primary-ny** that contains the line cards in slots 1 and 2 on member 0 and in slots 3 and 4 on member 1:

```
[edit chassis]
user@switch# set nssu upgrade-group primary-ny member 0 fpcs [1 2]

[edit chassis]
user@switch# set nssu upgrade-group primary-ny member 1 fpcs [3 4]
```

To configure line-card upgrade groups on a QFX Series Virtual Chassis or mixed Virtual Chassis, or a VCF:



NOTE: For Virtual Chassis or VCFs comprised of fixed-chassis switches that do not have separate line cards, you use the `upgrade-group` configuration statement with the `fpcs` option to specify the Virtual Chassis or VCF member IDs that you want to include in an upgrade group. The member hierarchy of the `upgrade-group` statement is not used.

- To create an upgrade group and add a Virtual Chassis or VCF member to the upgrade group:

```
[edit chassis]
user@switch# set nssu upgrade-group group-name fpcs value
```

For example, to create an upgrade group called **vcf** and add a line card member:

```
[edit chassis]
user@switch# set nssu upgrade-group vcf fpcs 2
```

If **vcf** already exists, this command adds member 2 to **vcf**.

- To create an upgrade group that contains multiple members in a Virtual Chassis or VCF:

```
[edit chassis]
user@switch# set nssu upgrade-group group-name fpcs [list-of-slot-numbers]
```

For example, to create an upgrade group called **vcf** that contains members 1 and 2:

```
[edit chassis]
user@switch# set nssu upgrade-group vcf fpcs [1 2]
```

Related Documentation

- [Example: Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade on EX Series Switches on page 494](#)
- [Understanding Nonstop Software Upgrade on EX Series Switches on page 471](#)
- [Understanding Nonstop Software Upgrade on a Virtual Chassis and Mixed Virtual Chassis](#)
- [Understanding Nonstop Software Upgrade on EX Series Switches on page 471](#)
- [Upgrading Software on an EX6200 or EX8200 Standalone Switch Using Nonstop Software Upgrade \(CLI Procedure\) on page 485](#)
- [Upgrading Software on an EX8200 Virtual Chassis Using Nonstop Software Upgrade \(CLI Procedure\) on page 650](#)

- [Upgrading Software on a Virtual Chassis and Mixed Virtual Chassis Using Nonstop Software Upgrade](#)
- [Upgrading Software on a Virtual Chassis Fabric Using Nonstop Software Upgrade](#)
- [Upgrading Software Using Nonstop Software Upgrade on EX Series Virtual Chassis and Mixed Virtual Chassis \(CLI Procedure\) on page 654](#)

Upgrading Software on an EX6200 or EX8200 Standalone Switch Using Nonstop Software Upgrade (CLI Procedure)

You can use nonstop software upgrade (NSSU) to upgrade the software on standalone EX6200 or EX8200 switches with redundant Routing Engines. NSSU upgrades the software running on the Routing Engines and line cards with minimal traffic disruption during the upgrade. NSSU is supported on EX8200 switches running Junos OS Release 10.4 or later and on EX6200 switches running Junos OS Release 12.2 or later.

This topic covers:

- [Preparing the Switch for Software Installation on page 485](#)
- [Upgrading Both Routing Engines Using NSSU on page 487](#)
- [Upgrading One Routing Engine Using NSSU \(EX8200 Switch Only\) on page 490](#)
- [Upgrading the Original Master Routing Engine \(EX8200 Switch Only\) on page 492](#)

Preparing the Switch for Software Installation

Before you begin software installation using NSSU:

- (Optional) Configure line-card upgrade groups as described in [“Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade” on page 482](#). By default, an NSSU upgrades line cards one at a time to allow aggregated Ethernet links that have members on different line cards to remain up through the upgrade process. Configuring line-card upgrade groups reduces the time an upgrade takes because the line cards in each upgrade group are upgraded at the same time rather than sequentially.
- Verify that the Routing Engines are running the same version of the software. Enter the following command:

```
{master}
user@switch> show version invoke-on all-routing-engines
re0:
-----
Hostname: switch
Model: ex8208
JUNOS Base OS boot [11.3-20110429.1]
JUNOS Base OS Software Suite [11.3-20110429.1]
JUNOS Kernel Software Suite [11.3-20110429.1]
JUNOS Crypto Software Suite [11.3-20110429.1]
JUNOS Online Documentation [11.3-20110429.1]
JUNOS Enterprise Software Suite [11.3-20110429.1]
LC JUNOS Installation Software [11.3-20110429.1]
```

```

JUNOS Routing Software Suite [11.3-20110429.1]
JUNOS Web Management [11.3-20110429.1]

rel:
-----
Hostname: switch
Model: ex8208
JUNOS Base OS boot [11.3-20110429.1]
JUNOS Base OS Software Suite [11.3-20110429.1]
JUNOS Kernel Software Suite [11.3-20110429.1]
JUNOS Crypto Software Suite [11.3-20110429.1]
JUNOS Online Documentation [11.3-20110429.1]
JUNOS Enterprise Software Suite [11.3-20110429.1]
LC JUNOS Installation Software [11.3-20110429.1]
JUNOS Routing Software Suite [11.3-20110429.1]
JUNOS Web Management [11.3-20110429.1]

```

If the Routing Engines are not running the same version of the software, use the **request system software add** command to upgrade the Routing Engine that is running the earlier software version.

- Ensure that nonstop active routing (NSR) and graceful Routing Engine switchover (GRES) are enabled. To verify that they are enabled, you need to check only the state of nonstop active routing—if nonstop active routing is enabled, then graceful Routing Engine switchover is enabled.

To verify that nonstop active routing is enabled, execute the following command:

```

{master}
user@switch> show task replication
    Stateful Replication: Enabled
    RE mode: Master

Protocol                Synchronization Status
-----
OSPF                    Complete
RIP                     Complete
PIM                     Complete
RSVP                    Complete

```

If nonstop active routing is not enabled (**Stateful Replication is Disabled**), see [“Configuring Nonstop Active Routing on Switches” on page 220](#) for information on how to enable it.

- (Optional) Enable nonstop bridging (NSB). Enabling NSB ensures that all NSB-supported Layer 2 protocols operate seamlessly during the Routing Engine switchover that is part of the NSSU.
- (Optional) Back up the system software on each Routing Engine to an external storage device with the **request system snapshot** command.

Upgrading Both Routing Engines Using NSSU

This procedure describes how to upgrade both Routing Engines using NSSU. When the upgrade completes, both Routing Engines are running the new version of the software, and the backup Routing Engine is the new master Routing Engine.

To upgrade both Routing Engines using NSSU:

1. Download the software package.
2. Copy the software package to the switch. We recommend that you use FTP to copy the file to the `/var/tmp` directory.
3. Log in to the master Routing Engine using the console connection. You can perform an NSSU from the management interface, but a console connection allows you to monitor the progress of the master Routing Engine reboot.
4. Install the new software package:

```
{master}
user@switch> request system software nonstop-upgrade reboot
/var/tmp/package-name-m.nZx-distribution.tgz
```

where `package-name-m.nZx-distribution.tgz` is, for example, `jinstall-ex-8200-10.4R1.5-domestic-signed.tgz`.

The switch displays the following status messages as the upgrade executes:

```
Chassis ISSU Check Done
ISSU: Validating Image
ISSU: Preparing Backup RE
Pushing bundle to re1
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
Backup upgrade done
Rebooting Backup RE

Rebooting re1
ISSU: Backup RE Prepare Done
Waiting for Backup RE reboot
GRES operational
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
```

Item	Status	Reason
FPC 0	Online (ISSU)	
FPC 1	Online (ISSU)	
FPC 2	Online (ISSU)	
FPC 3	Offline	Offlined by CLI command
FPC 4	Online (ISSU)	
FPC 5	Online (ISSU)	
FPC 6	Online (ISSU)	

```

FPC 7          Online (ISSU)
Resolving mastership...
Complete. The other routing engine becomes the master.
ISSU: RE switchover Done
ISSU: Upgrading Old Master RE
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
ISSU: Old Master Upgrade Done
ISSU: IDLE

*** FINAL System shutdown message from user@switch ***
System going down IMMEDIATELY

Shutdown NOW!
[pid 2635]

```



NOTE: If you omit the reboot option in this step when using an EX8200 switch, you must manually reboot the original master Routing Engine with the `request system reboot` command for the upgrade to complete.

The original master Routing Engine reboots automatically after updating the new master Routing Engine when an NSSU is used to upgrade an EX6200 switch with dual Routing Engines.

5. Log in after the reboot completes. To verify that both Routing Engines have been upgraded, enter the following command:

```

{backup}
user@switch> show version invoke-on all-routing-engines
re0:
-----
Hostname: switch
Model: ex8208
JUNOS Base OS boot [12.1-20111229.0]
JUNOS Base OS Software Suite [12.1-20111229.0]
JUNOS Kernel Software Suite [12.1-20111229.0]
JUNOS Crypto Software Suite [12.1-20111229.0]
JUNOS Online Documentation [12.1-20111229.0]
JUNOS Enterprise Software Suite [12.1-20111229.0]
LC JUNOS Installation Software [12.1-20111229.0]
JUNOS Routing Software Suite [12.1-20111229.0]
JUNOS Web Management [12.1-20111229.0]

re1:
-----
Hostname: switch
Model: ex8208
JUNOS Base OS boot [12.1-20111229.0]
JUNOS Base OS Software Suite [12.1-20111229.0]
JUNOS Kernel Software Suite [12.1-20111229.0]
JUNOS Crypto Software Suite [12.1-20111229.0]
JUNOS Online Documentation [12.1-20111229.0]
JUNOS Enterprise Software Suite [12.1-20111229.0]
LC JUNOS Installation Software [12.1-20111229.0]

```

```
JUNOS Routing Software Suite [12.1-20111229.0]
JUNOS Web Management [12.1-20111229.0]
```

6. To verify that the line cards that were online before the upgrade are online after the upgrade, log in to the master Routing Engine and enter the **show chassis nonstop-upgrade** command:

```
{backup}
user@switch> request routing-engine login master

{master}
user@switch> show chassis nonstop-upgrade
```

Item	Status	Reason
FPC 0	Online (ISSU)	
FPC 1	Online (ISSU)	
FPC 2	Online (ISSU)	
FPC 3	Offline	Offlined by CLI command
FPC 4	Online (ISSU)	
FPC 5	Online (ISSU)	
FPC 6	Online (ISSU)	
FPC 7	Online (ISSU)	

7. If you want to make **re0** the master Routing Engine again, enter the following command:

```
{master}
user@switch> request chassis routing-engine master switch
Toggle mastership between routing engines ? [yes,no] (no) yes
```

You can verify that **re0** is the master Routing Engine by executing the **show chassis routing-engine** command.

8. To ensure that the resilient dual-root partitions feature operates correctly, execute the following command to copy the new Junos OS image into the alternate root partition on each Routing Engine:

```
user@switch> request system snapshot slice alternate routing-engine both
```

Resilient dual-root partitions allow the switch to boot transparently from the alternate root partition if the system fails to boot from the primary root partition.

Upgrading One Routing Engine Using NSSU (EX8200 Switch Only)

This procedure describes how to upgrade one of the Routing Engines using NSSU on an EX8200 switch. When the upgrade completes, the backup Routing Engine is running the new software version and is the new master. The original master Routing Engine, now the backup Routing Engine, continues to run the previous software version.



NOTE: NSSU always upgrades the software on both Routing Engines on an EX6200 switch. Therefore, you cannot upgrade software on one Routing Engine using NSSU on an EX6200 switch.

To upgrade one Routing Engine using NSSU:

1. Download the software package.
2. Copy the software package to the switch. We recommend that you use FTP to copy the file to the `/var/tmp` directory.
3. Log in to the master Routing Engine.
4. Request an NSSU. On an EX8200 switch, specify the **no-old-master-upgrade** option when requesting the NSSU:

```
{master}
user@switch> request system software nonstop-upgrade
no-old-master-upgrade /var/tmp/package-name-m.nZx-distribution.tgz
```

where *package-name-m.nZx-distribution.tgz* is, for example, *jinstall-ex-8200-10.4R2.5-domestic-signed.tgz*.

The switch displays the following status messages as the upgrade executes:

```
Chassis ISSU Check Done
ISSU: Validating Image
ISSU: Preparing Backup RE
Pushing bundle to re1
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
Backup upgrade done
Rebooting Backup RE

Rebooting re1
ISSU: Backup RE Prepare Done
Waiting for Backup RE reboot
GRES operational
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
```

Item	Status	Reason
------	--------	--------


```

FPC 0      Online (ISSU)
FPC 1      Online (ISSU)
FPC 2      Online (ISSU)
FPC 3      Offline           Offlined by CLI command
FPC 4      Online (ISSU)
FPC 5      Online (ISSU)
FPC 6      Online (ISSU)
FPC 7      Online (ISSU)
Resolving mastership...
Complete. The other routing engine becomes the master.
ISSU: RE switchover Done
Skipping Old Master Upgrade
ISSU: IDLE

```

When the upgrade is complete, the original master Routing Engine (**re0**) becomes the backup Routing Engine.

- To verify that the original backup Routing Engine (**re1**) has been upgraded, enter the following command:

```

{backup}
user@switch> show version invoke-on all-routing-engines
re0:
-----
Hostname: switch
Model: ex8208
JUNOS Base OS boot [11.3-20110429.1]
JUNOS Base OS Software Suite [11.3-20110429.1]
JUNOS Kernel Software Suite [11.3-20110429.1]
JUNOS Crypto Software Suite [11.3-20110429.1]
JUNOS Online Documentation [11.3-20110429.1]
JUNOS Enterprise Software Suite [11.3-20110429.1]
LC JUNOS Installation Software [11.3-20110429.1]
JUNOS Routing Software Suite [11.3-20110429.1]
JUNOS Web Management [11.3-20110429.1]

re1:
-----
Hostname: switch
Model: ex8208
JUNOS Base OS boot [12.1-20111229.0]
JUNOS Base OS Software Suite [12.1-20111229.0]
JUNOS Kernel Software Suite [12.1-20111229.0]
JUNOS Crypto Software Suite [12.1-20111229.0]
JUNOS Online Documentation [12.1-20111229.0]
JUNOS Enterprise Software Suite [12.1-20111229.0]
LC JUNOS Installation Software [12.1-20111229.0]
JUNOS Routing Software Suite [12.1-20111229.0]
JUNOS Web Management [12.1-20111229.0]

```

- To verify that the line cards that were online before the upgrade are online after the upgrade, log in to the new master Routing Engine and enter the **show chassis nonstop-upgrade** command:

```

{backup}
user@switch> request routing-engine login master

```

```

--- JUNOS 12.1-20111229.0 built 2011-12-29 04:12:22 UTC
{master}
user@switch> show chassis nonstop-upgrade

```

Item	Status	Reason
FPC 0	Online	
FPC 1	Online	
FPC 2	Online	
FPC 3	Offline	Offlined by CLI command
FPC 4	Online	
FPC 5	Online	
FPC 6	Online	
FPC 7	Online	

7. To ensure that the resilient dual-root partitions feature operates correctly, copy the new Junos OS image into the alternate root partition of the Routing Engine:

```
user@switch> request system snapshot slice alternate
```

Resilient dual-root partitions allow the switch to boot transparently from the alternate root partition if the system fails to boot from the primary root partition.

Upgrading the Original Master Routing Engine (EX8200 Switch Only)

This procedure describes how to upgrade the original master Routing Engine after you have upgraded the original backup Routing Engine as described in [“Upgrading One Routing Engine Using NSSU \(EX8200 Switch Only\)” on page 490](#) for an EX8200 switch.

1. Log in to the current master Routing Engine (re1).
2. Enter configuration mode and disable nonstop active routing:

```

{master}[edit]
user@switch# delete routing-options nonstop-routing

```

3. Deactivate graceful Routing Engine switchover and commit the configuration:

```

{master}[edit]
user@switch# deactivate chassis redundancy graceful-switchover

{master}[edit]
user@switch# commit

```

4. Log in to the current backup Routing Engine (re0) using a console connection.
5. Request a software installation:

```

user@switch> request system software add reboot
/var/tmp/package-name-m.nZx-distribution.tgz

```



NOTE: When you use NSSU to upgrade only one Routing Engine, the installation package is not automatically deleted from `/var/tmp`, leaving the package available to be used to upgrade the original master Routing Engine.

6. After the upgrade completes, log in to the current master Routing Engine (**re1**) and enter CLI configuration mode.

7. Re-enable nonstop active routing and graceful Routing Engine switchover:

```
[edit]
user@switch# activate chassis redundancy graceful-switchover

[edit]
user@switch# set routing-options nonstop-routing

[edit]
user@switch# commit
```

8. To ensure that the resilient dual-root partitions feature operates correctly, exit the CLI configuration mode and copy the new Junos OS image into the alternate root partition of the Routing Engine:

```
user@switch> request system snapshot slice alternate
```

Resilient dual-root partitions allow the switch to boot transparently from the alternate root partition if the system fails to boot from the primary root partition.

9. (Optional) To return control to the original master Routing Engine (**re0**), enter the following command:

```
{master}
user@switch> request chassis routing-engine master switch
Toggle mastership between routing engines ? [yes,no] (no) yes
```

You can verify that **re0** is the master Routing Engine by executing the **show chassis routing-engine** command.

Related Documentation

- [Example: Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade on EX Series Switches on page 494](#)
- [Configuring Dual-Root Partitions](#)
- [Troubleshooting Software Installation](#)

Example: Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade on EX Series Switches

Nonstop software upgrade (NSSU) enables you to upgrade the software running on an EX Series switch with redundant Routing Engines or on most EX Series Virtual Chassis by using a single command and with minimal disruption to network traffic. By default, NSSU upgrades the software running on line cards one line card at a time.

To reduce the time an NSSU takes, you can configure line-card upgrade groups on an EX6200 or EX8200 switch with redundant Routing Engines or on an EX8200 Virtual Chassis.

This example shows how to configure NSSU to use line-card upgrade groups:

- [Requirements on page 494](#)
- [Overview and Topology on page 494](#)
- [Configuration on page 495](#)

Requirements

This example uses the following hardware and software components:

- An EX8200 switch with redundant Routing Engines
- Junos OS Release 10.4 or later for EX Series switches

Before you begin to configure line-card upgrade groups, ensure that you have configured the link aggregation groups (LAGs) as described in *Configuring Aggregated Ethernet Links (CLI Procedure)*. See [“Overview and Topology” on page 494](#) for details about the LAG configurations for this example.

Overview and Topology

In its default configuration, NSSU upgrades each line card in a switch or Virtual Chassis one at a time. Traffic continues to flow through the other line cards while a line card is being restarted as part of the upgrade. This behavior allows you minimize disruption to traffic by configuring link aggregation groups (LAGs) such that the member links of each LAG reside on different line cards. When one member link of a LAG is down, the remaining links are up, and traffic continues to flow through the LAG.

Because the default configuration upgrades each line card one at a time, the upgrade can take some time to complete. You can reduce the time it takes to perform an NSSU by configuring line-card upgrade groups. Instead of being upgraded sequentially, the line cards in an upgrade group are upgraded simultaneously. To achieve minimal traffic disruption, you must define the line-card upgrade groups such that the member links of the LAGs reside on line cards that are in different upgrade groups.

This example uses an EX8200 switch that has five line cards installed in slots 0 through 4. Two LAGs have been configured:

- **ae0**—Has two member links, one on the line card in slot 0 and one on the line card in slot 1.
- **ae1**—Has two member links, one on the line card in slot 2 and one on the line card in slot 3.

The interfaces on the line card in slot 4 are not part of either LAG.


To minimize the time an upgrade takes and to ensure that the member links of each LAG are in different upgrade groups, this example configures the following two line-card upgrade groups:

- **group1**—Contains the line cards in slots 0, 2, and 4.
- **group2**—Contains the line cards in slots 1 and 3.

The line card in slot 4 could be put in either group. It could also be left out of an upgrade group entirely, and it would be upgraded separately after the line cards in the upgrade groups have been upgraded. However, it is more efficient to include it in an upgrade group.

Figure 33 on page 495 illustrates the topology.

Figure 33: Example Line-Card Upgrade Group Topology



Configuration

To create line-card upgrade groups, perform these tasks:

CLI Quick Configuration

To quickly create the line-card upgrade groups, copy the following commands and paste them into the switch terminal window:

```
[edit]
set chassis nssu upgrade-group group1 fpcs [0 2 4]
set chassis nssu upgrade-group group2 fpcs [1 3]
```

Step-by-Step Procedure

To create the line-card upgrade groups for an NSSU:

1. Create the first line-card upgrade group:

```
[edit chassis]
user@switch# set nssu upgrade-group group1 fpcs [0 2 4]
```

2. Create the second line-card upgrade group:

```
[edit chassis]
user@switch# set nssu upgrade-group group2 fpcs (NSSU Upgrade Groups) [1 3]
```

Results Display the results of the configuration:

```
[edit chassis]
user@switch# show
nssu {
    upgrade-group group1 {
        fpcs [ 0 2 4 ];
    }
    upgrade-group group2 {
        fpcs [ 1 3 ];
    }
}
```

**Related
Documentation**

- [Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade on page 482](#)
- [Upgrading Software on an EX6200 or EX8200 Standalone Switch Using Nonstop Software Upgrade \(CLI Procedure\) on page 485](#)
- [Upgrading Software on an EX8200 Virtual Chassis Using Nonstop Software Upgrade \(CLI Procedure\) on page 650](#)

PART 13

Configuration Statements and Operational Commands

- [Configuration Statements: Adaptive Load Balancing on page 499](#)
- [Configuration Statements: Bidirectional Forwarding Detection on page 503](#)
- [Configuration Statements: Graceful Routing Engine Switchover on page 513](#)
- [Configuration Statements: Graceful Restart on page 517](#)
- [Configuration Statements: Nonstop Active Routing on page 545](#)
- [Configuration Statements: Nonstop Bridging on page 553](#)
- [Configuration Statements: NSSU on page 555](#)
- [Configuration Statements: Power Management on page 561](#)
- [Configuration Statements: Redundant Power System on page 565](#)
- [Configuration Statements: Routing Engine and Switching Control Board Redundancy on page 569](#)
- [Configuration Statements: Unified ISSU on page 593](#)
- [Configuration Statements: VRRP on page 597](#)
- [Administration on page 641](#)
- [Operational Commands on page 661](#)
- [Troubleshooting on page 797](#)

CHAPTER 33

Configuration Statements: Adaptive Load Balancing

- [adaptive on page 500](#)

adaptive

Syntax	<pre>adaptive { pps; scan-interval <i>multiple</i>; tolerance <i>tolerance-percentage</i>; }</pre>
Hierarchy Level	<pre>[edit dynamic-profiles <i>name</i> interfaces <i>name</i> aggregated-ether-options load-balance], [edit dynamic-profiles <i>name</i> interfaces <i>name</i> logical-tunnel-options load-balance], [edit dynamic-profiles <i>name</i> interfaces interface-range <i>name</i> aggregated-ether-options load-balance], [edit dynamic-profiles <i>name</i> interfaces interface-range <i>name</i> logical-tunnel-options load-balance], [edit dynamic-profiles <i>name</i> logical-systems <i>name</i> interfaces <i>name</i> aggregated-ether-options load-balance], [edit dynamic-profiles <i>name</i> logical-systems <i>name</i> interfaces <i>name</i> logical-tunnel-options load-balance], [edit dynamic-profiles <i>name</i> logical-systems <i>name</i> interfaces interface-range <i>name</i> aggregated-ether-options load-balance], [edit dynamic-profiles <i>name</i> logical-systems <i>name</i> interfaces interface-range <i>name</i> logical-tunnel-options load-balance], [edit interfaces <i>name</i> aggregated-ether-options load-balance], [edit interfaces <i>name</i> logical-tunnel-options load-balance], [edit interfaces interface-range <i>name</i> aggregated-ether-options load-balance], [edit interfaces interface-range <i>name</i> logical-tunnel-options load-balance]</pre>
Release Information	<p>Statement introduced in Junos OS Release 13.2R3 for MX Series Routers.</p> <p>Statement introduced in Junos OS Release 14.1 for PTX Series Packet Transport Routers.</p> <p>Statement introduced in Junos OS Release 15.1X53-D10 for the QFX Series.</p>
Description	Correct a genuine traffic imbalance by using a feedback mechanism to distribute the traffic across the links of an aggregated Ethernet bundle.
Options	<p>pps—(PTX Series only) The type of traffic rate among the members of the AE bundle is measured packets per second. The default rate type is bytes per second.</p> <p>scan-interval <i>multiple</i>—(PTX Series only) Scan interval, as a multiple of a 30-second interval.</p> <p>Range: 1 through 5</p> <p>Default: 1</p> <p>tolerance <i>tolerance-percentage</i>—(MX Series and PTX Series) Limit to the variance in the packet traffic flow to the aggregated Ethernet links in a percentage.</p> <p>Range: 1 through 100 percent</p> <p>Default: 20 percent</p>

Required Privilege Level	interface - To view this statement in the configuration. interface-control - To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Understanding Aggregated Ethernet Load Balancing</i>• <i>Example: Configuring Aggregated Ethernet Load Balancing</i>

CHAPTER 34

Configuration Statements: Bidirectional Forwarding Detection

- [dedicated-ukern-cpu \(BFD\) on page 504](#)
- [realtime-ukern-thread \(BFD\) on page 505](#)
- [authentication \(LAG\) on page 506](#)
- [bfd-liveness-detection \(LAG\) on page 507](#)
- [detection-time \(LAG\) on page 509](#)
- [traceoptions \(Protocols BFD\) on page 510](#)
- [transmit-interval \(LAG\) on page 512](#)

dedicated-ukern-cpu (BFD)

Syntax	dedicated-ukern-cpu;
Hierarchy Level	[edit chassis]
Release Information	Statement introduced in Junos OS Release 15.1X49-D100.
Description	Enable the dedicated Bidirectional Forwarding Detection (BFD) protocol. One dedicated CPU core is allocated for the flowd ukernel thread to handle the dedicated BFD. This ensures that the BFD packet processing does not compete with the Routing Engine daemons.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Enabling Dedicated and Real-Time BFD on page 106• show chassis dedicated-ukern-cpu on page 693• bfd-liveness-detection on page 507• authentication on page 506• detection-time on page 509• transmit-interval on page 512

realtime-ukern-thread (BFD)

Syntax	realtime-ukern-thread;
Hierarchy Level	[edit chassis]
Release Information	Command introduced in Junos OS Release 15.1X49-D100.
Description	Enable the real-time Bidirectional Forwarding Detection (BFD) protocol. After real-time BFD is enabled, the priority of the flowd ukernel thread is changed to the highest level and, therefore, the flowd ukernel thread gets more CPU cycles for processing the BFD packets.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Enabling Dedicated and Real-Time BFD on page 106• show chassis realtime-ukern-thread on page 694• bfd-liveness-detection on page 507• authentication on page 506• detection-time on page 509• transmit-interval on page 512

authentication (LAG)

Syntax	<pre>authentication { algorithm <i>algorithm-name</i>; key-chain <i>key-chain-name</i>; loose-check; }</pre>
Hierarchy Level	[edit interfaces aex aggregated-ether-options bfd-liveness-detection]
Release Information	Statement introduced in Junos OS Release 13.3.
Description	Configure the authentication criteria of the BFD session for aggregated Ethernet interfaces.
Options	<p>algorithm <i>algorithm-name</i>—Specify the algorithm to be used to authenticate the BFD session. You can use one of the following algorithms for authentication:</p> <ul style="list-style-type: none">• keyed-md5• keyed-sha-1• meticulous-keyed-md5• meticulous-keyed-sha-1• simple-password <p>key-chain <i>key-chain-name</i>—Specify the name that is associated with the security key for the BFD session. The name you specify must match one of the keychains configured in the authentication-key-chains key-chain statement at the [edit security] hierarchy level.</p> <p>loose-check—(Optional) Configure loose authentication checking on the BFD session. Use only for transitional periods when authentication might not be configured at both ends of the BFD session.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• bfd-liveness-detection on page 507• detection-time on page 509• transmit-interval on page 512• Configuring Independent Micro BFD Sessions for LAG on page 87• Example: Configuring Independent Micro BFD Sessions for LAG on page 93• Understanding Independent Micro BFD Sessions for LAG on page 44

bfd-liveness-detection (LAG)

Syntax

```

bfd-liveness-detection {
  authentication {
    algorithm algorithm-name;
    key-chain key-chain-name;
    loose-check;
  }
  detection-time {
    threshold milliseconds;
  }
  holddown-interval milliseconds;
  local-address bfd-local-address;
  minimum-interval milliseconds;
  minimum-receive-interval milliseconds;
  multiplier number;
  neighbor bfd-neighbor-address;
  no-adaptation;
  transmit-interval {
    minimum-interval milliseconds;
    threshold milliseconds;
  }
  version (1 | automatic);
}

```

Hierarchy Level [edit interfaces *aex* aggregated-ether-options]

Release Information Statement introduced in Junos OS Release 13.3.

Description Configure Bidirectional Forwarding Detection (BFD) timers and authentication for aggregated Ethernet interfaces.

Options **holddown-interval *milliseconds***— Specify a time limit, in milliseconds, indicating the time that a BFD session remains up before a state change notification is sent. If the BFD session goes down and then comes back up during the hold-down interval, the timer is restarted.

Range: 0 through 255,000

Default: 0

local-address *bfd-local-address*— Specify the loopback address or the AE interface address of the source of the BFD session.



NOTE: Beginning with Release 16.1R2, Junos OS checks and validates the configured micro BFD *local-address* against the interface or loopback IP address before the configuration commit. Junos OS performs this

check on both IPv4 and IPv6 micro BFD address configurations, and if they do not match, the commit fails.

minimum-interval *milliseconds*— Specify a minimum time interval after which the local routing device transmits a BFD packet and then expects to receive a reply from the BFD neighbor. Optionally, instead of using this statement, you can configure the minimum transmit and receive intervals separately using the **transmit-interval minimum-interval** statement.

Range: 1 through 255,000

minimum-receive-interval *milliseconds*— Specify the minimum time interval after which the routing device expects to receive a reply from the BFD neighbor.

Range: 1 through 255,000

multiplier *number*— Specify the number of BFD packets that were not received by the BFD neighbor before the originating interface is declared down.

Range: 1 through 255

neighbor *bfd-neighbor-address*— Specify the loopback address or the AE interface address of a remote destination to send BFD packets.

no-adaptation— Disable the BFD adaptation. Include this statement if you do not want the BFD sessions to adapt to changing network conditions. We recommend that you do not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

version— Configure the BFD version to detect (BFD version 1) or autodetect (the BFD version).



NOTE: The version option is not supported on the QFX Series.

Default: automatic

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
---------------------------------	---

- Related Documentation**
- [authentication on page 506](#)
 - [detection-time on page 509](#)
 - [transmit-interval on page 512](#)
 - [Configuring Independent Micro BFD Sessions for LAG on page 87](#)
 - [Example: Configuring Independent Micro BFD Sessions for LAG on page 93](#)
 - [Understanding Independent Micro BFD Sessions for LAG on page 44](#)

detection-time (LAG)

Syntax	<pre> detection-time { threshold <i>milliseconds</i>; } </pre>
Hierarchy Level	[edit interfaces <i>aex</i> aggregated-ether-options bfd-liveness-detection]
Release Information	Statement introduced in Junos OS Release 13.3.
Description	Configure BFD timers for aggregated Ethernet interfaces.
Options	threshold <i>milliseconds</i> — Specify the maximum time interval for detecting a BFD neighbor. If the transmit interval is greater than this value, the device triggers a trap.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • authentication on page 506 • bfd-liveness-detection on page 507 • transmit-interval on page 512 • Configuring Independent Micro BFD Sessions for LAG on page 87 • Example: Configuring Independent Micro BFD Sessions for LAG on page 93 • Understanding Independent Micro BFD Sessions for LAG on page 44

traceoptions (Protocols BFD)

Syntax	<pre>traceoptions { file <i>name</i> <size <i>size</i>> <files <i>number</i>> <world-readable no-world-readable>; flag <i>flag</i> <<i>flag-modifier</i>> <disable>; }</pre>
Hierarchy Level	[edit protocols bfd]
Release Information	Statement introduced before Junos OS Release 7.4. issu flag for BFD added in Junos OS Release 9.1.
Description	<p>Define tracing operations that track unified in-service software upgrade (ISSU) functionality in the router.</p> <p>To specify more than one tracing operation, include multiple flag statements.</p>
Default	If you do not include this statement, no global tracing operations are performed.
Options	<p>disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.</p> <p>file <i>name</i>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log. We recommend that you place global routing protocol tracing output in the file routing-log.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>Range: 2 through 1000 files</p> <p>Default: 2 files</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the size option.</p> <p>flag <i>flag</i>—Tracing operation to perform. The tracing options are as follows:</p> <ul style="list-style-type: none">• adjacency—Trace adjacency messages.• all—Trace everything.• error—Trace all errors.• events—Trace all events.

- **issu**—Trace ISSU packet activity.
- **nsr-packet**—Trace packet activity of NSR.
- **nsr-synchronization**—Trace NSR synchronization events.
- **packet**—Trace all packets.
- **pipe**—Trace pipe messages.
- **pipe-detail**—Trace pipe messages in detail.
- **ppm-packet**—Trace packet activity by periodic packet management.
- **state**—Trace state transitions.
- **timer**—Trace timer processing.

no-world-readable—Restrict users from reading the log file.

size size—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

Syntax: *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 128 KB

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

world-readable—Allow users to read the log file.

Required Privilege Level	routing and trace—To view this statement in the configuration. routing-control and trace-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Managing and Tracing BFD Sessions During Unified ISSU Procedures on page 462

transmit-interval (LAG)

Syntax	<pre>transmit-interval { minimum-interval <i>milliseconds</i>; threshold <i>milliseconds</i>; }</pre>
Hierarchy Level	[edit interfaces aex aggregated-ether-options bfd-liveness-detection]
Release Information	Statement introduced in Junos OS Release 13.3.
Description	Configure the minimum interval and the threshold for transmission of BFD packets for aggregated Ethernet interfaces.
Options	<p>minimum-interval <i>milliseconds</i>— Specify the minimum time interval between two transmissions of packets. Range: 1 through 255,000</p> <p>threshold <i>milliseconds</i>— Specify the maximum interval between transmission of packets. If the transmit interval is greater than this value, the device triggers a trap.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• authentication on page 506• bfd-liveness-detection on page 507• detection-time on page 509• Configuring Independent Micro BFD Sessions for LAG on page 87• Example: Configuring Independent Micro BFD Sessions for LAG on page 93• Understanding Independent Micro BFD Sessions for LAG on page 44

CHAPTER 35

Configuration Statements: Graceful Routing Engine Switchover

- [graceful-switchover](#) on page 513
- [graceful-switchover](#) on page 514
- [redundancy \(Graceful Switchover\)](#) on page 515

[graceful-switchover](#)

Syntax `graceful-switchover;`

Hierarchy Level [edit chassis [redundancy](#)]

Release Information Statement introduced before Junos OS Release 7.4.

Description For routing platforms with two Routing Engines, configure a master Routing Engine to switch over gracefully to a backup Routing Engine without interruption to packet forwarding.



NOTE: The `graceful-switchover` statement at the [edit chassis redundancy] hierarchy level is not supported for Junos OS Evolved. Graceful switchover is enabled on the Junos OS Evolved system by default.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation • [Configuring Graceful Routing Engine Switchover](#) on page 148

graceful-switchover

Syntax	graceful-switchover;
Hierarchy Level	[edit chassis (EX Series) redundancy]
Release Information	Statement introduced in Junos OS Release 9.2 for EX Series switches. Statement introduced in Junos OS Release 13.2 for the QFX Series.
Description	For switches with more than one Routing Engine, including those in a Virtual Chassis or a Virtual Chassis Fabric, configure the master Routing Engine to switch over gracefully to a backup Routing Engine without interruption to packet forwarding.
Default	Graceful Routing Engine switchover (GRES) is disabled.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Nonstop Active Routing on Switches on page 226• Configuring Graceful Routing Engine Switchover on page 148• Configuring Graceful Routing Engine Switchover in a Virtual Chassis on page 151• Configuring Nonstop Active Routing on Switches on page 220• <i>Installing Software on an EX Series Switch with Redundant Routing Engines (CLI Procedure)</i>

redundancy (Graceful Switchover)

Syntax	<pre> redundancy { failover { on-disk-failure; on-loss-of-keepalives; } graceful-switchover; } </pre>
Hierarchy Level	[edit chassis (EX Series)]
Release Information	<p>Statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Enable redundant Routing Engines on a Virtual Chassis with two or more member switches or on a Virtual Chassis Fabric, on a standalone EX6200 or EX8200 switch with more than one Routing Engine.</p> <p>The remaining statements are explained separately. See CLI Explorer.</p>
Default	Redundancy is enabled for the Routing Engines.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • graceful-switchover on page 514 • Configuring Graceful Routing Engine Switchover in a Virtual Chassis on page 151 • Configuring Graceful Routing Engine Switchover on page 148 • Installing Software on an EX Series Switch with Redundant Routing Engines (CLI Procedure) • High Availability Features for EX Series Switches Overview on page 9

Configuration Statements: Graceful Restart

- `disable` on page 518
- `disable` (BGP Graceful Restart) on page 519
- `dont-help-shared-fate-bfd-down` on page 520
- `graceful-restart` (Enabling Globally) on page 521
- `graceful-restart` (Multicast Snooping) on page 522
- `graceful-restart` (Protocols BGP) on page 523
- `graceful-restart` (Protocols OSPF) on page 525
- `helper-disable` (Multiple Protocols) on page 527
- `helper-disable` (OSPF) on page 528
- `kernel-replication` on page 529
- `maximum-helper-recovery-time` on page 530
- `maximum-helper-restart-time` (RSVP) on page 531
- `maximum-neighbor-reconnect-time` on page 532
- `maximum-neighbor-recovery-time` on page 533
- `no-strict-lsa-checking` on page 534
- `notify-duration` on page 535
- `not-on-disk-underperform` on page 536
- `reconnect-time` on page 537
- `recovery-time` on page 538
- `restart-duration` on page 539
- `restart-time` (BGP Graceful Restart) on page 541
- `stale-routes-time` on page 542
- `traceoptions` (Protocols) on page 543
- `warm-standby` on page 544


disable

Syntax	disable;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (bgp isis ldp ospf ospf3 pim rip ripng rsvp) graceful-restart],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (bgp ldp ospf ospf3 pim) graceful-restart],</p> <p>[edit protocols (bgp isis isis ospf ospf3 ldp pim rip ripng rsvp) graceful-restart],</p> <p>[edit protocols bgp group <i>group-name</i> graceful-restart],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>ip-address</i> graceful-restart],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (bgp ldp ospf ospf3 pim) graceful-restart],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options graceful-restart],</p> <p>[edit routing-options graceful-restart]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Disable graceful restart.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Enabling Graceful Restart on page 243 • Configuring Routing Protocols Graceful Restart on page 270 • Configuring Graceful Restart for MPLS-Related Protocols on page 277 • Configuring VPN Graceful Restart on page 280 • Configuring Logical System Graceful Restart on page 281 • Graceful Restart Configuration Statements • Configuring Graceful Restart for QFabric Systems on page 282

dont-help-shared-fate-bfd-down

Syntax	dont-help-shared-fate-bfd-down
Hierarchy Level	[edit protocols bgp graceful-restart]
Release Information	Statement introduced in Junos OS Release 18.3R1.
Description	<p>When BFD is control plane dependent and the device detects a BFD down event and is not already entering the graceful restart helper mode, this is treated as a regular BFD down event and the device enters the graceful restart helper mode. This behavior makes the control plane dependent BFD unusable in conjunction with graceful restart.</p> <p>Include the dont-help-shared-fate-bfd-down statement at the [edit protocols bgp graceful-restart] hierarchy to ensure that the device does not enter the graceful restart helper mode and data traffic continues to be forwarded to an alternate path even if there is an interface failure (without a control plane restart on the BGP neighbor).</p>
Default	By default, this option is not enabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Using Control Plane Dependent BFD along with Graceful Restart Helper Mode on page 272• <i>Understanding External BGP Peering Sessions</i>

graceful-restart (Enabling Globally)

Syntax	<pre> graceful-restart { disable; helper-disable; maximum-helper-recovery-time <i>seconds</i>; maximum-helper-restart-time <i>seconds</i>; notify-duration <i>seconds</i>; recovery-time <i>seconds</i>; restart-duration <i>seconds</i>; stale-routes-time <i>seconds</i>; }</pre>
Hierarchy Level	<pre> [edit logical-systems <i>logical-system-name</i> routing-options], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options], [edit routing-options], [edit routing-instances <i>routing-instance-name</i> routing-options]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>You configure the graceful restart routing option globally to enable the feature, but not to enable graceful restart for all routing protocols in a routing instance. To enable graceful restart globally, include the graceful-restart statement under the [edit routing options] hierarchy level. This enables graceful restart globally for all routing protocols. You can, optionally, modify the global settings at the individual protocol level.</p>
	<div>  <p>NOTE:</p> <ul style="list-style-type: none"> For VPNs, the graceful-restart statement allows a router whose VPN control plane is undergoing a restart to continue to forward traffic while recovering its state from neighboring routers. For BGP, if you configure graceful restart after a BGP session has been established, the BGP session restarts and the peers negotiate graceful restart capabilities. LDP sessions flap when graceful-restart configurations change. </div>
Default	Graceful restart is disabled by default.
Options	The remaining statements are explained separately. See CLI Explorer .

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Enabling Graceful Restart on page 243• Configuring Routing Protocols Graceful Restart on page 270• Configuring Graceful Restart for MPLS-Related Protocols on page 277• Configuring VPN Graceful Restart on page 280• Configuring Logical System Graceful Restart on page 281• Configuring Graceful Restart for QFabric Systems on page 282

graceful-restart (Multicast Snooping)

Syntax	<pre>graceful-restart { disable; restart-duration <i>seconds</i>; }</pre>
Hierarchy Level	[edit multicast-snooping-options]
Release Information	Statement introduced in Junos OS Release 9.2.
Description	Establish the graceful restart duration for multicast snooping. You can set this value between 0 and 300 seconds. If you set the duration to 0, graceful restart is effectively disabled. Set this value slightly larger than the IGMP query response interval.
Default	180 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring Multicast Snooping</i>• <i>query-response-interval (Bridge Domains)</i>

graceful-restart (Protocols BGP)

Syntax	<pre>graceful-restart { disable; restart-time seconds; stale-routes-time seconds; }</pre>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Configure graceful restart for BGP. Graceful restart allows a routing device undergoing a restart to inform its adjacent neighbors and peers of its condition. Graceful restart is disabled by default. However, helper mode, the ability to assist a neighboring router attempting a graceful restart, is enabled by default.</p> <p>To configure the duration of the BGP graceful restart period, include the restart-time statement at the [edit protocols bgp graceful-restart] hierarchy level. To set the length of time the router waits to receive messages from restarting neighbors before declaring them down, include the stale-routes-time statement at the [edit protocols bgp graceful-restart] hierarchy level.</p>



NOTE: If you configure graceful restart after a BGP session has been established, the BGP session restarts and the peers negotiate graceful restart capabilities.

Enable graceful restart mode for BGP (and other protocols) by configuring graceful-restart at the routing-options level. Note that you cannot enable graceful restart for specific protocols unless graceful restart is also enabled globally.

For example, this configuration is required to enable graceful restart:

```
routing-options {
  graceful-restart
}
```

If you want to disable graceful restart for some protocols, you can do this at the protocol's graceful-restart command. The following configuration along with the configuration above will keep graceful restart for all protocols but BGP.

```
protocols{
  bgp{
    graceful-restart; {
      disable;
    }
  }
}
```

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
---------------------------------	---

Related Documentation	<ul style="list-style-type: none">• Configuring Graceful Restart Options for BGP on page 271• Configuring Graceful Restart for QFabric Systems on page 282• <i>High Availability Feature Guide</i>
------------------------------	--

graceful-restart (Protocols OSPF)

Syntax	<pre> graceful-restart { disable; helper-disable (standard restart-signaling both); no-strict-lsa-checking; notify-duration <i>seconds</i>; restart-duration <i>seconds</i>; } </pre>
Hierarchy Level	<pre> [edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)], [edit protocols (ospf ospf3)], [edit routing-instances <i>routing-instance-name</i> protocols ospf] </pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Support for the no-strict-lsa-checking statement introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the helper mode standard, restart-signaling, and both options introduced in Junos OS Release 11.4.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Configure graceful restart for OSPF.</p> <p>Graceful restart allows a routing device to restart with minimal effects to the network, and is enabled for all routing protocols at the [edit routing-options] hierarchy level.</p>
Options	<p>disable—Disable graceful restart for OSPF.</p> <p>helper-disable (standard restart-signaling both)—Disable helper mode for graceful restart. When helper mode is disabled, a device cannot help a neighboring device that is attempting to restart. Beginning with Junos OS Release 11.4, you can configure restart signaling-based helper mode for OSPFv2 graceful restart configurations. The standard, restart-signaling, and both options are only supported for OSPFv2. Specify standard to disable helper mode for standard graceful restart (based on RFC 3623). Specify restart-signaling to disable helper mode for restart signaling-based graceful restart (based on RFC 4811, RFC 4812, and RFC 4813). Specify both to disable helper mode for both standard and restart signaling-based graceful restart. The last committed statement takes precedence over the previously configured statement.</p> <p>Default: Helper mode is enabled by default. For OSPFv2, both standard and restart-signaling based helper modes are enabled by default.</p> <p>no-strict-lsa-checking—Disable strict OSPF link-state advertisement (LSA) checking to prevent the termination of graceful restart by a helping router. LSA checking is enabled by default.</p>



NOTE: The `helper-disable` statement and the `no-strict-lsa-checking` statement cannot be configured at the same time. If you attempt to configure both statements at the same time, the routing device displays a warning message when you enter the `show protocols (ospf | ospf3)` command.

notify-duration seconds—Estimated time needed to send out purged grace LSAs over all the interfaces.

Range: 1 through 3600 seconds

Default: 30 seconds

restart-duration seconds—Estimated time needed to reacquire a full OSPF neighbor from each area.

Range: 1 through 3600 seconds

Default: 180 seconds


Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
---------------------------------	---

Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring Graceful Restart for OSPF</i>• <i>Example: Configuring the Helper Capability Mode for OSPFv2 Graceful Restart</i>• <i>Example: Configuring the Helper Capability Mode for OSPFv3 Graceful Restart</i>• <i>Example: Disabling Strict LSA Checking for OSPF Graceful Restart</i>
------------------------------	---

helper-disable (Multiple Protocols)

Syntax	helper-disable;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (isis ldp ospf ospf3 rsvp) graceful-restart],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ldp ospf ospf3) graceful-restart],</p> <p>[edit protocols (isis ldp ospf ospf3 rsvp) graceful-restart],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ldp ospf ospf3) graceful-restart]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>
Description	Disable helper mode for graceful restart. When helper mode is disabled, a router or switch cannot help a neighboring router that is attempting to restart.
Default	Helper mode is enabled by default for these supported protocols: IS-IS, LDP, OSPF/OSPFv3, and RSVP.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Routing Protocols Graceful Restart on page 270 • Configuring Graceful Restart for MPLS-Related Protocols on page 277

helper-disable (OSPF)

Syntax	helper-disable < both restart-signaling standard >;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ospf graceful-restart], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf graceful-restart], [edit protocols ospf graceful-restart], [edit routing-instances <i>routing-instance-name</i> protocols ospf graceful-restart]
Release Information	Statement introduced before Junos OS Release 7.4. Options both , restart-signaling , and standard introduced in Junos OS Release 11.4. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Disable helper mode for graceful restart. When helper mode is disabled, a router cannot help a neighboring router that is attempting to restart. The last committed statement takes precedence over the previously configured statement.
Default	Helper mode is enabled by default for OSPF.
Options	both —(Optional) Disable helper mode for both standard and restart signaling-based graceful restart. restart-signaling —(Optional) Disable helper mode for restart signaling-based graceful restart (based on RFC 4811, RFC 4812, and RFC 4813). <div> NOTE: Restart signaling-based helper mode is not supported for OSPFv3 configurations.</div> standard —(Optional) Disable helper mode for standard graceful restart (based on RFC 3623).
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Routing Protocols Graceful Restart on page 270• Configuring Graceful Restart for MPLS-Related Protocols on page 277

kernel-replication

Syntax	<pre>kernel-replication { no-multithreading; system-reboot recovery-failure; }</pre>
Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 17.2R1.
Description	Configure kernel replication. Use this configuration statement to debug the kernel synchronization process (ksyncd) and configure automatic recovery from ksyncd initialization errors.
Options	<p>no-multithreading—(Optional) Run ksyncd in single thread mode for debugging purposes.</p> <p>system-reboot recovery-failure—(Optional) Configure the backup RE to automatically reboot if a ksyncd initialization error is detected.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Understanding Graceful Routing Engine Switchover on page 135 • show system switchover on page 774

maximum-helper-recovery-time

Syntax	<code>maximum-helper-recovery-time <i>seconds</i>;</code>
Hierarchy Level	[edit protocols rsvp graceful-restart], [edit logical-systems <i>logical-system-name</i> protocols rsvp graceful-restart]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Specify the length of time the router or switch retains the state of its Resource Reservation Protocol (RSVP) neighbors while they undergo a graceful restart.
Options	<i>seconds</i> —Length of time that the router retains the state of its Resource Reservation Protocol (RSVP) neighbors while they undergo a graceful restart. Range: 1 through 3600 Default: 180
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Graceful Restart Options for RSVP, CCC, and TCC on page 278• maximum-helper-restart-time (RSVP) on page 531

maximum-helper-restart-time (RSVP)

Syntax	<code>maximum-helper-restart-time <i>seconds</i>;</code>
Hierarchy Level	[edit protocols rsvp graceful-restart], [edit logical-systems <i>logical-system-name</i> protocols rsvp graceful-restart]
Release Information	Statement introduced in Junos OS Release 8.3. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Specify the length of time the router or switch waits after it discovers that a neighboring router has gone down before it declares the neighbor down. This value is applied to all RSVP neighbor routers and should be based on the time that the slowest RSVP neighbor requires for restart.
Options	<i>seconds</i> —The time the router or switch waits after it discovers that a neighboring router has gone down before it declares the neighbor down. Range: 1 through 1800 Default: 60
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Graceful Restart Options for RSVP, CCC, and TCC on page 278 • maximum-helper-recovery-time on page 530

maximum-neighbor-reconnect-time

Syntax	maximum-neighbor-reconnect-time <i>seconds</i> ;
Hierarchy Level	[edit protocols ldp graceful-restart], [edit logical-systems <i>logical-system-name</i> protocols ldp graceful-restart], [edit routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart]
Release Information	Statement introduced in Junos OS Release 9.1.
Description	Specify the maximum length of time allowed to reestablish connection from a restarting neighbor.
Options	<i>seconds</i> —Maximum time allowed for reconnection. Range: 30 through 300
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Graceful Restart Options for LDP on page 279

maximum-neighbor-recovery-time

Syntax	<code>maximum-neighbor-recovery-time <i>seconds</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp graceful-restart], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart], [edit protocols ldp graceful-restart], [edit routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart]
Release Information	Statement introduced before Junos OS Release 7.4. Statement changed from maximum-recovery-time to maximum-neighbor-recovery-time in Junos OS Release 9.1. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Specify the maximum amount of time to wait before giving up an attempt to gracefully restart.
Options	<i>seconds</i> —Configure the maximum recovery time, in seconds. Range: 120 through 1800 seconds Default: 140 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Recovery Time and Maximum Recovery Time • Configuring Graceful Restart Options for LDP on page 279 • no-strict-lsa-checking on page 534 • recovery-time on page 538

no-strict-lsa-checking

Syntax	no-strict-lsa-checking;
Hierarchy Level	[edit protocols (ospf ospf3) graceful-restart]
Release Information	Statement introduced in Junos OS Release 8.5. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Disable strict OSPF link-state advertisement (LSA) checking to prevent the termination of graceful restart by a helping router or switch.
Default	By default, LSA checking is enabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Graceful Restart Options for OSPF and OSPFv3 on page 274• Configuring Graceful Restart for QFabric Systems on page 282• maximum-neighbor-recovery-time on page 533• recovery-time on page 538

notify-duration

Syntax	<code>notify-duration <i>seconds</i>;</code>
Hierarchy Level	<p>[edit protocols (ospf ospf3) graceful-restart], [edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) graceful-restart], [edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols (ospf ospf3) graceful-restart], [edit routing-instances <i>instance-name</i> protocols (ospf ospf3) graceful-restart]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.3. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Specify the length of time the router or switch notifies helper OSPF routers that it has completed graceful restart.
Options	<p><i>seconds</i>—Length of time in the router notifies helper OSPF routers that it has completed graceful restart. Range: 1 through 3600 Default: 30</p>
Required Privilege Level	<p>routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Graceful Restart Options for OSPF and OSPFv3 on page 274 • Configuring Graceful Restart for QFabric Systems on page 282 • restart-duration on page 539

not-on-disk-underperform

Syntax not-on-disk-underperform;

Hierarchy Level [edit chassis redundancy failover]

Release Information Statement introduced in Junos OS Release 13.3R6.

Description Prevent gstatd from causing failovers in dual Routing Engines set for graceful Routing Engine switchover (GRES). The gstatd log message is still generated. This is an optional configuration.



NOTE: Configure the `disk-write-threshold` and `disk-read-threshold` statements to customize the gstatd timeout threshold.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Documentation • [Preventing Graceful Routing Engine Switchover in the Case of Slow Disks on page 152](#)

reconnect-time

Syntax	<code>reconnect-time <i>seconds</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp graceful-restart], [edit protocols ldp graceful-restart], [edit routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart]
Release Information	Statement introduced in Junos OS Release 9.1. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Specify the length of time required to reestablish a Label Distribution Protocol (LDP) session after graceful restart.
Options	<i>seconds</i> —Time required for reconnection. Range: 30 through 300 Default: 60 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring LDP Graceful Restart on MPLS Applications Feature Guide • Configuring Graceful Restart Options for LDP on page 279

recovery-time

Syntax	<code>recovery-time seconds;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp graceful-restart], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart], [edit protocols ldp graceful-restart], [edit routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify the length of time a router or switch waits for Label Distribution Protocol (LDP) neighbors to assist it with a graceful restart.
Options	seconds —Time the router waits for LDP to restart gracefully. Range: 120 through 1800 Default: 160
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Graceful Restart Options for LDP on page 279• maximum-neighbor-recovery-time on page 533• no-strict-lsa-checking on page 534

restart-duration

Syntax	<code>restart-duration <i>seconds</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (isis ospf ospf3 pim) graceful-restart],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3 pim) graceful-restart],</p> <p>[edit protocols (esis isis ospf ospf3 pim) graceful-restart],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3 pim) graceful-restart],</p> <p>[edit routing-options graceful-restart]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Configure the grace period for graceful restart globally.</p> <p>Additionally, you can individually configure the duration of the graceful restart period for the End System-to-Intermediate System (ES-IS), Intermediate System-to-Intermediate System (IS-IS), Open Shortest Path First (OSPF), and OSPFv3 protocols and for Protocol Independent Multicast (PIM) sparse mode.</p>
Options	<p><i>seconds</i>—Time for the graceful restart period.</p> <p>Range:</p> <p>The range of values varies according to whether the graceful restart period is being set globally or for a particular protocol:</p> <ul style="list-style-type: none"> • [edit routing-options graceful-restart] (global setting)—120 through 900 • ES-IS—30 through 300 • IS-IS—30 through 300 • OSPF/OSPFv3—1 through 3600 • PIM—30 through 300 <p>Default:</p> <p>The default value varies according to whether the graceful restart period is being set globally or for a particular protocol:</p> <ul style="list-style-type: none"> • [edit routing-options graceful-restart] (global setting)—300 • ES-IS—180 • IS-IS—210 • OSPF/OSPFv3—180 • PIM—60

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Enabling Graceful Restart on page 243• Configuring Graceful Restart for MPLS-Related Protocols on page 277• Configuring VPN Graceful Restart on page 280• <i>Configuring Graceful Restart for VPNs</i>• Configuring Logical System Graceful Restart on page 281

restart-time (BGP Graceful Restart)

Syntax	<code>restart-time seconds;</code>
Hierarchy Level	<p>[edit protocols (bgp rip ripng) graceful-restart],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (bgp rip ripng) graceful-restart (Enabling Globally)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp graceful-restart],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp graceful-restart]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.3.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Configure the duration of the BGP, RIP, or next-generation RIP (RIPng) graceful restart period.
Options	<p>seconds—Length of time for the graceful restart period.</p> <p>Range: 1 through 600 seconds</p> <p>Default: Varies by protocol:</p> <ul style="list-style-type: none"> • BGP—120 seconds • RIP and RIPng—60 seconds
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Graceful Restart Options for BGP on page 271 • Configuring Graceful Restart Options for RIP and RIPng on page 275 • Configuring Graceful Restart for QFabric Systems on page 282 • stale-routes-time on page 542

stale-routes-time

Syntax	<code>stale-routes-time <i>seconds</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-routing-name</i> protocols bgp graceful-restart], [edit logical-systems <i>logical-routing-name</i> routing-instances <i>routing-instance-name</i> protocols bgp graceful-restart], [edit protocols bgp graceful-restart], [edit routing-instances <i>routing-instance-name</i> protocols bgp graceful-restart]
Release Information	Statement introduced in Junos OS Release 8.3. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.
Description	Specify the maximum time that stale routes are kept during a restart. The stale-routes-time statement allows you to set the length of time the routing device waits to receive messages from restarting neighbors before declaring them down.
Options	seconds —Time the router device waits to receive messages from restarting neighbors before declaring them down. Range: 1 through 600 seconds Default: 300 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Graceful Restart Options for BGP on page 271• Configuring Graceful Restart for QFabric Systems on page 282• restart-time (BGP Graceful Restart) on page 541

traceoptions (Protocols)

Syntax	<pre> traceoptions { file <i>name</i> <size <i>size</i>> <files <i>number</i>> <world-readable no-world-readable>; flag <i>flag</i> <<i>flag-modifier</i>> <disable>; } </pre>
Hierarchy Level	[edit protocols isis], [edit protocols (ospf ospf3)]
Release Information	Statement introduced before Junos OS Release 7.4. graceful-restart flag for IS-IS and OSPF/OSPFv3 added in Junos OS Release 8.4.
Description	<p>Define tracing operations that graceful restart functionality in the router or switch.</p> <p>To specify more than one tracing operation, include multiple flag statements.</p>
Default	If you do not include this statement, no global tracing operations are performed.
Options	<p>disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.</p> <p>file <i>name</i>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log. We recommend that you place global routing protocol tracing output in the file routing-log.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>Range: 2 through 1000 files</p> <p>Default: 2 files</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the size option.</p> <p>flag <i>flag</i>—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. The nonstop active routing tracing option is:</p> <ul style="list-style-type: none"> graceful-restart—Tracing operations for nonstop active routing <p>no-world-readable—Restrict users from reading the log file.</p> <p>size <i>size</i>—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named trace-file reaches this size, it is renamed</p>

trace-file.0. When the **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

Syntax: **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 128 KB

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

world-readable—Allow users to read the log file.

Required Privilege Level	routing and trace—To view this statement in the configuration. routing-control and trace-control—To add this statement to the configuration.
---------------------------------	---

Related Documentation	<ul style="list-style-type: none">• Tracking Graceful Restart Events on page 277
------------------------------	--

warm-standby

Syntax	warm-standby;
---------------	---------------

Hierarchy Level	[edit routing-options]
------------------------	------------------------

Release Information	Statement introduced in Junos OS Release 17.2R1.
----------------------------	--

Description	Set the routing protocols process (rpd) mode to warm standby. Warm standby mode helps the backup RE stay synchronized with the master RE, allowing for faster RE switchover during GRES.
--------------------	--

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
---------------------------------	---

Related Documentation	<ul style="list-style-type: none">• Understanding Graceful Routing Engine Switchover on page 135
------------------------------	--

CHAPTER 37

Configuration Statements: Nonstop Active Routing

- [nonstop-routing on page 546](#)
- [switchover-on-routing-crash on page 547](#)
- [synchronize on page 548](#)
- [traceoptions on page 550](#)

nonstop-routing

Syntax nonstop-routing;

Hierarchy Level [edit routing-options]



NOTE: Although `nonstop-routing` is also a valid keyword at the `logical-systems` hierarchy level, it is not supported.

Release Information Statement introduced in Junos OS Release 8.4.
Statement introduced in Junos OS Release 10.4 for EX Series switches.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Statement introduced in Junos OS Release 13.2X51-D20 for QFX Series switches.

Description For routing platforms with two Routing Engines, configure a master Routing Engine to switch over gracefully to a backup Routing Engine and to preserve routing protocol information.

Default disabled

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation • [Configuring Nonstop Active Routing on page 217](#)

switchover-on-routing-crash

Syntax	switchover-on-routing-crash;
Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 13.3 for M Series, MX Series, T Series, TX Matrix, PTX Series, EX Series, QFX Series.
Description	Prevent loss of traffic in the case of NSR being configured. With the switchover-on-routing-crash configuration statement enabled, when rpd on the master Routing Engine crashes with NSR configured, the Routing Engine will switch over immediately to the backup Routing Engine to preserve protocol state and adjacencies. Prior to having this statement, if NSR was configured and rpd on the master Routing Engine crashed, it would cause network impact (protocol neighbor and adjacency drops and traffic loss).
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Nonstop Active Routing on page 217

synchronize

Syntax	synchronize;
Hierarchy Level	[edit system commit]
Release Information	Statement introduced in Junos OS Release 7.4. Statement introduced in Junos OS Release 10.4 for EX Series switches.
Description	For devices with multiple Routing Engines only. Configure the commit command to automatically perform a commit synchronize action between dual Routing Engines within the same chassis. The Routing Engine on which you execute the commit command (the requesting Routing Engine) copies and loads its candidate configuration to the other (the responding) Routing Engine. Each Routing Engine then performs a syntax check on the candidate configuration file being committed. If no errors are found, the configuration is activated and becomes the current operational configuration on both Routing Engines.



NOTE: If you configure the **commit synchronize** statement at the [edit system] hierarchy level and issue a **commit** in the master Routing Engine, the master configuration is automatically synchronized with the backup. However, if the backup Routing Engine is down when you issue the **commit**, the Junos OS displays a warning and commits the candidate configuration in the master Routing Engine. When the backup Routing Engine comes up, its configuration will automatically be synchronized with the master. A newly inserted backup Routing Engine automatically synchronizes its configuration with the master Routing Engine configuration.



NOTE: When you configure nonstop active routing (NSR), you must configure the **commit synchronize** statement. Otherwise, the **commit** operation fails.

On the TX Matrix router, synchronization only occurs between the Routing Engines within the same chassis. When synchronization is complete, the new configuration is then distributed to the Routing Engines on the T640 routers. That is, the master Routing Engine on the TX Matrix router distributes the configuration to the master Routing Engine on each T640 router. Likewise, the backup Routing Engine on the TX Matrix router distributes the configuration to the backup Routing Engine on each T640 router.

On the TX Matrix Plus router, synchronization only occurs between the Routing Engines within the switch-fabric chassis and when synchronization is complete, the new configuration is then distributed to the Routing Engines on the line-card chassis (LCC). That is, the master Routing Engine on the TX Matrix Plus router distributes the configuration to the master Routing Engine on each LCC. Likewise, the backup Routing

Engine on the TX Matrix Plus router distributes the configuration to the backup Routing Engine on each LCC.

In EX Series Virtual Chassis configurations:

- On EX4200 switches in Virtual Chassis, synchronization occurs between the switch in the master role and the switch in the backup role.
- On EX8200 switches in a Virtual Chassis, synchronization occurs only between the master and backup XRE200 External Routing Engines.

Options **and-quit**—(Optional) Quit configuration mode if the commit synchronization succeeds.

at—(Optional) Time at which to activate configuration changes.

comment—(Optional) Write a message to the commit log.

force—(Optional) Force a commit synchronization on the other Routing Engine (ignore warnings).


scripts—(Optional) Push scripts to the other Routing Engine.

Required Privilege Level **system**—To view this statement in the configuration.
 system-control—To add this statement to the configuration.

Related Documentation

- [Synchronizing the Routing Engine Configuration on page 218](#)
- *Configuring Multiple Routing Engines to Synchronize Committed Configurations Automatically*

traceoptions

Syntax	<pre> traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <disable>; } </pre>
Hierarchy Level	<pre> [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast], [edit logical-systems <i>logical-system-name</i> routing-options], [edit logical-systems <i>logical-system-name</i> routing-options multicast], [edit routing-instances <i>routing-instance-name</i> routing-options], [edit routing-instances <i>routing-instance-name</i> routing-options multicast], [edit routing-options], [edit routing-options flow], [edit routing-options multicast] </pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>nsr-synchronization flag for BGP, IS-IS, LDP, and OSPF added in Junos OS Release 8.4.</p> <p>nsr-synchronization and nsr-packet flags for BFD sessions added in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>nsr-synchronization flag for RIP and RIPng added in Junos OS Release 9.0.</p> <p>nsr-synchronization flag for Layer 2 VPNs and VPLS added in Junos OS Release 9.1.</p> <p>nsr-synchronization flag for PIM added in Junos OS Release 9.3.</p> <p>nsr-synchronization flag for MPLS added in Junos OS Release 10.1.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>nsr-synchronization flag for MSDP added in Junos OS Release 12.1.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Define tracing operations that track all routing protocol functionality in the routing device.</p> <p>To specify more than one tracing operation, include multiple flag statements.</p>
	<div>  <p>NOTE: On Junos OS Evolved, traceoptions is hidden for op, event, and commit scripts. Instead, Junos OS Evolved enables default tracking and trace messages that are logged under /var/log/traces.</p> </div>
Default	If you do not include this statement, no global tracing operations are performed.
Options	Values:

disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file filename—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. We recommend that you place global routing protocol tracing output in the file **routing-log**.

files number—(Optional) Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. Note that if you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

Range: 2 through 1000 files

Default: 10 files

flag flag—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. These are the global routing protocol tracing options:

- **all**—All tracing operations
- **condition-manager**—Condition-manager events
- **config-internal**—Configuration internals
- **general**—All normal operations and routing table changes (a combination of the **normal** and **route** trace operations)
- **graceful-restart**—Graceful restart operations
- **normal**—All normal operations
- **nsr-packet**—Detailed trace information for BFD nonstop active routing only
- **nsr-synchronization**—Tracing operations for nonstop active routing
- **nsr-synchronization**—Nonstop active routing synchronization
- **parse**—Configuration parsing
- **policy**—Routing policy operations and actions
- **regex-parse**—Regular-expression parsing
- **route**—Routing table changes
- **state**—State transitions
- **task**—Interface transactions and processing
- **timer**—Timer usage

no-world-readable—(Optional) Prevent any user from reading the log file.

size size—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. Note that if you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 128 KB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege	routing and trace—To view this statement in the configuration.
Level	routing-control and trace-control—To add this statement to the configuration.

Related Documentation	<ul style="list-style-type: none">• <i>Example: Tracing Global Routing Protocol Operations</i>
------------------------------	--

CHAPTER 38

Configuration Statements: Nonstop Bridging

- [nonstop-bridging](#) on page 553
- [nonstop-bridging \(Ethernet Switching\)](#) on page 554

[nonstop-bridging](#)

Syntax	<code>nonstop-bridging;</code>
Hierarchy Level	[edit protocols layer2-control]
Release Information	Statement introduced in Junos OS Release 8.4.
Description	For platforms with two Routing Engines, configure a master Routing Engine to switch over gracefully to a backup Routing Engine and preserve Layer 2 Control Protocol (L2CP) information.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Synchronizing the Routing Engine Configuration on page 218• Configuring Nonstop Bridging on page 191• For information about configuring NSB on EX Series switches that do not support the Enhanced Layer 2 Software (ELS) CLI style, see Configuring Nonstop Bridging on EX Series Switches (CLI Procedure) on page 195• For information about configuring NSB on switches that support ELS, see Configuring Nonstop Bridging on Switches (CLI Procedure) on page 193

nonstop-bridging (Ethernet Switching)


Syntax	nonstop-bridging;
Hierarchy Level	[edit ethernet-switching-options]
Release Information	Statement introduced in Junos OS Release 11.3 for EX Series switches.
Description	For switches with two Routing Engines or for Virtual Chassis, configure a master Routing Engine to switch over gracefully to a backup Routing Engine and preserve Layer 2 protocol information for the Layer 2 protocols that support nonstop bridging (NSB). For a list of the EX Series switches and Layer 2 protocols that support nonstop bridging, see <i>EX Series Switch Software Features Overview</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Nonstop Bridging on EX Series Switches (CLI Procedure) on page 195

CHAPTER 39

Configuration Statements: NSSU

- `fpcs` (NSSU Upgrade Groups) on page 556
- `member` (NSSU Upgrade Groups) on page 558
- `nssu` on page 559
- `upgrade-group` on page 560

fpcs (NSSU Upgrade Groups)

Syntax	<code>fpcs (slot-number [list-of-slot-numbers]);</code>
Hierarchy Level	[edit chassis (EX Series) <code>nssu upgrade-group group-name</code>], [edit chassis (EX Series) <code>nssu upgrade-group group-name member member-id</code>]
Release Information	Statement introduced in Junos OS Release 10.4 for EX Series switches. Statement introduced in Junos OS Release 13.2X51-D20 for QFX Series switches.
Description	<p>Configure switch line cards, Virtual Chassis member switches, or Virtual Chassis Fabric (VCF) member switches as part of an NSSU upgrade group.</p> <p>To reduce the time an NSSU takes, you can configure line-card upgrade groups for an EX6200 or EX8200 switch with redundant Routing Engines; an EX8200 Virtual Chassis; QFX3500, QFX3600, and QFX5100 Virtual Chassis; or a Virtual Chassis Fabric (VCF).</p> <p>For switches that have separate line cards, use this statement to assign one or more line cards to an NSSU upgrade group by specifying their slot numbers with this statement.</p> <p>For Virtual Chassis or VCF member switches that do not have separate line cards, use this statement to assign one or more Virtual Chassis or VCF members to an NSSU upgrade group by specifying their member IDs.</p>
	<p> NOTE: You do not use this statement with the <code>member</code> keyword in this case. When to use the <code>member</code> statement hierarchy is explained next.</p>
	<p>To configure an upgrade group that includes line cards on different switches that support multiple line cards and comprise a Virtual Chassis, use this statement with the <code>member</code> statement hierarchy to specify the Virtual Chassis member ID and the desired line card slot number or numbers on that member switch to include in the upgrade group. Use multiple statements to add line cards from different Virtual Chassis members to the upgrade group.</p>
Options	<p><i>list-of-slot-numbers</i>—A list of slot numbers of multiple line cards or member IDs of Virtual Chassis or VCF members to be included in the upgrade group. Separate multiple slot numbers or member IDs with spaces and enclose the list in square brackets—for example: <code>[3 4 7]</code>.</p> <p><i>slot-number</i>—The slot number of a single line card or member ID of a Virtual Chassis or VCF member to be included in the upgrade group.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

**Related
Documentation**

- [Example: Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade on EX Series Switches on page 494](#)
- [Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade on page 482](#)

member (NSSU Upgrade Groups)

Syntax `member member-id {
 fpcs (slot-number | [list-of-slot-numbers]);
 }`

Hierarchy Level [edit chassis (EX Series) **nssu upgrade-group** *group-name*]

Release Information Statement introduced in Junos OS Release 11.1 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D20 for QFX Series switches.

Description Specify the Virtual Chassis member whose line-card slot numbers you are assigning to an NSSU upgrade group.



NOTE: This statement is not applicable to Virtual Chassis or VCF member switches that do not support separate line cards. To configure Virtual Chassis or VCF member switches that do not have separate line cards into an NSSU upgrade group, use the **fpcs** statement alone, and specify the Virtual Chassis or VCF member IDs to include in the upgrade group in place of line card slot numbers.

To reduce the time an NSSU takes, you can configure NSSU line-card upgrade groups on an EX6200 or EX8200 switch with redundant Routing Engines; EX8200 Virtual Chassis; QFX3500, QFX3600, and QFX5100 Virtual Chassis; and Virtual Chassis Fabric (VCF).

To configure an upgrade group that includes line cards on different switches that support multiple line cards and comprise a Virtual Chassis, use this statement hierarchy with the **fpcs** option to first specify the Virtual Chassis member ID and then desired line card slot number or numbers on that member switch to include in the upgrade group. Use multiple statements to add line cards from different Virtual Chassis members to the upgrade group.

Options **member-id**—The ID of the Virtual Chassis or VCF member switch containing one or more line cards to include in an NSSU upgrade group.

The remaining statement is explained separately. See [CLI Explorer](#).

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Documentation

- [Example: Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade on EX Series Switches on page 494](#)
- [Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade on page 482](#)

nssu

```
Syntax  nssu {
        rcp-count number;
        upgrade-group group-name {
            fpcs (slot-number | [list-of-slot-numbers]);
            member member-id {
                fpcs (slot-number | [list-of-slot-numbers]);
            }
        }
    }
```

Hierarchy Level [edit chassis (EX Series)]

Release Information Statement introduced in Junos OS Release 10.4 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D20 for QFX Series switches.
rcp-count statement introduced in Junos OS Release 14.1X53-D40 for QFX5100 switches.

Description Configure optional parameters used in the nonstop software upgrade (NSSU) process.



NOTE: The **request system software nonstop-upgrade** command is used to initiate NSSU.

For the **rcp-count** statement: (QFX5100 Virtual Chassis and Virtual Chassis Fabric (VCF) only) Configure the number of parallel **rcp** sessions NSSU uses to copy the new software to multiple Virtual Chassis or VCF member switches at a time. See *rcp-count* for details.

For **upgrade-group** statements: Define a line-card upgrade group for NSSU, for switch configurations that support upgrade groups. All line cards or Virtual Chassis or VCF members in an upgrade group are upgraded to the new software version at the same time when an NSSU is initiated and at least one upgrade group is configured. Line-card upgrade groups are not required to initiate an NSSU, and are not supported on some EX Series switches or EX Virtual Chassis that support NSSU. See **upgrade-group** for details.

The remaining statements are explained separately. See [CLI Explorer](#).

Default If **rcp-count** is not configured, NSSU uses a default algorithm to determine the number of parallel **rcp** sessions to use based on the number of members in the Virtual Chassis or VCF.

If no line-card upgrade groups are defined, NSSU upgrades line cards and members of a Virtual Chassis or VCF one at a time in ascending order by slot or member number.

Required Privilege interface—To view this statement in the configuration.
Level interface-control—To add this statement to the configuration.

Related Documentation

- [Example: Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade on EX Series Switches on page 494](#)
- [Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade on page 482](#)

upgrade-group

Syntax

```
upgrade-group group-name {  
    fpcs (slot-number | [list-of-slot-numbers]);  
    member member-id {  
        fpcs (slot-number | [list-of-slot-numbers]);  
    }  
}
```

Hierarchy Level [edit chassis (EX Series) [nssu](#)]

Release Information Statement introduced in Junos OS Release 10.4 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D20 for QFX Series switches.

Description Assign a name to a line-card upgrade group being created for nonstop software upgrade (NSSU).

To reduce the time an NSSU takes, you can configure line-card upgrade groups on an EX6200 or EX8200 switch with redundant Routing Engines; EX8200 Virtual Chassis; QFX3500, QFX3600, and QFX5100 Virtual Chassis; and Virtual Chassis Fabric (VCF).

Options *group-name*—Name of the upgrade group.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege interface—To view this statement in the configuration.
Level interface-control—To add this statement to the configuration.

Related Documentation


- [Example: Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade on EX Series Switches on page 494](#)
- [Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade on page 482](#)

CHAPTER 40

Configuration Statements: Power Management

- [power-budget-priority](#) on page 562
- [n-plus-n \(Power Management\)](#) on page 563
- [psu](#) on page 563
- [redundancy \(Power Management\)](#) on page 564

power-budget-priority

Syntax	<code>power-budget-priority <i>priority</i>;</code>
Hierarchy Level	[edit chassis (EX Series) fpc slot]
Release Information	Statement introduced in Junos OS Release 10.2 for EX Series switches.
Description	Assign a power priority to the specified line card slot on an EX6200 or EX8200 switch. <div><div></div><div>NOTE: On an EX6200 switch, you cannot change the power priority of a slot containing a Switch Fabric and Routing Engine (SRE) module. Although the CLI allows you to set a different power priority for the slot, your change does not go into effect, and the power priority remains 0. A message is sent to the system log to inform you that changing the power priority of the slot is unsupported.</div></div>
Default	All line card slots are initially assigned the lowest priority, with the exception of slot 4 and slot 5 on the EX6200 switch, which always are assigned a priority of 0.
Options	<i>priority</i> —Assigned power priority for the slot, with 0 being the highest priority: <ul style="list-style-type: none">• 0 through 9 for an EX6200 switch• 0 through 7 for an EX8208 switch• 0 through 15 for an EX8216 switch
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Power Priority of Line Cards (CLI Procedure) on page 302

n-plus-n (Power Management)

Syntax	n-plus-n;
Hierarchy Level	[edit chassis (EX Series) psu redundancy]
Release Information	Statement introduced in Junos OS Release 10.2 for EX Series switches.
Description	Configure N+N power supply redundancy for power management on an EX6200 or EX8200 switch.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Power Supply Redundancy (CLI Procedure) on page 303

psu

Syntax	<pre> psu { redundancy { n-plus-n (Power Management); } } </pre>
Hierarchy Level	[edit chassis (EX Series)]
Release Information	Statement introduced in Junos OS Release 10.2 for EX Series switches.
Description	<p>Configure N+N power supply redundancy for power management on an EX6200 or EX8200 switch.</p> <p>The remaining statements are explained separately. See CLI Explorer.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Power Supply Redundancy (CLI Procedure) on page 303

redundancy (Power Management)

Syntax	<pre>redundancy { n-plus-n (Power Management); }</pre>
Hierarchy Level	[edit chassis (EX Series) psu]
Release Information	Statement introduced in Junos OS Release 10.2 for EX Series switches.
Description	<p>Configure $N+N$ power supply redundancy for power management on an EX6200 or EX8200 switch.</p> <p>The remaining statement is explained separately. See CLI Explorer.</p>
Default	$N+1$ power supply redundancy is configured by default.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring Power Supply Redundancy (CLI Procedure) on page 303

CHAPTER 41

Configuration Statements: Redundant Power System

- [member \(Redundant Power System\)](#) on page 565
- [priority \(Redundant Power System\)](#) on page 566
- [redundant-power-system](#) on page 567

member (Redundant Power System)

Syntax `member vc-member-number {
 priority (0|1|2|3|4|5|6);
 }`

Hierarchy Level `[edit redundant-power-system]`

Release Information Statement introduced in Junos OS Release 12.1 for EX Series switches.

Description Specify the Virtual Chassis member ID of a switch connected to the Redundant Power System (RPS) for backup power supply. The member ID is required only for switches that can be configured in a Virtual Chassis. If the switch has never been configured in a Virtual Chassis, the value is always 0.

Options *member-number*—Member ID of a switch that has Virtual Chassis capability that is connected to the RPS.

Range: 0 through maximum members in the Virtual Chassis

Default: 0

The remaining statement is explained separately. See [CLI Explorer](#).

Required Privilege Level admin—To view this statement in the configuration.
 admin-control—To add this statement to the configuration.

Related Documentation • [Determining and Setting Priority for Switches Connected to an EX Series RPS](#) on page 310

priority (Redundant Power System)

Syntax	priority (0 1 2 3 4 5 6);
Hierarchy Level	[edit redundant-power-system member] [edit redundant-power-system member <i>member-number</i>]
Release Information	Statement introduced in Junos OS Release 12.1 for EX Series switches.
Description	<p>Configure the backup of any switch connected to the Redundant Power System (RPS) using the CLI on each switch. The determines the order in which the RPS supplies backup power to the switches connected to the RPS. 6 is the highest priority and 1 is lowest. Zero means off or no RPS backup.</p> <p>If the switch is not reconfigured from the CLI, the default priority is 1. In this case, priority is determined by connector location with the rightmost connector having the highest priority.</p> <p>For switches that can only be used as standalone switches, this hierarchy level is used for configuration:</p> <p>[edit redundant-power-system]</p> <p>For switches that can be used either as standalone switches or configured in a Virtual Chassis, this hierarchy level is used for configuration:</p> <p>[edit redundant-power-system member <i>vc-member-number</i>]</p> <p>If two or more connections are assigned the same , then the power of each connection is determined based on its switch connector port location, with the rightmost port receiving power first.</p>
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.

redundant-power-system

Syntax EX2200 switch:

```
redundant-power-system {
  priority (0|1|2|3|4|5|6)
}
```

EX3300 switch:

```
redundant-power-system {
  member vc-member-number {
    priority (0|1|2|3|4|5|6)
  }
}
```

Hierarchy Level [edit]

Release Information Statement introduced in Junos OS Release 12.1 for EX Series switches.

Description Configure Redundant Power System (RPS) member to ensure higher- switches always receive power backup.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- [Determining and Setting Priority for Switches Connected to an EX Series RPS on page 310](#)

CHAPTER 42

Configuration Statements: Routing Engine and Switching Control Board Redundancy

- [cfeb on page 570](#)
- [description \(Chassis Redundancy\) on page 571](#)
- [disk-failure-action on page 572](#)
- [failover \(Chassis\) on page 573](#)
- [failover \(Chassis\) on page 574](#)
- [failover \(System Process\) on page 575](#)
- [feb \(Creating a Redundancy Group\) on page 576](#)
- [feb \(Assigning a FEB to a Redundancy Group\) on page 577](#)
- [keepalive-time on page 578](#)
- [keepalive-time on page 579](#)
- [no-auto-failover on page 580](#)
- [on-disk-failure \(Chassis Redundancy Failover\) on page 580](#)
- [on-disk-failure on page 581](#)
- [on-loss-of-keepalives on page 582](#)
- [on-loss-of-keepalives on page 583](#)
- [redundancy on page 584](#)
- [redundancy-group on page 585](#)
- [routing-engine \(Chassis Redundancy\) on page 586](#)
- [routing-engine on page 587](#)
- [sfm \(Chassis Redundancy\) on page 588](#)
- [ssb on page 589](#)
- [vcp-no-hold-time on page 590](#)

cfeb

Syntax	<code>cfeb slot-number (always preferred);</code>
Hierarchy Level	[edit chassis redundancy]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	On M10i routers only, configure which Compact Forwarding Engine Board (CFEB) is the master and which is the backup.
Default	By default, the CFEB in slot 0 is the master and the CFEB in slot 1 is the backup.
Options	<p>slot-number—Specify which slot is the master and which is the backup.</p> <p>always—Define this CFEB as the sole device.</p> <p>preferred—Define this CFEB as the preferred device of at least two.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring CFEB Redundancy on the M10i Router on page 23

description (Chassis Redundancy)

Syntax `description description;`

Hierarchy Level [edit chassis [redundancy feb redundancy-group](#) *group-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Provide a description of the FEB redundancy group.

Options *description*—Provide a description for the FEB redundancy group.

Required Privilege interface—To view this statement in the configuration.

Level interface-control—To add this statement to the configuration.

Related Documentation • [Configuring FEB Redundancy on the M120 Router on page 24](#)

disk-failure-action

Syntax	disk-failure-action (halt reboot);
Hierarchy Level	[edit chassis redundancy on-disk-failure] [edit chassis routing-engine on-disk-failure]
Release Information	Statement introduced in Junos OS Release 9.2 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the Routing Engine to halt or reboot when the Routing Engine hard disk fails.
Options	halt —Specify the Routing Engine to halt. reboot —Specify the Routing Engine to reboot.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• graceful-switchover on page 514• <i>Enabling a Routing Engine to Reboot on Hard Disk Errors</i>• <i>Installing Software on an EX Series Switch with Redundant Routing Engines (CLI Procedure)</i>• High Availability Features for EX Series Switches Overview on page 9

failover (Chassis)

Syntax	<pre>failover { on-disk-failure; on-loss-of-keepalives; on-re-to-fpc-stale; }</pre>
Hierarchy Level	[edit chassis redundancy]
Release Information	Statement introduced before Junos OS Release 7.4. on-re-to-fpc-stale option introduced in Junos OS Release 15.2 on the MX240, MX480, MX960, MX2010, and MX2020.
Description	<p>Specify conditions on the master Routing Engine that cause the backup router to take mastership.</p> <p>The remaining statements are explained separately. See CLI Explorer.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• On Detection of a Hard Disk Error on the Master Routing Engine on page 120

failover (Chassis)

Syntax	<pre>failover { on-disk-failure; on-loss-of-keepalives; }</pre>
Hierarchy Level	[edit chassis redundancy]
Release Information	Statement introduced in Junos OS Release 9.2 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>Specify conditions on the master Routing Engine that cause the backup router to take mastership.</p> <p>The remaining statements are explained separately. See CLI Explorer.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• graceful-switchover on page 514• On Detection of a Hard Disk Error on the Master Routing Engine on page 120• <i>Installing Software on an EX Series Switch with Redundant Routing Engines (CLI Procedure)</i>• High Availability Features for EX Series Switches Overview on page 9

failover (System Process)

Syntax	failover (alternate-media other-routing-engine);
Hierarchy Level	[edit system processes <i>process-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure the router to reboot if the software process fails four times within 30 seconds, and specify the software to use during the reboot.
Options	<p>process-name—Junos OS process name. Some of the processes that support the failover statement are bootp, chassis-control, craft-control, ethernet-connectivity-fault-management, init, interface-control, neighbor-liveness, pfe, redundancy-interface-process, routing, smg-service, and vrrp.</p> <p>alternate-media—Use the Junos OS image on alternate media during the reboot.</p> <p>other-routing-engine—On routers with dual Routing Engines, use the Junos OS image on the other Routing Engine during the reboot. That Routing Engine assumes mastership; in the usual configuration, the other Routing Engine is the designated backup Routing Engine.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• When a Software Process Fails on page 122• <i>processes</i>

feb (Creating a Redundancy Group)

Syntax feb {
 [redundancy-group](#) *group-name* {
 [description](#) *description*;
 [feb](#) *slot-number* (backup | primary);
 [no-auto-failover](#);
 }
 }

Hierarchy Level [edit chassis [redundancy](#)]

Release Information Statement introduced in Junos OS Release 8.2.

Description On M120 routers only, configure a Forwarding Engine Board (FEB) redundancy group.

Options The remaining statements are described separately.

Required Privilege interface—To view this statement in the configuration.
Level interface-control—To add this statement to the configuration.

Related Documentation • [Configuring FEB Redundancy on the M120 Router on page 24](#)

feb (Assigning a FEB to a Redundancy Group)

Syntax	<code>feb slot-number (backup primary);</code>
Hierarchy Level	[edit chassis redundancy feb redundancy-group group-name]
Release Information	Statement introduced in Junos OS Release 8.2.
Description	On M120 routers only, configure a Forwarding Engine Board (FEB) as part of a FEB redundancy group.
Options	<p>slot-number—Slot number of the FEB. The range of values is from 0 to 5.</p> <p>backup—(Optional) For each redundancy group, you must configure exactly one backup FEB.</p> <p>primary—(Optional) For each redundancy group, you can optionally configure one primary FEB.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring FEB Redundancy on the M120 Router on page 24

keepalive-time

Syntax	keepalive-time <i>seconds</i> ;
Hierarchy Level	[edit chassis redundancy]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure the time period that must elapse before the backup router takes mastership when it detects loss of the keepalive signal.
Default	<p>The on-loss-of-keepalives statement at the [edit chassis redundancy failover] hierarchy level must be included for failover to occur.</p> <p>When the on-loss-of-keepalives statement is included and graceful Routing Engine switchover <i>is not</i> configured, failover occurs after 300 seconds (5 minutes).</p> <p>When the on-loss-of-keepalives statement is included and graceful Routing Engine switchover <i>is</i> configured, the keepalive signal is automatically enabled and the failover time is set to 2 seconds (4 seconds on M20 routers). You cannot manually reset the keepalive time.</p>
Options	seconds —Time before the backup router takes mastership when it detects loss of the keepalive signal. The range of values is 2 through 10,000.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• On Detection of a Loss of Keepalive Signal from the Master Routing Engine on page 121• failover (Chassis) on page 573• on-loss-of-keepalives on page 582

keepalive-time

Syntax	keepalive-time <i>seconds</i> ;
Hierarchy Level	[edit chassis redundancy]
Release Information	Statement introduced in Junos OS Release 9.2 for EX Series switches.
Description	Configure the time period that must elapse before the backup router takes mastership when it detects loss of the keepalive signal.
Default	<p>The on-loss-of-keepalives statement at the [edit chassis redundancy failover] hierarchy level must be included for failover to occur.</p> <p>When the on-loss-of-keepalives statement is included and graceful Routing Engine switchover <i>is not</i> configured, failover occurs after 300 seconds (5 minutes).</p> <p>When the on-loss-of-keepalives statement is included and graceful Routing Engine switchover <i>is</i> configured, the keepalive signal is automatically enabled and the failover time is set to 2 seconds.</p>
Options	seconds —Time before the backup router takes mastership when it detects loss of the keepalive signal. The range of values is 2 through 10,000.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• failover on page 574• graceful-switchover on page 514• on-loss-of-keepalives on page 583• High Availability Features for EX Series Switches Overview on page 9

no-auto-failover

Syntax	no-auto-failover;
Hierarchy Level	[edit chassis redundancy feb redundancy-group <i>group-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Disable automatic failover to a backup FEB when an active FEB in a redundancy group fails.
Default	Automatic failover is enabled by default.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring FEB Redundancy on the M120 Router on page 24

on-disk-failure (Chassis Redundancy Failover)

Syntax	on-disk-failure;
Hierarchy Level	[edit chassis redundancy failover]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Instruct the backup router to take mastership if it detects hard disk errors on the master Routing Engine.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• On Detection of a Hard Disk Error on the Master Routing Engine on page 120

on-disk-failure

Syntax	<pre>on-disk-failure { disk-failure-action (halt reboot); }</pre>
Hierarchy Level	[edit chassis redundancy] [edit chassis routing-engine]
Release Information	Statement introduced in Junos OS Release 9.2 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Instruct the router to halt or reboot if it detects hard disk errors on the Routing Engine.
Options	The remaining statement is explained separately. See CLI Explorer .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• graceful-switchover on page 514• <i>Enabling a Routing Engine to Reboot on Hard Disk Errors</i>• <i>Installing Software on an EX Series Switch with Redundant Routing Engines (CLI Procedure)</i>• High Availability Features for EX Series Switches Overview on page 9

on-loss-of-keepalives

Syntax	on-loss-of-keepalives;
Hierarchy Level	[edit chassis redundancy failover]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Instruct the backup router to take mastership if it detects a loss of keepalive signal from the master Routing Engine.
Default	<p>The on-loss-of-keepalives statement must be included at the [edit chassis redundancy failover] hierarchy level for failover to occur.</p> <p>When the on-loss-of-keepalives statement is included but graceful Routing Engine switchover <i>is not</i> configured, failover occurs after 300 seconds (5 minutes).</p> <p>When the on-loss-of-keepalives statement is included and graceful Routing Engine switchover <i>is</i> configured, the keepalive signal is automatically enabled and the failover time is set to 2 seconds (4 seconds on M20 routers) . The keepalive time is not configurable.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• On Detection of a Loss of Keepalive Signal from the Master Routing Engine on page 121• keepalive-time on page 578

on-loss-of-keepalives

Syntax	on-loss-of-keepalives;
Hierarchy Level	[edit chassis redundancy failover]
Release Information	Statement introduced in Junos OS Release 9.2 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Instruct the backup router to take mastership if it detects a loss of keepalive signal from the master Routing Engine.
Default	<p>The on-loss-of-keepalives statement must be included at the [edit chassis redundancy failover] hierarchy level for failover to occur.</p> <p>When the on-loss-of-keepalives statement is included but graceful Routing Engine switchover <i>is not</i> configured, failover occurs after 300 seconds (5 minutes).</p> <p>When the on-loss-of-keepalives statement is included and graceful Routing Engine switchover <i>is</i> configured, the keepalive signal is automatically enabled and the failover time is set to 2 seconds.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • graceful-switchover on page 514 • keepalive-time on page 579 • <i>Installing Software on an EX Series Switch with Redundant Routing Engines (CLI Procedure)</i> • High Availability Features for EX Series Switches Overview on page 9

redundancy

Syntax	<pre>redundancy { cfeb slot (always preferred); failover { on-disk-failure; on-loss-of-keepalives; on-re-to-fpc-stale; } feb { redundancy-group group-name { description description; feb slot-number (backup primary); no-auto-failover; } } graceful-switchover; keepalive-time seconds; routing-engine slot-number (backup disabled master); sfm slot-number (always preferred); ssb slot-number (always preferred); }</pre>
Hierarchy Level	[edit chassis]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure redundancy options.
Options	The remaining statements are explained separately. See CLI Explorer .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Routing Engine Redundancy on page 119• Configuring CFEB Redundancy on the M10i Router on page 23• Configuring FEB Redundancy on the M120 Router on page 24• Configuring SFM Redundancy on M40e and M160 Routers on page 26• Configuring SSB Redundancy on the M20 Router on page 27

redundancy-group

Syntax `redundancy-group group-name {
 description description;
 feb slot-number (backup | primary);
 no-auto-failover;
 }`

Hierarchy Level [edit chassis [redundancy feb](#)]

Release Information Statement introduced in Junos OS Release 8.2.

Description On M120 routers only, configure a Forwarding Engine Board (FEB) redundancy group.

Options *group-name* is the unique name for the redundancy group. The maximum length is 39 alphanumeric characters.

Other statements are explained separately.

Required Privilege interface—To view this statement in the configuration.
Level interface-control—To add this statement to the configuration.

Related Documentation • [Configuring FEB Redundancy on the M120 Router on page 24](#)

routing-engine (Chassis Redundancy)

Syntax	<code>routing-engine <i>slot-number</i> (backup disabled master);</code>
Hierarchy Level	[edit chassis redundancy]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure Routing Engine redundancy.
Default	By default, the Routing Engine in slot 0 is the master Routing Engine and the Routing Engine in slot 1 is the backup Routing Engine.
Options	<p><i>slot-number</i>—Specify the slot number (0 or 1).</p> <p>Set the function of the Routing Engine for the specified slot:</p> <ul style="list-style-type: none">• master—Routing Engine in the specified slot is the master.• backup—Routing Engine in the specified slot is the backup.• disabled—Routing Engine in the specified slot is disabled.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring Routing Engine Redundancy on page 119

routing-engine

Syntax	<pre>routing-engine { on-disk-failure { disk-failure-action (halt reboot); } }</pre>
Hierarchy Level	[edit chassis]
Release Information	Statement introduced in Junos OS Release 9.2 for EX Series switches.
Description	Configure a Routing Engine to halt or reboot automatically when a hard disk error occurs. A hard disk error may cause a Routing Engine to enter a state in which it responds to local pings and interfaces remain up, but no other processes are responding. Rebooting or halting prevents this.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• graceful-switchover on page 514• <i>Enabling a Routing Engine to Reboot on Hard Disk Errors</i>• High Availability Features for EX Series Switches Overview on page 9


sfm (Chassis Redundancy)

Syntax	<code>sfm slot-number (always preferred);</code>
Hierarchy Level	[edit chassis redundancy]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	On M40e and M160 routers, configure which Switching and Forwarding Module (SFM) is the master and which is the backup.
Default	By default, the SFM in slot 0 is the master and the SFM in slot 1 is the backup.
Options	<p>slot-number—Specify which slot is the master and which is the backup. On the M40e router, slot-number can be 0 or 1. On the M160 router, slot-number can be 0 through 3.</p> <p>always—Define this SFM as the sole device.</p> <p>preferred—Define this SFM as the preferred device of at least two.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring SFM Redundancy on M40e and M160 Routers on page 26

ssb

Syntax	<code>ssb slot-number (always preferred);</code>
Hierarchy Level	[edit chassis redundancy]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	On M20 routers, configure which System and Switch Board (SSB) is the master and which is the backup.
Default	By default, the SSB in slot 0 is the master and the SSB in slot 1 is the backup.
Options	<p><code>slot-number</code>—Specify which slot is the master and which is the backup.</p> <p><code>always</code>—Define this SSB as the sole device.</p> <p><code>preferred</code>—Define this SSB as the preferred device of at least two.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring SSB Redundancy on the M20 Router on page 27

vcp-no-hold-time

Syntax	vcp-no-hold-time;
Hierarchy Level	[edit virtual-chassis]
Release Information	Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches. Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.
Description	<p>Disable the Virtual Chassis port (VCP) holddown timer for all VCPs in the Virtual Chassis or Virtual Chassis Fabric (VCF).</p> <p>The VCP holddown timer is an internal mechanism that delays a Virtual Chassis reconvergence for several seconds when a VCP becomes inactive. The purpose of this delay is to provide the VCP time to return online without having to reconverge the Virtual Chassis to adjust to the inactive VCP. All traffic to the VCP is dropped while the VCP is inactive. If the VCP remains down for a time that exceeds the VCP holddown timer, a Virtual Chassis reconvergence occurs.</p> <p>When this statement is enabled, the VCP holddown timer is disabled and the Virtual Chassis reconvergence occurs when a VCP becomes inactive. The period of time where traffic is dropped waiting for the VCP to return online is avoided.</p> <p>We recommend enabling this statement after a Virtual Chassis is operational. We recommend disabling this statement when you are adding or removing member switches from your Virtual Chassis.</p> <p>The VCP holddown timer cannot be viewed and is not user-configurable. You can only control whether the VCP holddown timer is enabled or disabled by configuring this statement.</p> <div> NOTE: For the EX4300 Virtual Chassis, you should enable the <code>vcp-no-hold-time</code> statement before performing a software upgrade using NSSU. If you do not enable the <code>vcp-no-hold-time</code> statement, the Virtual Chassis may split during the upgrade. A split Virtual Chassis can cause disruptions to your network, and you may have to manually reconfigure your Virtual Chassis after the NSSU if the split and merge feature was disabled. For more information about a split Virtual Chassis, see <i>Understanding Split and Merge in a Virtual Chassis</i></div>
Default	The VCP holddown timer is enabled by default on all devices that support this statement.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.

- Related Documentation**
- *Understanding EX Series Virtual Chassis*
 - *Understanding QFX Series Virtual Chassis*
 - *Understanding Virtual Chassis Components*

CHAPTER 43

Configuration Statements: Unified ISSU

- [no-issu-timer-negotiation](#) on page 593
- [traceoptions \(Protocols BFD\)](#) on page 594

[no-issu-timer-negotiation](#)

Syntax `no-issu-timer-negotiation;`

Hierarchy Level `[edit protocols bfd],`
 `[edit logical-systems logical-system-name protocols bfd],`
 `[edit routing-instances routing-instance-name protocols bfd]`

Release Information Statement introduced in Junos OS Release 9.1.
 Statement introduced in Junos OS Release 13.2 for PTX5000 routers.

Description Disable unified ISSU timer negotiation for Bidirectional Forwarding Detection (BFD) sessions.



CAUTION: The sessions might flap during unified ISSU or Routing Engine switchover, depending on the detection intervals.

Required Privilege Level routing—To view this statement in the configuration.
 routing-control—To add this statement to the configuration.

Related Documentation • [Managing and Tracing BFD Sessions During Unified ISSU Procedures](#) on page 462
 • [Junos OS Routing Protocols Library](#)

traceoptions (Protocols BFD)

Syntax	<pre>traceoptions { file <i>name</i> <size <i>size</i>> <files <i>number</i>> <world-readable no-world-readable>; flag <i>flag</i> <<i>flag-modifier</i>> <disable>; }</pre>
Hierarchy Level	[edit protocols bfd]
Release Information	Statement introduced before Junos OS Release 7.4. issu flag for BFD added in Junos OS Release 9.1.
Description	<p>Define tracing operations that track unified in-service software upgrade (ISSU) functionality in the router.</p> <p>To specify more than one tracing operation, include multiple flag statements.</p>
Default	If you do not include this statement, no global tracing operations are performed.
Options	<p>disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.</p> <p>file <i>name</i>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log. We recommend that you place global routing protocol tracing output in the file routing-log.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>Range: 2 through 1000 files</p> <p>Default: 2 files</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the size option.</p> <p>flag <i>flag</i>—Tracing operation to perform. The tracing options are as follows:</p> <ul style="list-style-type: none">• adjacency—Trace adjacency messages.• all—Trace everything.• error—Trace all errors.• events—Trace all events.

- **issu**—Trace ISSU packet activity.
- **nsr-packet**—Trace packet activity of NSR.
- **nsr-synchronization**—Trace NSR synchronization events.
- **packet**—Trace all packets.
- **pipe**—Trace pipe messages.
- **pipe-detail**—Trace pipe messages in detail.
- **ppm-packet**—Trace packet activity by periodic packet management.
- **state**—Trace state transitions.
- **timer**—Trace timer processing.

no-world-readable—Restrict users from reading the log file.

size size—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

Syntax: *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 128 KB

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

world-readable—Allow users to read the log file.

Required Privilege Level	routing and trace—To view this statement in the configuration. routing-control and trace-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Managing and Tracing BFD Sessions During Unified ISSU Procedures on page 462

Configuration Statements: VRRP

- [accept-data](#) on page 599
- [advertise-interval](#) on page 600
- [asymmetric-hold-time](#) on page 601
- [asymmetric-hold-time](#) on page 602
- [authentication-key](#) on page 603
- [authentication-type](#) on page 604
- [bandwidth-threshold](#) on page 605
- [delegate-processing \(VRRP\)](#) on page 606
- [failover-delay](#) on page 607
- [fast-interval](#) on page 608
- [global-advertisements-threshold](#) on page 609
- [hold-time \(VRRP\)](#) on page 610
- [hold-time](#) on page 611
- [inherit-advertisement-interval](#) on page 611
- [inet6-advertise-interval](#) on page 612
- [inet6-advertise-interval](#) on page 613
- [interface](#) on page 614
- [preempt \(VRRP\)](#) on page 615
- [preempt](#) on page 616
- [priority \(Protocols VRRP\)](#) on page 617
- [priority](#) on page 618
- [priority-cost \(VRRP\)](#) on page 619
- [priority-hold-time](#) on page 620
- [route \(Interfaces\)](#) on page 621
- [skew-timer-disable](#) on page 622
- [startup-silent-period](#) on page 623
- [traceoptions \(Protocols VRRP\)](#) on page 624
- [traceoptions](#) on page 626

- [track \(VRRP\) on page 628](#)
- [version-3 on page 629](#)
- [virtual-address on page 630](#)
- [virtual-inet6-address on page 631](#)
- [virtual-inet6-address on page 632](#)
- [virtual-link-local-address on page 633](#)
- [virtual-link-local-address on page 634](#)
- [vrrp-group on page 635](#)
- [vrrp-inet6-group on page 637](#)
- [vrrp-inet6-group on page 638](#)
- [vrrp-inherit-from on page 639](#)

accept-data

Syntax (accept-data | no-accept-data);

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* *vrrp-group group-id*],
 [edit interfaces *interface-name* unit *logical-unit-number* family inet6 address *address* *vrrp-inet6-group group-id*],
 [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet address *address* *vrrp-group group-id*],
 [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet6 address *address* *vrrp-inet6-group group-id*]

Release Information Statement introduced before Junos OS Release 7.4.
 Statement introduced in Junos OS 11.3 for the QFX Series.
 Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.

Description In a Virtual Router Redundancy Protocol (VRRP) configuration, determine whether or not a router that is acting as the master router accepts all packets destined for the virtual IP address.

- **accept-data**—Enable the master router to accept all packets destined for the virtual IP address.
- **no-accept-data**—Prevent the master router from accepting packets other than the ARP packets destined for the virtual IP address.

Default If the router acting as the master router is the IP address owner or has its priority set to 255, the master router, by default, responds to all packets sent to the virtual IP address. However, if the router acting as the master router does not own the IP address or has its priority set to a value less than 255, the master router responds only to ARP requests.



NOTE:

- If you want to restrict the incoming IP packets to ICMP packets only, you must configure firewall filters to accept only ICMP packets.
- If you include the **accept-data** statement, your routing platform configuration does not comply with RFC 3768 (see section 6.4.3 of RFC 3768, *Virtual Router Redundancy Protocol (VRRP)*).

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

- Related Documentation**
- [Configuring an Interface to Accept All Packets Destined for the Virtual IP Address of a VRRP Group on page 371](#)

advertise-interval

Syntax `advertise-interval seconds;`

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* **vrrp-group** *group-id*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet address *address* **vrrp-group** *group-id*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS 11.3 for the QFX Series.
Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.

Description Configure the interval between Virtual Router Redundancy Protocol (VRRP) IPv4 advertisement packets.

All routers in the VRRP group must use the same advertisement interval.



NOTE: When VRRPv3 is enabled, the `advertise-interval` statement cannot be used to configure advertisement intervals. Instead, use the `fast-interval` statement to configure advertisement intervals.

Options ***seconds***—Interval between advertisement packets.
Range: 1 through 255 seconds
Default: 1 second

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

- Related Documentation**
- [Configuring the Advertisement Interval for the VRRP Master Router on page 350](#)
 - [fast-interval on page 608](#)
 - [inet6-advertise-interval on page 612](#)
 - [version-3 on page 629](#)


asymmetric-hold-time

Syntax	asymmetric-hold-time;
Hierarchy Level	[edit protocols vrrp]
Release Information	Statement introduced in Junos OS Release 9.5.
Description	Enable the VRRP master router to switch over to the backup router immediately, without waiting for the priority hold time to expire, when a tracked route or interface goes down. When the route or interface comes back online, the original master router that is now acting as the backup router waits for the priority hold time to expire before it reasserts mastership.
Default	asymmetric-hold-time is disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Asymmetric Hold Time for VRRP Routers on page 355


asymmetric-hold-time

Syntax	asymmetric-hold-time;
Hierarchy Level	[edit protocols vrrp]
Release Information	Statement introduced in Junos OS 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.
Description	<p>Configure a VRRP master to fail over to a backup immediately—without waiting for the priority hold time to expire—when a tracked route goes down. Otherwise, the master waits for the hold time to expire before it initiates a failover when a tracked interface or route goes down.</p> <p>When the tracked interface or route comes up again, the new backup (original master) router waits for the priority hold time to expire before it reasserts mastership.</p>
Default	asymmetric-hold-time is disabled.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring VRRP Preemption and Hold Time on page 349

authentication-key

Syntax	<code>authentication-key key;</code>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrrp-group <i>group-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrrp-group <i>group-id</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.</p> <p>Statement introduced in Junos OS Release 18.1R1 for the SRX Series devices.</p>
Description	<p>Configure a Virtual Router Redundancy Protocol (VRRP) IPv4 authentication key. You also must specify a VRRP authentication scheme by including the authentication-type statement.</p> <p>All devices in the VRRP group must use the same authentication scheme and password.</p>
	<p> NOTE: When VRRPv3 is enabled, the authentication-type and authentication-key statements cannot be configured for any VRRP groups.</p>
Options	<p>key—Authentication password. For simple authentication, it can be 1 through 8 characters long. For Message Digest 5 (MD5) authentication, it can be 1 through 16 characters long. If you include spaces, enclose all characters in quotation marks (" ").</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring VRRP Authentication (IPv4 Only) on page 348 • authentication-type on page 604 • version-3 on page 629 • <i>Understanding VRRP on SRX Series Devices</i> • <i>Example: Configuring VRRP/VRRPv3 on Chassis Cluster Redundant Ethernet Interfaces</i>

authentication-type

Syntax	<code>authentication-type (md5 simple);</code>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrp-group <i>group-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrp-group <i>group-id</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.</p> <p>Statement introduced in Junos OS Release 18.1R1 for the SRX Series devices.</p>
Description	<p>Enable Virtual Router Redundancy Protocol (VRRP) IPv4 authentication and specify the authentication scheme for the VRRP group. If you enable authentication, you must specify a password by including the authentication-key statement. The specific type of authentication used by OSPF is encoded in this field.</p> <p>All devices in the VRRP group must use the same authentication scheme and password.</p>
<div>  <p>NOTE: When VRRPv3 is enabled, the authentication-type and authentication-key statements cannot be configured for any VRRP groups.</p> </div>	
Options	<p>authentication—Authentication scheme:</p> <ul style="list-style-type: none"> simple—Use a simple password. The password is included in the transmitted packet, so this method of authentication is relatively insecure. md5—Use the MD5 algorithm to create an encoded checksum of the packet. The encoded checksum is included in the transmitted packet. The receiving routing platform uses the authentication key to verify the packet, discarding it if the digest does not match. This algorithm provides a more secure authentication scheme. <p>Default: none (no authentication is performed).</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring VRRP Authentication (IPv4 Only) on page 348 • authentication-key on page 603 • version-3 on page 629

- *Understanding VRRP on SRX Series Devices*
- *Example: Configuring VRRP/VRRPv3 on Chassis Cluster Redundant Ethernet Interfaces*

bandwidth-threshold

Syntax	<code>bandwidth-threshold <i>bits-per-second</i> priority-cost <i>priority</i>;</code>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrrp-group <i>group-id</i> track interface <i>interface-name</i>],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group <i>group-id</i> track interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrrp-group <i>group-id</i> track interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group <i>group-id</i> track interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.</p>
Description	Specify the bandwidth threshold for Virtual Router Redundancy Protocol (VRRP) logical interface tracking.
Options	<p><i>bits-per-second</i>—Bandwidth threshold for the tracked interface. When the bandwidth of the tracked interface drops below the specified value, the VRRP group uses the bandwidth threshold priority cost value. You can include up to five bandwidth threshold statements for each interface you track.</p> <p>Range: 1 through 100000000000000 bits per second</p> <p><i>priority-cost</i> <i>priority</i>—The value subtracted from the configured VRRP priority when the tracked interface or route is down to force a new master router election. The sum of all the costs for all interfaces or routes that are tracked must be less than or equal to the configured priority of the VRRP group.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring a Logical Interface to Be Tracked for a VRRP Group on page 358


delegate-processing (VRRP)

Syntax	<pre>delegate-processing { ae-irb; }</pre>
Hierarchy Level	[edit protocols vrrp]
Release Information	Statement introduced in Junos OS Release 9.6. ae-irb option introduced in Junos OS Release 15.1.
Description	<p>Configure the distributed periodic packet management process (ppmd) to send Virtual Router Redundancy Protocol (VRRP) advertisements .</p> <p>Using a hash logic based on iflIndex, the vrrp group ID, and the IP version, select one of the Flexible OIC Concentrators (FPCs) for distribution. The selected FPC is called the <i>anchor FPC</i>. All transmit instances and receive instances are from and to the anchor FPC. The anchor FPC is static, and VRRP is not guaranteed to get distributed to all available FPCs uniformly for all VRRP sessions.</p>
Options	<p>ae-irb—Enable distributed ppmmd for VRRP over aggregated Ethernet and integrated routing and bridging (IRB) interfaces.</p> <p>Using the ae-irb option is only for MPC line cards. ae-irb is not supported on small MX Series routing devices with built-in MPCs such as the MX104 and below. Using the ae-irb option requires use of the enhanced-ip mode.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Enabling the Distributed Periodic Packet Management Process for VRRP on page 373

failover-delay

Syntax	<code>failover-delay <i>milliseconds</i>;</code>
Hierarchy Level	[edit protocols vrrp]
Release Information	Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.
Description	<p>If you configure multiple VRRP groups on an interface (using multiple VLANs), traffic for some of the groups might be briefly dropped if a failover occurs. This can happen because the new master must send gratuitous ARP replies for each VRRP group to update the ARP tables in the connected devices, and there is a short delay between each gratuitous ARP reply. Traffic sent by devices that have not yet received the gratuitous ARP reply is dropped (until the device receives the reply and learns the MAC address of the new master).</p> <p>If you configure a failover delay, the new master delays sending gratuitous ARP replies for the period that you set. This allows the new master to send the ARP replies for all of the VRRP groups simultaneously.</p>
Options	<p><i>milliseconds</i>—Specify the failover delay time, in milliseconds.</p> <p>Range: 50 through 2000</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Troubleshooting VRRP on page 383 • <code>show vrrp</code>

fast-interval

Syntax	<code>fast-interval <i>milliseconds</i>;</code>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.</p>
Description	<p>Configure the interval, in milliseconds, between Virtual Router Redundancy Protocol (VRRP) advertisement packets.</p> <p>All routers in the VRRP group must use the same advertisement interval.</p>
Options	<p><i>milliseconds</i>—Interval between advertisement packets.</p> <p>Range: 10 through 40,950 milliseconds (range extended from 100–999 to 10–40,950 in Junos OS Release 12.2).</p>
<div>  <p>NOTE: When configuring VRRP for IPv4, if you have chosen not to enable VRRPv3, you cannot set a value less than 100 for <i>fast-interval</i>. Commit check fails if a value less than 100 is configured.</p> </div>	
Default:	1 second
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring the Advertisement Interval for the VRRP Master Router on page 350 • advertise-interval on page 600 • advertise-interval on page 600 • inet6-advertise-interval on page 612 • version-3 on page 629

global-advertisements-threshold

Syntax `global-advertisements-threshold advertisement-value;`

Hierarchy Level [edit protocols vrrp]

Release Information Statement introduced in Junos OS Release 12.2.

Description Configure the number of fast advertisements that can be missed by a backup router before the master router is declared as down.



NOTE:

- The advertisement value configured using the `global-advertisements-threshold` statement is applicable to all the Virtual Router Redundancy Protocol (VRRP) groups in the system.
- Setting the advertisement value of the `global-advertisements-threshold` configuration to 1 is not recommended for a scaled configuration with an aggressive advertisement interval. For example, if you have 1000 VRRP groups with an advertisement interval of 100 ms, then do not set the `global-advertisements-threshold` value to 1.
- Changing the advertisement value of the `global-advertisements-threshold` configuration during runtime can result in unpredictable behavior by the VRRP state machine. For example, momentary ownership change from the master router to the backup router and vice versa. Therefore, avoid changing the advertisement value of the `global-advertisements-threshold` statement during runtime.

Options `advertisement-value`—Number of VRRP advertisements missed before the master router is declared as down.

Range: 1 through 15

Default: 3

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [Improving the Convergence Time for VRRP on page 374](#)
- [Configuring VRRP to Improve Convergence Time on page 376](#)

hold-time (VRRP)

Syntax	<code>hold-time seconds;</code>
Hierarchy Level	<code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id preempt</i>],</code> <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id preempt</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id preempt</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id preempt</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS 11.3 for the QFX Series.
Description	In a Virtual Router Redundancy Protocol (VRRP) configuration, set the hold time before a higher-priority backup router preempts the master router.
Default	VRRP preemption is not timed.
Options	<i>seconds</i> —Hold-time period. Range: 0 through 3600 seconds Default: 0 seconds (VRRP preemption is not timed.)
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring a Backup Router to Preempt the VRRP Master Router on page 353• Configuring VRRP Preemption and Hold Time on page 349

hold-time

Syntax	<code>hold-time <i>seconds</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i> preempt]
Release Information	Statement introduced in Junos OS Release 10.0 for EX Series switches.
Description	Configure the time in seconds after which a backup router with the highest priority preempts the master router.
Options	<i>seconds</i> —Hold-time period.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring VRRP for IPv6 (CLI Procedure) on page 341

inherit-advertisement-interval

Syntax	<code>inherit-advertisement-interval <i>seconds</i>;</code>
Hierarchy Level	[edit protocols vrrp]
Release Information	Statement introduced in Junos OS Release 14.2R3.
Description	Set the time interval for advertisement for inherit sessions.
Options	inherit-advertisement-interval <i>seconds</i> —Time interval for inherit sessions advertisements in seconds. The default value is the recommended value. Default: 120 Range: 5 to 120
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> •

inet6-advertise-interval

Syntax `inet6-advertise-interval milliseconds;`

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family inet6 address *address* [vrrp-inet6-group group-id](#)],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet6 address *address* [vrrp-inet6-group group-id](#)]

Release Information Statement introduced in Junos OS Release 8.4R2.

Description Configure the interval between Virtual Router Redundancy Protocol (VRRP) IPv6 advertisement packets.

All routers in the VRRP group must use the same advertisement interval.



NOTE: When VRRPv3 is enabled, the `inet6-advertise-interval` statement cannot be used to configure advertisement intervals. Instead, use the `fast-interval` statement to configure advertisement intervals.

Options *milliseconds*—Interval, in milliseconds, between advertisement packets.
Range: 100 to 40,000 milliseconds (ms)
Default: 1 second

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring the Advertisement Interval for the VRRP Master Router on page 350](#)
- [advertise-interval on page 600](#)
- [fast-interval on page 608](#)
- [version-3 on page 629](#)

inet6-advertise-interval

Syntax	inet6-advertise-interval <i>milliseconds</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group <i>group-id</i>]
Release Information	Statement introduced in Junos OS Release 10.0 for EX Series switches.
Description	Configure the interval between Virtual Router Redundancy Protocol (VRRP) IPv6 advertisement packets.
Options	<i>milliseconds</i> —Interval, in milliseconds, between advertisement packets. Range: 100 to 40,000 ms Default: 1 second
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring VRRP for IPv6 (CLI Procedure) on page 341

interface

Syntax	<pre>interface <i>interface-name</i> { bandwidth-threshold <i>bits-per-second</i> priority-cost <i>priority</i>; priority-cost <i>priority</i>; }</pre>
Hierarchy Level	<pre>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrrp-group <i>group-id</i> track], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group <i>group-id</i> track], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrrp-group <i>group-id</i> track], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group <i>group-id</i> track]</pre>
Release Information	Statement introduced before Junos OS Release 7.4. bandwidth-threshold statement added in Junos OS Release 8.1. Statement introduced in Junos OS 11.3 for the QFX Series.
Description	Enable logical interface tracking for a Virtual Router Redundancy Protocol (VRRP) group.
Options	interface-name —Interface to be tracked for this VRRP group. Range: 1 through 10 interfaces The remaining statements are described separately.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring a Logical Interface to Be Tracked for a VRRP Group on page 358• <i>Junos OS Services Interfaces Library for Routing Devices</i>

preempt (VRRP)

Syntax (preempt | no-preempt) {
 hold-time seconds;
 }

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* *vrrp-group group-id*],
 [edit interfaces *interface-name* unit *logical-unit-number* family inet6 address *address* *vrrp-inet6-group group-id*],
 [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet address *address* *vrrp-group group-id*],
 [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet6 address *address* *vrrp-inet6-group group-id*]

Release Information Statement introduced before Junos OS Release 7.4.
 Statement introduced in Junos OS 11.3 for the QFX Series.
 Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.

Description In a Virtual Router Redundancy Protocol (VRRP) configuration, determine whether or not a backup router can preempt a master router:

- **preempt**—Allow the master router to be preempted.



NOTE: By default, a higher-priority backup router can preempt a lower-priority master router.

- **no-preempt**—Prohibit the preemption of the master router. When **no-preempt** is configured, the backup router cannot preempt the master router even if the backup router has a higher priority.

The remaining statement is explained separately. See [CLI Explorer](#).

Default By default the **preempt** statement is enabled, and a higher-priority backup router preempts a lower-priority master router even if the **preempt** statement is not explicitly configured.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring a Backup Router to Preempt the VRRP Master Router on page 353](#)
- [Configuring VRRP Preemption and Hold Time on page 349](#)


preempt

Syntax	(preempt no-preempt) { hold-time seconds; }
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group <i>group-id</i>]
Release Information	Statement introduced in Junos OS Release 10.0 for EX Series switches.
Description	<p>Configure whether a backup router can preempt a master router:</p> <ul style="list-style-type: none">• preempt—Allow the master router to be preempted.• no-preempt—Prohibit the preemption of the master router. <p>The remaining statement is explained separately. See CLI Explorer.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring VRRP for IPv6 (CLI Procedure) on page 341

priority (Protocols VRRP)

Syntax	<code>priority priority;</code>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.</p> <p>Statement introduced in Junos OS Release 18.1R1 for the SRX Series devices.</p>
Description	<p>Configure a Virtual Router Redundancy Protocol (VRRP) device's priority for becoming the master default device. The device with the highest priority within the group becomes the master. VRRP is designed to eliminate the single point of failure inherent in the static default routed environment. VRRP specifies an election protocol that dynamically assigns responsibility for a virtual router to one of the VRRP routers on a LAN. The VRRP router controlling the IP address(es) associated with a virtual router is called the Master, and forwards packets sent to these IP addresses. The election process provides dynamic fail-over in the forwarding responsibility when the Master become unavailable. Any of the virtual router's IP addresses on a LAN can then be used as the default first hop router by end-hosts. The advantage gained from using VRRP is a higher availability default path without requiring configuration of dynamic routing or router discovery protocols on every end-host.</p>
Options	<p>priority—Device's priority for being elected to be the master device in the VRRP group. A larger value indicates a higher priority for being elected.</p> <p>Range: 0 through 255</p> <p>Default: 100. If two or more devices have the highest priority in the VRRP group, the device with the VRRP interface that has the highest IP address becomes the master, and the others serve as backups.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Basic VRRP Support on page 332 • Understanding VRRP on SRX Series Devices • Example: Configuring VRRP/VRRPv3 on Chassis Cluster Redundant Ethernet Interfaces


priority

Syntax	<code>priority <i>number</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrrp-group <i>group-id</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group <i>group-id</i>]
Release Information	Statement introduced in Junos OS Release 10.0 for EX Series switches.
Description	Configure a switch's priority for becoming the master default routing platform. The routing platform with the highest priority within the group becomes the master.
Options	<i>number</i> —Routing platform's priority for being elected to be the master router in the VRRP group. A larger value indicates a higher priority for being elected. Range: 1 through 255 Default: 100 (for backup routers)
<div> NOTE: Priority 255 cannot be assigned to routed VLAN interfaces (RVIs).</div>	
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring VRRP for IPv6 (CLI Procedure) on page 341

priority-cost (VRRP)

Syntax	<code>priority-cost <i>priority</i>;</code>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id track interface interface-name</i>],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id track interface interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id track interface interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id track interface interface-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.</p> <p>Statement introduced in Junos OS Release 12.2 for ACX2000 Universal Metro Routers.</p>
Description	Configure a Virtual Router Redundancy Protocol (VRRP) router's priority cost for becoming the master default router. The router with the highest priority within the group becomes the master.
Options	<p><i>priority</i>—The value subtracted from the configured VRRP priority when the tracked interface or route is down to force a new master router election. The sum of all the costs for all interfaces or routes that are tracked must be less than or equal to the configured priority of the VRRP group.</p> <p>Range: 1 through 254</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring a Logical Interface to Be Tracked for a VRRP Group on page 358

priority-hold-time

Syntax	<code>priority-hold-time seconds;</code>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrrp-group group-id track],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group group-id track],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrrp-group group-id track],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group group-id track]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.</p>
Description	<p>Configure a Virtual Router Redundancy Protocol (VRRP) router's priority hold time to define the minimum length of time that must elapse between dynamic priority changes. If the dynamic priority changes because of a tracking event, the priority hold timer begins running. If another tracking event or manual configuration change occurs while the timer is running, the new dynamic priority update is postponed until the timer expires.</p>
	<p> NOTE: When the track feature is configured, and if VRRP should pre-empt due to the tracking interface or route transition, any configured pre-empt hold time will be ignored. VRRP master will pre-empt according to the configuration of the priority-hold time.</p>
Options	<p>seconds—Minimum length of time that must elapse between dynamic priority changes.</p> <p>Range: 0through 3600 seconds</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring a Logical Interface to Be Tracked for a VRRP Group on page 358

route (Interfaces)

Syntax	<code>route <i>prefix</i> routing-instance <i>instance-name</i> priority-cost <i>priority</i>;</code>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrrp-group <i>group-id</i> track],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group <i>group-id</i> track],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrrp-group <i>group-id</i> track],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group <i>group-id</i> track]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.0.</p> <p>Statement introduced in Junos OS 11.3 for QFX Series.</p> <p>Statement introduced in Junos OS 12.1 for EX Series switches.</p>
Description	Enable route tracking for a Virtual Router Redundancy Protocol (VRRP) group.
Options	<p><i>prefix</i>—Route to be tracked for this VRRP group.</p> <p><i>priority-cost <i>priority</i></i>—The value subtracted from the configured VRRP priority when the tracked interface or route is down, forcing a new master router election. The sum of all the costs for all interfaces or routes that are tracked must be less than or equal to the configured priority of the VRRP group.</p> <p><i>routing-instance <i>instance-name</i></i>—Routing instance in which the route is to be tracked. If the route is in the default, or global, routing instance, the value for <i>instance-name</i> must be default.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring a Route to Be Tracked for a VRRP Group on page 361

skew-timer-disable

Syntax skew-timer-disable;

Hierarchy Level [edit protocols vrrp]

Release Information Statement introduced in Junos OS Release 12.2.

Description Disable the skew timer, thereby reducing the time required to transition from the backup state to the master state.



NOTE: The skew-timer-disable statement is used when there is only one master router and one backup router in the network.

Default By default, the skew timer is enabled for all the VRRP groups.

Required Privilege Level routing—To view this statement in the configuration.
 routing-control—To add this statement to the configuration.

Related Documentation

- [Improving the Convergence Time for VRRP on page 374](#)
- [Configuring VRRP to Improve Convergence Time on page 376](#)

startup-silent-period

Syntax	<code>startup-silent-period <i>seconds</i>;</code>
Hierarchy Level	[edit protocols vrrp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.
Description	Instruct the system to ignore the Master Down Event when an interface transitions from the down state to the up state. This statement is used to avoid incorrect error alarms caused by the delay or interruption of incoming Virtual Router Redundancy Protocol (VRRP) advertisement packets during the interface startup phase.
Options	<i>seconds</i> —Number of seconds for the startup period. Default: 4 seconds Range: 1 through 2000 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Startup Period for VRRP Operations on page 353

traceoptions (Protocols VRRP)

Syntax	<pre>traceoptions { file <i>filename</i> <files <i>number</i>> <match <i>regular-expression</i>> <microsecond-stamp> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i>; no-remote-trace; }</pre>
Hierarchy Level	[edit protocols vrrp]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>Define tracing operations for the Virtual Router Redundancy Protocol (VRRP) process.</p> <p>To specify more than one tracing operation, include multiple flag statements.</p> <p>By default, VRRP logs the error, dcd configuration, and routing socket events in a file in the directory /var/log.</p>
Default	If you do not include this statement, no VRRP-specific tracing operations are performed.
Options	<p>file <i>filename</i>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log. By default, VRRP tracing output is placed in the file vrrpd.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files. When a trace file named <i>trace-file</i> reaches its maximum size, it is renamed <i>trace-file.0</i>, then <i>trace-file.1</i>, and so on, until the maximum number of trace files is reached. When the maximum number is reached, the oldest trace file is overwritten.</p> <p>Range: 0 through 4,294,967,296 files</p> <p>Default: 3 files</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the size option.</p> <p>flag <i>flag</i>—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. These are the VRRP-specific tracing options:</p> <ul style="list-style-type: none">• all—All VRRP tracing operations• database—Database changes• general—General events• interfaces—Interface changes

- **normal**—Normal events
- **packets**—Packets sent and received
- **state**—State transitions
- **timer**—Timer events

match *regular-expression*—(Optional) Refine the output to include only those lines that match the given regular expression.

microsecond-stamp—(Optional) Provide a timestamp with microsecond granularity.

no-world-readable—(Optional) Restrict users from reading the log file.

size *size*—(Optional) Maximum size of each trace file, in kilobytes, megabytes, or gigabytes. When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file.0***. When the ***trace-file*** again reaches its maximum size, ***trace-file.0*** is renamed ***trace-file.1*** and ***trace-file*** is renamed ***trace-file.0***. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

Syntax: *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB

Range: 10 KB through the maximum file size supported on your routing platform

Default: 1 MB

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

world-readable—(Optional) Allow users to read the log file.

Required Privilege Level	trace—To view this statement in the configuration.
	trace-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Tracing VRRP Operations on page 377

traceoptions

Syntax `traceoptions {
 file <filename> <files number> <match regular-expression> <microsecond-stamp>
 <size size> <world-readable | no-world-readable>;
 flag flag;
 no-remote-trace;
 }`

Hierarchy Level [edit protocols vrrp]

Release Information Statement introduced in Junos OS Release 11.3 for the QFX Series.
 Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.

Description Define tracing operations for the Virtual Router Redundancy Protocol (VRRP) process.

To specify more than one tracing operation, include multiple **flag** statements.

By default, VRRP logs the error, dcd configuration, and routing socket events in a file in the directory **/var/log**.



NOTE: The traceoptions statement is not supported on a QFabric system.

Default If you do not include this statement, no VRRP-specific tracing operations are performed.

Options **filename filename**—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. By default, VRRP tracing output is placed in the file **vrrpd**.

files number—(Optional) Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. When the maximum number is reached, the oldest trace file is overwritten.

Range: 0 through 4,294,967,296 files

Default: 3 files

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

flag flag—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. These are the VRRP-specific tracing options:

- **all**—All VRRP tracing operations

- **database**—Database changes
- **general**—General events
- **interfaces**—Interface changes
- **normal**—Normal events
- **packets**—Packets sent and received
- **state**—State transitions
- **timer**—Timer events

match *regex*—(Optional) Refine the output to include only those lines that match the given regular expression.

microsecond-stamp—(Optional) Provide a timestamp with microsecond granularity.

no-world-readable—Restrict users from reading the log file.

size *size*—(Optional) Maximum size of each trace file, in kilobytes, megabytes, or gigabytes. When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file.0***. When the ***trace-file*** again reaches its maximum size, ***trace-file.0*** is renamed ***trace-file.1*** and ***trace-file*** is renamed ***trace-file.0***. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

Syntax: *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB

Range: 10 KB through the maximum file size supported on your routing platform

Default: 1 MB

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

world-readable—Allow users to read the log file.

Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
---------------------------------	---

Related Documentation	• Tracing VRRP Operations on page 377
------------------------------	---

track (VRRP)

Syntax	<pre>track { interface <i>interface-name</i> { bandwidth-threshold <i>bits-per-second</i> priority-cost <i>priority</i>; priority-cost <i>priority</i>; } priority-hold-time <i>seconds</i>; route <i>prefix/prefix-length</i> routing-instance <i>instance-name</i> priority-cost <i>priority</i>; }</pre>
Hierarchy Level	<pre>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>priority-hold-time statement added in Junos OS Release 8.1.</p> <p>route statement added in Junos OS Release 9.0.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.</p>
Description	Enable logical interface tracking, route tracking, or both, for a Virtual Router Redundancy Protocol (VRRP) group.
Options	The remaining statements are described separately.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring a Logical Interface to Be Tracked for a VRRP Group on page 358• Configuring a Route to Be Tracked for a VRRP Group on page 361

version-3

Syntax `version-3;`

Hierarchy Level `[edit protocols vrrp]`

Release Information Statement introduced in Junos OS Release 12.2.

Description Enable Virtual Router Redundancy Protocol version 3 (VRRPv3).



NOTE:

- Even though the `version-3` statement can be configured only at the `[edit protocols vrrp]` hierarchy level, VRRPv3 is enabled on all the configured logical systems as well.
- When enabling VRRPv3, you must ensure that VRRPv3 is enabled on all the VRRP routers in the network. This is because VRRPv3 does not interoperate with the previous versions of VRRP.

Required Privilege Level `routing`—To view this statement in the configuration.
 `routing-control`—To add this statement to the configuration.

Related Documentation • [Junos OS Support for VRRPv3 on page 322](#)

virtual-address

Syntax	<code>virtual-address [<i>addresses</i>];</code>
Hierarchy Level	<code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series. Statement introduced in Junos OS Release 18.1R1 for the SRX Series devices.
Description	Configure the addresses of the devices in a Virtual Router Redundancy Protocol (VRRP) IPv4 or IPv6 group. You can configure up to eight addresses.
Options	addresses —Addresses of one or more devices. Do not include a prefix length. If the address is the same as the interface's physical address, the interface becomes the master device for the group.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Basic VRRP Support on page 332• Understanding VRRP on SRX Series Devices• Example: Configuring VRRP/VRRPv3 on Chassis Cluster Redundant Ethernet Interfaces

virtual-inet6-address

Syntax	<code>virtual-inet6-address [<i>addresses</i>];</code>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group group-id],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group group-id]</p>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure the addresses of the virtual routers in a Virtual Router Redundancy Protocol (VRRP) IPv6 group. You can configure up to eight addresses.
Options	<i>addresses</i> —Addresses of one or more virtual routers. Do not include a prefix length. If the address is the same as the interface's physical address, the interface becomes the master virtual router for the group.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Basic VRRP Support on page 332

virtual-inet6-address

Syntax `virtual-inet6-address [addresses];`

Hierarchy Level `[edit interfaces interface-name unit logical-unit-number family inet6 address address vrrp-inet6-group group-id]`

Release Information Statement introduced in Junos OS Release 10.0 for EX Series switches.

Description Configure the addresses of the virtual routers in a Virtual Router Redundancy Protocol (VRRP) IPv6 group. You can configure up to eight addresses.



NOTE: The address of an aggregated Ethernet interface (a LAG) or a routed VLAN interface (RVI) cannot be assigned as the virtual router address in a VRRP IPv6 group.

Options ***addresses***—Addresses of one or more virtual routers. Do not include a prefix length. If the address is the same as the interface's physical address, the interface becomes the master virtual router for the group.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Documentation • [Configuring VRRP for IPv6 \(CLI Procedure\) on page 341](#)

virtual-link-local-address

Syntax `virtual-link-local-address ipv6-address;`

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family inet6 address *address* [vrrp-inet6-group group-id](#)],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet6 address *address* [vrrp-inet6-group group-id](#)]

Release Information Statement introduced in Junos OS Release 8.4.
Statement introduced in Junos OS 11.3 for the QFX Series.

Description Configure a virtual link-local address for a Virtual Router Redundancy Protocol (VRRP) IPv6 group. You must explicitly define a virtual link-local address for each VRRP for IPv6 group. The virtual link-local address must be in the same subnet as the physical interface address.



NOTE: You do *not* need to configure link-local addresses and virtual link-local addresses when configuring VRRP for IPv6. Junos OS automatically generates link-local addresses and virtual link-local addresses. However, if link local addresses and virtual link-local addresses are configured, Junos OS considers the configured addresses.

Options *ipv6-address*—virtual link-local IPv6 address for VRRP for an IPv6 group.
Range: 0 through 255

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring Basic VRRP Support on page 332](#)
- [Junos OS Support for VRRPv3 on page 322](#)

virtual-link-local-address

Syntax	<code>virtual-link-local-address <i>ipv6-address</i>;</code>
Hierarchy Level	<code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrrp-inet6-group <i>group-id</i>]</code> <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group <i>group-id</i>]</code>
Release Information	Statement introduced in Junos OS Release 10.0 for EX Series switches.
Description	Configure a virtual link local address for a Virtual Router Redundancy Protocol (VRRP) IPv6 group. You must explicitly define a virtual link local address for each VRRP IPv6 group. The virtual link local address must be in the same subnet as the physical interface address.
Options	<i>ipv6-address</i> —Virtual link local IPv6 address for VRRP for an IPv6 group.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring VRRP for IPv6 (CLI Procedure) on page 341

vrrp-group

Syntax	<pre> vrrp-group <i>group-id</i> { (accept-data no-accept-data); advertise-interval <i>seconds</i>; global-advertisements-threshold <i>number</i>; authentication-key <i>key</i>; authentication-type <i>authentication</i>; fast-interval <i>milliseconds</i>; (preempt no-preempt) { hold-time <i>seconds</i>; } priority <i>number</i>; track { interface <i>interface-name</i> { bandwidth-threshold <i>bits-per-second</i> priority-cost <i>priority</i>; priority-cost <i>priority</i>; } priority-hold-time <i>seconds</i>; route <i>prefix/prefix-length</i> routing-instance <i>instance-name</i> priority-cost <i>priority</i>; } virtual-address [<i>addresses</i>]; vrrp-inherit-from <i>vrrp-group</i>; } </pre>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.</p>
Description	<p>Configure a Virtual Router Redundancy Protocol (VRRP) IPv4 group. As of Junos OS Release 13.2, VRRP nonstop active routing (NSR) is enabled only when you configure the nonstop-routing statement at the [edit routing-options] or [edit logical system <i>logical-system-name</i> routing-options] hierarchy level.</p>
Options	<p><i>group-id</i>—VRRP group identifier. If you enable MAC source address filtering on the interface, you must include the virtual MAC address in the list of source MAC addresses that you specify in the source-address-filter statement. MAC addresses ranging from 00:00:5e:00:53:00 through 00:00:5e:00:53:ff are reserved for VRRP, as defined in RFC 2338. The VRRP group number must be the decimal equivalent of the last hexadecimal byte of the virtual MAC address.</p> <p>Range: 0 through 255</p>



NOTE: Under certain circumstances, the group identifier that you enter must be different from any other group identifiers that you configured for logical units of this same physical interface. See [“Configuring Basic VRRP Support” on page 332](#) for more information.

The remaining statements are explained separately. Click a linked statement in the Syntax section for more information about that statement.

Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
---------------------------------	---

- | | |
|------------------------------|---|
| Related Documentation | <ul style="list-style-type: none">• Configuring Basic VRRP Support on page 332• Configuring VRRP on page 337• Example: Configuring VRRP for Load Sharing on page 378• vrrp-inet6-group on page 637• nonstop-routing on page 546 |
|------------------------------|---|

vrrp-inet6-group

Syntax	<pre> vrrp-inet6-group <i>group-id</i> { (accept-data no-accept-data); advertisements-threshold <i>number</i>; fast-interval <i>milliseconds</i>; inet6-advertise-interval <i>seconds</i>; (preempt no-preempt) { hold-time <i>seconds</i>; } priority <i>number</i>; track { interface <i>interface-name</i> { bandwidth-threshold <i>bits-per-second</i> priority-cost <i>priority</i>; priority-cost <i>priority</i>; } priority-hold-time <i>seconds</i>; route <i>prefix/prefix-length</i> routing-instance <i>instance-name</i> priority-cost <i>priority</i>; } virtual-inet6-address [<i>addresses</i>]; virtual-link-local-address <i>ipv6-address</i>; vrrp-inherit-from <i>vrrp-group</i>; } </pre>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i>]</p>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure a Virtual Router Redundancy Protocol (VRRP) IPv6 group.



NOTE: The group identifier that you enter must be different from any other group identifiers that you configured for logical units of this same physical interface.

Options *group-id*—VRRP group identifier. If you enable MAC source address filtering on the interface, you must include the virtual MAC address in the list of source MAC addresses that you specify in the **source-address-filter** statement. MAC addresses ranging from 00:00:5e:00:01:00 through 00:00:5e:00:01:ff are reserved for VRRP, as defined in RFC 2338. The VRRP group number must be the decimal equivalent of the last hexadecimal byte of the virtual MAC address.

Range: 0 through 255

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring Basic VRRP Support on page 332](#)

vrrp-inet6-group

Syntax `vrrp-inet6-group group-id {
 inet6-advertise-interval milliseconds;
 preempt{
 hold-time seconds;
 }
 priority number;
 virtual-inet6-address;
 virtual-link-local-address
}`

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family inet6 address *address*]

Release Information Statement introduced in Junos OS Release 10.0 for EX Series switches.

Description Configure a Virtual Router Redundancy Protocol (VRRP) IPv6 group.

Options *group-id*—VRRP group identifier. If you enable MAC source address filtering on the interface, you must include the virtual MAC address in the list of source MAC addresses that you specify in the **source-address-filter** statement. The MAC address **00-00-5E-00-02-{VRID}** is reserved for VRRP, as defined in RFC 5798. The VRRP group number must be the decimal equivalent of the last hexadecimal byte of the virtual MAC address.

Range: 0 through 255

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring VRRP for IPv6 \(CLI Procedure\) on page 341](#)

vrrp-inherit-from

Syntax	<pre>vrrp-inherit-from { active-group <i>group-index</i>; active-interface <i>active-interface-name</i>; }</pre>
Hierarchy Level	<pre>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 vrrp-inet6-group <i>group-id</i>] [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet vrrp-group <i>group-id</i>]</pre>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	VRRP group to follow for the vrrp-group or vrrp-inet6-group.
Options	<p><i>group-index</i>—Identifier for VRRP active group. Range: 0 through 255</p> <p><i>active-interface-name</i>—Interface name of VRRP active group.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Understanding VRRP on page 315

CHAPTER 45

Administration

- [Administration Tasks on page 641](#)
- [Verification Tasks on page 658](#)

Administration Tasks

- [Upgrading Software on an EX6200 or EX8200 Standalone Switch Using Nonstop Software Upgrade \(CLI Procedure\) on page 641](#)
- [Upgrading Software on an EX8200 Virtual Chassis Using Nonstop Software Upgrade \(CLI Procedure\) on page 650](#)
- [Upgrading Software Using Nonstop Software Upgrade on EX Series Virtual Chassis and Mixed Virtual Chassis \(CLI Procedure\) on page 654](#)

Upgrading Software on an EX6200 or EX8200 Standalone Switch Using Nonstop Software Upgrade (CLI Procedure)

You can use nonstop software upgrade (NSSU) to upgrade the software on standalone EX6200 or EX8200 switches with redundant Routing Engines. NSSU upgrades the software running on the Routing Engines and line cards with minimal traffic disruption during the upgrade. NSSU is supported on EX8200 switches running Junos OS Release 10.4 or later and on EX6200 switches running Junos OS Release 12.2 or later.

This topic covers:

- [Preparing the Switch for Software Installation on page 641](#)
- [Upgrading Both Routing Engines Using NSSU on page 643](#)
- [Upgrading One Routing Engine Using NSSU \(EX8200 Switch Only\) on page 646](#)
- [Upgrading the Original Master Routing Engine \(EX8200 Switch Only\) on page 648](#)

Preparing the Switch for Software Installation

Before you begin software installation using NSSU:

- (Optional) Configure line-card upgrade groups as described in [“Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade” on page 482](#). By default, an NSSU upgrades line cards one at a time to allow aggregated Ethernet links that have members on different line cards to remain up through the upgrade process. Configuring line-card upgrade groups reduces the time an upgrade takes because the line cards in each upgrade group are upgraded at the same time rather than sequentially.
- Verify that the Routing Engines are running the same version of the software. Enter the following command:

```
{master}
user@switch> show version invoke-on all-routing-engines
re0:
-----
Hostname: switch
Model: ex8208
JUNOS Base OS boot [11.3-20110429.1]
JUNOS Base OS Software Suite [11.3-20110429.1]
JUNOS Kernel Software Suite [11.3-20110429.1]
JUNOS Crypto Software Suite [11.3-20110429.1]
JUNOS Online Documentation [11.3-20110429.1]
JUNOS Enterprise Software Suite [11.3-20110429.1]
LC JUNOS Installation Software [11.3-20110429.1]
JUNOS Routing Software Suite [11.3-20110429.1]
JUNOS Web Management [11.3-20110429.1]

re1:
-----
Hostname: switch
Model: ex8208
JUNOS Base OS boot [11.3-20110429.1]
JUNOS Base OS Software Suite [11.3-20110429.1]
JUNOS Kernel Software Suite [11.3-20110429.1]
JUNOS Crypto Software Suite [11.3-20110429.1]
JUNOS Online Documentation [11.3-20110429.1]
JUNOS Enterprise Software Suite [11.3-20110429.1]
LC JUNOS Installation Software [11.3-20110429.1]
JUNOS Routing Software Suite [11.3-20110429.1]
JUNOS Web Management [11.3-20110429.1]
```

If the Routing Engines are not running the same version of the software, use the **request system software add** command to upgrade the Routing Engine that is running the earlier software version.

- Ensure that nonstop active routing (NSR) and graceful Routing Engine switchover (GRES) are enabled. To verify that they are enabled, you need to check only the state of nonstop active routing—if nonstop active routing is enabled, then graceful Routing Engine switchover is enabled.

To verify that nonstop active routing is enabled, execute the following command:

```
{master}
user@switch> show task replication
Stateful Replication: Enabled
RE mode: Master
```


Protocol	Synchronization Status
OSPF	Complete
RIP	Complete
PIM	Complete
RSVP	Complete

If nonstop active routing is not enabled (**Stateful Replication is Disabled**), see [“Configuring Nonstop Active Routing on Switches” on page 220](#) for information on how to enable it.

- (Optional) Enable nonstop bridging (NSB). Enabling NSB ensures that all NSB-supported Layer 2 protocols operate seamlessly during the Routing Engine switchover that is part of the NSSU.
- (Optional) Back up the system software on each Routing Engine to an external storage device with the **request system snapshot** command.

Upgrading Both Routing Engines Using NSSU

This procedure describes how to upgrade both Routing Engines using NSSU. When the upgrade completes, both Routing Engines are running the new version of the software, and the backup Routing Engine is the new master Routing Engine.

To upgrade both Routing Engines using NSSU:

1. Download the software package.
2. Copy the software package to the switch. We recommend that you use FTP to copy the file to the `/var/tmp` directory.
3. Log in to the master Routing Engine using the console connection. You can perform an NSSU from the management interface, but a console connection allows you to monitor the progress of the master Routing Engine reboot.
4. Install the new software package:

```
{master}
user@switch> request system software nonstop-upgrade reboot
/var/tmp/package-name-m.nZx-distribution.tgz
```

where **package-name-m.nZx-distribution.tgz** is, for example, **jinstall-ex-8200-10.4R1.5-domestic-signed.tgz**.

The switch displays the following status messages as the upgrade executes:

```
Chassis ISSU Check Done
ISSU: Validating Image
ISSU: Preparing Backup RE
Pushing bundle to re1
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
Backup upgrade done
Rebooting Backup RE
```

```

Rebooting re1
ISSU: Backup RE Prepare Done
Waiting for Backup RE reboot
GRES operational
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item          Status          Reason
  FPC 0         Online (ISSU)
  FPC 1         Online (ISSU)
  FPC 2         Online (ISSU)
  FPC 3         Offline          Offlined by CLI command
  FPC 4         Online (ISSU)
  FPC 5         Online (ISSU)
  FPC 6         Online (ISSU)
  FPC 7         Online (ISSU)
Resolving mastership...
Complete. The other routing engine becomes the master.
ISSU: RE switchover Done
ISSU: Upgrading Old Master RE
WARNING: A reboot is required to install the software
WARNING:      Use the 'request system reboot' command immediately
ISSU: Old Master Upgrade Done
ISSU: IDLE

*** FINAL System shutdown message from user@switch ***
System going down IMMEDIATELY

Shutdown NOW!
[pid 2635]

```



NOTE: If you omit the reboot option in this step when using an EX8200 switch, you must manually reboot the original master Routing Engine with the `request system reboot` command for the upgrade to complete.

The original master Routing Engine reboots automatically after updating the new master Routing Engine when an NSSU is used to upgrade an EX6200 switch with dual Routing Engines.

5. Log in after the reboot completes. To verify that both Routing Engines have been upgraded, enter the following command:

```

{backup}
user@switch> show version invoke-on all-routing-engines
re0:
-----
Hostname: switch
Model: ex8208
JUNOS Base OS boot [12.1-20111229.0]

```

```

JUNOS Base OS Software Suite [12.1-20111229.0]
JUNOS Kernel Software Suite [12.1-20111229.0]
JUNOS Crypto Software Suite [12.1-20111229.0]
JUNOS Online Documentation [12.1-20111229.0]
JUNOS Enterprise Software Suite [12.1-20111229.0]
LC JUNOS Installation Software [12.1-20111229.0]
JUNOS Routing Software Suite [12.1-20111229.0]
JUNOS Web Management [12.1-20111229.0]

```

```
rel:
```

```

-----
Hostname: switch
Model: ex8208
JUNOS Base OS boot [12.1-20111229.0]
JUNOS Base OS Software Suite [12.1-20111229.0]
JUNOS Kernel Software Suite [12.1-20111229.0]
JUNOS Crypto Software Suite [12.1-20111229.0]
JUNOS Online Documentation [12.1-20111229.0]
JUNOS Enterprise Software Suite [12.1-20111229.0]
LC JUNOS Installation Software [12.1-20111229.0]
JUNOS Routing Software Suite [12.1-20111229.0]
JUNOS Web Management [12.1-20111229.0]

```

- To verify that the line cards that were online before the upgrade are online after the upgrade, log in to the master Routing Engine and enter the `show chassis nonstop-upgrade` command:

```

{backup}
user@switch> request routing-engine login master

{master}
user@switch> show chassis nonstop-upgrade

```

Item	Status	Reason
FPC 0	Online (ISSU)	
FPC 1	Online (ISSU)	
FPC 2	Online (ISSU)	
FPC 3	Offline	Offlined by CLI command
FPC 4	Online (ISSU)	
FPC 5	Online (ISSU)	
FPC 6	Online (ISSU)	
FPC 7	Online (ISSU)	

- If you want to make **re0** the master Routing Engine again, enter the following command:

```

{master}
user@switch> request chassis routing-engine master switch
Toggle mastership between routing engines ? [yes,no] (no) yes

```

You can verify that **re0** is the master Routing Engine by executing the `show chassis routing-engine` command.

8. To ensure that the resilient dual-root partitions feature operates correctly, execute the following command to copy the new Junos OS image into the alternate root partition on each Routing Engine:

```
user@switch> request system snapshot slice alternate routing-engine both
```

Resilient dual-root partitions allow the switch to boot transparently from the alternate root partition if the system fails to boot from the primary root partition.

Upgrading One Routing Engine Using NSSU (EX8200 Switch Only)

This procedure describes how to upgrade one of the Routing Engines using NSSU on an EX8200 switch. When the upgrade completes, the backup Routing Engine is running the new software version and is the new master. The original master Routing Engine, now the backup Routing Engine, continues to run the previous software version.



NOTE: NSSU always upgrades the software on both Routing Engines on an EX6200 switch. Therefore, you cannot upgrade software on one Routing Engine using NSSU on an EX6200 switch.

To upgrade one Routing Engine using NSSU:

1. Download the software package.
2. Copy the software package to the switch. We recommend that you use FTP to copy the file to the `/var/tmp` directory.
3. Log in to the master Routing Engine.
4. Request an NSSU. On an EX8200 switch, specify the **no-old-master-upgrade** option when requesting the NSSU:

```
{master}
user@switch> request system software nonstop-upgrade
no-old-master-upgrade /var/tmp/package-name-m.nZx-distribution.tgz
```

where *package-name-m.nZx-distribution.tgz* is, for example, *jinstall-ex-8200-10.4R2.5-domestic-signed.tgz*.

The switch displays the following status messages as the upgrade executes:

```
Chassis ISSU Check Done
ISSU: Validating Image
ISSU: Preparing Backup RE
Pushing bundle to re1
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
Backup upgrade done
Rebooting Backup RE

Rebooting re1
ISSU: Backup RE Prepare Done
Waiting for Backup RE reboot
```

```

GRES operational
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item           Status           Reason
  FPC 0          Online (ISSU)
  FPC 1          Online (ISSU)
  FPC 2          Online (ISSU)
  FPC 3          Offline           Offlined by CLI command
  FPC 4          Online (ISSU)
  FPC 5          Online (ISSU)
  FPC 6          Online (ISSU)
  FPC 7          Online (ISSU)
Resolving mastership...
Complete. The other routing engine becomes the master.
ISSU: RE switchover Done
Skipping Old Master Upgrade
ISSU: IDLE

```

When the upgrade is complete, the original master Routing Engine (**re0**) becomes the backup Routing Engine.

5. To verify that the original backup Routing Engine (**re1**) has been upgraded, enter the following command:

```

{backup}
user@switch> show version invoke-on all-routing-engines
re0:
-----
Hostname: switch
Model: ex8208
JUNOS Base OS boot [11.3-20110429.1]
JUNOS Base OS Software Suite [11.3-20110429.1]
JUNOS Kernel Software Suite [11.3-20110429.1]
JUNOS Crypto Software Suite [11.3-20110429.1]
JUNOS Online Documentation [11.3-20110429.1]
JUNOS Enterprise Software Suite [11.3-20110429.1]
LC JUNOS Installation Software [11.3-20110429.1]
JUNOS Routing Software Suite [11.3-20110429.1]
JUNOS Web Management [11.3-20110429.1]

re1:
-----
Hostname: switch
Model: ex8208
JUNOS Base OS boot [12.1-20111229.0]
JUNOS Base OS Software Suite [12.1-20111229.0]
JUNOS Kernel Software Suite [12.1-20111229.0]
JUNOS Crypto Software Suite [12.1-20111229.0]
JUNOS Online Documentation [12.1-20111229.0]
JUNOS Enterprise Software Suite [12.1-20111229.0]
LC JUNOS Installation Software [12.1-20111229.0]

```

```
JUNOS Routing Software Suite [12.1-20111229.0]
JUNOS Web Management [12.1-20111229.0]
```

- To verify that the line cards that were online before the upgrade are online after the upgrade, log in to the new master Routing Engine and enter the **show chassis nonstop-upgrade** command:

```
{backup}
user@switch> request routing-engine login master

--- JUNOS 12.1-20111229.0 built 2011-12-29 04:12:22 UTC
{master}
user@switch> show chassis nonstop-upgrade
```

Item	Status	Reason
FPC 0	Online	
FPC 1	Online	
FPC 2	Online	
FPC 3	Offline	Offlined by CLI command
FPC 4	Online	
FPC 5	Online	
FPC 6	Online	
FPC 7	Online	

- To ensure that the resilient dual-root partitions feature operates correctly, copy the new Junos OS image into the alternate root partition of the Routing Engine:

```
user@switch> request system snapshot slice alternate
```

Resilient dual-root partitions allow the switch to boot transparently from the alternate root partition if the system fails to boot from the primary root partition.

Upgrading the Original Master Routing Engine (EX8200 Switch Only)

This procedure describes how to upgrade the original master Routing Engine after you have upgraded the original backup Routing Engine as described in [“Upgrading One Routing Engine Using NSSU \(EX8200 Switch Only\)” on page 490](#) for an EX8200 switch.

- Log in to the current master Routing Engine (re1).
- Enter configuration mode and disable nonstop active routing:

```
{master}[edit]
user@switch# delete routing-options nonstop-routing
```

- Deactivate graceful Routing Engine switchover and commit the configuration:

```
{master}[edit]
user@switch# deactivate chassis redundancy graceful-switchover

{master}[edit]
user@switch# commit
```

4. Log in to the current backup Routing Engine (**re0**) using a console connection.
5. Request a software installation:

```
user@switch> request system software add reboot
/var/tmp/package-name-m.nZx-distribution.tgz
```



NOTE: When you use NSSU to upgrade only one Routing Engine, the installation package is not automatically deleted from `/var/tmp`, leaving the package available to be used to upgrade the original master Routing Engine.

6. After the upgrade completes, log in to the current master Routing Engine (**re1**) and enter CLI configuration mode.
7. Re-enable nonstop active routing and graceful Routing Engine switchover:

```
[edit]
user@switch# activate chassis redundancy graceful-switchover

[edit]
user@switch# set routing-options nonstop-routing

[edit]
user@switch# commit
```

8. To ensure that the resilient dual-root partitions feature operates correctly, exit the CLI configuration mode and copy the new Junos OS image into the alternate root partition of the Routing Engine:

```
user@switch> request system snapshot slice alternate
```

Resilient dual-root partitions allow the switch to boot transparently from the alternate root partition if the system fails to boot from the primary root partition.

9. (Optional) To return control to the original master Routing Engine (**re0**), enter the following command:

```
{master}
user@switch> request chassis routing-engine master switch
Toggle mastership between routing engines ? [yes,no] (no) yes
```

You can verify that **re0** is the master Routing Engine by executing the **show chassis routing-engine** command.

- See Also**
- [Example: Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade on EX Series Switches on page 494](#)
 - [Configuring Dual-Root Partitions](#)
 - [Troubleshooting Software Installation](#)

Upgrading Software on an EX8200 Virtual Chassis Using Nonstop Software Upgrade (CLI Procedure)

You can use nonstop software upgrade (NSSU) to upgrade the software on an EX8200 Virtual Chassis. NSSU upgrades the software running on all Routing Engines with minimal traffic disruption during the upgrade. NSSU is supported on EX8200 Virtual Chassis with redundant XRE200 External Routing Engines running Junos OS Release 11.1 or later.



NOTE: NSSU upgrades all Routing Engines on all members of the Virtual Chassis and on the XRE200 External Routing Engines. Using NSSU, you cannot choose to upgrade the backup Routing Engines only, nor can you choose to upgrade a specific member of the Virtual Chassis. If you need to upgrade a specific member of the Virtual Chassis, see *Installing Software for a Single Device in an EX8200 Virtual Chassis*.

This topic covers:

- [Preparing the Switch for Software Installation on page 650](#)
- [Upgrading the Software Using NSSU on page 651](#)

Preparing the Switch for Software Installation

Before you begin software installation using NSSU:

- (Optional) Configure line-card upgrade groups as described in “[Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade](#)” on page 482. By default, NSSU upgrades line cards one at a time, starting with the line card in slot 0 of member 0. This permits aggregated Ethernet links that have members on different line cards remain up through the upgrade process. Configuring line-card upgrade groups reduces the time an upgrade takes because the line cards in each upgrade group are upgraded at the same time rather than sequentially.
- Verify that the members are running the same version of the software:

```
{master:8}  
user@external-routing-engine> show version all-members
```

If the Virtual Chassis members are not running the same version of the software, use the **request system software add** command to upgrade the software on the inconsistent members. For instructions, see *Installing Software for a Single Device in an EX8200 Virtual Chassis*.

- Ensure that nonstop active routing (NSR) and graceful Routing Engine switchover (GRES) are enabled. To verify that they are enabled, you need to check only the state of nonstop active routing—if nonstop active routing is enabled, then graceful Routing Engine switchover is enabled.

To verify that nonstop active routing is enabled:

```
{master:8}
user@switch> show task replication
Stateful Replication: Enabled
RE mode: Master
```

Protocol	Synchronization Status
PIM	Complete

If nonstop active routing is not enabled (**Stateful Replication is Disabled**), see [“Configuring Nonstop Active Routing on Switches” on page 220](#) for information on how to enable it.

Upgrading the Software Using NSSU

This procedure describes how to upgrade the software running on all Routing Engines using NSSU. When the upgrade completes, all Routing Engines are running the new version of the software. The backup external Routing Engine is now the master external Routing Engine, and the internal backup Routing Engines in the member switches are now the internal master Routing Engines in those member switches.

To upgrade all Routing Engines using NSSU:

1. Download the software package for the XRE200 External Routing Engine by following one of the procedures in *Downloading Software*. The name of the software package for the XRE200 External Routing Engine contains the term **xre200**.
2. Copy the software package to the switch. We recommend that you use FTP to copy the file to the **/var/tmp** directory.
3. Log in to the master external Routing Engine using the console connection. You can perform an NSSU from the management interface, but a console connection allows you to monitor the progress of the master Routing Engine reboot.
4. Install the new software package:

```
{master:8}
user@external-routing-engine> request system software nonstop-upgrade reboot
/var/tmp/package-name-m.nZx-distribution.tgz
```

where **package-name-m.nZx-distribution.tgz** is, for example, **jinstall-ex-xre200-11.1R2.5-domestic-signed.tgz**.



NOTE: You can omit reboot option. When you include the reboot option, NSSU automatically reboots the original master Routing Engines after the new image has been installed on them. If you omit the reboot option, you must manually reboot the original master Routing Engines (now the backup Routing Engines) to complete the upgrade. To perform the reboot, you must establish a connection to the console port on the Switch Fabric and Routing Engine (SRE) module or Routing Engine (RE) module.

The switch displays status messages similar to the following messages as the upgrade executes:

```
Chassis ISSU Check Done
ISSU: Validating Image
ISSU: Preparing LCC Backup REs
ISSU: Preparing Backup RE
Pushing bundle /var/tmp/jinstall-ex-xre200-11.1-20110208.0-domestic-signed.tgz
to member9
member9:
-----
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
VC Backup upgrade done
Rebooting VC Backup RE

Rebooting member9
ISSU: Backup RE Prepare Done
Waiting for VC Backup RE reboot
Pushing bundle to member0-backup
Pushing bundle to member1-backup
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately

Rebooting member0-backup
Rebooting LCC [member0-backup]

Rebooting member1-backup
Rebooting LCC [member1-backup]
ISSU: LCC Backup REs Prepare Done
GRES operational
Initiating Chassis Nonstop-Software-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking Nonstop-Upgrade status
member0:
-----
Item           Status           Reason
FPC 0          Online (ISSU)
FPC 1          Online (ISSU)
FPC 2          Online (ISSU)
FPC 5          Online (ISSU)
```

```

member1:
-----
  Item      Status      Reason
  FPC 0     Online (ISSU)
  FPC 1     Online (ISSU)
  FPC 2     Online (ISSU)
  FPC 5     Online (ISSU)

member0:
-----
  Item      Status      Reason
  FPC 0     Online (ISSU)
  FPC 1     Online (ISSU)
  FPC 2     Online (ISSU)
  FPC 5     Online (ISSU)

member1:
-----
  Item      Status      Reason
  FPC 0     Online (ISSU)
  FPC 1     Online (ISSU)
  FPC 2     Online (ISSU)
  FPC 5     Online (ISSU)
ISSU: Upgrading Old Master RE
Pushing bundle /var/tmp/incoming-package-8200.tgz to member0-master
Pushing bundle /var/tmp/incoming-package-8200.tgz to member1-master
ISSU: RE switchover Done
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
ISSU: Old Master Upgrade Done
ISSU: IDLE

*** FINAL System shutdown message from root@ ***
System going down IMMEDIATELY

Shutdown NOW!

```



NOTE: If you omit the reboot option in this step, you must complete the upgrade by separately rebooting the original master Routing Engine on each Virtual Chassis member and the original master external Routing Engine. To reboot the original master Routing Engine on a Virtual Chassis member, you must establish a connection to the console port on the Switch Fabric and Routing Engine (SRE) module or Routing Engine (RE) module.

5. Log in after the reboot completes. To verify that the software on all Routing Engines in the Virtual Chassis members has been upgraded, enter the following command:

```

{backup:8}
user@external-routing-engine> show version all-members

```

6. Verify that the line cards that were online before the upgrade are online after the upgrade by entering the `show chassis nonstop-upgrade` command:

```
{backup:8}
user@external-routing-engine> show chassis nonstop-upgrade
member0:
```

Item	Status	Reason
FPC 0	Online	
FPC 1	Online	
FPC 2	Online	
FPC 5	Online	

```
member1:
```

Item	Status	Reason
FPC 0	Online	
FPC 1	Online	
FPC 2	Online	
FPC 5	Online	

- See Also**
- [Upgrading Software Using Nonstop Software Upgrade on EX Series Virtual Chassis and Mixed Virtual Chassis \(CLI Procedure\) on page 654](#)
 - [Upgrading Software on an EX6200 or EX8200 Standalone Switch Using Nonstop Software Upgrade \(CLI Procedure\) on page 485](#)
 - [Example: Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade on EX Series Switches on page 494](#)
 - [Configuring Dual-Root Partitions](#)
 - [Troubleshooting Software Installation](#)
 - [Understanding Nonstop Software Upgrade on EX Series Switches on page 471](#)
 - [Understanding Software Installation on EX Series Switches](#)

Upgrading Software Using Nonstop Software Upgrade on EX Series Virtual Chassis and Mixed Virtual Chassis (CLI Procedure)

You can use nonstop software upgrade (NSSU) to upgrade the software running on all member switches in most EX Series Virtual Chassis with minimal traffic disruption during the upgrade.

NSSU is supported on the following EX Series Virtual Chassis platforms:

- EX3300 Virtual Chassis
- EX3400 Virtual Chassis
- EX4200 Virtual Chassis
- EX4300 Virtual Chassis
- EX4500 Virtual Chassis

- EX4550 Virtual Chassis
- All mixed Virtual Chassis composed of EX4200, EX4500, and EX4550 switches
- EX8200 Virtual Chassis

This topic covers:

- [Preparing the Switch for Software Installation on page 655](#)
- [Upgrading the Software Using NSSU on page 656](#)

Preparing the Switch for Software Installation

Before you begin software installation using NSSU:

- Ensure that the Virtual Chassis is configured correctly to support NSSU. Verify that:
 - The Virtual Chassis members are connected in a ring topology. A ring topology prevents the Virtual Chassis from splitting during an NSSU.
 - The Virtual Chassis master and backup are adjacent to each other in the ring topology. Adjacency permits the master and backup to always be in sync, even when the switches in linecard roles are rebooting.
 - The Virtual Chassis is preprovisioned so that the linecard role has been explicitly assigned to member switches acting in the linecard role. During an NSSU, the Virtual Chassis members must maintain their roles—the master and backup must maintain their master and backup roles (although mastership will change), and the other member switches must maintain their linecard roles.

For information on configuring a preprovisioned Virtual Chassis, see *Configuring an EX3300 Virtual Chassis (CLI Procedure)*, *Configuring an EX4200, EX4500, or EX4550 Virtual Chassis (CLI Procedure)*, *Configuring an EX2300, EX3400, or EX4300 Virtual Chassis*, and *Configuring an EX8200 Virtual Chassis (CLI Procedure)*.

- A two-member Virtual Chassis has **no-split-detection** configured so that the Virtual Chassis does not split when an NSSU upgrades a member.
- Verify that the members are running the same version of the software:

```
user@switch> show version
```

If the Virtual Chassis members are not running the same version of the software, use the **request system software add** command to upgrade the software on the inconsistent members.

- Ensure that nonstop active routing (NSR) and graceful Routing Engine switchover (GRES) are enabled. To verify that they are enabled, you need to check only the state of nonstop active routing—if nonstop active routing is enabled, then graceful Routing Engine switchover is enabled.

To verify that nonstop active routing is enabled:

```
user@switch> show task replication
Stateful Replication: Enabled
```

RE mode: Master	
Protocol	Synchronization Status
OSPF	Complete
BGP	Complete
PIM	Complete

If nonstop active routing is not enabled (**Stateful Replication is Disabled**), see [“Configuring Nonstop Active Routing on Switches” on page 220](#) for information on how to enable it.

- For the EX4300 Virtual Chassis, you should enable the **vcp-no-hold-time** statement at the **[edit virtual-chassis]** hierarchy level before performing a software upgrade using NSSU. If you do not enable the **vcp-no-hold-time** statement, the Virtual Chassis may split during the upgrade. A split Virtual Chassis can cause disruptions to your network, and you may have to manually reconfigure your Virtual Chassis after the NSSU if the split and merge feature was disabled. For more information about a split Virtual Chassis, see *Understanding Split and Merge in a Virtual Chassis*.
- (Optional) Enable nonstop bridging (NSB). Enabling NSB ensures that all NSB-supported Layer 2 protocols operate seamlessly during the Routing Engine switchover that is part of the NSSU.
- (Optional) Back up the system software—Junos OS, the active configuration, and log files—on each member to an external storage device with the **request system snapshot** command.

Upgrading the Software Using NSSU

This procedure describes how to upgrade the software running on all Virtual Chassis members using NSSU. When the upgrade completes, all members are running the new version of the software. Because a graceful Routing Engine switchover occurs during the upgrade, the original Virtual Chassis backup is the new master.

To upgrade all members using NSSU:

1. Download the software package. If you are upgrading the software running on a mixed Virtual Chassis, download the software packages for both switch types.
2. Copy the software package or packages to the Virtual Chassis. We recommend that you copy the file to the **/var/tmp** directory on the master.
3. Log in to the Virtual Chassis using the console connection or the virtual management Ethernet (VME) interface. Using a console connection allows you to monitor the progress of the master switch reboot.
4. Start the NSSU:
 - On an EX3300 Virtual Chassis, EX3400 Virtual Chassis, EX4200 Virtual Chassis, EX4300 Virtual Chassis, EX4500 Virtual Chassis, or EX4550 Virtual Chassis, enter:

```
user@switch> request system software nonstop-upgrade
/var/tmp/package-name.tgz
```

where *package-name.tgz* is, for example, *jinstall-ex4200-12.1R2.5-domestic-signed.tgz*.

- On a mixed Virtual Chassis, enter:

```
user@switch> request system software nonstop-upgrade set
[/var/tmp/package-name.tgz /var/tmp/package-name.tgz]
```

where *[/var/tmp/package-name.tgz /var/tmp/package-name.tgz]* specifies the EX4200 and EX4500 software packages.

The switch displays status messages similar to the following messages as the upgrade executes:

```
Chassis ISSU Check Done
ISSU: Validating Image
ISSU: Preparing Backup RE
Installing image on other FPC's along with the backup

Checking pending install on fpc1
Pushing bundle to fpc1
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
Completed install on fpc1

Checking pending install on fpc2
Pushing bundle to fpc2
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
Completed install on fpc2

Rebooting fpc1
ISSU: Backup RE Prepare Done
Waiting for Backup RE reboot
GRES operational
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item          Status          Reason
  FPC 0         Online
  FPC 1         Online
  FPC 2         Online (ISSU)
Going to install image on master
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
relinquish mastership
ISSU: IDLE

*** FINAL System shutdown message from user@switch ***

System going down IMMEDIATELY
```

```
Shutdown NOW!  
[pid 9336]
```

5. Log in after the reboot of the original master switch completes. To verify that the software on all Routing Engines in the Virtual Chassis members has been upgraded, enter the following command:

```
user@switch> show version
```

6. To ensure that the resilient dual-root partitions feature operates correctly, copy the new Junos OS image into the alternate root partitions of all members:

```
user@switch> request system snapshot slice alternate all-members
```

Resilient dual-root partitions allow the switch to boot transparently from the alternate root partition if the system fails to boot from the primary root partition.

- See Also**
- [Understanding Nonstop Software Upgrade on EX Series Switches on page 471](#)
 - [Configuring Dual-Root Partitions](#)
 - [Understanding Software Installation on EX Series Switches](#)
 - [Troubleshooting Software Installation](#)
 - [Understanding Nonstop Software Upgrade on EX Series Switches on page 471](#)

Verification Tasks

- [Verifying Power Configuration and Use on page 658](#)

Verifying Power Configuration and Use

Purpose Verify on an EX Series switch:

- The power redundancy and line card priority settings
- The PoE power budgets for line cards that support PoE
- Whether the $N+1$ or $N+N$ power requirements are being met
- Whether the switch has sufficient power for a new line card or an $N+N$ configuration

Action Enter the following command:

```
user@switch> show chassis power-budget-statistics
```

Example output for an EX6200 switch:

PSU	0	(EX6200-PWR-AC2500)	:	2500 W	Online
PSU	1	(EX6200-PWR-AC2500)	:	2500 W	Online

PSU 2	(EX6200-PWR-AC2500)	:	2500 W	Online		
PSU 3	(EX6200-PWR-AC2500)	:	2500 W	Online		
Total Power supplied by all Online PSUs		:	10000 W			
Power Redundancy Configuration		:	N+1			
Power Reserved for the Chassis		:	500 W			
Fan Tray Statistics		Base power	Power Used			
FTC 0		:	300 W	43.04 W		
FPC Statistics		Base power	Power Used	PoE power	Priority	
FPC 1	(EX6200-48P)	:	220 W	49.47 W	1440 W	1
FPC 2	(EX6200-48P)	:	220 W	47.20 W	800 W	2
FPC 3	(EX6200-48P)	:	220 W	1493.57 W	1440 W	0
FPC 4	(EX6200-SRE64-4XS)	:	100 W	51.38 W	0 W	0
FPC 5	(EX6200-SRE64-4XS)	:	100 W	50.28 W	0 W	0
FPC 6	(EX6200-48P)	:	220 W	49.38 W	800 W	6
FPC 8	(EX6200-48P)	:	220 W	61.41 W	1440 W	9
FPC 9	(EX6200-48T)	:	150 W	12.49 W	0 W	9
Total (non-PoE) Power allocated		:	1750 W			
Total Power allocated for PoE		:	5920 W			
Power Available (Redundant case)		:	5750 W			
Total Power Available		:	2515 W			

Example output for an EX8200 switch:

PSU 0	(EX8200-AC2K)	:	1200 W	Online	
PSU 1	(EX8200-AC2K)	:	1200 W	Online	
PSU 2	(EX8200-AC2K)	:	1200 W	Online	
PSU 3	(EX8200-AC2K)	:	1200 W	Online	
Total Power supplied by all Online PSUs		:	4800 W		
Power Redundancy Configuration		:	N+1		
Power Reserved for the Chassis		:	1600 W		
FPC Statistics					
		Base power	PoE power		Priority
FPC 0	(EX8200-48T)	:	350 W	0 W	2
FPC 1	(EX8200-2XS-40P)	:	387 W	300 W	0
FPC 2	(EX8200-48PL)	:	267 W	350 W	15
FPC 4	(EX8200-2XS-40P)	:	387 W	300 W	1
FPC 5	(EX8200-48TL)	:	230 W	0 W	15
FPC 6	(EX8200-48TL)	:	230 W	0 W	15
Total (non-PoE) Power allocated			:	3451 W	
Total Power allocated for PoE			:	950 W	
Power Available (Redundant case)			:	149 W	
Total Power Available			:	510 W	

- Meaning**
- Example output for an EX6200 switch —The online power supplies can supply a total of 10,000 W to the switch. The switch is configured for *N*+1 redundancy, which means 7500 W of redundant power can be supplied. The **Power Available (Redundant case)** field shows that the switch is meeting the *N*+1 power requirements, with an additional 5750 W available. This value is calculated by subtracting all power allocations except PoE power allocations from redundant power (7500 W).

The total amount of power available on the switch is 2515 W. This value is calculated by subtracting all power allocations, including PoE power allocations, from the total power (10,000 W). On a switch with PoE line cards, if **Total Power Available** is 0, some

or all of the PoE line cards might not be allocated their configured PoE power budgets, which means power to some or all PoE ports might be disabled.

The power priority order of the line cards, from highest priority line card to the lowest priority line card, is 4, 5, 3, 1, 2, 6, 8, 9. Slots 4 and 5, which contain the Switch Fabric and Routing Engine (SRE) modules, always have highest priority, even if a lower-numbered slot, such as slot 3 in this example, has a priority of 0. Should two or more 2500 W power supplies fail, power management will remove or reduce the PoE power allocations from the PoE line cards in the following order to balance the power budget: 8, 6, 2, 1, and 3.

The **Power Used** values for the fan tray and line cards shows the actual power being consumed for these components at the time the command was executed. These values are for your information only; power management uses allocated power, which is based on the maximum power the component might consume, and not actual power consumed, in determining its power budget.

- Example output for an EX8200 switch—The online power supplies can supply a total of 4800 W to the switch. The switch is configured for $N+1$ redundancy, which means 3600 W of redundant power can be supplied. The **Power Available (Redundant case)** field shows that the switch is meeting the $N+1$ power requirements, with an additional 149 W available. This value is calculated by subtracting all power allocations except PoE power allocations from redundant power (3600 W). Because 149 W is insufficient power for a line card, another line card cannot be added to the switch while maintaining $N+1$ redundancy.

The total amount of power available on the switch is 510 W. This value is calculated by subtracting all power allocations, including PoE power allocations, from the total power (4800 W). On a switch with PoE line cards, if **Total Power Available** is 0, some or all of the PoE line cards might not be allocated their configured PoE power budgets, which means power to some or all PoE ports might be disabled.

The power priority order of the line cards, from highest priority line card to the lowest priority line card, is 1, 4, 0, 2, 5, 6. Should one or more 1200 W power supplies fail, power management will remove or reduce the PoE power allocations from the PoE line cards in the following order to balance the power budget: 2, 4, and 1.

- See Also**
- [Configuring Power Supply Redundancy \(CLI Procedure\) on page 303](#)
 - [Configuring the Power Priority of Line Cards \(CLI Procedure\) on page 302](#)

CHAPTER 46

Operational Commands

- `show bgp neighbor`
- `show log`
- `show (ospf | ospf3) overview`
- `show chassis dedicated-ukern-cpu`
- `show chassis realtime-ukern-thread`
- `clear vrrp`
- `request chassis ssb master switch`
- `request redundant-power-system multi-backup`
- `request system software in-service-upgrade`
- `request system software in-service-upgrade (MX Series 5G Universal Routing Platforms and EX9200 Switches)`
- `request system software nonstop-upgrade`
- `request system software validate in-service-upgrade`
- `show chassis nonstop-upgrade`
- `show chassis nonstop-upgrade node-group`
- `show chassis power-budget-statistics`
- `show chassis redundant-power-system`
- `show redundant-power-system led`
- `show redundant-power-system multi-backup`
- `show redundant-power-system network`
- `show redundant-power-system power-supply`
- `show redundant-power-system status`
- `show redundant-power-system upgrade`
- `show redundant-power-system version`
- `show chassis ssb`
- `show nonstop-routing`
- `show pfe ssb`
- `show system switchover`

- [show task replication](#)
- [show vrrp](#)
- [show vrrp track](#)

show bgp neighbor

List of Syntax [Syntax on page 663](#)
 [Syntax \(EX Series Switch, QFX Series, and OCX Series\) on page 663](#)

Syntax show bgp neighbor
 <exact-instance *instance-name*>
 <instance *instance-name*>
 <logical-system (all | *logical-system-name*)>
 <neighbor-address>
 <output-queue>
 <orf (detail | *neighbor-address*)>

Syntax (EX Series Switch, QFX Series, and OCX Series) show bgp neighbor
 <instance *instance-name*>
 <exact-instance *instance-name*>
 <neighbor-address>
 <orf (*neighbor-address* | detail)

Release Information Command introduced before Junos OS Release 7.4.
 Command introduced in Junos OS Release 9.0 for EX Series switches.
 Command introduced in Junos OS Release 11.3 for the QFX Series.
 Command introduced in Junos OS Release 14.1x53-D20 for the OCX Series.
 orf option introduced in Junos OS Release 9.2.
 exact-instance option introduced in Junos OS Release 11.4.
 output-queue option introduced in Junos OS Release 16.1
 DontGRHelpFateSharingBfdDown is added to the **options** field of the command output in Junos OS Release 18.3R1.

Description Display information about BGP peers.

Options **none**—Display information about all BGP peers.

exact-instance *instance-name*—(Optional) Display information for the specified instance only.

instance *instance-name*—(Optional) Display information about BGP peers for all routing instances whose name begins with this string (for example, **cust1**, **cust11**, and **cust111** are all displayed when you run the **show bgp neighbor instance cust1** command).

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

neighbor-address—(Optional) Display information for only the BGP peer at the specified IP address.

orf (detail | neighbor-address)—(Optional) Display outbound route-filtering information for all BGP peers or only for the BGP peer at the specified IP address. The default is to display brief output. Use the **detail** option to display detailed output.

output-queue—(Optional) Display information regarding the number of routes currently queued in the 17 prioritized BGP output queues.

Additional Information For information about the **local-address**, **nlri**, **hold-time**, and **preference** statements, see the *Junos OS Routing Protocols Library*.

Required Privilege Level

view

Related Documentation

- *clear bgp neighbor*

List of Sample Output

[show bgp neighbor on page 672](#)
[show bgp neighbor \(dont-help-shared-fate-bfd-down is configured\) on page 673](#)
[show bgp neighbor \(CLNS\) on page 674](#)
[show bgp neighbor \(Layer 2 VPN\) on page 675](#)
[show bgp neighbor \(Layer 3 VPN\) \(Not supported on the OCX Series.\) on page 677](#)
[show bgp neighbor neighbor-address on page 678](#)
[show bgp neighbor neighbor-address on page 678](#)
[show bgp neighbor neighbor-address \(BGP Graceful Restart Enabled\) on page 679](#)
[show bgp neighbor neighbor-address \(BGP Long-Lived Graceful Restart\) on page 680](#)
[show bgp neighbor orf neighbor-address detail on page 680](#)
[show bgp neighbor logical-system on page 681](#)
[show bgp neighbor output-queue on page 681](#)
[show bgp neighbor \(Segment Routing Traffic Engineering\) on page 682](#)

Output Fields [Table 41 on page 664](#) describes the output fields for the **show bgp neighbor** command. Output fields are listed in the approximate order in which they appear.

Table 41: show bgp neighbor Output Fields

Field Name	Field Description
Peer	Address of the BGP neighbor. The address is followed by the neighbor port number.
AS	AS number of the peer.
Local	Address of the local routing device. The address is followed by the peer port number.
Type	Type of peer: Internal or External .

Table 41: show bgp neighbor Output Fields (continued)

Field Name	Field Description
State	<p>Current state of the BGP session:</p> <ul style="list-style-type: none"> • Active—BGP is initiating a transport protocol connection in an attempt to connect to a peer. If the connection is successful, BGP sends an Open message. • Connect—BGP is waiting for the transport protocol connection to be completed. • Established—The BGP session has been established, and the peers are exchanging update messages. • Idle—This is the first stage of a connection. BGP is waiting for a Start event. • OpenConfirm—BGP has acknowledged receipt of an open message from the peer and is waiting to receive a keepalive or notification message. • OpenSent—BGP has sent an open message and is waiting to receive an open message from the peer. • route reflector client—The BGP session is established with a route reflector client.
Flags	<p>Internal BGP flags:</p> <ul style="list-style-type: none"> • Aggregate Label—BGP has aggregated a set of incoming labels (labels received from the peer) into a single forwarding label. • CleanUp—The peer session is being shut down. • Delete—This peer has been deleted. • Idled—This peer has been permanently idled. • ImportEval—At the last commit operation, this peer was identified as needing to reevaluate all received routes. • Initializing—The peer session is initializing. • SendRtn—Messages are being sent to the peer. • Sync—This peer is synchronized with the rest of the peer group. • RSync—This peer in the backup Routing Engine is synchronized with the BGP peer in the master Routing Engine for nonstop active routing. • TryConnect—Another attempt is being made to connect to the peer. • Unconfigured—This peer is not configured. • WriteFailed—An attempt to write to this peer failed.
Last state	<p>Previous state of the BGP session:</p> <ul style="list-style-type: none"> • Active—BGP is initiating a transport protocol connection in an attempt to connect to a peer. If the connection is successful, BGP sends an Open message. • Connect—BGP is waiting for the transport protocol connection to be completed. • Established—The BGP session has been established, and the peers are exchanging update messages. • Idle—This is the first stage of a connection. BGP is waiting for a Start event. • OpenConfirm—BGP has acknowledged receipt of an open message from the peer and is waiting to receive a keepalive or notification message. • OpenSent—BGP has sent an open message and is waiting to receive an open message from the peer.

Table 41: show bgp neighbor Output Fields (continued)

Field Name	Field Description
Last event	<p>Last activity that occurred in the BGP session:</p> <ul style="list-style-type: none"> • Closed—The BGP session closed. • ConnectRetry—The transport protocol connection failed, and BGP is trying again to connect. • HoldTime—The session ended because the hold timer expired. • KeepAlive—The local routing device sent a BGP keepalive message to the peer. • Open—The local routing device sent a BGP open message to the peer. • OpenFail—The local routing device did not receive an acknowledgment of a BGP open message from the peer. • RecvKeepAlive—The local routing device received a BGP keepalive message from the peer. • RecvNotify—The local routing device received a BGP notification message from the peer. • RecvOpen—The local routing device received a BGP open message from the peer. • RecvUpdate—The local routing device received a BGP update message from the peer. • Start—The peering session started. • Stop—The peering session stopped. • TransportError—A TCP error occurred.
Last error	<p>Last error that occurred in the BGP session:</p> <ul style="list-style-type: none"> • Cease—An error occurred, such as a version mismatch, that caused the session to close. • Finite State Machine Error—In setting up the session, BGP received a message that it did not understand. • Hold Time Expired—The session's hold time expired. • Message Header Error—The header of a BGP message was malformed. • Open Message Error—A BGP open message contained an error. • None—No errors occurred in the BGP session. • Update Message Error—A BGP update message contained an error.
Export	Name of the export policy that is configured on the peer.
Import	Name of the import policy that is configured on the peer.

Table 41: show bgp neighbor Output Fields (continued)

Field Name	Field Description
Options	<p>Configured BGP options:</p> <ul style="list-style-type: none"> • AddressFamily—Configured address family: inet or inet-vpn. • AdvertiseBGPStatic—Configured BGP static routes are advertised. • AuthKeyChain—Authentication key change is enabled. • BfdEnabled—Status of BFD. • DontGRHelpFateSharingBfdDown—Status of the dont-help-shared-fate-bfd-down option. If this option is configured the device does not go into graceful restart helper mode. • DropPathAttributes—Certain path attributes are configured to be dropped from neighbor updates during inbound processing. • GracefulRestart—Graceful restart is configured. • HoldTime—Hold time configured with the hold-time statement. The hold time is three times the interval at which keepalive messages are sent. • IgnorePathAttributes—Certain path attributes are configured to be ignored in neighbor updates during inbound processing. • Local Address—Address configured with the local-address statement. • LLGR—BGP long-lived graceful restart capability is configured. • LLGRHelperDisabled—BGP long-lived graceful restart is completely disabled for a neighbor. • Multihop—Allow BGP connections to external peers that are not on a directly connected network. • NLRI—Configured MBGP state for the BGP group: multicast, unicast, or both if you have configured nlri any. • Peer AS—Configured peer autonomous system (AS). • Preference—Preference value configured with the preference statement. • Refresh—Configured to refresh automatically when the policy changes. • Rib-group—Configured routing table group. • RFC6514CompliantSafi129—Configured SAFI 129 according to RFC 6514 (BGP VPN multicast used to use SAFI 128).
Path-attributes dropped	Path attribute codes that are dropped from neighbor updates.
Path-attributes ignored	Path attribute codes that are ignored during neighbor updates.
Peer does not support LLGR Restarter or Receiver functionality	BGP neighbor does not support long-lived graceful restart (LLGR) restarter mode completely.
Peer does not support LLGR Restarter functionality	BGP neighbor does not support long-lived graceful restart (LLGR) restarter mode for any family.
Authentication key change	(Appears only if the authentication-keychain statement has been configured) Name of the authentication keychain enabled.
Authentication algorithm	(Appears only if the authentication-algorithm statement has been configured) Type of authentication algorithm enabled: hmac or md5 .

Table 41: show bgp neighbor Output Fields (continued)

Field Name	Field Description
Address families configured	Names of configured address families for the VPN.
BGP-Static Advertisement Policy	Name of the BGP static policy that is configured on the peer.
Local Address	Address of the local routing device.
Remove-private options	Options associated with the remove-private statement.
Holdtime	Hold time configured with the hold-time statement. The hold time is three times the interval at which keepalive messages are sent.
Flags for NLRI inet-label-unicast	Flags related to labeled-unicast: <ul style="list-style-type: none"> • TrafficStatistics—Collection of statistics for labeled-unicast traffic is enabled.
Traffic statistics	Information about labeled-unicast traffic statistics: <ul style="list-style-type: none"> • Options—Options configured for collecting statistics about labeled-unicast traffic. • File—Name and location of statistics log files. • size—Size of all the log files, in bytes. • files—Number of log files.
Traffic Statistics Interval	Time between sample periods for labeled-unicast traffic statistics, in seconds.
Preference	Preference value configured with the preference statement.
Outbound Timer	Time for which the route is available in Junos OS routing table before it is exported to BGP. This field is displayed in the output only if the out-delay parameter is configured to a non-zero value.
Number of flaps	Number of times the BGP session has gone down and then come back up.
Peer ID	Router identifier of the peer.
Group index	Index number for the BGP peer group. The index number differentiates between groups when a single BGP group is split because of different configuration options at the group and peer levels.
Peer index	Index that is unique within the BGP group to which the peer belongs.
Local ID	Router identifier of the local routing device.
Local Interface	Name of the interface on the local routing device.
Active holdtime	Hold time that the local routing device negotiated with the peer.
Keepalive Interval	Keepalive interval, in seconds.

Table 41: show bgp neighbor Output Fields (continued)

Field Name	Field Description
BFD	Status of BFD failure detection.
Local Address	Name of directly connected interface over which direct EBGP peering is established.
NLRI and times for LLGR configured on peer	<p>Names of address families and stale time for BGP long-lived graceful restart configured on the BGP peer or neighbor.</p> <p>Times are displayed using the routing protocol daemon (rpd) %#OT format:</p> <p><weeks>w<days>d <hours>:<minutes>:<seconds></p> <p>Zero leading elements are omitted, for example, a value less than one week do not include the weeks.</p>
NLRI and times that peer supports LLGR Restarter for	<p>Names of address families and stale time that the BGP peer supports for restarter mode for BGP long-lived graceful restart.</p> <p>Times are displayed using the routing protocol daemon (rpd) %#OT format:</p> <p><weeks>w<days>d <hours>:<minutes>:<seconds></p> <p>Zero leading elements are omitted, for example, a value less than one week do not include the weeks.</p>
NLRI that peer saved LLGR forwarding for	Name of the address family for which the BGP peer saved BGP long-lived graceful restart forwarding.
Graceful Restart Details	Amount of time that is remaining until LLGR expires and the time remaining on the GR stale timer, along with RIB details, are displayed while LLGR receiver mode is active (a peer that negotiated LLGR has disconnected and not yet reconnected).
NLRI we are holding stale routes for	Names of address families (NLRIs) for which that stale routes are held or preserved when BGP graceful restart receiver mode is active for a neighbor.
Time until end-of-rib is assumed for stale routes	<p>Amount of time remaining on the stale timer until which end-of-RIB (EoR) markers are assumed when BGP graceful restart receiver mode is active for a neighbor.</p> <p>Time is displayed in Coordinated Universal Time (UTC) format (YYYY-MM-DD-HH:MM:SS). Note that the stale timer display ('Time until end-of-rib is assumed') is also present when a session is active, but the neighbor as not yet sent all of the end-of-rib indications.</p>
Time until stale routes are deleted or become long-lived stale	Amount of time up to which stale routes are deleted or become long-lived stale routes when BGP graceful restart receiver mode is active for a neighbor.
NLRI for restart configured on peer	Names of address families configured for restart.
NLRI advertised by peer	Address families supported by the peer: unicast or multicast .
NLRI for this session	Address families being used for this session.
Peer supports Refresh capability	Remote peer's ability to send and request full routing table readvertisement (route refresh capability). For more information, see RFC 2918, <i>Route Refresh Capability for BGP-4</i> .

Table 41: show bgp neighbor Output Fields (continued)

Field Name	Field Description
Restart time configured on peer	Configured time allowed for restart on the neighbor.
Stale routes from peer are kept for	When graceful restart is negotiated, the maximum time allowed to hold routes from neighbors after the BGP session has gone down.
Peer does not support Restarter functionality	Graceful restart restarter-mode is disabled on the peer.
Peer does not support Receiver functionality	Graceful restart helper-mode is disabled on the peer.
Restart time requested by this peer	Restart time requested by this neighbor during capability negotiation.
Restart flag received from the peer	When this field appears, the BGP speaker has restarted (Restarting), and this peer should not wait for the end-of-rib marker from the speaker before advertising routing information to the speaker.
NLRI that peer supports restart for	Neighbor supports graceful restart for this address family.
NLRI peer can save forwarding state	Neighbor supporting this address family saves all forwarding states.
NLRI that peer saved forwarding for	Neighbor saves all forwarding states for this address family.
NLRI that restart is negotiated for	Router supports graceful restart for this address family.
NLRI of received end-of-rib markers	Address families for which end-of-routing-table markers are received from the neighbor.
NLRI of all end-of-rib markers sent	Address families for which end-of-routing-table markers are sent to the neighbor.
Peer supports 4 byte AS extension (peer-as 1)	Peer understands 4-byte AS numbers in BGP messages. The peer is running Junos OS Release 9.1 or later.
NLRIs for which peer can receive multiple paths	Appears in the command output of the local router if the downstream peer is configured to receive multiple BGP routes to a single destination, instead of only receiving the active route. Possible value is inet-unicast.
NLRIs for which peer can send multiple paths: inet-unicast	Appears in the command output of the local router if the upstream peer is configured to send multiple BGP routes to a single destination, instead of only sending the active route. Possible value is inet-unicast.

Table 41: show bgp neighbor Output Fields (continued)

Field Name	Field Description
Table inet.number	<p>Information about the routing table:</p> <ul style="list-style-type: none"> • RIB State—BGP is in the graceful restart process for this routing table: restart is complete or restart in progress. • Bit—Number that represents the entry in the routing table for this peer. • Send state—State of the BGP group: in sync, not in sync, or not advertising. • Active prefixes—Number of prefixes received from the peer that are active in the routing table. • Received prefixes—Total number of prefixes from the peer, both active and inactive, that are in the routing table. • Accepted prefixes—Total number of prefixes from the peer that have been accepted by a routing policy. • Suppressed due to damping—Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols.
Last traffic (seconds)	Last time any traffic was received from the peer or sent to the peer, and the last time the local routing device checked.
Input messages	Messages that BGP has received from the receive socket buffer, showing the total number of messages, number of update messages, number of times a policy is changed and refreshed, and the buffer size in octets. The buffer size is 16 KB.
Output messages	Messages that BGP has written to the transmit socket buffer, showing the total number of messages, number of update messages, number of times a policy is changed and refreshed, and the buffer size in octets. The buffer size is 16 KB.
Input dropped path attributes	<p>Information about dropped path attributes:</p> <ul style="list-style-type: none"> • Code—Path attribute code. • Count—Path attribute count.
Input ignored path attributes	<p>Information about ignored path attributes:</p> <ul style="list-style-type: none"> • Code—Path attribute code. • Count—Path attribute count.
Output queue	<p>Number of BGP packets that are queued to be transmitted to a particular neighbor for a particular routing table. Output queue 0 is for unicast NLRIs, and queue 1 is for multicast NLRIs.</p> <p>It also specifies the routing table name and the NLRI that the table was advertised through, in the format (<i>routing table name</i>, <i>NLRI</i>).</p> <p>NOTE: The output queue of routing tables that are not advertised, will only show up at extensive output level.</p>
Trace options	Configured tracing of BGP protocol packets and operations.
Trace file	Name of the file to receive the output of the tracing operation.

Table 41: show bgp neighbor Output Fields (continued)

Field Name	Field Description
Filter Updates rcv	(orf option only) Number of outbound-route filters received for each configured address family. NOTE: The counter is cumulative. For example, the counter is increased after the remote peer either resends or clears the outbound route filtering prefix list.
Immediate	(orf option only) Number of route updates received with the immediate flag set. The immediate flag indicates that the BGP peer should readvertise the updated routes. NOTE: The counter is cumulative. For example, the counter is increased after the remote peer either resends or clears the outbound route filtering prefix list.
Filter	(orf option only) Type of prefix filter received: prefix-based or extended-community .
Received filter entries	(orf option only) List of received filters displayed.
seq	(orf option only) Numerical order assigned to this prefix entry among all the received outbound route filter prefix entries.
prefix	(orf option only) Address for the prefix entry that matches the filter.
minlength	(orf option only) Minimum prefix length, in bits, required to match this prefix.
maxlength	(orf option only) Maximum prefix length, in bits, required to match this prefix.
match	(orf option only) For this prefix match, whether to permit or deny route updates.

Sample Output

show bgp neighbor

```
user@host > show bgp neighbor
```

For M Series, MX Series, and T Series routers running Junos OS Release 16.1 or later, the **show bgp neighbor** output includes the BGP group the peer belongs to, the routing instance (if any) that the peer is configured in, and the routing instance that the peer is using for the forwarding context (if applicable). An example follows.

```
Peer: 10.255.7.250+179 AS 10   Local: 10.255.7.248+63740 AS 10
  Group: toAsbr2              Routing-Instance: master
  Forwarding routing-instance: toAs2
    Type: Internal   State: Established   Flags: <Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Export: [ redist_static ]
  Options: <Preference LocalAddress PeerAS Refresh>
  Options: <AdvertiseBGPStatic>
  Local Address: 10.255.7.248 Holdtime: 90 Preference: 170 Outbound Timer: 50
  Number of flaps: 0
  Peer ID: 10.255.7.250   Local ID: 10.255.7.248   Active Holdtime: 90
  Keepalive Interval: 30   Group index: 0   Peer index: 0
  BFD: disabled, down
```

```

NLRI for restart configured on peer: inet-unicast
NLRI advertised by peer: inet-unicast
NLRI for this session: inet-unicast
Peer supports Refresh capability (2)
Stale routes from peer are kept for: 300
Peer does not support Restarter functionality
NLRI that restart is negotiated for: inet-unicast
NLRI of received end-of-rib markers: inet-unicast
NLRI of all end-of-rib markers sent: inet-unicast
Peer supports 4 byte AS extension (peer-as 10)
Peer does not support Addpath
NLRI that we support extended nexthop encoding for: inet-unicast
NLRI that peer supports extended nexthop encoding for: inet-unicast

```

```

Table inet.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          1
  Received prefixes:       1
  Accepted prefixes:       1
  Suppressed due to damping: 0
  Advertised prefixes:     1
Last traffic (seconds): Received 9   Sent 5   Checked 5
Input messages: Total 36   Updates 2   Refreshes 0   Octets 718
Output messages: Total 37   Updates 1   Refreshes 0   Octets 796
Output Queue[0]: 0 (inet.0, inet-unicast)

```

```

Peer: 10.255.162.214+52193 AS 100 Local: 10.255.167.205+179 AS 100
  Type: Internal   State: Established (route reflector client)Flags: <Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress Cluster AddressFamily Rib-group Refresh>
  Address families configured: inet-unicast inet-vpn-unicast route-target
  Local Address: 10.255.167.205 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 10.255.162.214   Local ID: 10.255.167.205   Active Holdtime: 90
  Keepalive Interval: 30    Group index: 0    Peer index: 1

```

show bgp neighbor (dont-help-shared-fate-bfd-down is configured)

```
user@host> show bgp neighbor
```

```

Peer: 10.1.1.1 AS 200           Local: unspecified AS 17
  Group: one                     Routing-Instance: master
  Forwarding routing-instance: master
  Type: External   State: Idle   Flags: <PeerInterfaceError>
  Last State: NoState   Last Event: NoEvent
  Last Error: None
  Options: <Preference PeerAS Refresh>
  Options: <BfdEnabled>
  Options: <DontGRHelpFateSharingBfdDown>
  Holdtime: 90 Preference: 170
  Number of flaps: 0
  Trace options: bridge
  Trace file: /var/log/bgp-log size 131072 files 10

Peer: 20.1.1.1 AS 200           Local: unspecified AS 17
  Group: one                     Routing-Instance: master

```

```

Forwarding routing-instance: master
Type: External      State: Idle      Flags: <PeerInterfaceError>
Last State: NoState      Last Event: NoEvent
Last Error: None
Options: <Preference PeerAS Refresh>
Options: <BfdEnabled>
Options: <DontGRHelpFateSharingBfdDown>
Holdtime: 90 Preference: 170
Number of flaps: 0
Trace options: bridge
Trace file: /var/log/bgp-log size 131072 files 10

Peer: 30.1.1.1 AS 200      Local: unspecified AS 17
Group: two      Routing-Instance: master
Forwarding routing-instance: master
Type: External      State: Idle      Flags: <PeerInterfaceError>
Last State: NoState      Last Event: NoEvent
Last Error: None
Options: <Preference PeerAS Refresh>
Options: <BfdEnabled>
Options: <DontGRHelpFateSharingBfdDown>
Holdtime: 90 Preference: 170
Number of flaps: 0
Trace options: bridge
Trace file: /var/log/bgp-log size 131072 files 10

```

show bgp neighbor (CLNS)

```

user@host> show bgp neighbor

Peer: 10.245.245.1+179 AS 200 Local: 10.245.245.3+3770 AS 100
Type: External      State: Established      Flags: <ImportEval Sync>
Last State: OpenConfirm      Last Event: RecvKeepAlive
Last Error: None
Options: <Multihop Preference LocalAddress HoldTime AddressFamily PeerAS
Rib-group Refresh>
Address families configured: iso-vpn-unicast
Local Address: 10.245.245.3 Holdtime: 90 Preference: 170
Number of flaps: 0
Peer ID: 10.245.245.1      Local ID: 10.245.245.3      Active Holdtime: 90
Keepalive Interval: 30      Peer index: 0
NLRI advertised by peer: iso-vpn-unicast
NLRI for this session: iso-vpn-unicast
Peer supports Refresh capability (2)
Table bgp.isovpn.0 Bit: 10000
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: in sync
  Active prefixes:          3
  Received prefixes:        3
  Suppressed due to damping: 0
  Advertised prefixes:      3
Table aaaa.iso.0
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: not advertising
  Active prefixes:          3
  Received prefixes:        3
  Suppressed due to damping: 0
Last traffic (seconds): Received 6      Sent 5      Checked 5

```



```

Input messages: Total 1736    Updates 4        Refreshes 0        Octets 33385
Output messages: Total 1738    Updates 3        Refreshes 0        Octets 33305
Output Queue[0]: 0 (bgp.isovpn.0, iso-vpn-unicast)
Output Queue[1]: 0 (aaaa.iso.0, iso-vpn-unicast)

```

show bgp neighbor (Layer 2 VPN)

```

user@host> show bgp neighbor
Peer: 10.69.103.2      AS 65536 Local: 10.69.103.1      AS 65539
Type: External      State: Active      Flags: <ImportEval>
Last State: Idle      Last Event: Start
Last Error: None
Export: [ BGP-INET-import ]
Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily PeerAS
Refresh>
Address families configured: inet-unicast
Local Address: 10.69.103.1 Holdtime: 90 Preference: 170
Number of flaps: 0
Peer: 10.69.104.2      AS 65539 Local: 10.69.104.1      AS 65539
Type: External      State: Active      Flags: <ImportEval>
Last State: Idle      Last Event: Start
Last Error: None
Export: [ BGP-L-import ]
Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily PeerAS
Refresh>
Address families configured: inet-labeled-unicast
Local Address: 10.69.104.1 Holdtime: 90 Preference: 170
Number of flaps: 0
Peer: 10.255.14.182+179 AS 69      Local: 10.255.14.176+2131 AS 69
Type: Internal      State: Established  Flags: <ImportEval>
Last State: OpenConfirm Last Event: RecvKeepAlive
Last Error: None
Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily
Rib-group Refresh>
Address families configured: inet-vpn-unicast l2vpn
Local Address: 10.255.14.176 Holdtime: 90 Preference: 170
Number of flaps: 0
Peer ID: 10.255.14.182      Local ID: 10.255.14.176      Active Holdtime: 90
Keepalive Interval: 30
NLRI for restart configured on peer: inet-vpn-unicast l2vpn
NLRI advertised by peer: inet-vpn-unicast l2vpn
NLRI for this session: inet-vpn-unicast l2vpn
Peer supports Refresh capability (2)
Restart time configured on the peer: 120
Stale routes from peer are kept for: 300
Restart time requested by this peer: 120
NLRI that peer supports restart for: inet-vpn-unicast l2vpn
NLRI peer can save forwarding state: inet-vpn-unicast l2vpn
NLRI that peer saved forwarding for: inet-vpn-unicast l2vpn
NLRI that restart is negotiated for: inet-vpn-unicast l2vpn
NLRI of received end-of-rib markers: inet-vpn-unicast l2vpn
Table bgp.l3vpn.0 Bit: 10000
RIB State: BGP restart in progress
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:          10
Received prefixes:        10
Suppressed due to damping: 0
Table bgp.l2vpn.0 Bit: 20000

```

```

RIB State: BGP restart in progress
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:          1
Received prefixes:        1
Suppressed due to damping: 0
Table BGP-INET.inet.0 Bit: 30000
RIB State: BGP restart in progress
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:          2
Received prefixes:        2
Suppressed due to damping: 0
Table BGP-L.inet.0 Bit: 40000
RIB State: BGP restart in progress
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:          2
Received prefixes:        2
Suppressed due to damping: 0
Table LDP.inet.0 Bit: 50000
RIB State: BGP restart is complete
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:          1
Received prefixes:        1
Suppressed due to damping: 0
Table OSPF.inet.0 Bit: 60000
RIB State: BGP restart is complete
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:          2
Received prefixes:        2
Suppressed due to damping: 0
Table RIP.inet.0 Bit: 70000
RIB State: BGP restart is complete
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:          2
Received prefixes:        2
Suppressed due to damping: 0
Table STATIC.inet.0 Bit: 80000
RIB State: BGP restart is complete
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:          1
Received prefixes:        1
Suppressed due to damping: 0
Table L2VPN.l2vpn.0 Bit: 90000
RIB State: BGP restart is complete
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:          1
Received prefixes:        1
Suppressed due to damping: 0
Last traffic (seconds): Received 0    Sent 0    Checked 0
Input messages: Total 14    Updates 13    Refreshes 0    Octets 1053
Output messages: Total 3    Updates 0    Refreshes 0    Octets 105
Output Queue[0]: 0 (bgp.l3vpn.0, inet-vpn-unicast)
Output Queue[1]: 0 (bgp.l2vpn.0, inet-vpn-unicast)

```

```

Output Queue[2]: 0 (BGP-INET.inet.0, inet-vpn-unicast)
Output Queue[3]: 0 (BGP-L.inet.0, inet-vpn-unicast)
Output Queue[4]: 0 (LDP.inet.0, inet-vpn-unicast)
Output Queue[5]: 0 (OSPF.inet.0, inet-vpn-unicast)
Output Queue[6]: 0 (RIP.inet.0, inet-vpn-unicast)
Output Queue[7]: 0 (STATIC.inet.0, inet-vpn-unicast)
Output Queue[8]: 0 (L2VPN.l2vpn.0, inet-vpn-unicast)

```

show bgp neighbor (Layer 3 VPN) (Not supported on the OCX Series.)

```

user@host> show bgp neighbor
Peer: 192.0.2.0.179      AS 10045 Local: 192.0.2.1+1214      AS 10045
  Type: Internal      State: Established      Flags: <ImportEval>
  Last State: OpenConfirm      Last Event: RecvKeepAlive
  Last Error: None
  Export: [ match-all ] Import: [ match-all ]
  Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily
    Rib-group Refresh>
  Address families configured: inet-vpn-unicast
  Local Address: 192.0.2.1 Holdtime: 90 Preference: 170
  Flags for NLRI inet-labeled-unicast: TrafficStatistics
  Traffic Statistics: Options: all File: /var/log/bstat.log
                                size 131072 files 10
  Traffic Statistics Interval: 60
  Number of flaps: 0
  Peer ID: 192.168.1.110      Local ID: 192.168.1.111      Active Holdtime: 90
  Keepalive Interval: 30
  NLRI for restart configured on peer: inet-vpn-unicast
  NLRI advertised by peer: inet-vpn-unicast
  NLRI for this session: inet-vpn-unicast
  Peer supports Refresh capability (2)
  Restart time configured on the peer: 120
  Stale routes from peer are kept for: 300
  Restart time requested by this peer: 120
  NLRI that peer supports restart for: inet-vpn-unicast
  NLRI peer can save forwarding state: inet-vpn-unicast
  NLRI that peer saved forwarding for: inet-vpn-unicast
  NLRI that restart is negotiated for: inet-vpn-unicast
  NLRI of received end-of-rib markers: inet-vpn-unicast
  NLRI of all end-of-rib markers sent: inet-vpn-unicast
  Table bgp.l3vpn.0 Bit: 10000
    RIB State: BGP restart is complete
    RIB State: VPN restart is complete
    Send state: in sync
    Active prefixes:          2
    Received prefixes:        2
    Suppressed due to damping: 0
  Table vpn-green.inet.0 Bit: 20001
    RIB State: BGP restart is complete
    RIB State: VPN restart is complete
    Send state: in sync
    Active prefixes:          2
    Received prefixes:        2
    Suppressed due to damping: 0
  Last traffic (seconds): Received 15      Sent 20      Checked 20
  Input messages: Total 40      Updates 2      Refreshes 0      Octets 856
  Output messages: Total 44      Updates 2      Refreshes 0      Octets 1066
  Output Queue[0]: 0 (bgp.l3vpn.0, inet-vpn-unicast)
  Output Queue[1]: 0 (vpn-green.inet.0, inet-vpn-unicast)

```

```
Trace options: detail packets
Trace file: /var/log/bgpr.log size 131072 files 10
```

show bgp neighbor neighbor-address

```
user@host> show bgp neighbor 192.168.1.111
Peer: 10.255.245.12+179 AS 35 Local: 10.255.245.13+2884 AS 35
  Type: Internal State: Established (route reflector client)Flags: <Sync>
  Last State: OpenConfirm Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress HoldTime Cluster AddressFamily Rib-group
  Refresh>
  Options: RFC6514CompliantSafi129
  Address families configured: inet-vpn-unicast inet-labeled-unicast
  Local Address: 10.255.245.13 Holdtime: 90 Preference: 170
  Flags for NLRI inet-vpn-unicast: AggregateLabel
  Flags for NLRI inet-labeled-unicast: AggregateLabel
  Number of flaps: 0
  Peer ID: 10.255.245.12 Local ID: 10.255.245.13 Active Holdtime: 90
  Keepalive Interval: 30
  BFD: disabled
  NLRI advertised by peer: inet-vpn-unicast inet-labeled-unicast
  NLRI for this session: inet-vpn-unicast inet-labeled-unicast
  Peer supports Refresh capability (2)
  Restart time configured on the peer: 300
  Stale routes from peer are kept for: 60
  Restart time requested by this peer: 300
  NLRI that peer supports restart for: inet-unicast inet6-unicast
  NLRI that restart is negotiated for: inet-unicast inet6-unicast
  NLRI of received end-of-rib markers: inet-unicast inet6-unicast
  NLRI of all end-of-rib markers sent: inet-unicast inet6-unicast
  Table inet.0 Bit: 10000
    RIB State: restart is complete
    Send state: in sync
    Active prefixes: 4
    Received prefixes: 6
    Suppressed due to damping: 0
  Table inet6.0 Bit: 20000
    RIB State: restart is complete
    Send state: in sync
    Active prefixes: 0
    Received prefixes: 2
    Suppressed due to damping: 0
  Last traffic (seconds): Received 3 Sent 3 Checked 3
  Input messages: Total 9 Updates 6 Refreshes 0 Octets 403
  Output messages: Total 7 Updates 3 Refreshes 0 Octets 365
  Output Queue[0]: 0 (inet.0, inet-unicast)
  Output Queue[1]: 0 (inet6.0, inet6-unicast)
  Trace options: detail packets
  Trace file: /var/log/bgpr size 131072 files 10
```

show bgp neighbor neighbor-address

```
user@host> show bgp neighbor 192.168.4.222
Peer: 192.168.4.222+4902 AS 65501 Local: 192.168.4.221+179 AS 65500
  Type: External State: Established Flags: <Sync>
  Last State: OpenConfirm Last Event: RecvKeepAlive
```

```

Last Error: Cease
Export: [ export-policy ] Import: [ import-policy ]
Options: <Preference HoldTime AddressFamily PeerAS PrefixLimit Refresh>
Address families configured: inet-unicast inet-multicast
Holdtime: 60000 Preference: 170
Number of flaps: 4
Last flap event: RecvUpdate
Error: 'Cease' Sent: 5 Recv: 0
Peer ID: 10.255.245.6      Local ID: 10.255.245.5      Active Holdtime: 60000
Keepalive Interval: 20000      Peer index: 0
BFD: disabled, down
Local Interface: fxp0.0
NLRI advertised by peer: inet-unicast inet-multicast
NLRI for this session: inet-unicast inet-multicast
Peer supports Refresh capability (2)
Table inet.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          8
  Received prefixes:        10
  Accepted prefixes:        10
  Suppressed due to damping: 0
  Advertised prefixes:      3
Table inet.2 Bit: 20000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          0
  Received prefixes:        0
  Accepted prefixes:        0
  Suppressed due to damping: 0
  Advertised prefixes:      0
Last traffic (seconds): Received 357 Sent 357 Checked 357
Input messages: Total 4 Updates 2 Refreshes 0 Octets 211
Output messages: Total 4 Updates 1 Refreshes 0 Octets 147
Output Queue[0]: 0 (inet.0, inet-unicast)
Output Queue[1]: 0 (inet.2, inet-multiicast)
Trace options: all
Trace file: /var/log/bgp size 10485760 files 10

```

show bgp neighbor neighbor-address (BGP Graceful Restart Enabled)

```
user@router> show bgp neighbor 10.255.255.16
```

```

Peer: 10.255.255.16 AS 100      Local: 10.255.255.12 AS 100
Type: Internal      State: Active      Flags: <>
Last State: Idle      Last Event: Start
Last Error: None
Options: <Preference LocalAddress AddressFamily Rib-group Refresh>
Options: <LLGR>
Address families configured: 12vpn
Local Address: 10.255.255.12 Holdtime: 90 Preference: 170
NLRI 12vpn:
Number of flaps: 6
Last flap event: Restart
NLRI we are holding stale routes for: inet-vpn-unicast
Time until stale routes are deleted or become long-lived stale: 00:01:57
Time until end-of-rib is assumed for stale routes: 00:04:43
Table bgp.13vpn.0
  RIB State: BGP restart is complete

```

```

RIB State: VPN restart is complete
Send state: not advertising
Active prefixes:          0
Received prefixes:        7
Accepted prefixes:        7
Suppressed due to damping: 0
Table foo.inet.0 Bit: 30000
RIB State: BGP restart is complete
RIB State: VPN restart is complete
Send state: not in sync
Active prefixes:          0
Received prefixes:        7
Accepted prefixes:        7
Suppressed due to damping: 0

```

show bgp neighbor neighbor-address (BGP Long-Lived Graceful Restart)

```
user@router> show bgp neighbor 10.4.12.11
```

```

Peer: 10.4.12.11 AS 100      Local: 10.6.128.225 AS 100
Type: Internal      State: Active      Flags: <>
Last State: Idle      Last Event: Start
Last Error: None
Export: [ foo ]
Options: <Preference LocalAddress Refresh GracefulRestart>
Options: <LLGR>
Local Address: 10.6.128.225 Holdtime: 90 Preference: 170
Number of flaps: 3
Last flap event: Restart
Error: 'Cease' Sent: 0 Recv: 1
Time until long-lived stale routes deleted: inet-vpn-unicast 10:00:22
route-target 10:00:22
Table bgp.l3vpn.0
RIB State: BGP restart is complete
RIB State: VPN restart is complete
Send state: not advertising
Active prefixes:          0
Received prefixes:        7
Accepted prefixes:        7
Suppressed due to damping: 0
Table foo.inet.0 Bit: 30000
RIB State: BGP restart is complete
RIB State: VPN restart is complete
Send state: not in sync
Active prefixes:          0
Received prefixes:        7
Accepted prefixes:        7
Suppressed due to damping: 0

```

show bgp neighbor orf neighbor-address detail

```
user@host > show bgp neighbor orf 192.168.165.56 detail
```

```

Peer: 192.168.165.56+179 Type: External
Group: ext1

inet-unicast
Filter updates rcv:          1 Immediate:          1

```

```

Filter: prefix-based receive
Received filter entries:
  seq 1: prefix 2.2.2.2/32: minlen 32: maxlen 32: match deny:

inet6-unicast
Filter updates rcv:          0 Immediate:          1
Filter: prefix-based receive
Received filter entries:
  *.*

```

show bgp neighbor logical-system

```

user@host > show bgp neighbor logical-system ITR1
Peer: 10.79.8.2+179 AS 65536 Local: 10.79.8.1+50891 AS 65500
Description: MX1
Type: External State: Established Flags: <ImportEval Sync>
Last State: OpenConfirm Last Event: RecvKeepAlive
Last Error: None
...
Table inet.0 Bit: 10000
RIB State: BGP restart is complete
Send state: in sync
Active prefixes:          1
Received prefixes:        1
Accepted prefixes:        1
Suppressed due to damping: 0
Advertised prefixes:      10
Stale prefixes:           4: <=new, line only appears if count is non-0
It is the Number of prefixes marked as stale;
LLGR-stale prefixes:      5: <=new, line only appears if count is non-0
It is the Number of prefixes marked as LLGR-stale

```

show bgp neighbor output-queue

```

user@host > show bgp neighbor output-queue
Peer: 192.0.2.2+179 AS 103 Local: 192.0.2.1+50799 AS 102
Output Queue[0]: 0 (inet.0, inet-unicast)
Priority 1 : 0
Priority 2 : 0
Priority 3 : 0
Priority 4 : 0
Priority 5 : 0
Priority 6 : 0
Priority 7 : 0
Priority 8 : 0
Priority 9 : 0
Priority 10: 0
Priority 11: 0
Priority 12: 0
Priority 13: 0
Priority 14: 0
Priority 15: 0
Priority 16: 0
Expedited : 0

```

show bgp neighbor (Segment Routing Traffic Engineering)

```
user@host > show bgp neighbor

run show bgp neighbor 1.1.1.254
  Peer: 1.1.1.254+60180 AS 100   Local: 1.1.1.1+179 AS 100
  Group: toB                    Routing-Instance: master
  Forwarding routing-instance: master
  Type: Internal   State: Established   Flags: <Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress>
  Address families configured: inet-segment-routing-te
  Local Address: 1.1.1.1 Holdtime: 90 Preference: 170 Local AS: 100 Local System
  AS: 0
  Number of flaps: 0
  Peer ID: 128.9.150.15   Local ID: 128.9.150.110   Active Holdtime: 90
  Keepalive Interval: 30   Group index: 0   Peer index: 0
  I/O Session Thread: bgpio-0 State: Enabled
  BFD: disabled, down
  NLRI for restart configured on peer: inet-segment-routing-te
  NLRI advertised by peer: inet-segment-routing-te
  NLRI for this session: inet-segment-routing-te
  Peer supports Refresh capability (2)
  Stale routes from peer are kept for: 300
  Peer does not support Restarter functionality
  Restart flag received from the peer: Notification
  NLRI that restart is negotiated for: inet-segment-routing-te
  Peer does not support LLGR Restarter functionality
  Peer supports 4 byte AS extension (peer-as 100)
  Peer does not support Addpath
  Last traffic (seconds): Received 17628 Sent 25   Checked 17628
  Input messages:   Total 2   Updates 0   Refreshes 0   Octets 82
  Output messages: Total 1   Updates 0   Refreshes 0   Octets 19
  Trace options: all
  Trace file: /var/log/bgp.log size 10485760 files 10
```


show log

List of Syntax [Syntax on page 683](#)
 [Syntax \(QFX Series and OCX Series\) on page 683](#)
 [Syntax \(TX Matrix Router\) on page 683](#)

Syntax `show log`
 `<filename | user <username>>`

Syntax (QFX Series and OCX Series) `show log filename`
 `<device-type (device-id | device-alias)>`

Syntax (TX Matrix Router) `show log`
 `<all-lcc | lcc number | scc>`
 `<filename | user <username>>`

Release Information Command introduced before Junos OS Release 7.4.
 Command introduced in Junos OS Release 9.0 for EX Series switches.
 Command introduced in Junos OS Release 11.1 for the QFX Series.
 Option *device-type (device-id | device-alias)* is introduced in Junos OS Release 13.1 for the QFX Series.
 Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description List log files, display log file contents, or display information about users who have logged in to the router or switch.



NOTE: On MX Series routers, modifying a configuration to replace a service interface with another service interface is treated as a catastrophic event. When you modify a configuration, the entire configuration associated with the service interface—including NAT pools, rules, and service sets—is deleted and then re-created for the newly specified service interface. If there are active sessions associated with the service interface that is being replaced, these sessions are deleted and the NAT pools are then released, which leads to the generation of the NAT_POOL_RELEASE system log messages. However, because NAT pools are already deleted as a result of the catastrophic configuration change and no longer exist, the NAT_POOL_RELEASE system log messages are not generated for the changed configuration.

Options **none**—List all log files.

<all-lcc | lcc number | scc>—(Routing matrix only)(Optional) Display logging information about all T640 routers (or line-card chassis) or a specific T640 router (replace

number with a value from 0 through 3) connected to a TX Matrix router. Or, display logging information about the TX Matrix router (or switch-card chassis).

device-type—(QFabric system only) (Optional) Display log messages for only one of the following device types:

- **director-device**—Display logs for Director devices.
- **infrastructure-device**—Display logs for the logical components of the QFabric system infrastructure, including the diagnostic Routing Engine, fabric control Routing Engine, fabric manager Routing Engine, and the default network Node group and its backup (NW-NG-0 and NW-NG-0-backup).
- **interconnect-device**—Display logs for Interconnect devices.
- **node-device**—Display logs for Node devices.



NOTE: If you specify the **device-type** optional parameter, you must also specify either the **device-id** or **device-alias** optional parameter.

(device-id | device-alias)—If a device type is specified, display logs for a device of that type. Specify either the device ID or the device alias (if configured).

filename—(Optional) Display the log messages in the specified log file. For the routing matrix, the filename must include the chassis information.



NOTE: The **filename** parameter is mandatory for the QFabric system. If you did not configure a syslog filename, specify the default filename of messages.

user <username>—(Optional) Display logging information about users who have recently logged in to the router or switch. If you include **username**, display logging information about the specified user.

Required Privilege Level

trace

Related Documentation

- [syslog \(System\)](#)

List of Sample Output

[show log on page 685](#)
[show log filename on page 685](#)
[show log filename \(QFabric System\) on page 685](#)
[show log user on page 686](#)

Sample Output

show log

```
user@host> show log
total 57518
-rw-r--r-- 1 root bin      211663 Oct  1 19:44 dcd
-rw-r--r-- 1 root bin      999947 Oct  1 19:41 dcd.0
-rw-r--r-- 1 root bin      999994 Oct  1 17:48 dcd.1
-rw-r--r-- 1 root bin      238815 Oct  1 19:44 rpd
-rw-r--r-- 1 root bin     1049098 Oct  1 18:00 rpd.0
-rw-r--r-- 1 root bin     1061095 Oct  1 12:13 rpd.1
-rw-r--r-- 1 root bin     1052026 Oct  1 06:08 rpd.2
-rw-r--r-- 1 root bin     1056309 Sep 30 18:21 rpd.3
-rw-r--r-- 1 root bin     1056371 Sep 30 14:36 rpd.4
-rw-r--r-- 1 root bin     1056301 Sep 30 10:50 rpd.5
-rw-r--r-- 1 root bin     1056350 Sep 30 07:04 rpd.6
-rw-r--r-- 1 root bin     1048876 Sep 30 03:21 rpd.7
-rw-rw-r-- 1 root bin       19656 Oct  1 19:37 wtmp
```

show log filename

```
user@host> show log rpd
Oct  1 18:00:18 trace_on: Tracing to ?/var/log/rpd? started
Oct  1 18:00:18 EVENT <MTU> ds-5/2/0.0 index 24 <Broadcast PointToPoint Multicast
Oct  1 18:00:18
Oct  1 18:00:19 KRT rcv len 56 V9 seq 148 op add Type route/if af 2 addr
192.0.2.21 nhop type local nhop 192.0.2.21
Oct  1 18:00:19 KRT rcv len 56 V9 seq 149 op add Type route/if af 2 addr
192.0.2.22 nhop type unicast nhop 192.0.2.22
Oct  1 18:00:19 KRT rcv len 48 V9 seq 150 op add Type ifaddr index 24 devindex
43
Oct  1 18:00:19 KRT rcv len 144 V9 seq 151 op chnge Type ifdev devindex 44
Oct  1 18:00:19 KRT rcv len 144 V9 seq 152 op chnge Type ifdev devindex 45
Oct  1 18:00:19 KRT rcv len 144 V9 seq 153 op chnge Type ifdev devindex 46
Oct  1 18:00:19 KRT rcv len 1272 V9 seq 154 op chnge Type ifdev devindex 47
...
```

show log filename (QFabric System)

```
user@qfabric> show log messages
Mar 28 18:00:06 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:06 ED1486
chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 2159)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1486
chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 2, jnxFruL3Index 0,
jnxFruName PIC: @ 0/1/*, jnxFruType 11, jnxFruSlot 0, jnxFruOfflineReason 2,
jnxFruLastPowerOff 0, jnxFruLastPowerOn 2191)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1492
chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 242726)
```

```

Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1492
chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 2, jnxFruL3Index 0,
jnxFruName PIC: @ 0/1/*, jnxFruType 11, jnxFruSlot 0, jnxFruOfflineReason 2,
jnxFruLastPowerOff 0, jnxFruLastPowerOn 242757)
Mar 28 18:00:16 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:16 ED1486
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:00:27 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:27 ED1486
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:00:50 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:50
_DCF_default__NW-INE-0_REO_ file: UI_COMMIT: User 'root' requested 'commit'
operation (comment: none)
Mar 28 18:00:50 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:50
_DCF_default__NW-INE-0_REO_ file: UI_COMMIT: User 'root' requested 'commit'
operation (comment: none)
Mar 28 18:00:55 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:55 ED1492
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:01:10 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:01:10 ED1492
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:02:37 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:02:37 ED1491
chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 33809)

```

show log user

```
user@host> show log user
```

usera	mg2546		Thu Oct 1 19:37	still logged in
usera	mg2529		Thu Oct 1 19:08 - 19:36	(00:28)
usera	mg2518		Thu Oct 1 18:53 - 18:58	(00:04)
root	mg1575		Wed Sep 30 18:39 - 18:41	(00:02)
root	ttyp2	aaa.bbbb.com	Wed Sep 30 18:39 - 18:41	(00:02)
userb	ttyp1	192.0.2.0	Wed Sep 30 01:03 - 01:22	(00:19)

show (ospf | ospf3) overview

List of Syntax	Syntax on page 687 Syntax (EX Series Switch and QFX Series) on page 687
Syntax	<pre>show (ospf ospf3) overview <brief extensive> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)> <realm (ipv4-multicast ipv4-unicast ipv6-multicast)></pre>
Syntax (EX Series Switch and QFX Series)	<pre>show (ospf ospf3) overview <brief extensive> <instance <i>instance-name</i>></pre>
Release Information	<p>Command introduced in Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>realm option introduced in Junos OS Release 9.2.</p> <p>Database protection introduced in Junos 10.2.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Display Open Shortest Path First (OSPF) overview information.
Options	<p>none—Display standard information about all OSPF neighbors for all routing instances.</p> <p>brief extensive—(Optional) Display the specified level of output.</p> <p>instance <i>instance-name</i>—(Optional) Display all OSPF interfaces under the named routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>realm (ipv4-multicast ipv4-unicast ipv6-multicast)—(Optional) (OSPFv3 only) Display information about the specified OSPFv3 realm, or address family. Use the realm option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.</p>
Required Privilege Level	view
List of Sample Output	show ospf overview (without SRGB) on page 689 show ospf overview (with SRGB) on page 690 show ospf overview (With Database Protection) on page 691 show ospf3 overview (With Database Protection) on page 691 show ospf overview extensive on page 691

Output Fields Table 42 on page 688 lists the output fields for the **show ospf overview** command. Output fields are listed in the approximate order in which they appear.

Table 42: show ospf overview Output Fields

Field name	Field Description	Level of Output
Instance	OSPF routing instance.	All levels
Router ID	Router ID of the routing device.	All levels
Route table index	Route table index.	All levels
Configured overload	Overload capability is enabled. If the overload timer is also configured, display the time that remains before it is set to expire. This field is not displayed after the timer expires.	All levels
Topology	Topology identifier.	All levels
Prefix export count	Number of prefixes exported into OSPF.	All levels
Full SPF runs	Number of complete Shortest Path First calculations.	All levels
SPF delay	Delay before performing consecutive Shortest Path First calculations.	All levels
SPF holddown	Delay before performing additional Shortest Path First (SPF) calculations after the maximum number of consecutive SPF calculations is reached.	All levels
SPF rapid runs	Maximum number of Shortest Path First calculations that can be performed in succession before the hold-down timer begins.	All levels
LSA refresh time	Refresh period for link-state advertisement (in minutes).	All levels
SPRING	Source protocol routing in networking: enable or disable .	All levels
Node Segments	Nodes of source protocol routing in networking: enable or disable .	All levels
Ipv4 Index	Ipv4 Index.	All levels
Index Range	Ipv4 Index range.	All levels
Node Segment Blocks Allocated	Details about node segment blocks.	All levels
Database protection state	Current state of database protection.	All levels
Warning threshold	Threshold at which a warning message is logged (percentage of maximum LSA count).	All levels
Non self-generated LSAs	Number of LSAs whose router ID is not equal to the local router ID: Current , Warning (threshold), and Allowed .	All levels

Table 42: show ospf overview Output Fields (continued)

Field name	Field Description	Level of Output
Ignore time	How long the database has been in the ignore state.	All levels
Reset time	How long the database must stay out of the ignore or isolated state before it returns to normal operations.	All levels
Ignore count	Number of times the database has been in the ignore state: Current and Allowed .	All levels
Restart	Graceful restart capability: enabled or disabled .	All levels
Restart duration	Time period for complete reacquisition of OSPF neighbors.	All levels
Restart grace period	Time period for which the neighbors should consider the restarting routing device as part of the topology.	All levels
Graceful restart helper mode	(OSPFv2) Standard graceful restart helper capability (based on RFC 3623): enabled or disabled .	All levels
Restart-signaling helper mode	(OSPFv2) Restart signaling-based graceful restart helper capability (based on RFC 4811, RFC 4812, and RFC 4813): enabled or disabled .	All levels
Helper mode	(OSPFv3) Graceful restart helper capability: enabled or disabled .	All levels
Trace options	OSPF-specific trace options.	extensive
Trace file	Name of the file to receive the output of the tracing operation.	extensive
Area	Area number. Area 0.0.0.0 is the backbone area.	All levels
Stub type	Stub type of area: Normal Stub , Not Stub , or Not so Stubby Stub .	All levels
Authentication Type	Type of authentication: None , Password , or MD5 . NOTE: The Authentication Type field refers to the authentication configured at the <code>[edit protocols ospf area area-id]</code> level. Any authentication configured for an interface in this area will not affect the value of this field.	All levels
Area border routers	Number of area border routers.	All levels
Neighbors	Number of autonomous system boundary routers.	All levels

Sample Output

show ospf overview (without SRGB)

```
user@host> show ospf overview
```

```
Instance: master
Router ID: 10.255.245.6
Route table index: 0
```

```

Configured overload, expires in 118 seconds
LSA refresh time: 50 minutes
SPRING: Enabled
Node Segments: Enabled
Ipv4 Index : 10, Index Range: 2048
Node Segment Blocks Allocated:
Start Index : 0, Size : 256, Label-Range: [ 802048, 802303 ]
Start Index : 256, Size : 256, Label-Range: [ 802304, 802559 ]
Start Index : 512, Size : 256, Label-Range: [ 802560, 802815 ]
Start Index : 768, Size : 256, Label-Range: [ 802816, 803071 ]
Start Index : 1024, Size : 256, Label-Range: [ 803072, 803327 ]
Start Index : 1280, Size : 256, Label-Range: [ 803328, 803583 ]
Start Index : 1536, Size : 256, Label-Range: [ 803584, 803839 ]
Start Index : 1792, Size : 256, Label-Range: [ 803840, 804095 ]
Restart: Enabled
Restart duration: 20 sec
Restart grace period: 40 sec
Helper mode: enabled
Area: 0.0.0.0
Stub type: Not Stub
Authentication Type: None
Area border routers: 0, AS boundary routers: 0
Neighbors
Up (in full state): 0
Topology: default (ID 0)
Prefix export count: 0
Full SPF runs: 1
SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3

```

show ospf overview (with SRGB)

```

user@host> show ospf overview

Instance: master
Router ID: 10.10.10.10
Route table index: 0
LSA refresh time: 50 minutes
Traffic engineering
SPRING: Enabled
SRGB Config Range :
SRGB Start-Label : 1000, SRGB Index-Range : 2000
SRGB Block Allocation: Success
SRGB Start Index : 1000, SRGB Size : 2000, Label-Range: [ 1000, 2999 ]
Node Segments: Enabled
Ipv4 Index : 1000
Post Convergence Backup: Disabled
Area: 0.0.0.0
Stub type: Not Stub
Authentication Type: None
Area border routers: 0, AS boundary routers: 0
Neighbors
Up (in full state): 3
Topology: default (ID 0)
Prefix export count: 0
Full SPF runs: 5
SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3
Backup SPF: Enabled, Remote Backup calculation enabled

```


show ospf overview (With Database Protection)

```

user@host> show ospf overview

Instance: master
  Router ID: 10.255.112.218
  Route table index: 0
  LSA refresh time: 50 minutes
  Traffic engineering
  Restart: Enabled
    Restart duration: 180 sec
    Restart grace period: 210 sec
    Graceful restart helper mode: Enabled
    Restart-signaling helper mode: Enabled
  Database protection state: Normal
    Warning threshold: 70 percent
    Non self-generated LSAs: Current 582, Warning 700, Allowed 1000
    Ignore time: 30, Reset time: 60
    Ignore count: Current 0, Allowed 1
  Area: 0.0.0.0
    Stub type: Not Stub
    Authentication Type: None
    Area border routers: 0, AS boundary routers: 0
  Neighbors
    Up (in full state): 160
  Topology: default (ID 0)
    Prefix export count: 0
    Full SPF runs: 70
    SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3
    Backup SPF: Not Needed

```

show ospf3 overview (With Database Protection)

```

user@host> show ospf3 overview

Instance: master
  Router ID: 10.255.112.128
  Route table index: 0
  LSA refresh time: 50 minutes
  Database protection state: Normal
    Warning threshold: 80 percent
    Non self-generated LSAs: Current 3, Warning 8, Allowed 10
    Ignore time: 30, Reset time: 60
    Ignore count: Current 0, Allowed 2
  Area: 0.0.0.0
    Stub type: Not Stub
    Area border routers: 0, AS boundary routers: 0
  Neighbors
    Up (in full state): 1
  Topology: default (ID 0)
    Prefix export count: 0
    Full SPF runs: 7
    SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3
    Backup SPF: Not Needed

```

show ospf overview extensive

```

user@host> show ospf overview extensive

```

```
Instance: master
Router ID: 1.1.1.103
Route table index: 0
Full SPF runs: 13, SPF delay: 0.200000 sec
LSA refresh time: 50 minutes
Restart: Disabled
Trace options: lsa
Trace file: /var/log/ospf size 131072 files 10
Area: 0.0.0.0
  Stub type: Not Stub
  Authentication Type: None
  Area border routers: 0, AS boundary routers: 0
  Neighbors
    Up (in full state): 1
```

show chassis dedicated-ukern-cpu

Syntax	show chassis dedicated-ukern-cpu
Release Information	Command introduced in Junos OS Release 15.1X49-D100.
Description	Display whether dedicated Bidirectional Forwarding Detection (BFD) is enabled or disabled. If dedicated BFD is enabled, the output of the show command displays the value of the Dedicated Ukern CPU Status field as Enabled .
Options	This command has no options.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Enabling Dedicated and Real-Time BFD on page 106 • dedicated-ukern-cpu (BFD) on page 504 • Understanding BFD for BGP on page 36 • Understanding Distributed BFD on page 47
List of Sample Output	show chassis dedicated-ukern-cpu on page 693
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

show chassis dedicated-ukern-cpu

```
user@host> show chassis dedicated-ukern-cpu
Dedicated Ukern CPU Status: Enabled
```

show chassis realtime-ukern-thread

Syntax	show chassis realtime-ukern-thread
Release Information	Command introduced in Junos OS Release 15.1X49-D100.
Description	Display whether real-time Bidirectional Forwarding Detection (BFD) is enabled or disabled. If real-time BFD is enabled, the output of the show command displays the value of the realtime Ukern thread Status field as Enabled .
Options	This command has no options.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• Enabling Dedicated and Real-Time BFD on page 106• realtime-ukern-thread (BFD) on page 505• Understanding BFD for BGP on page 36• Understanding Distributed BFD on page 47
List of Sample Output	show chassis realtime-ukern-thread on page 694
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

show chassis realtime-ukern-thread

```
user@host> show chassis realtime-ukern-thread
realtime Ukern thread Status: Enabled
```

clear vrrp

Syntax	<code>clear vrrp (all interface <i>interface-name</i>)</code>
Release Information	Command introduced before Junos OS Release 7.4.
Description	Set Virtual Router Redundancy Protocol (VRRP) interface statistics to zero.
Options	all —Clear statistics on all interfaces. interface <i>interface-name</i> —Clear statistics on the specified interface only.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show vrrp on page 782
List of Sample Output	clear vrrp all on page 695
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear vrrp all

```
user@host> clear vrrp all
```

request chassis ssb master switch

Syntax	request chassis ssb master switch <no-confirm>
Release Information	Command introduced before Junos OS Release 7.4.
Description	(M20 router only) Control which System and Switch Board (SSB) is master.
Options	no-confirm —(Optional) Do not request confirmation for the switch.
Additional Information	<p>By default, the SSB in slot 0 (SSB0) is the master and the SSB in slot 1 (SSB1) is the backup. If you use this command to change the master, and then restart the chassis software for any reason, the master reverts to the default setting. To change the default master SSB, include the ssb statement at the [edit chassis redundancy] hierarchy level in the configuration. For more information, see the <i>Junos OS Administration Library</i>.</p> <p>The configurations on the two SSBs do not have to be the same, and they are not automatically synchronized. If you configure both SSBs as masters, when the chassis software restarts for any reason, the SSB in slot 0 becomes the master and the one in slot 1 becomes the backup.</p> <p>The switchover from the primary SSB to the backup SSB is immediate. The SSB takes several seconds to reinitialize the Flexible PIC Concentrators (FPCs) and restart the PICs. The interior gateway protocol (IGP) and BGP convergence times depend on the specific network environment.</p>
Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none">• show chassis ssb on page 763
List of Sample Output	request chassis ssb master switch on page 696 request chassis ssb master switch no-confirm on page 697
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

request chassis ssb master switch

```
user@host> request chassis ssb master switch
warning: Traffic will be interrupted while the PFE is re-initialized
Toggle mastership between system switch boards ? [yes,no] (no) yes
```

Switch initiated, use “show chassis ssb” to verify

request chassis ssb master switch no-confirm

```
user@host> request chassis ssb master switch no-confirm
```

Switch initiated, use “show chassis ssb” to verify

request redundant-power-system multi-backup

Syntax EX2200 switch:

```
request redundant-power-system multi-backup
request redundant-power-system no-multi-backup
```

EX3300 switch:

```
request redundant-power-system multi-backup member member-number
request redundant-power-system no-multi-backup member member-number
```

Release Information Command introduced in Junos OS Release 12.1 for EX Series switches.

Description Configure a redundant power system (RPS) to back up six non-Power-over-Ethernet (PoE) powered switches instead of the default which is to back up three PoE-powered switches.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- [EX Series Redundant Power System Hardware Overview on page 305](#)

List of Sample Output [request redundant-power-system multi-backup on page 698](#)

Sample Output

request redundant-power-system multi-backup

```
user@switch> request redundant-power-system multi-backup member 1
Sending multi-backup setting to RPS
```


request system software in-service-upgrade

Syntax	<code>request system software in-service-upgrade <i>package-name</i></code> <code><no-old-master-upgrade></code> <code><reboot></code>
Syntax (QFX Series)	<code>request system software in-service-upgrade <i>package-name</i></code>
Release Information	<p>Command introduced in Junos OS Release 9.0.</p> <p>Command introduced in Junos OS Release 12.3R2, 13.1R2, and 13.2R1 for TX Matrix Plus routers.</p> <p>Command introduced in Junos OS Release 13.2 for PTX5000 routers.</p> <p>Command introduced in Junos OS Release 13.2 X51-D15 for the QFX Series.</p> <p>Command introduced in Junos OS Release 15.1X54-D60 for the ACX5000 line of routers.</p>
Description	<p>Perform a unified in-service software upgrade (ISSU). A unified ISSU enables you to upgrade from one Junos OS Release to another with no disruption on the control plane and with minimal disruption of traffic.</p> <p>On QFX5100 switches, enable nonstop active routing (NSR) and nonstop bridging (NSB).</p> <p>On QFX5200 switches, non-stop active routing (NSR) and nonstop bridging (NSB).</p>
Options	<p><i>package-name</i>—Location from which the software package or bundle is to be installed. For example:</p> <ul style="list-style-type: none"> • <i>/var/tmp/package-name</i>—For a software package or bundle that is being installed from a local directory on the router. • <i>protocol://hostname/pathname/package-name</i>—For a software package or bundle that is to be downloaded and installed from a remote location. Replace <i>protocol</i> with one of the following: <ul style="list-style-type: none"> • ftp—File Transfer Protocol • http—Hypertext Transfer Protocol • scp—Secure copy (available only for Canada and U.S. version) <p>no-old-master-upgrade—(Optional) When the no-old-master-upgrade option is included, after the backup Routing Engine is rebooted with the new software package and a switchover occurs to make it the new master Routing Engine, the former master (new backup) Routing Engine will not be upgraded to the new software. In this case, you must manually upgrade the former master (new backup) Routing Engine. If you do not include the no-old-master-upgrade option, the system will automatically upgrade the former master Routing Engine.</p>



NOTE: This option is not available on the QFX5100 and QFX5200 switches.

reboot—(Optional) When the **reboot** option is included, the former master (new backup) Routing Engine is automatically rebooted after being upgraded to the new software. When the **reboot** option is not included, you must manually reboot the former master (new backup) Routing Engine using the **request system reboot** command.



NOTE: This option is not available on the QFX series.

Additional Information The following conditions apply to unified ISSUs:

- Unified ISSU is not supported on every platform. For a list of supported platforms, see [“Unified ISSU System Requirements” on page 401](#).
- Unsupported PICs are restarted during a unified ISSU on certain routing devices. For information about supported PICs, see the *High Availability Feature Guide*.
- Unsupported protocols will experience packet loss during a unified ISSU. For information about supported protocols, see the *High Availability Feature Guide*.
- During a unified ISSU, you cannot bring any PICs online or offline on certain routing devices.

For more information, see the *High Availability Feature Guide*.

Required Privilege Level

view

Related Documentation

- *request system software abort*
- *show chassis in-service-upgrade*
- [Getting Started with Unified In-Service Software Upgrade on page 387](#)
- [Performing an In-Service Software Upgrade \(ISSU\) with Non-Stop Routing on page 452](#)
- *Understanding In-Service Software Upgrade with Graceful Restart*
- *Performing an In-Service Software Upgrade (ISSU) with Graceful Restart*
- [Example: Performing a Unified ISSU on page 424](#)

List of Sample Output

[request system software-in-service upgrade reboot on page 701](#)
[request system software-in-service upgrade reboot \(TX Matrix Plus Router\) on page 703](#)

[request system software-in-service upgrade \(QFX5100 Switch\) on page 711](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

[request system software-in-service upgrade reboot](#)

```
{master}

user@host> request system software in-service-upgrade
/var/tmp/jinstall-9.0-20080114.2-domestic-signed.tgz reboot

ISSU: Validating Image
PIC 0/3 will be offlined (In-Service-Upgrade not supported)
Do you want to continue with these actions being taken ? [yes,no] (no) yes

ISSU: Preparing Backup RE
Pushing bundle to re1
Checking compatibility with configuration
Initializing...
Using jbase-9.0-20080114.2
Verified manifest signed by PackageProduction_9_0_0
Using /var/tmp/jinstall-9.0-20080114.2-domestic-signed.tgz
Verified jinstall-9.0-20080114.2-domestic.tgz signed by PackageProduction_9_0_0
Using jinstall-9.0-20080114.2-domestic.tgz
Using jbundle-9.0-20080114.2-domestic.tgz
Checking jbundle requirements on /
Using jbase-9.0-20080114.2.tgz
Verified manifest signed by PackageProduction_9_0_0
Using jkernel-9.0-20080114.2.tgz
Verified manifest signed by PackageProduction_9_0_0
Using jcrypto-9.0-20080114.2.tgz
Verified manifest signed by PackageProduction_9_0_0
Using jpfe-9.0-20080114.2.tgz
Using jdocs-9.0-20080114.2.tgz
Verified manifest signed by PackageProduction_9_0_0
Using jroute-9.0-20080114.2.tgz
Verified manifest signed by PackageProduction_9_0_0
Hardware Database regeneration succeeded
Validating against /config/juniper.conf.gz
mgd: commit complete
Validation succeeded
Installing package '/var/tmp/jinstall-9.0-20080114.2-domestic-signed.tgz' ...
Verified jinstall-9.0-20080114.2-domestic.tgz signed by PackageProduction_9_0_0
Adding jinstall...
Verified manifest signed by PackageProduction_9_0_0

WARNING: This package will load JUNOS 9.0-20080114.2 software.
WARNING: It will save JUNOS configuration files, and SSH keys
WARNING: (if configured), but erase all other files and information
WARNING: stored on this machine. It will attempt to preserve dumps
WARNING: and log files, but this can not be guaranteed. This is the
WARNING: pre-installation stage and all the software is loaded when
WARNING: you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in
/var/db/config/juniper.conf.pre-install
```

Installing the bootstrap installer ...

WARNING: A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
 WARNING: 'request system reboot' command when software installation is
 WARNING: complete. To abort the installation, do not reboot your system,
 WARNING: instead use the 'request system software delete jinstall'
 WARNING: command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall-9.0-20080114.2-domestic-signed.tgz

...

Saving state for rollback ...

Backup upgrade done

Rebooting Backup RE

Rebooting re1

ISSU: Backup RE Prepare Done

Waiting for Backup RE reboot

GRES operational

Initiating Chassis In-Service-Upgrade

Chassis ISSU started

ISSU: Backup RE Prepare Done

ISSU: Preparing Daemons

ISSU: Daemons Ready for ISSU

ISSU: Starting Upgrade for FRUs

ISSU: Preparing for Switchover

ISSU: Ready for Switchover

Checking In-Service-Upgrade status

Item	Status	Reason
FPC 0	Online (ISSU)	
FPC 1	Online (ISSU)	
FPC 2	Online (ISSU)	
FPC 6	Online (ISSU)	
FPC 7	Online (ISSU)	

Resolving mastership...

Complete. The other routing engine becomes the master.

ISSU: RE switchover Done

ISSU: Upgrading Old Master RE

Installing package '/var/tmp/paKEuy' ...

Verified jinstall-9.0-20080114.2-domestic.tgz signed by PackageProduction_9_0_0

Adding jinstall...

Verified manifest signed by PackageProduction_9_0_0

WARNING: This package will load JUNOS 9.0-20080114.2 software.
 WARNING: It will save JUNOS configuration files, and SSH keys
 WARNING: (if configured), but erase all other files and information
 WARNING: stored on this machine. It will attempt to preserve dumps
 WARNING: and log files, but this can not be guaranteed. This is the
 WARNING: pre-installation stage and all the software is loaded when
 WARNING: you reboot the system.

Saving the config files ...

NOTICE: uncommitted changes have been saved in

/var/db/config/juniper.conf.pre-install

Installing the bootstrap installer ...

WARNING: A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
 WARNING: 'request system reboot' command when software installation is
 WARNING: complete. To abort the installation, do not reboot your system,
 WARNING: instead use the 'request system software delete jinstall'
 WARNING: command as soon as this operation completes.

```

Saving package file in /var/sw/pkg/jinstall-9.0-20080114.2-domestic-signed.tgz
...
cp: /var/tmp/paKEuy is a directory (not copied).
Saving state for rollback ...
ISSU: Old Master Upgrade Done
ISSU: IDLE
Shutdown NOW!
Reboot consistency check bypassed - jinstall 9.0-20080114.2 will complete
installation upon reboot
[pid 30227]

*** FINAL System shutdown message from root@host ***

System going down IMMEDIATELY

Connection to host closed.

```

request system software-in-service upgrade reboot (TX Matrix Plus Router)

```

{master}

user@host> request system software in-service-upgrade
/var/tmp/jinstall-12.3R2-domestic-signed.tgz

Chassis ISSU Check Done
ISSU: Validating Image
PIC 8/1 will be offlined (In-Service-Upgrade not supported)
PIC 19/2 will be offlined (In-Service-Upgrade not supported)
PIC 15/3 will be offlined (In-Service-Upgrade not supported)
Do you want to continue with these actions being taken ? [yes,no] (no) yes

Checking compatibility with configuration
Initializing...
Using jbase-12.3R2
Verified manifest signed by PackageProduction_12_3_0
Using /var/tmp/jinstall-12.3R2-domestic-signed.tgz
Verified jinstall-12.3R2-domestic.tgz signed by PackageProduction_12_3_0
Using jinstall-12.3R2-domestic.tgz
Using jbundle-12.3R2-domestic.tgz
Checking jbundle requirements on /
Using jbase-12.3R2.tgz
Verified manifest signed by PackageProduction_12_3_0
Verified jbase-12.3R2 signed by PackageProduction_12_3_0
Using /var/validate/chroot/tmp/jbundle/jboot-12.3R2.tgz
Using jcrypto-12.3R2.tgz
Verified manifest signed by PackageProduction_12_3_0
Verified jcrypto-12.3R2 signed by PackageProduction_12_3_0
Using jdocs-12.3R2.tgz
Verified manifest signed by PackageProduction_12_3_0
Verified jdocs-12.3R2 signed by PackageProduction_12_3_0
Using jkernel-12.3R2.tgz
Verified manifest signed by PackageProduction_12_3_0
Verified jkernel-12.3R2 signed by PackageProduction_12_3_0
Using jpfe-12.3R2.tgz
WARNING: jpfe-12.3R2.tgz: not a signed package
WARNING: jpfe-common-12.3R2.tgz: not a signed package
Verified jpfe-common-12.3R2 signed by PackageProduction_12_3_0
WARNING: jpfe-T-12.3R2.tgz: not a signed package

```

```
Verified jpfe-T-12.3R2 signed by PackageProduction_12_3_0
Using jplatform-12.3R2.tgz
Verified manifest signed by PackageProduction_12_3_0
Verified jplatform-12.3R2 signed by PackageProduction_12_3_0
Using jroute-12.3R2.tgz
Verified manifest signed by PackageProduction_12_3_0
Verified jroute-12.3R2 signed by PackageProduction_12_3_0
Using jruntime-12.3R2.tgz
Verified manifest signed by PackageProduction_12_3_0
Verified jruntime-12.3R2 signed by PackageProduction_12_3_0
Using jservices-12.3R2.tgz
Using jservices-crypto-12.3R2.tgz
Hardware Database regeneration succeeded
Validating against /config/juniper.conf.gz
mgd: commit complete
Validation succeeded
ISSU: Preparing LCC Backup REs
Pushing bundle to lcc0-re1
Pushing bundle to lcc1-re1
Pushing bundle to lcc2-re1
Pushing bundle to lcc3-re1
Pushing bundle to sfc0-re1
Installing package '/var/tmp/jinstall-12.3R2-domestic-signed.tgz' ...
Verified jinstall-12.3R2-domestic.tgz signed by PackageProduction_12_3_0
Adding jinstall...
Verified manifest signed by PackageProduction_12_3_0

WARNING: This package will load JUNOS 12.3R2 software.
WARNING: It will save JUNOS configuration files, and SSH keys
WARNING: (if configured), but erase all other files and information
WARNING: stored on this machine. It will attempt to preserve dumps
WARNING: and log files, but this can not be guaranteed. This is the
WARNING: pre-installation stage and all the software is loaded when
WARNING: you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in
/var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

WARNING: A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING: 'request system reboot' command when software installation is
WARNING: complete. To abort the installation, do not reboot your system,
WARNING: instead use the 'request system software delete jinstall'
WARNING: command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall-12.3R2-domestic-signed.tgz ...
Saving state for rollback ...
Installing package '/var/tmp/jinstall-12.3R2-domestic-signed.tgz' ...
Verified jinstall-12.3R2-domestic.tgz signed by PackageProduction_12_3_0
Adding jinstall...
Verified manifest signed by PackageProduction_12_3_0

WARNING: This package will load JUNOS 12.3R2 software.
WARNING: It will save JUNOS configuration files, and SSH keys
WARNING: (if configured), but erase all other files and information
WARNING: stored on this machine. It will attempt to preserve dumps
WARNING: and log files, but this can not be guaranteed. This is the
WARNING: pre-installation stage and all the software is loaded when
WARNING: you reboot the system.
```

```

Saving the config files ...
NOTICE: uncommitted changes have been saved in
/var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

WARNING:  A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING:  'request system reboot' command when software installation is
WARNING:  complete. To abort the installation, do not reboot your system,
WARNING:  instead use the 'request system software delete jinstall'
WARNING:  command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall-12.3R2-domestic-signed.tgz ...
Saving state for rollback ...
Installing package '/var/tmp/jinstall-12.3R2-domestic-signed.tgz' ...
Verified jinstall-12.3R2-domestic.tgz signed by PackageProduction_12_3_0
Adding jinstall...
Verified manifest signed by PackageProduction_12_3_0

WARNING:  This package will load JUNOS 12.3R2 software.
WARNING:  It will save JUNOS configuration files, and SSH keys
WARNING:  (if configured), but erase all other files and information
WARNING:  stored on this machine. It will attempt to preserve dumps
WARNING:  and log files, but this can not be guaranteed. This is the
WARNING:  pre-installation stage and all the software is loaded when
WARNING:  you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in
/var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

WARNING:  A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING:  'request system reboot' command when software installation is
WARNING:  complete. To abort the installation, do not reboot your system,
WARNING:  instead use the 'request system software delete jinstall'
WARNING:  command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall-12.3R2-domestic-signed.tgz ...
Saving state for rollback ...
Installing package '/var/tmp/jinstall-12.3R2-domestic-signed.tgz' ...
Verified jinstall-12.3R2-domestic.tgz signed by PackageProduction_12_3_0
Adding jinstall...
Verified manifest signed by PackageProduction_12_3_0

WARNING:  This package will load JUNOS 12.3R2 software.
WARNING:  It will save JUNOS configuration files, and SSH keys
WARNING:  (if configured), but erase all other files and information
WARNING:  stored on this machine. It will attempt to preserve dumps
WARNING:  and log files, but this can not be guaranteed. This is the
WARNING:  pre-installation stage and all the software is loaded when
WARNING:  you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in
/var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

WARNING:  A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING:  'request system reboot' command when software installation is

```

```
WARNING: complete. To abort the installation, do not reboot your system,
WARNING: instead use the 'request system software delete jinstall'
WARNING: command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall-12.3R2-domestic-signed.tgz ...
Saving state for rollback ...
ISSU: Preparing SFC Backup RE
NOTICE: Validating configuration against jinstall-12.3R2-domestic-signed.tgz.
NOTICE: Use the 'no-validate' option to skip this if desired.
Checking compatibility with configuration
Initializing...
Using jbase-12.3R2
Verified manifest signed by PackageProduction_12_3_0
Using /var/tmp/jinstall-12.3R2-domestic-signed.tgz
Verified jinstall-12.3R2-domestic.tgz signed by PackageProduction_12_3_0
Using jinstall-12.3R2-domestic.tgz
Using jbundle-12.3R2-domestic.tgz
Checking jbundle requirements on /
Using jbase-12.3R2.tgz
Verified manifest signed by PackageProduction_12_3_0
Verified jbase-12.3R2 signed by PackageProduction_12_3_0
Using /var/validate/chroot/tmp/jbundle/jboot-12.3R2.tgz
Using jcrypto-12.3R2.tgz
Verified manifest signed by PackageProduction_12_3_0
Verified jcrypto-12.3R2 signed by PackageProduction_12_3_0
Using jdocs-12.3R2.tgz
Verified manifest signed by PackageProduction_12_3_0
Verified jdocs-12.3R2 signed by PackageProduction_12_3_0
Using jkernel-12.3R2.tgz
Verified manifest signed by PackageProduction_12_3_0
Verified jkernel-12.3R2 signed by PackageProduction_12_3_0
Using jpfe-12.3R2.tgz
WARNING: jpfe-12.3R2.tgz: not a signed package
WARNING: jpfe-common-12.3R2.tgz: not a signed package
Verified jpfe-common-12.3R2 signed by PackageProduction_12_3_0
WARNING: jpfe-T-12.3R2.tgz: not a signed package
Verified jpfe-T-12.3R2 signed by PackageProduction_12_3_0
Using jplatform-12.3R2.tgz
Verified manifest signed by PackageProduction_12_3_0
Verified jplatform-12.3R2 signed by PackageProduction_12_3_0
Using jroute-12.3R2.tgz
Verified manifest signed by PackageProduction_12_3_0
Verified jroute-12.3R2 signed by PackageProduction_12_3_0
Using jruntime-12.3R2.tgz
Verified manifest signed by PackageProduction_12_3_0
Verified jruntime-12.3R2 signed by PackageProduction_12_3_0
Using jservices-12.3R2.tgz
Using jservices-crypto-12.3R2.tgz
Hardware Database regeneration succeeded
Validating against /config/juniper.conf.gz
mgd: commit complete
Validation succeeded
Installing package '/var/tmp/jinstall-12.3R2-domestic-signed.tgz' ...
Verified jinstall-12.3R2-domestic.tgz signed by PackageProduction_12_3_0
Adding jinstall...
Verified manifest signed by PackageProduction_12_3_0

WARNING: This package will load JUNOS 12.3R2 software.
WARNING: It will save JUNOS configuration files, and SSH keys
WARNING: (if configured), but erase all other files and information
```



```

WARNING:      stored on this machine. It will attempt to preserve dumps
WARNING:      and log files, but this can not be guaranteed. This is the
WARNING:      pre-installation stage and all the software is loaded when
WARNING:      you reboot the system.

```

```

Saving the config files ...
NOTICE: uncommitted changes have been saved in
/var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

```

```

WARNING:      A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING:      'request system reboot' command when software installation is
WARNING:      complete. To abort the installation, do not reboot your system,
WARNING:      instead use the 'request system software delete jinstall'
WARNING:      command as soon as this operation completes.

```

```

Saving package file in /var/sw/pkg/jinstall-12.3R2-domestic-signed.tgz ...
Saving state for rollback ...
SFC Backup upgrade done
Rebooting SFC Backup RE

```

```

Rebooting sfc0-re1
ISSU: SFC Backup RE Prepare Done
Waiting for SFC Backup RE reboot

```

```

Rebooting lcc0-re1
Rebooting LCC [lcc0-re1]

```

```

Rebooting lcc1-re1
Rebooting LCC [lcc1-re1]

```

```

Rebooting lcc2-re1
Rebooting LCC [lcc2-re1]

```

```

Rebooting lcc3-re1
Rebooting LCC [lcc3-re1]
LCC Backup REs have rebooted
Waiting for LCC Backup REs come back online
ISSU: LCC Backup REs Prepare Done
GRES operational
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
lcc0-re0:

```

Item	Status	Reason
FPC 1	Online (ISSU)	
PIC 0	Online (ISSU)	
FPC 2	Online (ISSU)	
FPC 3	Online (ISSU)	
PIC 1	Online (ISSU)	
FPC 4	Online (ISSU)	
FPC 6	Online (ISSU)	
FPC 7	Online (ISSU)	

```
lcc1-re0:
```

```
-----
Item      Status      Reason
FPC 0     Online (ISSU)
  PIC 3    Online (ISSU)
FPC 1     Online (ISSU)
FPC 2     Online (ISSU)
FPC 4     Online (ISSU)
FPC 6     Online (ISSU)
FPC 7     Online (ISSU)
```

```
lcc2-re0:
```

```
-----
Item      Status      Reason
FPC 0     Online (ISSU)
FPC 2     Online (ISSU)
FPC 3     Online (ISSU)
  PIC 0    Online (ISSU)
FPC 4     Online (ISSU)
FPC 6     Online (ISSU)
FPC 7     Online (ISSU)
  PIC 1    Online (ISSU)
```

```
lcc3-re0:
```

```
-----
Item      Status      Reason
FPC 0     Online (ISSU)
  PIC 0    Online (ISSU)
FPC 1     Online (ISSU)
FPC 2     Online (ISSU)
FPC 3     Online (ISSU)
  PIC 2    Online (ISSU)
FPC 4     Online (ISSU)
FPC 5     Online (ISSU)
FPC 6     Online (ISSU)
FPC 7     Online (ISSU)
  PIC 1    Online (ISSU)
```

```
lcc0-re0:
```

```
-----
Resolving mastership...
Complete. The other routing engine becomes the master.
```

```
lcc1-re0:
```

```
-----
Resolving mastership...
Complete. The other routing engine becomes the master.
```

```
lcc2-re0:
```

```
-----
Resolving mastership...
Complete. The other routing engine becomes the master.
```

```
lcc3-re0:
```

```
-----
Resolving mastership...
Complete. The other routing engine becomes the master.
Resolving mastership...
Complete. The other routing engine becomes the master.
ISSU: RE switchover Done
```

ISSU: Upgrading SFC Old Master RE

```
lcc0-re0:
Installing package '/var/tmp/jinstall-12.3R2-domestic-signed.tgz' ...
Verified jinstall-12.3R2-domestic.tgz signed by PackageProduction_12_3_0
Adding jinstall...
Verified manifest signed by PackageProduction_12_3_0
```

```
WARNING: This package will load JUNOS 12.3R2 software.
WARNING: It will save JUNOS configuration files, and SSH keys
WARNING: (if configured), but erase all other files and information
WARNING: stored on this machine. It will attempt to preserve dumps
WARNING: and log files, but this can not be guaranteed. This is the
WARNING: pre-installation stage and all the software is loaded when
WARNING: you reboot the system.
```

```
Saving the config files ...
NOTICE: uncommitted changes have been saved in
/var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...
```

```
WARNING: A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING: 'request system reboot' command when software installation is
WARNING: complete. To abort the installation, do not reboot your system,
WARNING: instead use the 'request system software delete jinstall'
WARNING: command as soon as this operation completes.
```

```
Saving package file in /var/sw/pkg/jinstall-12.3R2-domestic-signed.tgz ...
Saving state for rollback ...
```

```
lcc1-re0:
Installing package '/var/tmp/jinstall-12.3R2-domestic-signed.tgz' ...
Verified jinstall-12.3R2-domestic.tgz signed by PackageProduction_12_3_0
Adding jinstall...
Verified manifest signed by PackageProduction_12_3_0
```

```
WARNING: This package will load JUNOS 12.3R2 software.
WARNING: It will save JUNOS configuration files, and SSH keys
WARNING: (if configured), but erase all other files and information
WARNING: stored on this machine. It will attempt to preserve dumps
WARNING: and log files, but this can not be guaranteed. This is the
WARNING: pre-installation stage and all the software is loaded when
WARNING: you reboot the system.
```

```
Saving the config files ...
NOTICE: uncommitted changes have been saved in
/var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...
```

```
WARNING: A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING: 'request system reboot' command when software installation is
WARNING: complete. To abort the installation, do not reboot your system,
WARNING: instead use the 'request system software delete jinstall'
WARNING: command as soon as this operation completes.
```

```
Saving package file in /var/sw/pkg/jinstall-12.3R2-domestic-signed.tgz ...
Saving state for rollback ...
```

```
lcc2-re0:
Installing package '/var/tmp/jinstall-12.3R2-domestic-signed.tgz' ...
```

```
Verified jinstall-12.3R2-domestic.tgz signed by PackageProduction_12_3_0
Adding jinstall...
Verified manifest signed by PackageProduction_12_3_0

WARNING: This package will load JUNOS 12.3R2 software.
WARNING: It will save JUNOS configuration files, and SSH keys
WARNING: (if configured), but erase all other files and information
WARNING: stored on this machine. It will attempt to preserve dumps
WARNING: and log files, but this can not be guaranteed. This is the
WARNING: pre-installation stage and all the software is loaded when
WARNING: you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in
/var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

WARNING: A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING: 'request system reboot' command when software installation is
WARNING: complete. To abort the installation, do not reboot your system,
WARNING: instead use the 'request system software delete jinstall'
WARNING: command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall-12.3R2-domestic-signed.tgz ...
Saving state for rollback ...

lcc3-re0:
Installing package '/var/tmp/jinstall-12.3R2-domestic-signed.tgz' ...
Verified jinstall-12.3R2-domestic.tgz signed by PackageProduction_12_3_0
Adding jinstall...
Verified manifest signed by PackageProduction_12_3_0

WARNING: This package will load JUNOS 12.3R2 software.
WARNING: It will save JUNOS configuration files, and SSH keys
WARNING: (if configured), but erase all other files and information
WARNING: stored on this machine. It will attempt to preserve dumps
WARNING: and log files, but this can not be guaranteed. This is the
WARNING: pre-installation stage and all the software is loaded when
WARNING: you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in
/var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

WARNING: A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING: 'request system reboot' command when software installation is
WARNING: complete. To abort the installation, do not reboot your system,
WARNING: instead use the 'request system software delete jinstall'
WARNING: command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall-12.3R2-domestic-signed.tgz ...
Saving state for rollback ...
Installing package '/var/tmp/paBWTg' ...
Verified jinstall-12.3R2-domestic.tgz signed by PackageProduction_12_3_0
Adding jinstall...
Verified manifest signed by PackageProduction_12_3_0

WARNING: This package will load JUNOS 12.3R2 software.
WARNING: It will save JUNOS configuration files, and SSH keys
```

```

WARNING:      (if configured), but erase all other files and information
WARNING:      stored on this machine. It will attempt to preserve dumps
WARNING:      and log files, but this can not be guaranteed. This is the
WARNING:      pre-installation stage and all the software is loaded when
WARNING:      you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in
/var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

WARNING:      A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING:      'request system reboot' command when software installation is
WARNING:      complete. To abort the installation, do not reboot your system,
WARNING:      instead use the 'request system software delete jinstall'
WARNING:      command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall-12.3R2-domestic-signed ...
cp: /var/tmp/paBWTg is a directory (not copied).
Saving state for rollback ...
ISSU: SFC Old Master Upgrade Done
ISSU: IDLE

```

request system software-in-service upgrade (QFX5100 Switch)

```

{master}

user@switch> request system software in-service-upgrade
/var/tmp/jinstall-qfx-132_x51_vjunos.0-domestic.tgz

ISSU: Validating Image
Prepare for ISSU
spawn the backup VM
ISSU: Preparing Backup RE
Backup upgrade done
ISSU: Backup RE Prepare Done
waiting for backup RE switchover ready
GRES operational
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: FPC Warm Booting
ISSU: FPC Warm Booted
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item          Status          Reason
  FPC 0         Online (ISSU)
send ISSU done to chassisd on backup VM
Chassis ISSU Completed
ISSU: IDLE
mgd_package_opus_issu: Initiate em0 device handoff

```

request system software in-service-upgrade (MX Series 5G Universal Routing Platforms and EX9200 Switches)

Syntax `request system software in-service-upgrade package-name`
 `<no-copy>`
 `<no-old-master-upgrade>`
 `<reboot>`
 `<unlink>`

Release Information Command introduced in Junos OS Release 11.2.
 Command introduced in Junos OS Release 14.1 for MX Series Virtual Chassis.
 Command introduced in Junos OS Release 14.2 for EX Series switches.

Description Perform a unified in-service software upgrade (unified ISSU). Unified ISSU enables you to upgrade from one Junos OS release to another with no disruption on the control plane and with minimal disruption of traffic. Unified ISSU is supported only by dual Routing Engine platforms. In addition, graceful Routing Engine switchover (GRES) and nonstop active routing (NSR) must be enabled.

Options *package-name*—Location from which the software package or bundle is to be installed.
 For example:

- */var/tmp/package-name*—For a software package or bundle that is being installed from a local directory on the router.
- *protocol://hostname/pathname/package-name*—For a software package or bundle that is to be downloaded and installed from a remote location. Replace *protocol* with one of the following:
 - **ftp**—File Transfer Protocol
 - **http**—Hypertext Transfer Protocol
 - **scp**—Secure copy (available only for Canada and U.S. version)

no-copy—(Optional) When the **no-copy** option is included, copies of package files are not saved on the Packet Forwarding Engine.

The **no-copy** option is not available for an MX Series Virtual Chassis or an EX9200 Virtual Chassis.

no-old-master-upgrade—(Optional) When the **no-old-master-upgrade** option is included, after the backup Routing Engine is rebooted with the new software package and a switchover occurs to make it the new master Routing Engine, the former master (new backup) Routing Engine is not upgraded to the new software. In this case, you must manually upgrade the former master (new backup) Routing Engine. If you do not include the **no-old-master-upgrade** option, the system automatically upgrades the former master Routing Engine.

The **no-old-master-upgrade** option is not available for an MX Series Virtual Chassis or an EX9200 Virtual Chassis.

reboot—(Optional) When the **reboot** option is included, the former master (new backup) Routing Engine is automatically rebooted after being upgraded to the new software. When the **reboot** option is not included, you must manually reboot the former master (new backup) Routing Engine using the **request system reboot** command.

The **reboot** option is accepted but ignored for an MX Series Virtual Chassis or an EX9200 Virtual Chassis. A unified ISSU in an MX Series Virtual Chassis or EX9200 Virtual Chassis always reboots all Routing Engines in the member routers or switches.

unlink—(Optional) When the **unlink** option is included, the package is removed from **/var/home** whether the installation is successful or unsuccessful.

The **unlink** option is not available for an MX Series Virtual Chassis or an EX9200 Virtual Chassis.

Additional Information The following conditions apply to unified ISSUs:

- Unified ISSUs are supported on MX Series 5G Universal Routing Platforms and EX9200 switches.
- Unsupported PICs (on EX9200, PICs are known as “line cards”) are restarted during a unified ISSU. For information about supported PICs, see the *High Availability Feature Guide*. For information about supported EX9200 line cards, see “Unified ISSU System Requirements” on page 401.
- Unsupported protocols will experience packet loss during a unified ISSU. For information about supported protocols, see the *High Availability Feature Guide* or, for EX9200, see “Unified ISSU System Requirements” on page 401.
- During a unified ISSU, you cannot bring any PICs online or offline.

For more information, see the *High Availability Feature Guide*.

Required Privilege Level view

Related Documentation

- *request system software abort*
- *show chassis in-service-upgrade*

List of Sample Output [request system software in-service-upgrade reboot on page 714](#)
[request system software in-service-upgrade \(MX Series Virtual Chassis\) on page 725](#)

Output Fields When you enter this command, you are provided feedback about the status of your request.

Sample Output

request system software in-service-upgrade reboot

```
{master}
```

```
user@host> request system software in-service-upgrade
/var/tmp/jinstall-11.2B2.1-domestic-signed.tgz reboot

Chassis ISSU Check Done
ISSU: Validating Image
Checking compatibility with configuration
Initializing...
Using jbase-11.2B1.5
Verified manifest signed by PackageProduction_11_2_0
Verified jbase-11.2B1.5 signed by PackageProduction_11_2_0
Using /var/tmp/jinstall-11.2B2.1-domestic-signed.tgz
Verified jinstall-11.2B2.1-domestic.tgz signed by PackageProduction_11_2_0
Using jinstall-11.2B2.1-domestic.tgz
Using jbundle-11.2B2.1-domestic.tgz
Checking jbundle requirements on /
Using jbase-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jbase-11.2B2.1 signed by PackageProduction_11_2_0
Using /var/validate/chroot/tmp/jbundle/jboot-11.2B2.1.tgz
Using jcrypto-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jcrypto-11.2B2.1 signed by PackageProduction_11_2_0
Using jdocs-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jdocs-11.2B2.1 signed by PackageProduction_11_2_0
Using jkernel-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jkernel-11.2B2.1 signed by PackageProduction_11_2_0
Using jpfe-11.2B2.1.tgz
Using jroute-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jroute-11.2B2.1 signed by PackageProduction_11_2_0
Using jruntime-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jruntime-11.2B2.1 signed by PackageProduction_11_2_0
Using jservices-11.2B2.1.tgz
Auto-deleting old jservices-voice ...
Removing /opt/sdk/service-packages/jservices-voice ...
Removing jservices-voice-bsg-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-voice ...
Verified jservices-voice-bsg-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /var/sw/pkg ...
Creating /opt/sdk/service-packages/jservices-voice ...
Storing jservices-voice-bsg-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-voice/jservices-voice-bsg ->
/var/sw/pkg/jservices-voice-bsg-11.2B2.1.tgz...
Auto-deleting old jservices-bgf ...
Removing /opt/sdk/service-packages/jservices-bgf ...
Removing jservices-bgf-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-bgf ...
Verified jservices-bgf-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-bgf ...
```



```

Storing jservices-bgf-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-bgf/jservices-bgf-pic ->
/var/sw/pkg/jservices-bgf-pic-11.2B2.1.tgz...
Auto-deleting old jservices-aac1 ...
Removing /opt/sdk/service-packages/jservices-aac1 ...
Removing jservices-aac1-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-aac1 ...
Verified jservices-aac1-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-aac1 ...
Storing jservices-aac1-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-aac1/jservices-aac1-pic ->
/var/sw/pkg/jservices-aac1-pic-11.2B2.1.tgz...
Auto-deleting old jservices-llpdf ...
Removing /opt/sdk/service-packages/jservices-llpdf ...
Removing jservices-llpdf-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-llpdf ...
Verified jservices-llpdf-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-llpdf ...
Storing jservices-llpdf-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-llpdf/jservices-llpdf-pic ->
/var/sw/pkg/jservices-llpdf-pic-11.2B2.1.tgz...
Auto-deleting old jservices-ptsp ...
Removing /opt/sdk/service-packages/jservices-ptsp ...
Removing jservices-ptsp-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-ptsp ...
Verified jservices-ptsp-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-ptsp ...
Storing jservices-ptsp-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-ptsp/jservices-ptsp-pic ->
/var/sw/pkg/jservices-ptsp-pic-11.2B2.1.tgz...
Auto-deleting old jservices-sfw ...
Removing /opt/sdk/service-packages/jservices-sfw ...
Removing jservices-sfw-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-sfw ...
Verified jservices-sfw-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-sfw ...
Storing jservices-sfw-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-sfw/jservices-sfw-pic ->
/var/sw/pkg/jservices-sfw-pic-11.2B2.1.tgz...
Auto-deleting old jservices-nat ...
Removing /opt/sdk/service-packages/jservices-nat ...
Removing jservices-nat-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-nat ...
Verified jservices-nat-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-nat ...
Storing jservices-nat-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-nat/jservices-nat-pic ->
/var/sw/pkg/jservices-nat-pic-11.2B2.1.tgz...
Auto-deleting old jservices-alg ...
Removing /opt/sdk/service-packages/jservices-alg ...
Removing jservices-alg-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-alg ...
Verified jservices-alg-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-alg ...

```

```
Storing jservices-alg-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-alg/jservices-alg-pic ->
/var/sw/pkg/jservices-alg-pic-11.2B2.1.tgz...
Auto-deleting old jservices-cpcd ...
Removing /opt/sdk/service-packages/jservices-cpcd ...
Removing jservices-cpcd-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-cpcd ...
Verified jservices-cpcd-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-cpcd ...
Storing jservices-cpcd-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-cpcd/jservices-cpcd-pic ->
/var/sw/pkg/jservices-cpcd-pic-11.2B2.1.tgz...
Auto-deleting old jservices-rpm ...
Removing /opt/sdk/service-packages/jservices-rpm ...
Removing jservices-rpm-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-rpm ...
Verified jservices-rpm-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-rpm ...
Storing jservices-rpm-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-rpm/jservices-rpm-pic ->
/var/sw/pkg/jservices-rpm-pic-11.2B2.1.tgz...
Auto-deleting old jservices-hcm ...
Removing /opt/sdk/service-packages/jservices-hcm ...
Removing jservices-hcm-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-hcm ...
Verified jservices-hcm-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-hcm ...
Storing jservices-hcm-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-hcm/jservices-hcm-pic ->
/var/sw/pkg/jservices-hcm-pic-11.2B2.1.tgz...
Auto-deleting old jservices-appid ...
Removing /opt/sdk/service-packages/jservices-appid ...
Removing jservices-appid-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-appid ...
Verified jservices-appid-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-appid ...
Storing jservices-appid-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-appid/jservices-appid-pic ->
/var/sw/pkg/jservices-appid-pic-11.2B2.1.tgz...
Auto-deleting old jservices-idp ...
Removing /opt/sdk/service-packages/jservices-idp ...
Removing jservices-idp-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-idp ...
Verified jservices-idp-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-idp ...
Storing jservices-idp-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-idp/jservices-idp-pic ->
/var/sw/pkg/jservices-idp-pic-11.2B2.1.tgz...
Using jservices-crypto-11.2B2.1.tgz
Auto-deleting old jservices-crypto-base ...
Removing /opt/sdk/service-packages/jservices-crypto-base ...
Removing jservices-crypto-base-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-crypto-base ...
Verified jservices-crypto-base-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
```

```

Creating /opt/sdk/service-packages/jservices-crypto-base ...
Storing jservices-crypto-base-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-crypto-base/jservices-crypto-base-pic
-> /var/sw/pkg/jservices-crypto-base-pic-11.2B2.1.tgz...
Auto-deleting old jservices-ssl ...
Removing /opt/sdk/service-packages/jservices-ssl ...
Removing jservices-ssl-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-ssl ...
Verified jservices-ssl-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-ssl ...
Storing jservices-ssl-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-ssl/jservices-ssl-pic ->
/var/sw/pkg/jservices-ssl-pic-11.2B2.1.tgz...
Hardware Database regeneration succeeded
Validating against /config/juniper.conf.gz
mgd: commit complete
Validation succeeded
ISSU: Preparing Backup RE
Pushing bundle to re1
NOTICE: Validating configuration against jinstall-11.2B2.1-domestic-signed.tgz.
NOTICE: Use the 'no-validate' option to skip this if desired.
Checking compatibility with configuration
Initializing...
Using jbase-11.2B1.5
Verified manifest signed by PackageProduction_11_2_0
Verified jbase-11.2B1.5 signed by PackageProduction_11_2_0
Using /var/tmp/jinstall-11.2B2.1-domestic-signed.tgz
Verified jinstall-11.2B2.1-domestic.tgz signed by PackageProduction_11_2_0
Using jinstall-11.2B2.1-domestic.tgz
Using jbundle-11.2B2.1-domestic.tgz
Checking jbundle requirements on /
Using jbase-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jbase-11.2B2.1 signed by PackageProduction_11_2_0
Using /var/validate/chroot/tmp/jbundle/jboot-11.2B2.1.tgz
Using jcrypto-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jcrypto-11.2B2.1 signed by PackageProduction_11_2_0
Using jdocs-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jdocs-11.2B2.1 signed by PackageProduction_11_2_0
Using jkernel-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jkernel-11.2B2.1 signed by PackageProduction_11_2_0
Using jpfe-11.2B2.1.tgz
Using jroute-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jroute-11.2B2.1 signed by PackageProduction_11_2_0
Using jruntime-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jruntime-11.2B2.1 signed by PackageProduction_11_2_0
Using jservices-11.2B2.1.tgz
Auto-deleting old jservices-voice ...
Removing /opt/sdk/service-packages/jservices-voice ...
Removing jservices-voice-bsg-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-voice ...
Verified jservices-voice-bsg-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /var/sw/pkg ...

```

```
Creating /opt/sdk/service-packages/jservices-voice ...
Storing jservices-voice-bsg-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-voice/jservices-voice-bsg ->
/var/sw/pkg/jservices-voice-bsg-11.2B2.1.tgz...
Auto-deleting old jservices-bgf ...
Removing /opt/sdk/service-packages/jservices-bgf ...
Removing jservices-bgf-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-bgf ...
Verified jservices-bgf-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-bgf ...
Storing jservices-bgf-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-bgf/jservices-bgf-pic ->
/var/sw/pkg/jservices-bgf-pic-11.2B2.1.tgz...
Auto-deleting old jservices-aac1 ...
Removing /opt/sdk/service-packages/jservices-aac1 ...
Removing jservices-aac1-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-aac1 ...
Verified jservices-aac1-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-aac1 ...
Storing jservices-aac1-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-aac1/jservices-aac1-pic ->
/var/sw/pkg/jservices-aac1-pic-11.2B2.1.tgz...
Auto-deleting old jservices-llpdf ...
Removing /opt/sdk/service-packages/jservices-llpdf ...
Removing jservices-llpdf-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-llpdf ...
Verified jservices-llpdf-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-llpdf ...
Storing jservices-llpdf-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-llpdf/jservices-llpdf-pic ->
/var/sw/pkg/jservices-llpdf-pic-11.2B2.1.tgz...
Auto-deleting old jservices-ptsp ...
Removing /opt/sdk/service-packages/jservices-ptsp ...
Removing jservices-ptsp-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-ptsp ...
Verified jservices-ptsp-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-ptsp ...
Storing jservices-ptsp-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-ptsp/jservices-ptsp-pic ->
/var/sw/pkg/jservices-ptsp-pic-11.2B2.1.tgz...
Auto-deleting old jservices-sfw ...
Removing /opt/sdk/service-packages/jservices-sfw ...
Removing jservices-sfw-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-sfw ...
Verified jservices-sfw-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-sfw ...
Storing jservices-sfw-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-sfw/jservices-sfw-pic ->
/var/sw/pkg/jservices-sfw-pic-11.2B2.1.tgz...
Auto-deleting old jservices-nat ...
Removing /opt/sdk/service-packages/jservices-nat ...
Removing jservices-nat-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-nat ...
Verified jservices-nat-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
```

```

Creating /opt/sdk/service-packages/jservices-nat ...
Storing jservices-nat-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-nat/jservices-nat-pic ->
/var/sw/pkg/jservices-nat-pic-11.2B2.1.tgz...
Auto-deleting old jservices-alg ...
Removing /opt/sdk/service-packages/jservices-alg ...
Removing jservices-alg-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-alg ...
Verified jservices-alg-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-alg ...
Storing jservices-alg-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-alg/jservices-alg-pic ->
/var/sw/pkg/jservices-alg-pic-11.2B2.1.tgz...
Auto-deleting old jservices-cpcd ...
Removing /opt/sdk/service-packages/jservices-cpcd ...
Removing jservices-cpcd-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-cpcd ...
Verified jservices-cpcd-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-cpcd ...
Storing jservices-cpcd-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-cpcd/jservices-cpcd-pic ->
/var/sw/pkg/jservices-cpcd-pic-11.2B2.1.tgz...
Auto-deleting old jservices-rpm ...
Removing /opt/sdk/service-packages/jservices-rpm ...
Removing jservices-rpm-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-rpm ...
Verified jservices-rpm-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-rpm ...
Storing jservices-rpm-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-rpm/jservices-rpm-pic ->
/var/sw/pkg/jservices-rpm-pic-11.2B2.1.tgz...
Auto-deleting old jservices-hcm ...
Removing /opt/sdk/service-packages/jservices-hcm ...
Removing jservices-hcm-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-hcm ...
Verified jservices-hcm-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-hcm ...
Storing jservices-hcm-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-hcm/jservices-hcm-pic ->
/var/sw/pkg/jservices-hcm-pic-11.2B2.1.tgz...
Auto-deleting old jservices-appid ...
Removing /opt/sdk/service-packages/jservices-appid ...
Removing jservices-appid-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-appid ...
Verified jservices-appid-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-appid ...
Storing jservices-appid-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-appid/jservices-appid-pic ->
/var/sw/pkg/jservices-appid-pic-11.2B2.1.tgz...
Auto-deleting old jservices-idp ...
Removing /opt/sdk/service-packages/jservices-idp ...
Removing jservices-idp-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-idp ...
Verified jservices-idp-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0

```

```
Creating /opt/sdk/service-packages/jservices-idp ...
Storing jservices-idp-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-idp/jservices-idp-pic ->
/var/sw/pkg/jservices-idp-pic-11.2B2.1.tgz...
Using jservices-crypto-11.2B2.1.tgz
Auto-deleting old jservices-crypto-base ...
Removing /opt/sdk/service-packages/jservices-crypto-base ...
Removing jservices-crypto-base-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-crypto-base ...
Verified jservices-crypto-base-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-crypto-base ...
Storing jservices-crypto-base-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-crypto-base/jservices-crypto-base-pic
-> /var/sw/pkg/jservices-crypto-base-pic-11.2B2.1.tgz...
Auto-deleting old jservices-ssl ...
Removing /opt/sdk/service-packages/jservices-ssl ...
Removing jservices-ssl-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-ssl ...
Verified jservices-ssl-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-ssl ...
Storing jservices-ssl-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-ssl/jservices-ssl-pic ->
/var/sw/pkg/jservices-ssl-pic-11.2B2.1.tgz...
Hardware Database regeneration succeeded
Validating against /config/juniper.conf.gz
mgd: commit complete
Validation succeeded
Installing package '/var/tmp/jinstall-11.2B2.1-domestic-signed.tgz' ...
Verified jinstall-11.2B2.1-domestic.tgz signed by PackageProduction_11_2_0
Adding jinstall...
Verified manifest signed by PackageProduction_11_2_0

WARNING: This package will load JUNOS 11.2B2.1 software.
WARNING: It will save JUNOS configuration files, and SSH keys
WARNING: (if configured), but erase all other files and information
WARNING: stored on this machine. It will attempt to preserve dumps
WARNING: and log files, but this can not be guaranteed. This is the
WARNING: pre-installation stage and all the software is loaded when
WARNING: you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in
/var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

WARNING: A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING: 'request system reboot' command when software installation is
WARNING: complete. To abort the installation, do not reboot your system,
WARNING: instead use the 'request system software delete jinstall'
WARNING: command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall-11.2B2.1-domestic-signed.tgz ...
Saving state for rollback ...
Backup upgrade done
Rebooting Backup RE

Rebooting re1
ISSU: Backup RE Prepare Done
```

```

Waiting for Backup RE reboot
GRES operational
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item             Status             Reason
  FPC 1            Online (ISSU)
  FPC 4            Online (ISSU)
  FPC 8            Online (ISSU)
  FPC 10           Online (ISSU)
Resolving mastership...
Complete. The other routing engine becomes the master.
ISSU: RE switchover Done
ISSU: Upgrading Old Master RE
NOTICE: Validating configuration against jinstall-11.2B2.1-domestic-signed.tgz.
NOTICE: Use the 'no-validate' option to skip this if desired.
Checking compatibility with configuration
Initializing...
Using jbase-11.2B1.5
Verified manifest signed by PackageProduction_11_2_0
Verified jbase-11.2B1.5 signed by PackageProduction_11_2_0
Using /var/tmp/jinstall-11.2B2.1-domestic-signed.tgz
Verified jinstall-11.2B2.1-domestic.tgz signed by PackageProduction_11_2_0
Using jinstall-11.2B2.1-domestic.tgz
Using jbundle-11.2B2.1-domestic.tgz
Checking jbundle requirements on /
Using jbase-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jbase-11.2B2.1 signed by PackageProduction_11_2_0
Using /var/validate/chroot/tmp/jbundle/jboot-11.2B2.1.tgz
Using jcrypto-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jcrypto-11.2B2.1 signed by PackageProduction_11_2_0
Using jdocs-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jdocs-11.2B2.1 signed by PackageProduction_11_2_0
Using jkernel-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jkernel-11.2B2.1 signed by PackageProduction_11_2_0
Using jpfe-11.2B2.1.tgz
Using jroute-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jroute-11.2B2.1 signed by PackageProduction_11_2_0
Using jruntime-11.2B2.1.tgz
Verified manifest signed by PackageProduction_11_2_0
Verified jruntime-11.2B2.1 signed by PackageProduction_11_2_0
Using jservices-11.2B2.1.tgz
Auto-deleting old jservices-voice ...
Removing /opt/sdk/service-packages/jservices-voice ...
Removing jservices-voice-bsg-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-voice ...
Verified jservices-voice-bsg-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /var/sw/pkg ...
Creating /opt/sdk/service-packages/jservices-voice ...

```

```
Storing jservices-voice-bsg-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-voice/jservices-voice-bsg ->
/var/sw/pkg/jservices-voice-bsg-11.2B2.1.tgz...
Auto-deleting old jservices-bgf ...
Removing /opt/sdk/service-packages/jservices-bgf ...
Removing jservices-bgf-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-bgf ...
Verified jservices-bgf-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-bgf ...
Storing jservices-bgf-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-bgf/jservices-bgf-pic ->
/var/sw/pkg/jservices-bgf-pic-11.2B2.1.tgz...
Auto-deleting old jservices-aac1 ...
Removing /opt/sdk/service-packages/jservices-aac1 ...
Removing jservices-aac1-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-aac1 ...
Verified jservices-aac1-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-aac1 ...
Storing jservices-aac1-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-aac1/jservices-aac1-pic ->
/var/sw/pkg/jservices-aac1-pic-11.2B2.1.tgz...
Auto-deleting old jservices-llpdf ...
Removing /opt/sdk/service-packages/jservices-llpdf ...
Removing jservices-llpdf-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-llpdf ...
Verified jservices-llpdf-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-llpdf ...
Storing jservices-llpdf-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-llpdf/jservices-llpdf-pic ->
/var/sw/pkg/jservices-llpdf-pic-11.2B2.1.tgz...
Auto-deleting old jservices-ptsp ...
Removing /opt/sdk/service-packages/jservices-ptsp ...
Removing jservices-ptsp-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-ptsp ...
Verified jservices-ptsp-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-ptsp ...
Storing jservices-ptsp-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-ptsp/jservices-ptsp-pic ->
/var/sw/pkg/jservices-ptsp-pic-11.2B2.1.tgz...
Auto-deleting old jservices-sfw ...
Removing /opt/sdk/service-packages/jservices-sfw ...
Removing jservices-sfw-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-sfw ...
Verified jservices-sfw-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-sfw ...
Storing jservices-sfw-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-sfw/jservices-sfw-pic ->
/var/sw/pkg/jservices-sfw-pic-11.2B2.1.tgz...
Auto-deleting old jservices-nat ...
Removing /opt/sdk/service-packages/jservices-nat ...
Removing jservices-nat-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-nat ...
Verified jservices-nat-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-nat ...
```



```

Storing jservices-nat-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-nat/jservices-nat-pic ->
/var/sw/pkg/jservices-nat-pic-11.2B2.1.tgz...
Auto-deleting old jservices-alg ...
Removing /opt/sdk/service-packages/jservices-alg ...
Removing jservices-alg-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-alg ...
Verified jservices-alg-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-alg ...
Storing jservices-alg-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-alg/jservices-alg-pic ->
/var/sw/pkg/jservices-alg-pic-11.2B2.1.tgz...
Auto-deleting old jservices-cpcd ...
Removing /opt/sdk/service-packages/jservices-cpcd ...
Removing jservices-cpcd-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-cpcd ...
Verified jservices-cpcd-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-cpcd ...
Storing jservices-cpcd-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-cpcd/jservices-cpcd-pic ->
/var/sw/pkg/jservices-cpcd-pic-11.2B2.1.tgz...
Auto-deleting old jservices-rpm ...
Removing /opt/sdk/service-packages/jservices-rpm ...
Removing jservices-rpm-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-rpm ...
Verified jservices-rpm-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-rpm ...
Storing jservices-rpm-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-rpm/jservices-rpm-pic ->
/var/sw/pkg/jservices-rpm-pic-11.2B2.1.tgz...
Auto-deleting old jservices-hcm ...
Removing /opt/sdk/service-packages/jservices-hcm ...
Removing jservices-hcm-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-hcm ...
Verified jservices-hcm-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-hcm ...
Storing jservices-hcm-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-hcm/jservices-hcm-pic ->
/var/sw/pkg/jservices-hcm-pic-11.2B2.1.tgz...
Auto-deleting old jservices-appid ...
Removing /opt/sdk/service-packages/jservices-appid ...
Removing jservices-appid-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-appid ...
Verified jservices-appid-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-appid ...
Storing jservices-appid-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-appid/jservices-appid-pic ->
/var/sw/pkg/jservices-appid-pic-11.2B2.1.tgz...
Auto-deleting old jservices-idp ...
Removing /opt/sdk/service-packages/jservices-idp ...
Removing jservices-idp-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-idp ...
Verified jservices-idp-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-idp ...

```

```

Storing jservices-idp-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-idp/jservices-idp-pic ->
/var/sw/pkg/jservices-idp-pic-11.2B2.1.tgz...
Using jservices-crypto-11.2B2.1.tgz
Auto-deleting old jservices-crypto-base ...
Removing /opt/sdk/service-packages/jservices-crypto-base ...
Removing jservices-crypto-base-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-crypto-base ...
Verified jservices-crypto-base-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-crypto-base ...
Storing jservices-crypto-base-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-crypto-base/jservices-crypto-base-pic
-> /var/sw/pkg/jservices-crypto-base-pic-11.2B2.1.tgz...
Auto-deleting old jservices-ssl ...
Removing /opt/sdk/service-packages/jservices-ssl ...
Removing jservices-ssl-pic-11.2B1.5.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-ssl ...
Verified jservices-ssl-pic-11.2B2.1.tgz signed by PackageProduction_11_2_0
Creating /opt/sdk/service-packages/jservices-ssl ...
Storing jservices-ssl-pic-11.2B2.1.tgz in /var/sw/pkg ...
Link: /opt/sdk/service-packages/jservices-ssl/jservices-ssl-pic ->
/var/sw/pkg/jservices-ssl-pic-11.2B2.1.tgz...
Hardware Database regeneration succeeded
Validating against /config/juniper.conf.gz
mgd: commit complete
Validation succeeded
Installing package '/var/tmp/jinstall-11.2B2.1-domestic-signed.tgz' ...
Verified jinstall-11.2B2.1-domestic.tgz signed by PackageProduction_11_2_0
Adding jinstall...
Verified manifest signed by PackageProduction_11_2_0

WARNING: This package will load JUNOS 11.2B2.1 software.
WARNING: It will save JUNOS configuration files, and SSH keys
WARNING: (if configured), but erase all other files and information
WARNING: stored on this machine. It will attempt to preserve dumps
WARNING: and log files, but this can not be guaranteed. This is the
WARNING: pre-installation stage and all the software is loaded when
WARNING: you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in
/var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

WARNING: A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING: 'request system reboot' command when software installation is
WARNING: complete. To abort the installation, do not reboot your system,
WARNING: instead use the 'request system software delete jinstall'
WARNING: command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall-11.2B2.1-domestic-signed.tgz ...
Saving state for rollback ...
ISSU: Old Master Upgrade Done
ISSU: IDLE
Shutdown NOW!
Reboot consistency check bypassed - jinstall 11.2B2.1 will complete installation
upon reboot
[pid 66780]

```

```
*** FINAL System shutdown message from user@host> ***
System going down IMMEDIATELY
```

request system software in-service-upgrade (MX Series Virtual Chassis)

```
{master:member0-re0}
```

```
user@host> request system software in-service-upgrade
jinstall-14.1-20140114.2-domestic-signed.tgz
```

```
[Jan 30 10:45:32]:ISSU: IDLE
```

```
Beginning in-service-upgrade at Jan 30, 2014; 10:45:34
```

```
[Jan 30 10:45:34]:ISSU: Validating Image
```

```
Validating VC readiness...
```

```
Validating required configuration...
```

```
Validating release compatibility...
```

```
Validation successful
```

```
Initiating chassis in-service-upgrade
```

```
[Jan 30 10:46:56]:ISSU: Preparing LCC Backup REs
```

```
Copying new release to all RE's
```

```
Pushing bundle to member0-re0
```

```
Pushing bundle to member1-re0
```

```
Pushing bundle to member1-re1
```

```
[Jan 30 10:51:11]:ISSU: Preparing Backup RE
```

```
Arming new release on all RE's
```

```
member0-re0:
```

```
-----
Installing package
```

```
'/var/tmp/jinstall-14.1-20140114_ib_14_1_psd.1-domestic-signed.tgz' ...
```

```
Verified jinstall-14.1-20140114_ib_14_1_psd.1-domestic.tgz signed by
```

```
PackageDevelopmentEc_2014
```

```
Adding jinstall...
```

```
WARNING: The software that is being installed has limited support.
```

```
WARNING: Run 'file show /etc/notices/unsupported.txt' for details.
```

```
verixec: accepting signer: PackageDevelopmentEc_2014
```

```
Verified manifest signed by PackageDevelopmentEc_2014
```

```
WARNING: This package will load JUNOS 14.1-20140114_ib_14_1_psd.1 software.
```

```
WARNING: It will save JUNOS configuration files, and SSH keys
```

```
WARNING: (if configured), but erase all other files and information
```

```
WARNING: stored on this machine. It will attempt to preserve dumps
```

```
WARNING: and log files, but this can not be guaranteed. This is the
```

```
WARNING: pre-installation stage and all the software is loaded when
```

```
WARNING: you reboot the system.
```

```
Saving the config files ...
```

```
NOTICE: uncommitted changes have been saved in
```

```
/var/db/config/juniper.conf.pre-install
```

```
Installing the bootstrap installer ...
```

```
WARNING: A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
```

```
WARNING: 'request system reboot' command when software installation is
```

```
WARNING: complete. To abort the installation, do not reboot your system,
```

```
WARNING: instead use the 'request system software delete jinstall'
```

```
WARNING: command as soon as this operation completes.
```

```

Saving package file in
/var/sw/pkg/jinstall-14.1-20140114_ib_14_1_psd.1-domestic-signed.tgz ...
Saving state for rollback ...

member1-re0:
-----
Installing package
'/var/tmp/jinstall-14.1-20140114_ib_14_1_psd.1-domestic-signed.tgz' ...
Verified jinstall-14.1-20140114_ib_14_1_psd.1-domestic.tgz signed by
PackageDevelopmentEc_2014
Adding jinstall...

WARNING:    The software that is being installed has limited support.
WARNING:    Run 'file show /etc/notices/unsupported.txt' for details.

verixec: accepting signer: PackageDevelopmentEc_2014
Verified manifest signed by PackageDevelopmentEc_2014

WARNING:    This package will load JUNOS 14.1-20140114_ib_14_1_psd.1 software.
WARNING:    It will save JUNOS configuration files, and SSH keys
WARNING:    (if configured), but erase all other files and information
WARNING:    stored on this machine. It will attempt to preserve dumps
WARNING:    and log files, but this can not be guaranteed. This is the
WARNING:    pre-installation stage and all the software is loaded when
WARNING:    you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in
/var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

WARNING:    A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING:    'request system reboot' command when software installation is
WARNING:    complete. To abort the installation, do not reboot your system,
WARNING:    instead use the 'request system software delete jinstall'
WARNING:    command as soon as this operation completes.

Saving package file in
/var/sw/pkg/jinstall-14.1-20140114_ib_14_1_psd.1-domestic-signed.tgz ...
Saving state for rollback ...

member1-re1:
-----
Installing package
'/var/tmp/jinstall-14.1-20140114_ib_14_1_psd.1-domestic-signed.tgz' ...
Verified jinstall-14.1-20140114_ib_14_1_psd.1-domestic.tgz signed by
PackageDevelopmentEc_2014
Adding jinstall...

WARNING:    The software that is being installed has limited support.
WARNING:    Run 'file show /etc/notices/unsupported.txt' for details.

verixec: accepting signer: PackageDevelopmentEc_2014
Verified manifest signed by PackageDevelopmentEc_2014

WARNING:    This package will load JUNOS 14.1-20140114_ib_14_1_psd.1 software.
WARNING:    It will save JUNOS configuration files, and SSH keys
WARNING:    (if configured), but erase all other files and information
WARNING:    stored on this machine. It will attempt to preserve dumps

```

```

WARNING:    and log files, but this can not be guaranteed. This is the
WARNING:    pre-installation stage and all the software is loaded when
WARNING:    you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in
/var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

WARNING:    A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING:    'request system reboot' command when software installation is
WARNING:    complete. To abort the installation, do not reboot your system,
WARNING:    instead use the 'request system software delete jinstall'
WARNING:    command as soon as this operation completes.

Saving package file in
/var/sw/pkg/jinstall-14.1-20140114_ib_14_1_psd.1-domestic-signed.tgz ...
Saving state for rollback ...
Installing package
'/var/tmp/jinstall-14.1-20140114_ib_14_1_psd.1-domestic-signed.tgz' ...
Verified jinstall-14.1-20140114_ib_14_1_psd.1-domestic.tgz signed by
PackageDevelopmentEc_2014
Adding jinstall...

WARNING:    The software that is being installed has limited support.
WARNING:    Run 'file show /etc/notices/unsupported.txt' for details.

verixec: accepting signer: PackageDevelopmentEc_2014
Verified manifest signed by PackageDevelopmentEc_2014

WARNING:    This package will load JUNOS 14.1-20140114_ib_14_1_psd.1 software.
WARNING:    It will save JUNOS configuration files, and SSH keys
WARNING:    (if configured), but erase all other files and information
WARNING:    stored on this machine. It will attempt to preserve dumps
WARNING:    and log files, but this can not be guaranteed. This is the
WARNING:    pre-installation stage and all the software is loaded when
WARNING:    you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in
/var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

WARNING:    A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING:    'request system reboot' command when software installation is
WARNING:    complete. To abort the installation, do not reboot your system,
WARNING:    instead use the 'request system software delete jinstall'
WARNING:    command as soon as this operation completes.

Saving package file in
/var/sw/pkg/jinstall-14.1-20140114_ib_14_1_psd.1-domestic-signed.tgz ...
Saving state for rollback ...
[Jan 30 11:03:12]:ISSU: Backup RE Prepare Done
Rebooting standby RE's
Sending Reboot Command to member0-re0
Shutdown NOW!
Reboot consistency check bypassed - jinstall 14.1-20140114_ib_14_1_psd.1 will
complete installation upon reboot
[pid 2757]
Sending Reboot Command to member1-re1

```

```

Shutdown NOW!
Reboot consistency check bypassed - jinstall 14.1-20140114_ib_14_1_psd.1 will
complete installation upon reboot
[pid 2670]
Waiting for standby RE's to boot
[Jan 30 11:18:26]:ISSU: LCC Backup REs Prepare Done
Waiting for standby RE's to have the correct ISSU state
Waiting for protocol backup to be ready to switch mastership
Switching mastership on the protocol backup chassis to slot 1
Waiting for protocol backup chassis master switch to complete
Globally updating ISSU state
Waiting for protocol backup chassis to become GRES ready
[Jan 30 11:19:18]:ISSU: VC Protocol Backup has Switched
Passing ISSU control to chassisd
Chassis ISSU Started
[Jan 30 11:21:01]:ISSU: Preparing Daemons
[Jan 30 11:22:02]:ISSU: Daemons Ready for ISSU
[Jan 30 11:22:06]:ISSU: Starting Upgrade for FRUs
[Jan 30 11:25:42]:ISSU: Preparing for Switchover
[Jan 30 11:26:06]:ISSU: Ready for Switchover
[Jan 30 11:26:20]:ISSU: All VC Members Ready for Switchover
Waiting for master chassis to be switch ready
Switching mastership locally
Resolving mastership...
Complete. The other routing engine becomes the master.
Waiting for virtual chassis roles to switch
Globally updating ISSU state to IDLE
[Jan 30 11:26:33]:ISSU: IDLE
Rebooting protocol backup standby RE.
Sending Reboot Command to member1-re0

member1-re0:
-----
Shutdown NOW!
Reboot consistency check bypassed - jinstall 14.1-20140114_ib_14_1_psd.1 will
complete installation upon reboot
[pid 10462]
Rebooting locally to complete the in service upgrade.
Shutdown NOW!
Reboot consistency check bypassed - jinstall 14.1-20140114_ib_14_1_psd.1 will
complete installation upon reboot
[pid 13458]

{local:member0-re1}
user@host>
*** FINAL System shutdown message from user@host ***

System going down IMMEDIATELY

Connection closed by foreign host.

```

request system software nonstop-upgrade

Syntax	<pre>request system software nonstop-upgrade (<i>package-name</i> set [<i>package-name</i> <i>package-name</i>]) <force-host> <no-copy> <no-old-master-upgrade> <reboot > <unlink></pre>
Release Information	<p>Command introduced in Junos OS Release 10.4 for EX Series switches.</p> <p>Option set [<i>package-name package-name</i>] added in Junos OS Release 12.1 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2X50-D20 for the QFX Series.</p> <p>Command introduced in Junos OS Release 15.1X53-D55 for EX3400 switches.</p> <p>Command introduced in Junos OS Release 19.3R1 for EX4650 and QFX5120 switches.</p>
Description	<p>Perform a nonstop software upgrade (NSSU) on a switch with redundant Routing Engines or on a Virtual Chassis or Virtual Chassis Fabric (VCF). The behavior of this command depends on the type of switch, Virtual Chassis, or VCF where you run it, as follows:</p> <ul style="list-style-type: none"> When you run this command on any of the following Virtual Chassis or VCF configurations, NSSU upgrades all members of the Virtual Chassis: <ul style="list-style-type: none"> EX3300, EX3400, EX4200, EX4300, EX4500, EX4550, EX4600, or EX4650 Virtual Chassis Mixed Virtual Chassis composed of any combination of EX4200, EX4500, and EX4550 switches, or EX4300 and EX4600 switches QFX3500 and QFX3600 Virtual Chassis QFX5100 Virtual Chassis QFX5120-48Y Virtual Chassis Fixed configuration of switches in a VCF (QFX3500/QFX3600 and QFX5100 switches) Mixed VCF composed of any combination of QFX3500/QFX3600, QFX5100, and EX4300 switches <p>The original Virtual Chassis or VCF backup becomes the master. The new master automatically upgrades and reboots the original master, which then rejoins the Virtual Chassis or VCF as the backup.</p> When you run this command on an EX6200 or EX8200 switch, NSSU upgrades both the backup and master Routing Engines. The original backup Routing Engine becomes the new master at the end of the upgrade. <ul style="list-style-type: none"> On an EX6200 switch, NSSU automatically reboots the original master Routing Engine.

- On an EX8200 switch, NSSU does not automatically reboot the original master Routing Engine unless you specify the **reboot** option.
- When you run this command on an EX8200 Virtual Chassis, NSSU upgrades all master and backup Routing Engines in the Virtual Chassis, including the external Routing Engines. The original backup Routing Engines become the new master Routing Engines. NSSU does not automatically reboot the original master Routing Engines unless you specify the **reboot** option.

This command has the following requirements:

- All Virtual Chassis members, VCF members, and all Routing Engines must be running the same Junos OS release.
- Graceful Routing Engine switchover (GRES) must be enabled.
- Nonstop active routing (NSR) must be enabled.



NOTE: Although not required, we recommend you enable nonstop bridging (NSB). NSB ensures that all NSB-supported Layer 2 protocols operate seamlessly during the Routing Engine switchover during NSSU. See [“Configuring Nonstop Bridging on EX Series Switches \(CLI Procedure\)” on page 195](#).

- You must run the command from the master Routing Engine on a standalone switch or from the master on a Virtual Chassis.
- For minimal traffic disruption, you must define link aggregation groups (LAGs) such that the member links reside on:
 - Different Virtual Chassis members for EX3300, EX3400, EX4200, EX4300, EX4500, EX4550, EX4600, EX4650, QFX3500, QFX3600, QFX5100, and QFX5120 Virtual Chassis, mixed Virtual Chassis, or VCF.
 - Different line cards for EX6200 and EX8200 switches and EX8200 Virtual Chassis.
- For EX3300, EX3400, EX4200, EX4300, EX4500, EX4550, EX4600, EX4650, QFX3500, QFX3600, QFX5100, and QFX5120 Virtual Chassis, and mixed Virtual Chassis:
 - The Virtual Chassis members must be connected in a ring topology. A ring topology prevents the Virtual Chassis from splitting during an NSSU.
 - The Virtual Chassis master and backup must be adjacent to each other in the ring topology. With adjacent placement, the master and backup are always in sync while the switches in line-card roles are rebooting.
 - The Virtual Chassis must be preprovisioned so the line-card role is explicitly assigned to member switches acting in a line-card role. During an NSSU, the master and backup member switches must maintain their Routing Engine roles (although

mastership switches to the backup), and the remaining switches must maintain their line-card roles.

- You must configure **no-split-detection** in a two-member Virtual Chassis so the Virtual Chassis doesn't split during NSSU.
- For Virtual Chassis Fabric:
 - You can only have two members preprovisioned in the Routing Engine role. If more than two Routing Engines are configured, NSSU issues a warning message and the NSSU process stops.
 - The VCF members should be connected in a spine and leaf topology. A spine and leaf topology prevents the VCF from splitting during NSSU. Each leaf device must be connected to both spine devices.
 - The VCF must be preprovisioned so that the line-card role has been explicitly assigned to member switches acting in a line-card role, and likewise the Routing Engine role has been explicitly assigned to the member switches acting in a Routing Engine role. During an NSSU, the master and backup member switches must maintain their Routing Engine roles (although mastership switches to the backup), and the remaining switches must maintain their line-card roles.
 - You must configure **no-split-detection** in a two-member VCF so the VCF does not split during NSSU.

Options *package-name*—Location of the software package or bundle to be installed. For example:

- */var/tmp/package-name*—For a software package or bundle installed from a local directory on the switch.
- *protocol://hostname/pathname/package-name*—For a software package or bundle downloaded and installed from a remote location. Replace *protocol* with one of the following:
 - **ftp**—File Transfer Protocol.
Use **ftp://hostname/pathname/package-name**.
To specify authentication credentials, use **ftp://<username>:<password>@hostname/pathname/package-name**.
To have the system prompt you for the password, specify **prompt** in place of the password.
The command displays an error message if a password is required and you do not specify the password or **prompt**.
 - **http**—Hypertext Transfer Protocol.
Use **http://hostname/pathname/package-name**.
To specify authentication credentials, use **http://<username>:<password>@hostname/pathname/package-name**.
The command prompts you for a password if one is required and you didn't include it.
 - **scp**—Secure copy (available only for Canada and U.S. version).
Use **scp://hostname/pathname/package-name**.

To specify authentication credentials, use
`scp://<username>:<password>@hostname/pathname/package-name`.



NOTE: The *pathname* in the protocol is the relative path to the user home directory on the remote system and not the root directory.

set [*package-name package-name*]—(Mixed Virtual Chassis only) Locations of the different installation packages required by the different types of member switches. These packages must be for the same Junos OS release. See this command's *package-name* option for information about how to specify the installation packages.

force-host—(Optional) Force adding the host software package or bundle (and ignore warnings) on EX4650, QFX5100, or QFX5120 devices.

no-copy—(Optional) Install a software package or bundle, but do not save copies of the package or bundle files.

no-old-master-upgrade—(Optional) (EX8200 switches only) Upgrade the backup Routing Engine only. After the upgrade completes, the original master Routing Engine becomes the backup Routing Engine and continues running the previous software version.

reboot—(Optional) (EX8200 switches and EX8200 Virtual Chassis only) When you include the **reboot** option, NSSU automatically reboots the original master (new backup) Routing Engine after being upgraded to the new software. When you omit the **reboot** option, you must manually reboot the original master (new backup) Routing Engine using the **request system reboot** command.



NOTE: If you do not use the **reboot** option on an EX8200 Virtual Chassis, you must establish a connection to the console port on the Switch Fabric and Routing Engine (SRE) module or Routing Engine (RE) module to manually reboot the backup Routing Engines.

unlink—(Optional) Remove the software package after a successful upgrade.

Required Privilege Level maintenance

Related Documentation

- [show chassis nonstop-upgrade on page 743](#)
- [Upgrading Software on an EX6200 or EX8200 Standalone Switch Using Nonstop Software Upgrade \(CLI Procedure\) on page 485](#)

- [Upgrading Software on an EX8200 Virtual Chassis Using Nonstop Software Upgrade \(CLI Procedure\) on page 650](#)
- [Upgrading Software on a Virtual Chassis and Mixed Virtual Chassis Using Nonstop Software Upgrade](#)
- [Upgrading Software on a Virtual Chassis Fabric Using Nonstop Software Upgrade](#)

List of Sample Output	request system software nonstop-upgrade (EX4200 Virtual Chassis) on page 733 request system software nonstop-upgrade (EX6200 Switch) on page 735 request system software nonstop-upgrade reboot (EX8200 Switch) on page 736 request system software nonstop-upgrade no-old-master-upgrade (EX8200 Switch) on page 736 request system software nonstop-upgrade reboot (EX8200 Virtual Chassis) on page 737
Output Fields	This command reports feedback on the status of the request. Some functions are shared between NSSU and the in-service software upgrade (ISSU) feature, so you might see what appear to be ISSU messages as well as NSSU messages in the output from this command.

Sample Output

[request system software nonstop-upgrade \(EX4200 Virtual Chassis\)](#)

```

user@switch> request system software nonstop-upgrade
/var/tmp/jinstall-ex-4200-12.1R5.5-domestic-signed.tgz
Chassis ISSU Check Done
ISSU: Validating Image
ISSU: Preparing Backup RE
Installing image on other FPC's along with the backup

Checking pending install on fpc1
Pushing bundle to fpc1
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
Completed install on fpc1

Checking pending install on fpc2
Pushing bundle to fpc2
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
Completed install on fpc2

Checking pending install on fpc3
Pushing bundle to fpc3
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
Completed install on fpc3

Checking pending install on fpc4
Pushing bundle to fpc4
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
Completed install on fpc4

```

```

Checking pending install on fpc5
Pushing bundle to fpc5
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
Completed install on fpc5

Checking pending install on fpc6
Pushing bundle to fpc6
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
Completed install on fpc6

Checking pending install on fpc7
Pushing bundle to fpc7
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
Completed install on fpc7
Backup upgrade done
Rebooting Backup RE

Rebooting fpc1
ISSU: Backup RE Prepare Done
Waiting for Backup RE reboot
GRES operational
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item      Status      Reason
  FPC 0      Online
  FPC 1      Online
  FPC 2      Online (ISSU)
  FPC 3      Online (ISSU)
  FPC 4      Online (ISSU)
  FPC 5      Online (ISSU)
  FPC 6      Online (ISSU)
  FPC 7      Online (ISSU)
Going to install image on master
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
relinquish mastership
ISSU: IDLE

*** FINAL System shutdown message from root@switch ***

System going down IMMEDIATELY

Shutdown NOW!
[pid 9336]

```

request system software nonstop-upgrade (EX6200 Switch)

```
{master}
user@switch> request system software nonstop-upgrade
/var/tmp/jinstall-ex-6200-12.2R5.5-domestic-signed.tgz
Chassis ISSU Check Done
ISSU: Validating Image
ISSU: Preparing Backup RE
Pushing bundle to re0
NOTICE: Validating configuration against
jinstall-ex-6200-12.2R5.5-domestic-signed.tgz.
NOTICE: Use the 'no-validate' option to skip this if desired.
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
Backup upgrade done
Rebooting Backup RE

Rebooting re0
ISSU: Backup RE Prepare Done
Waiting for Backup RE reboot
GRES operational
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item           Status           Reason
  FPC 0          Online (ISSU)
  FPC 1          Online (ISSU)
  FPC 2          Online (ISSU)
  FPC 3          Online (ISSU)
  FPC 4          Online
  FPC 5          Online
  FPC 6          Online (ISSU)
  FPC 7          Online (ISSU)
  FPC 8          Online (ISSU)
  FPC 9          Online (ISSU)
Going to install image on master
NOTICE: Validating configuration against
jinstall-ex-6200-12.2R5.5-domestic-signed.tgz.
NOTICE: Use the 'no-validate' option to skip this if desired.
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
relinquish mastership
ISSU: IDLE
Trying to relinquish mastership before rebooting...
Resolving mastership...
Complete. The other routing engine becomes the master.

*** FINAL System shutdown message from user@switch ***

System going down IMMEDIATELY
```

request system software nonstop-upgrade reboot (EX8200 Switch)

```
{master}
user@switch> request system software nonstop-upgrade reboot
/var/tmp/install-ex-8200-10.4R1.5-domestic-signed.tgz
Chassis ISSU Check Done
ISSU: Validating Image
ISSU: Preparing Backup RE
Pushing bundle to re1
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
Backup upgrade done
Rebooting Backup RE

Rebooting re1
ISSU: Backup RE Prepare Done
Waiting for Backup RE reboot
GRES operational
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item          Status          Reason
  FPC 0         Online (ISSU)
  FPC 2         Offline
  FPC 3         Online (ISSU)
                Offlined by CLI command
Resolving mastership...
Complete. The other routing engine becomes the master.
ISSU: RE switchover Done
ISSU: Upgrading Old Master RE
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
ISSU: Old Master Upgrade Done
ISSU: IDLE
Shutdown NOW!
[pid 2635]

*** FINAL System shutdown message from user@switch ***
System going down IMMEDIATELY
```

request system software nonstop-upgrade no-old-master-upgrade (EX8200 Switch)

```
{master}
user@switch> request system software nonstop-upgrade no-old-master-upgrade
/var/tmp/install-ex-8200-10.4R1.5-domestic-signed.tgz
Chassis ISSU Check Done
ISSU: Validating Image
ISSU: Preparing Backup RE
Pushing bundle to re1
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
Backup upgrade done
Rebooting Backup RE
```

```

Rebooting re1
ISSU: Backup RE Prepare Done
Waiting for Backup RE reboot
GRES operational
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item          Status          Reason
  FPC 0         Online (ISSU)
  FPC 1         Online (ISSU)
  FPC 2         Online (ISSU)
  FPC 3         Offline         Offlined by CLI command
  FPC 4         Online (ISSU)
  FPC 5         Online (ISSU)
  FPC 6         Online (ISSU)
  FPC 7         Online (ISSU)
Resolving mastership...
Complete. The other routing engine becomes the master.
ISSU: RE switchover Done
Skipping Old Master Upgrade
ISSU: IDLE

```

request system software nonstop-upgrade reboot (EX8200 Virtual Chassis)

```

{master:9}
user@external-routing-engine> request system software nonstop-upgrade reboot
/var/tmp/jinstall-ex-xre200-11.1-20101130.0-domestic-signed.tgz
Chassis ISSU Check Done
ISSU: Validating Image
ISSU: Preparing LCC Backup REs
ISSU: Preparing Backup RE
Pushing bundle /var/tmp/jinstall-ex-xre200-11.1-20101130.0-domestic-signed.tgz
to member8
-----
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
VC Backup upgrade done
Rebooting VC Backup RE

Rebooting member8
ISSU: Backup RE Prepare Done
Waiting for VC Backup RE reboot
Pushing bundle to member0-backup
Pushing bundle to member1-backup
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately

Rebooting member0-backup
Rebooting LCC [member0-backup]

Rebooting member1-backup

```

```

Rebooting LCC [member1-backup]
ISSU: LCC Backup REs Prepare Done
GRES operational
Initiating Chassis Nonstop-Software-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking Nonstop-Upgrade status
member0:
-----
  Item      Status      Reason
  FPC 0     Online (ISSU)
  FPC 1     Online (ISSU)
  FPC 2     Online (ISSU)
  FPC 5     Online (ISSU)

member1:
-----
  Item      Status      Reason
  FPC 0     Online (ISSU)
  FPC 1     Offline      Offlined due to config
  FPC 2     Online (ISSU)
  FPC 3     Online (ISSU)
  FPC 4     Online (ISSU)
  FPC 5     Online (ISSU)
  FPC 7     Online (ISSU)

member0:
-----
  Item      Status      Reason
  FPC 0     Online (ISSU)
  FPC 1     Online (ISSU)
  FPC 2     Online (ISSU)
  FPC 5     Online (ISSU)

member1:
-----
  Item      Status      Reason
  FPC 0     Online (ISSU)
  FPC 1     Offline      Offlined due to config
  FPC 2     Online (ISSU)
  FPC 3     Online (ISSU)
  FPC 4     Online (ISSU)
  FPC 5     Online (ISSU)
  FPC 7     Online (ISSU)
ISSU: Upgrading Old Master RE
Pushing bundle /var/tmp/incoming-package-8200.tgz to member0-master
Pushing bundle /var/tmp/incoming-package-8200.tgz to member1-master

ISSU: RE switchover Done
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
Rebooting ...
shutdown: [pid 2188]
Shutdown NOW!
ISSU: Old Master Upgrade Done
ISSU: IDLE

```


Shutdown NOW!

*** FINAL System shutdown message from root@ ***
System going down IMMEDIATELY

request system software validate in-service-upgrade

Syntax	<code>request system software validate in-service-upgrade <i>package-name</i></code>
Release Information	<p>Command introduced in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 13.2 for PTX5000 routers.</p> <p>Command introduced in Junos OS Release 14.2 for EX Series switches.</p>
Description	<p>Perform a compatibility check to ensure that the software and hardware components and the configuration on the device support unified ISSU. The request system software validate in-service-upgrade command enables you to detect any compatibility issues before actually issuing the request system software in-service-upgrade command to initiate unified ISSU.</p>
Options	<p><i>package-name</i>—Location from which the software package or bundle is to be installed. For example:</p> <ul style="list-style-type: none">• <i>/var/tmp/package-name</i>—For a software package or bundle that is being installed from a local directory on the router.• <i>protocol://hostname/pathname/package-name</i>—For a software package or bundle that is to be downloaded and installed from a remote location. Replace <i>protocol</i> with one of the following:<ul style="list-style-type: none">• ftp—File Transfer Protocol• http—Hypertext Transfer Protocol• scp—Secure copy (available only for Canada and U.S. version)
Additional Information	<p>Unified ISSU is not supported on every platform. For a list of supported platforms, see “Unified ISSU System Requirements” on page 401.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• request system software validate• request system software in-service-upgrade on page 699• request system software abort• show chassis in-service-upgrade• Getting Started with Unified In-Service Software Upgrade on page 387• Example: Performing a Unified ISSU on page 424

List of Sample Output [request system software validate in-service-upgrade on page 741](#)

Output Fields When you enter this command, Junos OS displays the status of your request.

Sample Output

[request system software validate in-service-upgrade](#)

```
{master}

user@host> request system software validate in-service-upgrade
/var/tmp/jinstall-9.0-20080114.2-domestic-signed.tgz reboot

Checking compatibility with configuration
Initializing...
Using jbase-9.5-20090127.0
Verified manifest signed by PackageProduction_9_5_0
Using /var/tmp/jinstall-9.6-daily-domestic-signed.tgz
Verified jinstall-9.6-20090706.0-domestic.tgz signed by PackageProduction_9_6_0
Using jinstall-9.6-20090706.0-domestic.tgz
Using jbundle-9.6-20090706.0-domestic.tgz
Checking jbundle requirements on /
Using jbase-9.6-20090706.0.tgz
Verified manifest signed by PackageProduction_9_6_0
Using jkernel-9.6-20090706.0.tgz
Verified manifest signed by PackageProduction_9_6_0
Using jcrypto-9.6-20090706.0.tgz
Verified manifest signed by PackageProduction_9_6_0
Using jpfe-9.6-20090706.0.tgz
Using jdocs-9.6-20090706.0.tgz
Verified manifest signed by PackageProduction_9_6_0
Using jroute-9.6-20090706.0.tgz
Verified manifest signed by PackageProduction_9_6_0
Using jservices-9.6-20090706.0.tgz
[: /var/validate/chroot/tmp/jservices/packages/jservices-voice-9.6-20090706.0.tgz:
unexpected operator
Auto-deleting old jservices-voice ...
Removing /opt/sdk/jservices-voice ...
Removing jservices-voice-bsg-9.5-20090127.0.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-voice ...
Verified jservices-voice-bsg-9.6-20090706.0.tgz signed by PackageProduction_9_6_0
Creating /var/sw/pkg ...
Creating /opt/sdk/jservices-voice ...
Storing jservices-voice-bsg-9.6-20090706.0.tgz in /var/sw/pkg ...
Link: /opt/sdk/jservices-voice/jservices-voice-bsg ->
/var/sw/pkg/jservices-voice-bsg-9.6-20090706.0.tgz...
Installing new jservices-bgf ...
Verified jservices-bgf-pic-9.6-20090706.0.tgz signed by PackageProduction_9_6_0
Creating /opt/sdk/jservices-bgf ...
Storing jservices-bgf-pic-9.6-20090706.0.tgz in /var/sw/pkg ...
Link: /opt/sdk/jservices-bgf/jservices-bgf-pic ->
/var/sw/pkg/jservices-bgf-pic-9.6-20090706.0.tgz...
Auto-deleting old jservices-aac1 ...
Removing /opt/sdk/jservices-aac1 ...
Removing jservices-aac1-pic-9.5-20090127.0.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-aac1 ...
Verified jservices-aac1-pic-9.6-20090706.0.tgz signed by PackageProduction_9_6_0
```

```
Creating /opt/sdk/jservices-aac1 ...
Storing jservices-aac1-pic-9.6-20090706.0.tgz in /var/sw/pkg ...
Link: /opt/sdk/jservices-aac1/jservices-aac1-pic ->
/var/sw/pkg/jservices-aac1-pic-9.6-20090706.0.tgz...
Auto-deleting old jservices-llpdf ...
Removing /opt/sdk/jservices-llpdf ...
Removing jservices-llpdf-pic-9.5-20090127.0.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-llpdf ...
Verified jservices-llpdf-pic-9.6-20090706.0.tgz signed by PackageProduction_9_6_0
Creating /opt/sdk/jservices-llpdf ...
Storing jservices-llpdf-pic-9.6-20090706.0.tgz in /var/sw/pkg ...
Link: /opt/sdk/jservices-llpdf/jservices-llpdf-pic ->
/var/sw/pkg/jservices-llpdf-pic-9.6-20090706.0.tgz...
Auto-deleting old jservices-sfw ...
Removing /opt/sdk/jservices-sfw ...
Removing jservices-sfw-pic-9.5-20090127.0.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-sfw ...
Verified jservices-sfw-pic-9.6-20090706.0.tgz signed by PackageProduction_9_6_0
Creating /opt/sdk/jservices-sfw ...
Storing jservices-sfw-pic-9.6-20090706.0.tgz in /var/sw/pkg ...
Link: /opt/sdk/jservices-sfw/jservices-sfw-pic ->
/var/sw/pkg/jservices-sfw-pic-9.6-20090706.0.tgz...
Auto-deleting old jservices-appid ...
Removing /opt/sdk/jservices-appid ...
Removing jservices-appid-pic-9.5-20090127.0.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-appid ...
Verified jservices-appid-pic-9.6-20090706.0.tgz signed by PackageProduction_9_6_0
Creating /opt/sdk/jservices-appid ...
Storing jservices-appid-pic-9.6-20090706.0.tgz in /var/sw/pkg ...
Link: /opt/sdk/jservices-appid/jservices-appid-pic ->
/var/sw/pkg/jservices-appid-pic-9.6-20090706.0.tgz...
Auto-deleting old jservices-idp ...
Removing /opt/sdk/jservices-idp ...
Removing jservices-idp-pic-9.5-20090127.0.tgz from /var/sw/pkg ...
Notifying mspd ...
Installing new jservices-idp ...
Verified jservices-idp-pic-9.6-20090706.0.tgz signed by PackageProduction_9_6_0
Creating /opt/sdk/jservices-idp ...
Storing jservices-idp-pic-9.6-20090706.0.tgz in /var/sw/pkg ...
Link: /opt/sdk/jservices-idp/jservices-idp-pic ->
/var/sw/pkg/jservices-idp-pic-9.6-20090706.0.tgz...
Hardware Database regeneration succeeded
Validating against /config/juniper.conf.gz
mgd: commit complete
Validation succeeded
PIC 7/0 will be offlined (In-Service-Upgrade not supported)
PIC 7/1 will be offlined (In-Service-Upgrade not supported)
PIC 4/2 will be offlined (In-Service-Upgrade not supported)
PIC 4/3 will be offlined (In-Service-Upgrade not supported)
```

show chassis nonstop-upgrade

Syntax `show chassis nonstop-upgrade`

Release Information Command introduced in Junos OS Release 10.4 for EX Series switches.
Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.

Description (EX6200 switches, EX8200 switches, EX8200 Virtual Chassis, QFX3500 and QFX3600 Virtual Chassis, and Virtual Chassis Fabric only) Display the status of the line cards or Virtual Chassis members in the linecard role after the most recent nonstop software upgrade (NSSU). This command must be issued on the master Routing Engine.

Required Privilege Level view

Related Documentation

- [request system software nonstop-upgrade on page 729](#)
- [Upgrading Software on an EX6200 or EX8200 Standalone Switch Using Nonstop Software Upgrade \(CLI Procedure\) on page 485](#)
- [Upgrading Software on a Virtual Chassis and Mixed Virtual Chassis Using Nonstop Software Upgrade](#)
- [Upgrading Software on a Virtual Chassis Fabric Using Nonstop Software Upgrade](#)
- [Upgrading Software on an EX8200 Virtual Chassis Using Nonstop Software Upgrade \(CLI Procedure\) on page 650](#)

List of Sample Output [show chassis nonstop-upgrade \(EX8200 Switch\) on page 744](#)
[show chassis nonstop-upgrade \(EX8200 Virtual Chassis\) on page 744](#)
[show chassis nonstop-upgrade \(Virtual Chassis Fabric\) on page 744](#)

Output Fields [Table 43 on page 743](#) lists the output fields for the **show chassis nonstop-upgrade** command. Output fields are listed in the approximate order in which they appear.

Table 43: show chassis nonstop-upgrade Output Fields

Field Name	Field Description
Item	Line card slot number.
Status	State of line card: <ul style="list-style-type: none"> • Error—Line card is in an error state. • Offline—Line card is powered down. • Online—Line card is online and running.
Reason	Reason for the state (if the line card is offline).

Sample Output

show chassis nonstop-upgrade (EX8200 Switch)

```
user@switch> show chassis nonstop-upgrade
```

Item	Status	Reason
FPC 0	Online	
FPC 1	Online	
FPC 2	Online	
FPC 3	Offline	Offlined by CLI command
FPC 4	Online	
FPC 5	Online	
FPC 6	Online	
FPC 7	Online	

show chassis nonstop-upgrade (EX8200 Virtual Chassis)

```
user@external-routing-engine> show chassis nonstop-upgrade  
member0:
```

Item	Status	Reason
FPC 0	Online	
FPC 1	Online	
FPC 2	Online	
FPC 5	Online	

```
member1:
```

Item	Status	Reason
FPC 0	Online	
FPC 1	Offline	Offlined due to config
FPC 2	Online	
FPC 3	Online	
FPC 4	Online	
FPC 5	Online	
FPC 7	Online	

show chassis nonstop-upgrade (Virtual Chassis Fabric)

Item	Status	Reason
FPC 0	Online	
FPC 1	Online	
FPC 2	Online	
FPC 3	Online	
FPC 4	Online	
FPC 5	Online	

show chassis nonstop-upgrade node-group

Syntax `show chassis nonstop-upgrade node-group node-group-name`

Release Information Command introduced in Junos OS Release 12.2 for the QFX Series.

Description Display the status of the Node group after the most recent nonstop software upgrade (NSSU).

Required Privilege Level view

Related Documentation

- *Performing a Nonstop Software Upgrade on the QFabric System*
- *request system software nonstop-upgrade*

List of Sample Output [show chassis nonstop-upgrade node-group on page 745](#)

Output Fields [Table 43 on page 743](#) lists the output fields for the **show chassis nonstop-upgrade node-group** command. Output fields are listed in the approximate order in which they appear.

Table 44: show chassis nonstop-upgrade node-group Output Fields

Field Name	Field Description
Item	Node device slot number.
Status	State of Node device: <ul style="list-style-type: none"> • Error—Node device is in an error state. • Offline—Node device is powered down. • Online—Node device is online and running.
Reason	Reason for the state (if the line card is offline).

Sample Output

show chassis nonstop-upgrade node-group

```
user@qfabric> show chassis nonstop-upgrade node-group NW-NG-0
```

Item	Status	Reason
P1550-C	Online	

show chassis power-budget-statistics

Syntax	<code>show chassis power-budget-statistics</code>
Release Information	Command introduced in Junos OS Release 10.2 for EX Series switches.
Description	Display the power budget of an EX Series switch.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Verifying Power Configuration and Use on page 658 • Configuring the Power Priority of Line Cards (CLI Procedure) on page 302 • Configuring Power Supply Redundancy (CLI Procedure) on page 303
List of Sample Output	show chassis power-budget-statistics (EX6200 Switch) on page 748 show chassis power-budget-statistics (EX8200 Switch) on page 749
Output Fields	Table 45 on page 746 lists the output fields for the show chassis power-budget-statistics command. Output fields are listed in the approximate order in which they appear.

Table 45: show chassis power-budget-statistics Output Fields

Field Name	Field Description
PSU <i>n</i> (<i>supply type</i>)	Capacity rating of the power supply and whether the power supply is currently operating (Online) or not (Offline). If a power supply is offline, the capacity is shown as 0 W.
Total Power supplied by all Online PSUs	Total number of watts supplied by all currently operating power supplies.
Power Redundancy Configuration	Configured power redundancy setting, either <i>N+1</i> or <i>N+N</i> .
Base power reserved	Total number of watts reserved for the switch.
Non-PoE power being consumed	The amount of power, in W, currently being consumed for PoE.

Table 45: show chassis power-budget-statistics Output Fields (continued)

Field Name	Field Description
Power Reserved for the Chassis	<p>Power reserved for the chassis:</p> <ul style="list-style-type: none"> For an EX6200 switch, 500 W. For an EX8208 switch: 1600 W in an <i>N</i>+1 configuration; 1200 W in an <i>N</i>+<i>N</i> configuration For an EX8216 switch: 2400 W in an <i>N</i>+1 configuration; 1800 W in an <i>N</i>+<i>N</i> configuration <p>The power reserved for the chassis includes the maximum power requirements for the fan tray and Switch Fabric and Routing Engine (SRE), Routing Engine (RE), and Switch Fabric (SF) modules in both base and redundant configurations.</p>
Fan Tray Statistics	<p>(EX6200 switch only) Information about the fan tray:</p> <ul style="list-style-type: none"> Base power—Power allocated to the fan tray in the power budget. This allocation is included in Power Reserved for the Chassis. Power Used—Actual power being used by the fan tray. This value is for informational purposes only: the power budget for the switch is based on allocated power (the theoretical maximum the fan tray might use) rather than used power.
FPC <i>n</i> (card type)	<p>Information about the line card installed in slot <i>n</i>. For EX6200 switches, information about the SRE modules in slot 4 and slot 5 is also shown.</p> <ul style="list-style-type: none"> Base power—For line cards without PoE ports, the total power allocated to the line card. For line cards with PoE ports, the power allocated to the line card before the PoE power budget is allocated. The base power includes 37 W of PoE power that is always allocated to line cards that support PoE. Power Used—(EX6200 switch only) The actual power being consumed by the line card or SRE module, including PoE power. This value is for informational purposes only: the power budget for the switch is based on allocated power (the theoretical maximum the line card might use) rather than used power. PoE power—For line cards with PoE ports, the PoE power budget allocated to the line card. This value includes the 37 W of PoE power that is always part of the base power allocation for line cards that support PoE. For line cards without PoE ports, the value is always 0 W. The power priority assigned to the line card slot.
Total (non-PoE) Power allocated	<p>Power budgeted for all the components in the switch, excluding the PoE power budget allocated to line cards. This value is equal to the power reserved for the chassis plus the base power allocations of all online line cards.</p>

Table 45: show chassis power-budget-statistics Output Fields (continued)

Field Name	Field Description
Total Power allocated for PoE	The total of the PoE power budgets allocated to the line cards in the switch. This figure includes the 37 W of PoE power always included in the base allocation for each line card that supports PoE.
Total PoE power consumed	The amount of power that has been consumed by PoE.
Total PoE power remaining	The amount of available power remaining that can be used for PoE.
Power Available (Redundant case)	Unused power available to the switch in the power budget, not including the power reserved for redundancy. If power is insufficient to meet the $N+1$ or $N+N$ redundancy requirements, this value is 0. PoE power allocations are not included in the calculation of this value.
Total Power Available	Unused power available to the switch in the power budget. This value is derived by subtracting all power allocations, including PoE power allocations, from the total power available on the switch (the Total Power supplied by all Online PSUs value).

Sample Output

show chassis power-budget-statistics (EX6200 Switch)

```

user@switch> show chassis power-budget-statistics
PSU 0      (EX6200-PWR-AC2500)      :    2500 W   Online
PSU 1      (EX6200-PWR-AC2500)      :    2500 W   Online
PSU 2      (EX6200-PWR-AC2500)      :    2500 W   Online
PSU 3      (EX6200-PWR-AC2500)      :    2500 W   Online
Total Power supplied by all Online PSUs :   10000 W
Power Redundancy Configuration         :      N+1
Power Reserved for the Chassis         :      500 W
Fan Tray Statistics                    Base power  Power Used
FTC 0                                     :    300 W    43.04 W
FPC Statistics                          Base power  Power Used  PoE power  Priority
FPC 1  (EX6200-48P)                    :    220 W    49.47 W    1440 W     1
FPC 2  (EX6200-48P)                    :    220 W    47.20 W     800 W     2
FPC 3  (EX6200-48P)                    :    220 W   1493.57 W    1440 W     0
FPC 4  (EX6200-SRE64-4XS)              :    100 W    51.38 W       0 W     0
FPC 5  (EX6200-SRE64-4XS)              :    100 W    50.28 W       0 W     0
FPC 6  (EX6200-48P)                    :    220 W    49.38 W     800 W     6
FPC 8  (EX6200-48P)                    :    220 W    61.41 W    1440 W     9
FPC 9  (EX6200-48T)                    :    150 W    12.49 W       0 W     9

Total (non-PoE) Power allocated        :    1750 W
Total Power allocated for PoE          :    5920 W
Power Available (Redundant case)       :    5750 W
Total Power Available                  :    2515 W

```

show chassis power-budget-statistics (EX8200 Switch)

```

user@switch> show chassis power-budget-statistics
PSU 0      (EX8200-AC2K)      :    2000 W  Online
PSU 1      (EX8200-AC2K)      :    2000 W  Online
PSU 2      (EX8200-AC2K)      :    2000 W  Online
PSU 3      (EX8200-AC2K)      :    2000 W  online
PSU 4      (EX8200-AC2K)      :    2000 W  Online
Total Power supplied by all Online PSUs :   10000 W
Power Redundancy Configuration          :    N+1
Power Reserved for the Chassis          :    2400 W
FPC Statistics                          Base power    PoE power    Priority

FPC 1      (EX8200-48T)          :    350 W          0 W      15
FPC 5      (EX8200-2XS-40P)      :    387 W          792 W      0
FPC 9      (EX8200-48PL)         :    267 W          915 W     15
FPC 10     (EX8200-2XS-40T)      :    350 W          0 W      1
FPC 12     (EX8200-48T)          :    350 W          0 W     15

Total (non-PoE) Power allocated         :    4104 W
Total Power allocated for PoE           :    1707 W
Power Available (Redundant case)         :    3896 W
Total Power Available                    :    4263 W

```

show chassis redundant-power-system

Syntax	<code>show chassis redundant-power-system</code>
Release Information	Command introduced in Junos OS Release 12.1 for EX Series switches.
Description	Display information about the Redundant Power Systems (RPS) connected to the switch.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> Determining and Setting Priority for Switches Connected to an EX Series RPS on page 310
List of Sample Output	show chassis redundant-power-system (Standalone Switch) on page 750 show chassis redundant-power-system (Virtual Chassis member) on page 751
Output Fields	Table 46 on page 750 lists the output fields for the <code>show chassis redundant-power-system</code> command. Output fields are listed in the approximate order in which they appear.

Table 46: show chassis redundant-power-system Output Fields

Field Name	Field Description	Level of Output
Member	Member number of the switch connected to the RPS—For a switch that has never been configured in a Virtual Chassis, the value is always zero. For a Virtual Chassis member, the range is zero through the maximum number of members in the Virtual Chassis.	All levels
Status	Status of the RPS: <ul style="list-style-type: none"> ARMED—The switch is ready to get backup power from the RPS if power supply fails on the switch. OFF—The switch has zero and is not configured to receive backup power from the RPS. BACKED-UP—The switch is receiving power backup from the RPS. OVER-SUBSCRIBED—The switch cannot receive backup power from the RPS even if you set the . 	All levels
RPS	Serial number of the RPS.	
Port	Number of the switch connector on the RPS that is connected to a switch.	All levels

Sample Output

show chassis redundant-power-system (Standalone Switch)

```
user@switch> show chassis redundant-power-system
```

Member	Status	RPS	Port
0	Armed	CG0209121807	0

show chassis redundant-power-system (Virtual Chassis member)

```
user@switch> show chassis redundant-power-system
```

Member	Status	RPS	Port
0	Armed	CG0209121814	5
2	Armed	CG0209121815	4

show redundant-power-system led

Syntax `show redundant-power-system led`

Release Information Command introduced in Junos OS Release 12.1 for EX Series switches.

Description Display information about fan status, Redundant Power System (RPS) status, and the switch connectors as displayed by the corresponding LEDs on the RPS.

Required Privilege Level view

Related Documentation

- [LEDs on an EX Series Redundant Power System](#)

List of Sample Output [show redundant-power-system led \(Standalone Switch\) on page 753](#)
[show redundant-power-system led \(EX3300 Virtual Chassis\) on page 753](#)

Output Fields [Table 47 on page 752](#) lists the output fields for the `show redundant-power-system led` command. Output fields are listed in the approximate order in which they appear.

Table 47: show redundant-power-system led Output Fields

Field Name	Field Description	Level of Output
RPS	The serial number of the RPS.	
RPS Fan	Status of the RPS power supply fans as displayed by the LED: <ul style="list-style-type: none"> • Green—All RPS power supply fans are operating fine. • Amber—A fan has failed in at least one RPS power supply. 	All levels
RPS System Status	Status of the RPS system as displayed by the LED: <ul style="list-style-type: none"> • Green—The RPS is active. • Blinking green—The RPS is booting. • Amber—An RPS power supply has failed. • Off—The RPS is off. 	All levels
RPS Port LED Status	Status of the RPS switch connectors as displayed by the LEDs. These LEDs indicate whether the redundant power source is being used. <ul style="list-style-type: none"> • Green—The RPS connector is enabled and connected to a switch but the RPS is not actively backing up the switch. • Blinking green—The RPS is backing up the switch connected to the port. • Off—The RPS connector is not connected to a switch. • Amber—The RPS is oversubscribed and the backup power to the switch has failed. 	All levels
Port	Number of one of the six switch connectors on the RPS.	All levels

Table 47: show redundant-power-system led Output Fields (continued)

Field Name	Field Description	Level of Output
Status	Status of each switch connector on the RPS.	All levels

Sample Output

show redundant-power-system led (Standalone Switch)

```
user@switch> show redundant-power-system led
```

```
Gathering requested information.
RPS-CG0209121807
  RPS Fan: GREEN
  RPS System Status: GREEN
RPS Port LED Status
Port  Status
0    GREEN
1    OFF
2    OFF
3    OFF
4    OFF
5    OFF
```

show redundant-power-system led (EX3300 Virtual Chassis)

```
user@switch> show redundant-power-system led
```

```
Gathering requested information.
RPS-CG0209121814
  RPS Fan: GREEN
  RPS System Status: GREEN
RPS Port LED Status
Port  Status
0    OFF
1    OFF
2    OFF
3    OFF
4    OFF
5    GREEN
RPS-CG0209121815
  RPS Fan: GREEN
  RPS System Status: GREEN
RPS Port LED Status
Port  Status
0    OFF
1    OFF
2    OFF
3    OFF
4    GREEN
5    OFF
```

show redundant-power-system multi-backup

Syntax `show redundant-power-system multi-backup`

`show redundant-power-system multi-backup member member-number`

Release Information Command introduced in Junos OS Release 12.1 for EX Series switches.

Description Display the current status of the Redundant Power System's (RPS's) ability to back up two switches per power supply when enough power to support Power over Ethernet (PoE) is not needed. This ability is referred to as the RPS's multi-backup ability.

Required Privilege Level view

Related Documentation • [request redundant-power-system multi-backup on page 698](#)

List of Sample Output [show redundant-power-system multi-backup on page 754](#)

Sample Output

show redundant-power-system multi-backup

```
User@switch> show redundant-power-system multi-backup
```

```
Requesting information from redundant-power-system..      Multi-Backup: enabled
```

show redundant-power-system network

Syntax show redundant-power-system network

Release Information Command introduced in Junos OS Release 12.1 for EX Series switches.

Description Display the Redundant Power Supply (RPS) IP address, netmask address, and gateway address required for firmware backup.

Required Privilege Level view

Related Documentation • *Upgrading Firmware on an EX Series Redundant Power System*

List of Sample Output [show redundant-power-system network on page 755](#)

Sample Output

show redundant-power-system network

```
user@switch> show redundant-power-system network
Requesting information from redundant-power-system..
IP Address:  10.93.2.38
Netmask: 255.255.254.0
Gateway:  10.93.3.254
```

show redundant-power-system power-supply

Syntax	show redundant-power-system power-supply
Release Information	Command introduced in Junos OS Release 12.1 for EX Series switches.
Description	Display information about the power supplies installed in the Redundant Power System (RPS). After installing a power supply, we recommend that you use this command to be sure that the power supply installed correctly.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • <i>Installing a Power Supply in the EX Series Redundant Power System</i>
List of Sample Output	show redundant-power-system power-supply (Standalone Switch) on page 756 show redundant-power-system power-supply (EX3300 Virtual Chassis) on page 757
Output Fields	Table 48 on page 756 lists the output fields for the show redundant-power-system power-supply command. Output fields are listed in the approximate order in which they appear.

Table 48: show redundant-power-system power-supply Output Fields

Field Name	Field Description	Level of Output
RPS	Serial number of the RPS.	All levels
PSU Slot	Number of the power supply slot. Slots are numbered 1 through 3.	All levels
Status	Status of the power supply slots: <ul style="list-style-type: none"> • Present—The slot contains an RPS power supply. • Empty—The slot is empty. 	All levels
Description	Description of the RPS power supply installed in the slot.	All levels

Sample Output

show redundant-power-system power-supply (Standalone Switch)

```
user@switch> show redundant-power-system power-supply
```

```
Gathering requested information.
RPS-CG0209121807
PSU Slot Status    Description
      1 Online    930W AC
```

```
2 offline    ---
3 Online    930W AC
```

show redundant-power-system power-supply (EX3300 Virtual Chassis)

```
user@switch> show redundant-power-system power-supply
```

```
Gathering requested information.
```

```
RPS-CG0209121814
```

PSU Slot	Status	Description
1	Online	930W AC
2	offline	---
3	Online	930W AC

```
RPS-CG0209121815
```

PSU Slot	Status	Description
1	Online	930W AC
2	Online	930W AC
3	Online	930W AC

show redundant-power-system status

Syntax `show redundant-power-system status`

Release Information Command introduced in Junos OS Release 12.1 for EX Series switches.

Description Display the status information for the switch connectors on the Redundant Power System (RPS).

Required Privilege Level view

Related Documentation

- [Determining and Setting Priority for Switches Connected to an EX Series RPS on page 310](#)
- [Installing a Power Supply in the EX Series Redundant Power System](#)

List of Sample Output [show redundant-power-system status \(Standalone Switch\) on page 759](#)
[show redundant-power-system status \(EX3300 Virtual Chassis\) on page 759](#)

Output Fields [Table 49 on page 758](#) lists the output fields for the **show redundant-power-system status** command. Output fields are listed in the approximate order in which they appear.

Table 49: show redundant-power-system status Output Fields

Field Name	Field Description	Level of Output
RPS	Serial number of the RPS.	All levels
Port	Number of the switch connector.	All levels
Status	Status of the switch connector: <ul style="list-style-type: none"> • ARMED—The switch is ready to get backup power from RPS if power supply fails on the switch. • OFF—The switch has zero and is not configured to receive backup power from RPS. • BACKED-UP—The switch is receiving power backup from RPS. • OVER-SUBSCRIBED—The switch cannot receive backup power from RPS even if you set the . 	All levels
Priority	Priority value of the switch connector.	All levels
Power-Requested	Power requested by the switch on the corresponding switch connector.	All levels

Sample Output

show redundant-power-system status (Standalone Switch)

```
user@switch> show redundant-power-system status
```

```
Gathering requested information.
RPS-CG0209121807
Port Status      Power-requested
0  Armed          3      930W
1  Off             1      ---
2  Off             1      ---
3  Off             1      ---
4  Off             1      ---
5  Off             1      ---
```

show redundant-power-system status (EX3300 Virtual Chassis)

```
user@switch> show redundant-power-system status
```

```
Gathering requested information.
RPS-CG0209121814
Port Status      Power-requested
0  OFF             1      ---
1  OFF             1      ---
2  OFF             1      ---
3  OFF             1      ---
4  OFF             1      ---
5  Armed           5      930W
RPS-CG0209121815
Port Status      Power-requested
0  OFF             1      ---
1  OFF             1      ---
2  OFF             1      ---
3  OFF             1      ---
4  Armed           4      930W
5  OFF             1      ---
```

show redundant-power-system upgrade

Syntax	show redundant-power-system upgrade
Release Information	Command introduced in Junos OS Release 12.1 for EX Series switches.
Description	Display RPS firmware upgrade status (pass or fail), previous RPS firmware version, and current RPS firmware version.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• request redundant-power-system multi-backup on page 698
List of Sample Output	show redundant-power-system upgrade on page 760
Output Fields	Table 50 on page 760 lists the output fields for the show redundant-power-system status command. Output fields are listed in the approximate order in which they appear.

Table 50: show redundant-power-system upgrade Output Fields

Field Name	Field Description	Level of Output
Firmware Upgrade Status	Indicates whether the upgrade passed or failed	All levels
Previous Firmware Version	Firmware version before the upgrade	All levels
Current Firmware Version	Firmware version after the upgrade	

Sample Output

show redundant-power-system upgrade

```
user@switch> show redundant-power-system upgrade
Requesting information from redundant-power-system..
Firmware Upgrade Status:  Pass
Previous Firmware Version: 1.0
Current Firmware Version: 1.0
```

show redundant-power-system version

Syntax	show redundant-power-system version
Release Information	Command introduced in Junos OS Release 12.1 for EX Series switches.
Description	Display version information about the Redundant Power System (RPS).
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • <i>Installing a Power Supply in the EX Series Redundant Power System</i> • <i>Packing an EX Series Redundant Power System or Redundant Power System Components for Shipping</i>
List of Sample Output	show redundant-power-system version (Standalone Switch) on page 761 show redundant-power-system version (EX3300 Virtual Chassis) on page 761
Output Fields	Table 51 on page 761 lists the output fields for the show redundant-power-system version command. Output fields are listed in the approximate order in which they appear.

Table 51: show redundant-power-system version Output Fields

Field Name	Field Description	Level of Output
RPS	Serial number of the RPS.	All levels
Model	Model name of the RPS.	All levels
RPS Firmware Version	Version number of the firmware installed on the RPS.	All levels
RPS U-Boot Version	Version of the bootup software installed on the RPS.	All levels

Sample Output

show redundant-power-system version (Standalone Switch)

```
user@switch> show redundant-power-system version
```

```
RPS-CG0209121807
Model: EX-PWR_RPS200
RPS Firmware Version [1.0]
RPS U-Boot Version [1.1.6]
```

show redundant-power-system version (EX3300 Virtual Chassis)

```
user@switch> show redundant-power-system version
```

```
RPS-CG0209121814
  Model: EX-PWR_RPS200
  RPS Firmware Version [1.0]
  RPS U-Boot Version [1.1.6]
RPS-CG0209121815
  Model: EX-PWR_RPS200
  RPS Firmware Version [1.0]
  RPS U-Boot Version [1.1.6]
```


show chassis ssb

Syntax	<code>show chassis ssb</code> <code><slot></code>
Release Information	Command introduced before Junos OS Release 7.4.
Description	(M20 routers only) Display status information about the System and Switch Board (SSB).
Options	<p>none—Display information about all SSBs.</p> <p>slot—(Optional) Display information about the SSB in the specified slot. Replace slot with 0 or 1.</p>
Required Privilege Level	view
Related Documentation	
List of Sample Output	show chassis ssb on page 764
Output Fields	Table 52 on page 763 lists the output fields for the <code>show chassis ssb</code> command. Output fields are listed in the approximate order in which they appear.

Table 52: show chassis ssb Output Fields

Field Name	Field Description
Failover	Number of times mastership has changed.
Slot	SSB slot number.
State	<p>Current state of the SSB in this slot. State can be any one of the following:</p> <ul style="list-style-type: none"> • Master—SSB is online, operating as master. • Backup—SSB running as backup. • Empty—No SSB is present.
Temperature	Temperature of the air passing by the SSB, in degrees Celsius.
CPU utilization	Total percentage of the CPU being used by the SSB's processor.
Interrupt utilization	Of the total CPU being used by the SSB's processor, the percentage being used for interrupts.
Heap utilization	Percentage of heap space being used by the SSB's processor.

Table 52: show chassis ssb Output Fields (continued)

Field Name	Field Description
Buffer utilization	Percentage of buffer space being used by the SSB's processor.
DRAM	Total DRAM available to the SSB's processor.
Start time	Time when the SSB started running.
Uptime	How long the SSB has been up and running.

Sample Output

show chassis ssb

```
user@host> show chassis ssb
```

```
SSB status:
  Failover:                0 time
  Slot 0:
    State:                  Master
    Temperature:            33 Centigrade
    CPU utilization:         0 percent
    Interrupt utilization:   0 percent
    Heap utilization:        0 percent
    Buffer utilization:       6 percent
    DRAM:                   64 Mbytes
    Start time:              1999-01-15 22:05:36 UTC
    Uptime:                  21 hours, 21 minutes, 22 seconds
...
```

show nonstop-routing

Syntax	show nonstop-routing
Release Information	Command introduced in Junos OS Release 13.3.
Description	Display the status of nonstop active routing that includes the automerger statistics and state.
Required Privilege Level	View
Related Documentation	<ul style="list-style-type: none"> • nonstop-routing on page 546
List of Sample Output	show nonstop-routing (MX Series Router) on page 766 show nonstop-routing (MX Series Router) on page 767
Output Fields	Table 53 on page 765 describes the output fields for the show nonstop-routing command. Output fields are listed in the approximate order in which they appear.

Table 53: show nonstop-routing Output Fields

Field Name	Field Description
Nonstop Routing	State of NSR.
Precision Timers state	<p>State of precision timer feature in the kernel.</p> <ul style="list-style-type: none"> • Enabled—By default, autokeepalive precision timers are enabled on the kernel after switchover. • Disabled—Autokeepalive precision timers are disabled. • Inactive—Precision timer is inactive if it is disabled. • Ready—Kernel precision timer is ready but is never activated. • InProcess—Kernel precision timer is operational and is generating keepalives on behalf of the RPD after switchover. The / count indicates the number of sessions being serviced against the total sessions. • Completed—Kernel has completed keepalive generation for all the sessions after switchover, and RPD has taken over all of them successfully. • Error—Error while retrieving the precision timer status of the kernel.
Precision Timers max period	Maximum period, in seconds, after the switchover from standby to master event for which the kernel autogenerates keepalives on behalf of BGP.

Table 53: show nonstop-routing Output Fields (continued)

Field Name	Field Description
Automerge	<p>Status of the automerge.</p> <ul style="list-style-type: none"> • Active—Automerge of sockets by the kernel after switchover is active. • Inactive—Automerge of sockets by the kernel after switchover is inactive.
Batching	<p>Status of Batching.</p> <ul style="list-style-type: none"> • Yes—Automerge of sockets by the kernel after a switchover. • No—Automerge of sockets by the kernel after switchover is inactive.
Batch count	Number of sockets merged per batch.
Batch count adjust	<p>Speed at which the batch count is adjusted.</p> <ul style="list-style-type: none"> • Slow—Number of sockets merged per batch is incremented additively. • Exp—Number of sockets merged per batch is incremented exponentially. • None—Number of sockets merged per batch remains constant.
Batch interval	Time interval between batches of automerge operation.
Batch interval adjust	<p>Speed at which the batch interval is adjusted.</p> <ul style="list-style-type: none"> • Exp—Time interval between automerge of batches is increased exponentially. • None—Time interval between automerge of batches is not adjusted.
Automerge State	<p>State of the automerge</p> <ul style="list-style-type: none"> • Ready—Ready to automerge socket pairs from secondary to primary routing engine • InProgress—Kernel is performing automerge after switchover • Switchover Completed—Sessions merged after switchover
Sessions Processed	Count of sessions that are automergerd.

Sample Output

show nonstop-routing (MX Series Router)

```
user@host show nonstop-routing
```

```
Nonstop Routing : Enabled
Precision Timers state: Enabled: Completed - 0/0
Precision Timers max period: 200
```

```
Automerge : Active
Batching: No
Batch count: 200
Batch count adjust: Exponential
Batch interval: 20 msec
Batch interval adjust: None
Automerge State: Ready
Sessions Processed: 0
```

show nonstop-routing (MX Series Router)

```
user@host> show nonstop-routing
```

```
Nonstop Routing : Enabled
  Automerge : Active
  Batching: Yes
  Batch count: 500
  Batch count adjust: Slow
  Batch interval: 50 msec
  Batch interval adjust: None
  Automerge State: Ready
  Sessions Processed: 0
```

show pfe ssb

Syntax `show pfe ssb`

Release Information Command introduced before Junos OS Release 7.4.

Description (M20 routers only) Display Packet Forwarding Engine System and Switch Board (SSB) status and statistics information.

Options This command has no options.

Required Privilege Level admin

List of Sample Output [show pfe ssb on page 770](#)

Output Fields [Table 54 on page 768](#) lists the output fields for the **show pfe ssb** command. Output fields are listed in the approximate order in which they appear.

Table 54: show pfe ssb Output Fields

Field Name	Field Description
Uptime (total)	SSB uptime.
Failures	Number of failures .
Pending	Number of pending.
Peer message type receive qualifiers	Information about Peer message type receive qualifiers.
Message Type	Peer message type.
Receive Qualifier	Peer receive qualifier.
TTP	Peer message type TTP.
IFD	Peer message type IFD.
IFL	Peer message type IFL.
Nexthop	Peer message type Nexthop.
COS	Peer message type COS.
Route	Peer message type Route.

Table 54: show pfe ssb Output Fields (continued)

Field Name	Field Description
SW Firewall	Peer message type SW Firewall.
HW Firewall	Peer message type HW Firewall.
PFE Statistics	Peer message type PFE Statistics.
PIC Statistics	Peer message type PIC Statistics.
Sampling	Peer message type Sampling .
Monitoring	Peer message type Monitoring.
ASP	Peer message type ASP.
L2TP	Peer message type L2TP.
Collector	Peer message type Collector.
PIC Configuration	Peer message type PIC Configuration.
Queue Statistics	Peer message type Queue Statistics.
PFE Listener statistics	Information about Packet Forwarding Engine listener statistics: <ul style="list-style-type: none"> • Open—Number of PFE listeners in the “open” state. • Close—Number of PFE listeners in the “close” state. • Sleep—Number of PFE listeners in the “sleep” state. • Wakeup—Number of PFE listeners in the “wakeup” state. • Resync Request—Number of PFE listeners in the “resync request” state. • Resync Done—Number of PFE listeners in the “resync done” state. • Resync Fail—Number of PFE listeners in the “resync fail” state • Resync Time—Number of PFE listeners in the resync time state.

Table 54: show pfe ssb Output Fields (continued)

Field Name	Field Description
PFE IPC statistics	<p>Information about Packet Forwarding Engine IPC statistics.</p> <ul style="list-style-type: none"> • type—Type of IPC message. <ul style="list-style-type: none"> • Header—IPC message type Header. • Test—IPC message type Test. • Interface—IPC message type Interface. • Chassis—IPC message type Chassis. • Boot—IPC message type Boot • Next-hop—IPC message type Next-hop. • Jtree—IPC message type Jtree. • Cprod—IPC message type Cprod. • Route—IPC message type Route. • Pfe—IPC message type PFE. • Dfw—IPC message type Dfw. • Mastership—IPC message type Mastership. • Sampling—IPC message type Sampling. • GUCP—IPC message type GUCP. • CoS—IPC message type CoS. • GCCP—IPC message type GCCP. • GHCP—IPC message type GHCP. • IRSD—IPC message type IRSD. • Monitoring—IPC message type Monitoring. • RE—IPC message type RE. • PIC—IPC message type PIC. • ASP cfg—IPC message type ASP configuration. • ASP cmd—IPC message type ASP command.. • L2TP cfg—IPC message type L2TP configuration. • Collector—IPC message type Collector. • PIC state—IPC message type PIC state. • Aggregator—IPC message type Aggregate. • Empty—IPC message type Empty. • PFE socket-buffer mbuf depth—Information about Packet Forwarding Engine socket-buffer depth • bucket—mbuf bucket value. • count—mbuf count value.
PFE socket-buffer bytes pending transmit—	<p>Information about Packet Forwarding Engine socket-buffer bytes pending for transmit.</p> <ul style="list-style-type: none"> • TX Messages—Number of transmitted messages. • RX messages—Number of received messages.

Sample Output

show pfe ssb

```
user@host> show pfe ssb
```


SSB status:

```

Slot:          Present
State:         Online
Last State Change: 2005-03-06 03:10:28 PST
Uptime (total): 11:23:27
Failures:      0
Pending:       0

```

Peer message type receive qualifiers:

Message Type	Receive Qualifier
-----	-----
TTP	Slot only
IFD	All
IFL	All
Nexthop	All
COS	All
Route	All
SW Firewall	All
HW Firewall	All
PFE Statistics	All
PIC Statistics	None
Sampling	All
Monitoring	None
ASP	None
L2TP	None
Collector	None
PIC Configuration	None
Queue Statistics	None
(null)	None

PFE listener statistics:

```

Open:          1
Close:         0
Sleep:         0
Wakeup:        0
Resync Request: 0
Resync Done:   1
Resync Fail:   0
Resync Time:   0

```

PFE IPC statistics:

type	TX Messages	RX messages
-----	-----	-----
Header	0	0
Test	0	0
Interface	737	9911
Chassis	0	0
Boot	0	0
Next-hop	48	0
Jtree	0	0
Cprod	0	0
Route	94	0
Pfe	2034	683
Dfw	8	0
Mastership	0	0
Sampling	0	0
GUCP	0	0
CoS	73	0

GCCP	0	0
GHCP	0	0
IRSD	0	0
Monitoring	0	0
RE	0	0
PIC	0	0
ASP cfg	0	0
ASP cmd	0	0
L2TP cfg	0	0
Collector	0	0
PIC state	0	0
Aggregator	0	0
Empty	0	0

PFE socket-buffer mbuf depth:

bucket	count
-----	-----
0	0
1	0
2	0
3	0
4	0
5	0
6	0
7	0
8	0
9	0
10	0
11	0
12	0
13	0
14	0
15	0
16	0
17	0
18	0
19	0
20	0
21	0

PFE socket-buffer bytes pending transmit:

bucket	count
-----	-----
0	0
1	0
2	0
3	0
4	0
5	0
6	0
7	0
8	0
9	0
10	0
11	0
12	0
13	0
14	0
15	0
16	0

17	0
18	0
19	0
20	0
21	0

show system switchover

List of Syntax [Syntax on page 774](#)
 [Syntax \(TX Matrix Router\) on page 774](#)
 [Syntax \(TX Matrix Plus Router\) on page 774](#)
 [Syntax \(MX Series Router\) on page 774](#)

Syntax show system switchover

Syntax (TX Matrix Router) show system switchover
 <all-chassis | all-lcc | lcc *number* | scc>

Syntax (TX Matrix Plus Router) show system switchover
 <all-chassis | all-lcc | lcc *number* | sfc *number*>

Syntax (MX Series Router) show system switchover
 <all-members>
 <local>
 <member *member-id*>

Release Information Command introduced before Junos OS Release 7.4.
 Command introduced in Junos OS Release 9.0 for EX Series switches.
 sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6.
 Command introduced in Junos OS Release 13.2X51-D20 for QFX Series switches.

Description Display whether graceful Routing Engine switchover is configured, the state of the kernel replication (ready or synchronizing), any replication errors, and whether the primary and standby Routing Engines are using compatible versions of the kernel database.



NOTE: Issue the **show system switchover** command *only* on the backup Routing Engine. This command is *not* supported on the master Routing Engine because the kernel-replication process daemon does not run on the master Routing Engine. This process runs only on the backup Routing Engine.

Beginning Junos OS Release 9.6, the **show system switchover** command has been deprecated on the master Routing Engine on all routers other than a TX Matrix (switch-card chassis) or a TX Matrix Plus (switch-fabric chassis) router.

However, in a routing matrix, if you issue the **show system switchover** command on the master Routing Engine of the TX Matrix router (or switch-card chassis), the CLI displays graceful switchover information for the master Routing Engine of the T640 routers (or line-card chassis) in the routing matrix. Likewise, if you issue the **show system switchover** command on the master Routing Engine of a TX Matrix Plus router (or switch-fabric

chassis), the CLI displays output for the master Routing Engine of T1600 or T4000 routers in the routing matrix.

Options **all-chassis**—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display graceful Routing Engine switchover information for all Routing Engines on the TX Matrix router and the T640 routers configured in the routing matrix. On a TX Matrix Plus router, display graceful Routing Engine switchover information for all Routing Engines on the TX Matrix Plus router and the T1600 or T4000 routers configured in the routing matrix.

all-lcc—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display graceful Routing Engine switchover information for all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus router, display graceful Routing Engine switchover information for all connected T1600 or T4000 LCCs.

Note that in this instance, packets get dropped. The LCCs perform GRES on their own chassis (GRES cannot be handled by one particular chassis for the entire router) and synchronization is not possible as the LCC plane bringup time varies for each LCC. Therefore, when there is traffic on these planes, there may be a traffic drop.

all-members—(MX Series routers only) (Optional) Display graceful Routing Engine switchover information for all Routing Engines on all members of the Virtual Chassis configuration.

lcc *number*—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display graceful Routing Engine switchover information for a specific T640 router connected to the TX Matrix router. On a TX Matrix Plus router, display graceful Routing Engine switchover information for a specific router connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(MX Series routers only) (Optional) Display graceful Routing Engines switchover information for all Routing Engines on the local Virtual Chassis member.

member *member-id*—(MX Series routers only) (Optional) Display graceful Routing Engine switchover information for all Routing Engines on the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

scc—(TX Matrix router only) (Optional) Display graceful Routing Engine switchover information for the TX Matrix router (or switch-card chassis).

sfc—(TX Matrix Plus routers only) (Optional) Display graceful Routing Engine switchover information for the TX Matrix Plus router.

Additional Information If you issue the **show system switchover** command on a TX Matrix backup Routing Engine, the command is broadcast to all the T640 backup Routing Engines that are connected to it.

Likewise, if you issue the **show system switchover** command on a TX Matrix Plus backup Routing Engine, the command is broadcast to all the T1600 or T4000 backup Routing Engines that are connected to it.

If you issue the **show system switchover** command on the active Routing Engine in the master router of an MX Series Virtual Chassis, the router displays a message that this command is not applicable on this member of the Virtual Chassis.

Required Privilege Level view

Related Documentation

- [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

List of Sample Output [show system switchover \(Backup Routing Engine - Ready\) on page 777](#)
[show system switchover \(Backup Routing Engine - Not Ready\) on page 777](#)
[show system switchover all-lcc \(Routing Matrix and Routing Matrix Plus\) on page 778](#)

Output Fields [Table 55 on page 776](#) describes the output fields for the **show system switchover** command. Output fields are listed in the approximate order in which they appear.

Table 55: show system switchover Output Fields

Field Name	Field Description
Graceful switchover	Display graceful Routing Engine switchover status: <ul style="list-style-type: none"> • On—Indicates graceful-switchover is specified for the routing-options configuration command. • Off—Indicates graceful-switchover is not specified for the routing-options configuration command.
Configuration database	State of the configuration database: <ul style="list-style-type: none"> • Ready—Configuration database has synchronized. • Synchronizing—Configuration database is synchronizing. Displayed when there are updates within the last 5 seconds. • Synchronize failed—Configuration database synchronize process failed.

Table 55: show system switchover Output Fields (continued)

Field Name	Field Description
Kernel database	<p>State of the kernel database:</p> <ul style="list-style-type: none"> • Ready—Kernel database has synchronized. This message implies that the system is ready for GRES. • Synchronizing—Kernel database is synchronizing. Displayed when there are updates within the last 5 seconds. • Version incompatible—The primary and standby Routing Engines are running incompatible kernel database versions. • Replication error—An error occurred when the state was replicated from the primary Routing Engine. Inspect Steady State for possible causes, or notify Juniper Networks customer support.
Peer state	<p>Routing Engine peer state:</p> <p>This field is displayed only when ksyncd is running in multichassis mode (LCC master).</p> <ul style="list-style-type: none"> • Steady State—Peer completed switchover transition. • Peer Connected—Peer in switchover transition.
Switchover Status	<p>Switchover Status:</p> <ul style="list-style-type: none"> • Ready—Message for system being switchover ready. • Not Ready—Message for system not being ready for switchover.

Sample Output

show system switchover (Backup Routing Engine - Ready)

```
user@host> show system switchover
```

```
Graceful switchover: On
Configuration database: Ready
Kernel database: Ready
Peer state: Steady State
Switchover Status: Ready
```

Switchover Status: Ready is the way the last line of the output reads if you are running Junos OS Release 16.1R1 or later. If you are running Junos OS Release 15.x, the last line of the output reads as Switchover Ready, for example:

```
user@host> show system switchover
```

```
Graceful switchover: On
Configuration database: Ready
Kernel database: Ready
Switchover Ready
```

show system switchover (Backup Routing Engine - Not Ready)

```
user@host> show system switchover
```

```
Graceful switchover: On
Configuration database: Ready
Kernel database: Ready
```

```
Peer state: Steady State
Switchover Status: Not Ready
```

Switchover Status: Not Ready is the way the last line of the output reads if you are running Junos OS Release 16.1R1 or later. If you are running Junos OS Release 15.x, the last line of the output reads as Not ready for mastership switch, try after xxx secs, for example:

```
user@host> show system switchover

Graceful switchover: On
Configuration database: Ready
Kernel database: Ready
Not ready for mastership switch, try after xxx secs.
```

show system switchover all-lcc (Routing Matrix and Routing Matrix Plus)

```
user@host> show system switchover all-lcc
```

```
lcc0-re0:
-----
Multichassis replication: On
Configuration database: Ready
Kernel database: Ready
Peer state: Steady State
Switchover Status: Ready
```

```
lcc2-re0:
-----
Multichassis replication: On
Configuration database: Ready
Kernel database: Ready
Peer state: Steady State
Switchover Status: Ready
```


show task replication

Syntax `show task replication`

Release Information Command introduced in Junos OS Release 8.5.
 Command introduced in Junos OS Release 9.0 for EX Series switches.
 Command introduced in Junos OS Release 13.2X51-D20 for QFX Series switches.
 Support for logical systems introduced in Junos OS Release 13.3

Description Displays nonstop active routing (NSR) status. When you issue this command on the master Routing Engine, the status of nonstop active routing synchronization is also displayed.



CAUTION: If BGP is configured, before attempting nonstop active routing switchover, check the output of `show bgp replication` to confirm that BGP routing table synchronization has completed on the backup Routing Engine. The complete status in the output of `show task replication` only indicates that the socket replication has completed and the BGP synchronization is in progress.

To determine whether BGP synchronization is complete, you must check the Protocol state and Synchronization state fields in the output of `show bgp replication` on the master Routing Engine. The Protocol state must be idle and the Synchronization state must be complete. If you perform NSR switchover before the BGP synchronization has completed, the BGP session might flap.

Options This command has no options.

Required Privilege Level view

Related Documentation • [Example: Configuring Nonstop Active Routing on Switches on page 226](#)

List of Sample Output [show task replication \(Issued on the Master Routing Engine\) on page 780](#)
[show task replication \(Issued on the Backup Routing Engine\) on page 780](#)
[show task replication \(Junos OS Evolved\) on page 780](#)

Output Fields [Table 56 on page 780](#) lists the output fields for the `show task replication` command. Output fields are listed in the approximate order in which they appear.

Table 56: show task replication Output Fields

Field Name	Field Description
Stateful replication	Displays whether or not graceful Routing Engine switchover is configured. The status can be Enabled or Disabled .
RE mode	Displays the Routing Engine on which the command is issued: Master , Backup , or Not applicable (when the router has only one Routing Engine).
Protocol	Protocols that are supported by nonstop active routing.
Synchronization Status	Nonstop active routing synchronization status for the supported protocols. States are NotStarted , InProgress , and Complete . Synchronization states are shown for each of the supported protocols that are running on the device at that moment.

Sample Output

show task replication (Issued on the Master Routing Engine)

```
user@host> show task replication

Stateful Replication: Enabled
RE mode: Master

Protocol      Synchronization Status
-----
OSPF          NotStarted
BGP           Complete
IS-IS        NotStarted
LDP           Complete
PIM           Complete
```

show task replication (Issued on the Backup Routing Engine)

```
user@host> show task replication

Stateful Replication: Enabled
RE mode: Backup
```

show task replication (Junos OS Evolved)

In Junos OS Evolved, both the master and backup Routings have the same CLI output. If you configured any protocol, you should see the synchronization state for the same.

```
user@host> show task replication

Stateful Replication: Enabled
RE mode: Master

Protocol      Synchronization Status
-----
OSPF          NotStarted
BGP           Complete
IS-IS        NotStarted
```

LDP	Complete
PIM	Complete

show vrrp

Syntax	<pre>show vrrp <brief detail extensive summary> <interface <i>interface-name</i> <group number>> <logical-system <i>logical-system-name</i> > <nsr></pre>
Release Information	Command introduced before Junos OS Release 7.4. nsr option added in Junos OS Release 13.2.
Description	Display status information about Virtual Router Redundancy Protocol (VRRP) groups.
Options	<p>none—(Same as brief) Display brief status information about all VRRP interfaces.</p> <p>brief detail extensive summary—(Optional) Display the specified level of output.</p> <p>interface <i>interface-name</i> <group number>—(Optional) Display information and status about the specified VRRP interface and, optionally, the group number.</p> <p>logical-system <i>logical-system-name</i>—(Optional) Perform this operation on a particular logical system.</p> <p>nsr—(Optional) Display state replication information when graceful Routing Engine switchover (GRES) with nonstop active routing (NSR) is configured. Use only on the backup Routing Engine.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• show vrrp track on page 793• clear vrrp on page 695
List of Sample Output	<ul style="list-style-type: none">• show vrrp on page 788• show vrrp brief on page 788• show vrrp detail (IPv6) on page 788• show vrrp detail (Route Track) on page 789• show vrrp detail (Route Track) on page 789• show vrrp extensive on page 790• show vrrp interface on page 790• show vrrp nsr on page 791• show vrrp summary on page 792
Output Fields	Table 57 on page 783 lists the output fields for the show vrrp command. Output fields are listed in the approximate order in which they appear

Table 57: show vrrp Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the logical interface.	brief extensive none summary
Interface index	Physical interface index number, which reflects its initialization sequence.	extensive
Groups	Total number of VRRP groups configured on the interface.	extensive
Active	Total number of VRRP groups that are active (that is, whose interface state is either up or down).	extensive
Interface VRRP PDU statistics	Non-errored statistics for the logical interface: <ul style="list-style-type: none"> • Advertisement sent—Number of VRRP advertisement protocol data units (PDUs) that the interface has transmitted. • Advertisement received—Number of VRRP advertisement PDUs received by the interface. • Packets received—Number of VRRP packets received for VRRP groups on the interface. • No group match received—Number of VRRP packets received for VRRP groups that do not exist on the interface. 	extensive
Interface VRRP PDU error statistics	Errored statistics for the logical interface: <ul style="list-style-type: none"> • Invalid IPAH next type received—Number of packets received that use the IP Authentication Header protocol (IPAH) and that do not encapsulate VRRP packets. • Invalid VRRP ttl value received—Number of packets received whose IP time-to-live (TTL) value is not 255. • Invalid VRRP version received—Number of packets received whose VRRP version is not 2. • Invalid VRRP pdu type received—Number of packets received whose VRRP PDU type is not 1. • Invalid VRRP authentication type received—Number of packets received whose VRRP authentication is not none, simple, or md5. • Invalid VRRP IP count received—Number of packets received whose VRRP IP count exceeds 8. • Invalid VRRP checksum received—Number of packets received whose VRRP checksum does not match the calculated one. 	extensive
Physical interface	Name of the physical interface.	detail extensive
Unit	Logical unit number.	All levels
Address	Address of the physical interface.	brief detail extensive none
Index	Physical interface index number, which reflects its initialization sequence.	detail extensive
SNMP ifIndex	SNMP index number for the physical interface.	detail extensive

Table 57: show vrrp Output Fields (continued)

Field Name	Field Description	Level of Output
VRRP-Traps	Status of VRRP traps: Enabled or Disabled .	detail extensive
VRRP-Version	VRRP version: 2 or 3 .	detail extensive
Type and Address	Identifier for the address and the address itself: <ul style="list-style-type: none"> lcl—Configured local interface address. mas—Address of the master virtual router. This address is displayed only when the local interface is acting as a backup router. vip—Configured virtual IP addresses. 	brief none summary
Interface state/Int state/State	State of the physical interface: <ul style="list-style-type: none"> down—The device is present and the link is unavailable. not present—The interface is configured, but no physical device is present. unknown—The VRRP process has not had time to query the kernel about the state of the interface. up—The device is present and the link is established. 	brief extensive none summary
Group	VRRP group number.	brief extensive none summary
State	The state of the interface on which VRRP is running: <ul style="list-style-type: none"> backup—The interface is acting as the backup router interface. bringup—VRRP is just starting and the physical device is not yet present. idle—VRRP is configured on the interface and is disabled. This can occur when VRRP is first enabled on an interface whose link is established. init—VRRP is initializing. master—The interface is acting as the master router interface. master(ISSU)—The master router interface is going through a unified in-service software upgrade. transition—The interface is changing between being the backup and being the master router. 	extensive
VRRP Mode	If the interface inherits its state and configuration from the active VRRP group, or if it is part of the active VRRP group. <ul style="list-style-type: none"> Active—Part of the active VRRP group Inherit—Inherits state and configuration from the active VRRP group. 	detail extensive
Priority	Configured VRRP priority for the interface.	detail extensive
Advertisement interval	Configured VRRP advertisement interval.	detail extensive
Authentication type	Configured VRRP authentication type: none , simple , or md5 .	detail extensive

Table 57: show vrrp Output Fields (continued)

Field Name	Field Description	Level of Output
Advertisement Threshold	A value from 1 through 15, used for setting the time when a peer should be considered down. <ul style="list-style-type: none"> The time a peer is considered down is equal to the advertisement-threshold multiplied by the advertisement-interval. (advertisement-threshold *advertisement-interval) = Peer down. 	detail extensive
Computed Send Rate	How many protocol data units (PDUs) are generated per second. Based on the number of instances and the advertisement interval.	detail extensive
Preempt	Whether preemption is allowed on the interface: yes or no .	detail extensive
Accept-data mode	Whether the interface is configured to accept packets destined for the virtual IP address: yes or no .	detail extensive
VIP count	Number of virtual IP addresses that have been configured on the interface.	detail extensive
VIP	List of virtual IP addresses configured on the interface.	detail extensive
Advertisement timer	How long, in seconds, until the advertisement timer expires.	detail extensive
Master router	IP address of the interface that is acting as the master. If the VRRP interface is down, the output is N/A .	detail extensive
Virtual router uptime	How long, in seconds, that the virtual router has been up.	detail extensive
Master router uptime	How long, in seconds, that the master route has been up.	detail extensive
Virtual MAC	MAC address associated with the virtual IP address.	detail extensive
Tracking	Whether tracking is enabled or disabled .	detail extensive
Current priority	Current operational priority for being the VRRP master.	detail extensive
Configured priority	Configured base priority for being the VRRP master.	detail extensive
Priority hold-time	Minimum time interval, in seconds, between successive changes to the current priority. Disabled indicates no minimum interval.	detail extensive
Remaining-time	(track option only) Displays the time remaining in the priority hold-time interval.	detail
Interface tracking	Whether interface tracking is enabled or disabled. When enabled, the output also displays the number of tracked interfaces.	detail extensive
Interface/Tracked interface/Track Int	Name of the tracked interface.	detail extensive

Table 57: show vrrp Output Fields (continued)

Field Name	Field Description	Level of Output
Int state/Interface state/State	Current operational state of the tracked interface: up or down .	detail extensive
Int speed/Speed	Current operational speed, in bits per second, of the tracked interface.	detail extensive
Incurred priority cost	Operational priority cost incurred due to the state and speed of this tracked interface. This cost is applied to the configured priority to obtain the current priority.	detail extensive
Threshold	Speed below which the corresponding priority cost is incurred. In other words, when the speed of the interface drops below the threshold speed, the corresponding priority cost is incurred. An entry of down means that the corresponding priority cost is incurred when the interface is down.	detail extensive
Route tracking	Whether route tracking is enabled or disabled. When enabled, the output also displays the number of tracked routes.	detail extensive
Route count	The number of routes being tracked.	detail extensive
Route	The IP address of the route being tracked.	detail extensive
VRF name	The VPN routing and forwarding (VRF) routing instance that the tracked route is in.	detail extensive
Route state	The state of the route being tracked: up , down , or unknown .	detail extensive
Priority cost	Configured priority cost. This value is incurred when the interface speed drops below the corresponding threshold or when the tracked route goes down.	detail extensive
Active	Whether the threshold is active (*). If the threshold is active, the corresponding priority cost is incurred.	detail extensive
Group VRRP PDU statistics	Number of VRRP advertisements sent and received by the group.	extensive

Table 57: show vrrp Output Fields (continued)

Field Name	Field Description	Level of Output
Group VRRP PDU error statistics	<p>Errored statistics for the VRRP group:</p> <ul style="list-style-type: none"> • Bad authentication type received—Number of VRRP PDUs received with an invalid authentication type. The received authentication can be none, simple, or md5 and must be the same for all routers in the VRRP group. • Bad password received—Number of VRRP PDUs received with an invalid key (password). The password for simple authentication must be the same for all routers in the VRRP group • Bad MD5 digest received—Number of VRRP PDUs received for which the MD5 digest computed from the VRRP PDU differs from the digest expected by the VRRP instance configured on the router. • Bad advertisement timer received—Number of VRRP PDUs received with an advertisement time interval that is inconsistent with the one in use among the routers in the VRRP group. • Bad VIP count received—Number of VRRP PDUs whose virtual IP address counts differ from the count that has been configured on the VRRP instance. • Bad VIPADDR received—Number of VRRP PDUs whose virtual IP addresses differ from the list of virtual IP addresses configured on the VRRP instance. 	extensive
Group state transition statistics	<p>State transition statistics for the VRRP group:</p> <ul style="list-style-type: none"> • Idle to master transitions—Number of times that the VRRP instance transitioned from the idle state to the master state. • Idle to backup transitions—Number of times that the VRRP instance transitioned from the idle state to the backup state. • Backup to master transitions—Number of times that the VRRP instance transitioned from the backup state to the master state. • Master to backup transitions—Number of times that the VRRP instance transitioned from the master state to the backup state. 	extensive
VR state	<p>The state of the VRRP:</p> <ul style="list-style-type: none"> • backup—The interface is acting as the backup router interface. • bringup—VRRP is just starting, and the physical device is not yet present. • idle—VRRP is configured on the interface and is disabled. This can occur when VRRP is first enabled on an interface whose link is established. • init—VRRP is initializing. • master—The interface is acting as the master router interface. • transition—The interface is changing between being the backup and being the master router. <p>NOTE: When show vrrp nsr is used on the backup Routing Engine, it displays the current VRRP state on the master Routing Engine, which is the future VRRP state for the backup Routing Engine. Do not use on the master Routing Engine.</p>	brief none summary

Table 57: show vrrp Output Fields (continued)

Field Name	Field Description	Level of Output
NSR	<p>VRRP nonstop active routing is enabled for the configured VRRP group: yes or no.</p> <p>NOTE: A yes value means that the new master Routing Engine will immediately start with the VRRP State value from the original master Routing Engine.</p> <p>A no value means that the VRRP session will:</p> <ul style="list-style-type: none"> • Start afresh. • Go through asilent startup period. • Move to a backup state. • Wait for the D Timer to run out before becoming the master (only if the master has not been configured already). 	brief none
RPD-NSR	The routing options have been set to nonstop active routing: yes or no .	brief none
Timer	<p>VRRP timer information:</p> <ul style="list-style-type: none"> • A—How long, in seconds, until the advertisement timer expires. • D—How long, in seconds, until the Master is Down timer expires. 	brief none

Sample Output

show vrrp

```
user@host> show vrrp
```

```

Interface      State      Group  VR state  Timer  Type  Address
fe-0/0/0.121   up         1      master   A 1.052 1c1  fec0::12:1:1:1
                                     vip  fe80::12:1:1:99
                                     vip  fec0::12:1:1:99
fe-0/0/2.131   up         1      master   A 0.364 1c1  fec0::13:1:1:1
                                     vip  fe80::13:1:1:99
                                     vip  fec0::13:1:1:99

```

show vrrp brief

The output for the **show vrrp brief** command is identical to that for the **show vrrp** command. For sample output, see [show vrrp on page 788](#).

show vrrp detail (IPv6)

```
user@host> show vrrp detail
```

```
Physical interface: fe-0/0/0, Unit: 121, Vlan-id: 212, Address: fec0::12:1:1:1/120
```

```

Index: 67, SNMP ifIndex: 45, VRRP-Traps: enabled
Interface state: up, Group: 1, State: master, VRRP Mode: Active
Priority: 200, Advertisement interval: 1, Authentication type: none
Advertisement threshold: 3, Computed send rate: 0
Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: fe80::12:1:1:99,
fec0::12:1:1:99
Advertisement timer: 1.121s, Master router: fe80::12:1:1:1
Virtual router uptime: 00:03:47, Master router uptime: 00:03:41
Virtual MAC: 00:00:5e:00:02:01
Tracking: disabled

Physical interface: fe-0/0/2, Unit: 131, Vlan-id: 213, Address: fec0::13:1:1:120

Index: 69, SNMP ifIndex: 47, VRRP-Traps: enabled
Interface state: up, Group: 1, State: master
Priority: 200, Advertisement interval: 1, Authentication type: none
Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: fe80::13:1:1:99,
fec0::13:1:1:99
Advertisement timer: 0.327s, Master router: fe80::13:1:1:1
Virtual router uptime: 00:03:47, Master router uptime: 00:03:41
Virtual MAC: 00:00:5e:00:02:01
Tracking: disabled

```

show vrrp detail (Route Track)

```

user@host> show vrrp detail

Physical interface: ge-0/0/0, Unit: 1, Vlan-id: 1, Address: 101.1.1.1/24
Index: 324, SNMP ifIndex: 623, VRRP-Traps: enabled, VRRP-Version: 2
Interface state: up, Group: 1, State: master(ISSU), VRRP Mode: Active
Priority: 200, Advertisement interval: 1, Authentication type: none
Advertisement threshold: 3, Computed send rate: 0
Preempt: yes, Accept-data mode: no, VIP count: 1, VIP: 101.1.1.3
Advertisement Timer: 0.469s, Master router: 101.1.1.1
Virtual router uptime: 00:02:10, Master router uptime: 00:02:05
Virtual Mac: 00:00:5e:00:01:01
Tracking: disabled

```

show vrrp detail (Route Track)

```

user@host> show vrrp detail

Physical interface: ge-1/2/0, Unit: 0, Address: 30.30.30.30/24
Index: 67, SNMP ifIndex: 379, VRRP-Traps: enabled, VRRP-Version: 2
Interface state: up, Group: 100, State: master
Priority: 150, Advertisement interval: 1, Authentication type: none
Preempt: yes, Accept-data mode: no, VIP count: 1, VIP: 30.30.30.100
Advertisement timer: 1.218s, Master router: 30.30.30.30
Virtual router uptime: 00:04:28, Master router uptime: 00:00:13
Virtual MAC: 00:00:5e:00:01:64
Tracking: enabled
  Current priority: 150, Configured priority: 150
  Priority hold-time: disabled
  Interface tracking: disabled
  Route tracking: enabled, Route count: 1
    Route          VRF name    Route state  Priority cost
    192.168.40.0/22 default      up           30

```

show vrrp extensive

```

user@host> show vrrp extensive

Interface: ge-2/0/0.0, Interface index :65539, Groups: 1, Active :1
  Interface VRRP PDU statistics
    Advertisement sent                :0
    Advertisement received            :0
    Packets received                  :0
    No group match received            :0
  Interface VRRP PDU error statistics
    Invalid IPAH next type received    :0
    Invalid VRRP TTL value received    :0
    Invalid VRRP version received      :0
    Invalid VRRP PDU type received     :0
    Invalid VRRP authentication type received:0
    Invalid VRRP IP count received     :0
    Invalid VRRP checksum received     :0

Physical interface: ge-2/0/0, Unit: 0, Address: 10.10.10.1/24
  Index: 65539, SNMP ifIndex: 648, VRRP-Traps: enabled, VRRP-Version: 3
  Interface state: up, Group: 1, State: backup, VRRP Mode: Active
  Priority: 100, Advertisement interval: 1, Authentication type: none
  Advertisement threshold: 3, Computed send rate: 0
  Preempt: yes, Accept-data mode: no, VIP count: 1, VIP: 10.10.10.2
  Dead timer: 3.078s, Master priority: 0, Master router: 10.10.10.1
  Virtual router uptime: 00:00:04
  Tracking: disabled
  Group VRRP PDU statistics
    Advertisement sent                :0
    Advertisement received            :0
  Group VRRP PDU error statistics
    Bad authentication Type received    :0
    Bad password received               :0
    Bad MD5 digest received             :0
    Bad advertisement timer received    :0
    Bad VIP count received              :0
    Bad VIPADDR received               :0
  Group state transition statistics
    Idle to master transitions          :0
    Idle to backup transitions          :1
    Backup to master transitions        :0
    Master to backup transitions        :0

```

show vrrp interface

```

user@host> show vrrp interface ge-0/0/0.1

Interface: ge-0/0/0.1, Interface index :324, Groups: 2, Active :2
  Interface VRRP PDU statistics
    Advertisement sent                :39
    Advertisement received            :0
    Packets received                  :0
    No group match received            :0
  Interface VRRP PDU error statistics
    Invalid IPAH next type received    :0
    Invalid VRRP TTL value received    :0
    Invalid VRRP version received      :0
    Invalid VRRP PDU type received     :0
    Invalid VRRP authentication type received:0

```

```

Invalid VRRP IP count received      :0
Invalid VRRP checksum received     :0

Physical interface: ge-0/0/0, Unit: 1, Vlan-id: 1, Address: 101.1.1.1/24
Index: 324, SNMP ifIndex: 623, VRRP-Traps: enabled, VRRP-Version: 2
Interface state: up, Group: 1, State: master(ISSU), VRRP Mode: Active
Advertisement threshold: 3, Computed send rate: 0
Priority: 200, Advertisement interval: 1, Authentication type: none
Advertisement threshold: 3, Computed send rate: 0
Preempt: yes, Accept-data mode: no, VIP count: 1, VIP: 101.1.1.3
Advertisement Timer: 0.619s, Master router: 101.1.1.1
Virtual router uptime: 00:00:22, Master router uptime: 00:00:17
Virtual Mac: 00:00:5e:00:01:01
Tracking: disabled
Group VRRP PDU statistics
  Advertisement sent                :20
  Advertisement received            :0
Group VRRP PDU error statistics
  Bad authentication Type received  :0
  Bad password received             :0
  Bad MD5 digest received           :0
  Bad advertisement timer received  :0
  Bad VIP count received            :0
  Bad VIPADDR received             :0
Group state transition statistics
  Idle to master transitions        :0
  Idle to backup transitions        :1
  Backup to master transitions      :1
  Master to backup transitions      :0
Interface: fe-0/0/0.121, Interface index: 67, Groups: 1, Active : 1
Interface VRRP PDU statistics
  Advertisement sent                :      205
  Advertisement received            :         0
  Packets received                  :         0
  No group match received           :         0
Interface VRRP PDU error statistics
  Invalid IPAH next type received   :         0
  Invalid VRRP TTL value received   :         0
  Invalid VRRP version received    :         0
  Invalid VRRP PDU type received   :         0
  Invalid VRRP authentication type received:         0
  Invalid VRRP IP count received   :         0
  Invalid VRRP checksum received   :         0

```

show vrrp nsr

This command is similar to **show vrrp**. Here, the **VR state** column displays the current VRRP state on the master Routing Engine, which is the future VRRP state for the backup Routing Engine. Do not use on the master Routing Engine.

NSR is yes if VRRP nonstop active routing is enabled for the configured VRRP group.

RPD-NSR is yes if the routing options have been set to nonstop active routing.

```
user@host>show vrrp nsr
```

Interface	State	Group	VR state	VR Mode	Type	NSR	RPD-NSR	Address
ge-1/0/1.0	up	1	master	Active	1cl	yes	yes	10.0.0.1

```

                                vip                10.0.0.3
ge-1/0/1.0 up 2 master Active lcl yes yes 20.0.0.1
                                vip                20.0.0.3
ge-1/0/1.0 up 3 master Active lcl yes yes 30.0.0.1
                                vip                30.0.0.3
ge-1/0/1.0 up 4 master Active lcl yes yes 40.0.0.1
                                vip                40.0.0.3
ge-1/0/1.0 up 5 master Active lcl yes yes 50.0.0.1
                                vip                50.0.0.3
ge-1/0/1.0 up 1 master Active lcl yes yes 1000::1
fe80::200:5eff:fe00:1
                                vip
                                vip                1000::3
ge-1/0/1.0 up 2 master Active lcl yes yes 2000::1
fe80::200:5eff:fe00:2
                                vip
                                vip                2000::3
ge-1/0/1.0 up 3 master Active lcl yes yes 3000::1
fe80::200:5eff:fe00:3
                                vip
                                vip                3000::3
ge-1/0/1.0 up 4 master Active lcl yes yes 4000::1
fe80::200:5eff:fe00:4
                                vip
                                vip                4000::3
ge-1/0/1.0 up 5 master Active lcl yes yes 5000::1
fe80::200:5eff:fe00:5
                                vip
                                vip                5000::3

```

show vrrp summary

```
user@host> show vrrp summary
```

Interface	State	Group	VR state	Type	Address
ge-4/2/0.0	up	1	backup	lcl	10.57.0.2
				vip	10.57.0.100

show vrrp track

Syntax	<pre>show vrrp track <all interfaces routes> <detail summary> <logical-system <i>logical-system-name</i>></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>all and routes options added in Junos OS Release 17.1.</p>
Description	Display status information about Virtual Router Redundancy Protocol (VRRP) tracked routes and tracked interfaces.
Options	<p>none—(Same as summary) Display summarized status information of tracked routes and tracked interfaces.</p> <p>all interfaces routes—(Optional) These options display the following information:</p> <ul style="list-style-type: none"> all—Output is the same as for the show vrrp track command. interfaces—Show summary of VRRP tracked interfaces. routes—Show summary of VRRP tracked routes <p>detail summary—(Optional) Display detailed or summarized information.</p> <p>logical-system <i>logical-system-name</i>—(Optional) Perform this operation on a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> Configuring a Logical Interface to Be Tracked for a VRRP Group on page 358 Configuring a Route to Be Tracked for a VRRP Group on page 361 show vrrp on page 782
List of Sample Output	<p>show vrrp track summary on page 795</p> <p>show vrrp track detail on page 795</p> <p>show vrrp track interfaces summary on page 795</p> <p>show vrrp track interfaces detail on page 795</p> <p>show vrrp track routes summary on page 796</p> <p>show vrrp track routes detail on page 796</p>
Output Fields	<p>Table 58 on page 794 lists the output fields for the show vrrp track command. Output fields are listed in the approximate order in which they appear.</p>

Table 58: show vrrp track Output Fields

Fields	Description	Level
Tracked interface/Track Int	Name of the tracked interface.	detail or summary
State	Current operational state of the tracked interface: up or down .	detail or summary
Speed	Current operational speed, in bits per second, of the tracked interface.	detail or summary
Incurred priority cost	Operational priority cost incurred resulting from the state and speed of this tracked interface. This cost is applied to the configured priority to obtain the current priority cost.	detail
VRRP Int/Tracking VRRP interface	Name of the VRRP interface.	detail or summary
Group	VRRP group number.	detail or summary
VR state	<p>The state of the VRRP:</p> <ul style="list-style-type: none"> • backup—The interface is acting as the backup router interface. • bringup—VRRP is just starting, and the physical device is not yet present. • idle—VRRP is configured on the interface and is disabled. This can occur when VRRP is first enabled on an interface whose link is established. • init—VRRP is initializing. • master—The interface is acting as the master router interface. • transition—The interface is changing between being the backup and being the master router. <p>NOTE: When the show vrrp nsr command is used on the backup Routing Engine, it displays the current VRRP state on the master Routing Engine, which is the future VRRP state for the backup Routing Engine. Do not use the show vrrp nsr command on the master Routing Engine.</p>	detail or summary
Current priority	Current operational priority for being the VRRP master.	detail or summary
Priority hold-time	Minimum time interval, in seconds, between successive changes to the current priority cost. Disabled indicates no minimum interval.	detail
Track route	IP address of route.	detail or summary
State	State of route. Possible values are unknown , up , and down .	detail or summary
Cost	Priority cost. When the route state is not up , the cost will be deducted from the configured priority of the VRRP session.	detail or summary
Interface	Name of the logical interface (for example, ge-0/0/1.0) on which the corresponding VRRP session is configured.	detail or summary
Cfg	Configured priority.	detail or summary

Table 58: show vrrp track Output Fields (continued)

Fields	Description	Level
Run	Current (or running) priority cost.	detail or summary

Sample Output

show vrrp track summary

```
user@host> show vrrp track summary
```

Track Int	State	Speed	VRRP Int	Group	VR State	Current prio
ge-0/0/2.0	up	1g	ge-0/0/1.0	1	master	80
ge-0/0/8.0	up	1g	ge-0/0/1.0	1	master	80

Track route	State	Cost	Interface	Group	Cfg	Run	VR State
44.44.44.0/24	unknown	10	ge-0/0/1.0	1	100	80	master
55.55.55.0/24	unknown	10	ge-0/0/1.0	1	100	80	master

show vrrp track detail

```
user@host> show vrrp track detail
```

```
Tracked interface: ge-0/0/2.0
State: up, Speed: 1g
Incurred priority cost: 0
Tracking VRRP interface: ge-0/0/1.0, Group: 1
VR State: master
Current priority: 80, Configured priority: 100
Priority hold-time: disabled

Tracked interface: ge-0/0/8.0
State: up, Speed: 1g
Incurred priority cost: 0
Tracking VRRP interface: ge-0/0/1.0, Group: 1
VR State: master
Current priority: 80, Configured priority: 100
Priority hold-time: disabled
```


Track route	State	Cost	Interface	Group	Cfg	Run	VR State
44.44.44.0/24	unknown	10	ge-0/0/1.0	1	100	80	master
55.55.55.0/24	unknown	10	ge-0/0/1.0	1	100	80	master

show vrrp track interfaces summary

```
user@host> show vrrp track interfaces summary
```

Track Int	State	Speed	VRRP Int	Group	VR State	Current prio
ge-0/0/2.0	up	1g	ge-0/0/1.0	1	master	80
ge-0/0/8.0	up	1g	ge-0/0/1.0	1	master	80

show vrrp track interfaces detail

```
user@host> show vrrp track interfaces detail
```

```
Tracked interface: ge-0/0/2.0
State: up, Speed: 1g
Incurred priority cost: 0
Tracking VRRP interface: ge-0/0/1.0, Group: 1
VR State: master
Current priority: 80, Configured priority: 100
Priority hold-time: disabled

Tracked interface: ge-0/0/8.0
State: up, Speed: 1g
Incurred priority cost: 0
Tracking VRRP interface: ge-0/0/1.0, Group: 1
VR State: master
Current priority: 80, Configured priority: 100
Priority hold-time: disabled
```

show vrrp track routes summary

```
user@host> show vrrp track routes summary
```

Track route	State	Cost	Interface	Group	Cfg	Run	VR State
44.44.44.0/24	unknown	10	ge-1/0/0.0	1	100	60	bringup
55.55.55.0/24	unknown	10	ge-1/0/0.0	1	100	60	bringup

show vrrp track routes detail

The output for **show vrrp track routes detail** is the same as that for **show vrrp track routes summary**.

Troubleshooting

- [Tracing Nonstop Active Routing Synchronization Events on page 797](#)
- [Troubleshooting the EX Series Redundant Power System Power On and Power Backup Issues on page 799](#)

Tracing Nonstop Active Routing Synchronization Events

To track the progress of nonstop active routing synchronization between Routing Engines, you can configure nonstop active routing trace options flags for each supported protocol and for BFD sessions and record these operations to a log file.

To configure nonstop active routing trace options for supported routing protocols, include the **nsr-synchronization** statement at the **[edit protocols *protocol-name* traceoptions flag]** hierarchy level and optionally specify one or more of the **detail**, **disable**, **receive**, and **send** options:

```
[edit protocols]
bgp {
  traceoptions {
    flag nsr-synchronization <detail> <disable> <receive> <send>;
  }
}
isis {
  traceoptions {
    flag nsr-synchronization <detail> <disable> <receive> <send>;
  }
}
ldp {
  traceoptions {
    flag nsr-synchronization <detail> <disable> <receive> <send>;
  }
}
mpls {
  traceoptions {
    flag nsr-synchronization;
    flag nsr-synchronization-detail;
  }
}
msdp {
  traceoptions {
    flag nsr-synchronization <detail> <disable> <receive> <send>;
```

```
    }  
  }  
  (ospf | ospf3) {  
    traceoptions {  
      flag nsr-synchronization <detail> <disable> <receive> <send>;  
    }  
  }  
  rip {  
    traceoptions {  
      flag nsr-synchronization <detail> <disable> <receive> <send>;  
    }  
  }  
  ripng {  
    traceoptions {  
      flag nsr-synchronization <detail> <disable> <receive> <send>;  
    }  
  }  
  pim {  
    traceoptions {  
      flag nsr-synchronization <detail> <disable> <receive> <send>;  
    }  
  }  
}
```

To configure nonstop active routing trace options for BFD sessions, include the **nsr-synchronization** and **nsr-packet** statements at the **[edit protocols bfd traceoptions flag]** hierarchy level.

```
[edit protocols]  
bfd {  
  traceoptions {  
    flag nsr-synchronization;  
    flag nsr-packet;  
  }  
}
```

To trace the Layer 2 VPN signaling state replicated from routes advertised by BGP, include the **nsr-synchronization** statement at the **[edit routing-options traceoptions flag]** hierarchy level. This flag also traces the label and logical interface association that VPLS receives from the kernel replication state.

```
[edit routing-options]  
traceoptions {  
  flag nsr-synchronization;  
}
```

**Related
Documentation**

- [Configuring Nonstop Active Routing on page 217](#)
- [Configuring Nonstop Active Routing on Switches on page 220](#)
- [Example: Configuring Nonstop Active Routing on Switches on page 226](#)
- [Example: Configuring Nonstop Active Routing on page 222](#)

Troubleshooting the EX Series Redundant Power System Power On and Power Backup Issues

This topic provides troubleshooting information for problems related to the EX Series Redundant Power System (RPS).

1. [The EX Series RPS Is Not Powering On on page 799](#)
2. [A Switch Is Not Recognized by the RPS on page 799](#)
3. [An Error Message Indicates That an RPS Power Supply is Not Supported on page 800](#)
4. [The EX Series Redundant Power System Is Not Providing Power Backup to a Connected Switch on page 800](#)
5. [The Wrong Switches Are Being Backed Up on page 801](#)
6. [Six Switches That Do Not Require PoE Are Not All Being Backed Up on page 801](#)

The EX Series RPS Is Not Powering On

Problem **Description:** The RPS does not power on even though it has a power supply installed and is connected to an AC power source outlet.

Environment: The RPS with one EX-PWR3-930-AC power supply installed in it is connected to a switch.

Symptoms: The SYS LED on the power supply side of the RPS is off, and when you check the RPS status using the CLI command `show chassis redundant-power-system`, the message **No RPS connected** is displayed.

Cause A power supply must be installed in the middle slot on the RPS to power on the RPS.

Solution Install a power supply in the middle slot on the power supply side of the RPS and verify that the AC power source outlet is properly connected to it. See *Installing a Power Supply in the EX Series Redundant Power System*.

Verify that the **AC OK** LED and the **DC OK** LED on the power supply in the RPS are lit green.

A Switch Is Not Recognized by the RPS

Problem **Description:** I cannot set up the RPS.

Cause A switch must be active to be recognized by the RPS.

Solution Activate the switch by configuring it and issuing a commit statement.

An Error Message Indicates That an RPS Power Supply is Not Supported

- Problem** **Description:** An RPS error message indicates that an RPS power supply is not supported.
- Cause** RPS supports only one power supply, the EX-PWR3-930-AC. If you install another similar power supply, it may fit in the slot but it is not compatible with RPS.
- Solution** The power supply shipped with your RPS (in a separate box) is an EX-PWR3-930-AC. If you installed more power supplies, you ordered them separately. Replace any other power supply model (such as the EX-PWR2-930-AC) with an EX-PWR3-930-AC model.

The EX Series Redundant Power System Is Not Providing Power Backup to a Connected Switch

- Problem** **Description:** The RPS does not provide power backup to a connected switch.
- Environment:** The RPS has an EX-PWR3-930-AC power supply installed in the middle power supply slot and is connected to two switches with power loss, one connected to RPS switch connector port 1 and the other on port 2.
- Symptoms:** The status LED on the associated switch connector port is not blinking green—it is either solid green (connected) or not lit (off).
- Cause** The RPS provides backup power based on the power priority assigned to each switch.
- Solution** If the status LED on a switch connector port is off, ensure that the RPS cable is properly connected to both the RPS and the switch, and ensure that the priority configured for the switch is not 0. See [show redundant-power-system status](#).
- If the status LED on switch connector port 1 is on and is steadily green, check the backup priority configured for the switch and assign it a higher priority. See [“Determining and Setting Priority for Switches Connected to an EX Series RPS” on page 310](#)
- If the status LED on switch connector port 1 is amber, check if the RPS has enough power supplies installed in it to provide backup power. If it does not, install a power supply in an empty power supply slot on the RPS. See *Installing a Power Supply in the EX Series Redundant Power System*.
- If the status LED on switch connector port 1 is still off, check the priority configured for the switch. Ensure that the is not set to 0, which means off. See [show redundant-power-system status](#). The priority assigned must be from 1 through 6. See [“Determining and Setting Priority for Switches Connected to an EX Series RPS” on page 310](#).
- Verify that a dedicated power supply is installed in the switch. The RPS cannot boot a switch that does not have a dedicated power supply. See *Installing a Power Supply in the EX Series Redundant Power System*.

Also keep in mind that when the command `request redundant-power-system multi-backup` has been set, support for switches that supply PoE is not guaranteed. To reverse this setting, use the command `request redundant-power-system no-multi-backup`.

The Wrong Switches Are Being Backed Up

Problem **Description:** Four or more switches are connected to an RPS with three power supplies. When all four switches fail, the wrong three switches have .

Environment: Four or more switches are connected to an RPS with three power supplies. One or more switches provide PoE to other devices.

Symptoms: When all four switches fail, the wrong three switches have .

Cause The RPS provides backup power based on the power priority assigned to each switch. This is derived from two configurations, one of which has precedence over the other one. Initial is derived from the location of the port used to attach a switch—the leftmost connector has lowest priority and the rightmost connector has highest priority. The second, dominant priority configuration is derived from a CLI priority setting on the switch itself. With this CLI configuration, 6 is highest priority and 1 is the lowest priority.

Solution Connect the three switches to the three rightmost connectors on the RPS. Then, using the CLI on each switch, set each switch's priority to 1 using the `redundant-power-system` configuration command `redundant-power-system 1`. Now, physical connection location is determining .

If you do not want to change the cabling on the switches, you can use the configuration statement `redundant-power-system` on all four switches, assigning priority **6** (highest), **5**, **4** and **3** to the appropriate switches. Priority configuration on the switch always overcomes set by connector location.

Six Switches That Do Not Require PoE Are Not All Being Backed Up

Problem **Description:** Only three switches out of six are simultaneously backed up when all switches experience power supply failure. None of these switches supply PoE power to any device.

Environment: The RPS with three EX-PWR3-930-AC power supplies installed in it is connected to six switches, none of which is connected to a non-PoE device.

Symptoms: Only three switches out of six are simultaneously backed up when all switches experience power supply failure. None of these switches supply PoE power to any device.

Cause Each power supply can support two switches that do not need enough power for PoE, as long as you configure the RPS to do so.

Solution From any of the attached switches, issue the `request redundant-power-system multi-backup` command from the operational mode. Now standard power will be supplied to two non-PoE switches per power supply.