



Junos[®] OS

High Availability Feature Guide for Routing Devices



Modified: 2017-03-15

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

Juniper Networks, Junos, Steel-Belted Radius, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. The Juniper Networks Logo, the Junos logo, and JunosE are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, 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 for Routing Devices
Copyright © 2017, 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 <http://www.juniper.net/support/eula.html>. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

| | | |
|------------------|--|-----------|
| | About the Documentation | xv |
| | Documentation and Release Notes | xv |
| | Supported Platforms | xv |
| | Using the Examples in This Manual | xv |
| | Merging a Full Example | xvi |
| | Merging a Snippet | xvi |
| | Documentation Conventions | xvii |
| | Documentation Feedback | xix |
| | Requesting Technical Support | xix |
| | Self-Help Online Tools and Resources | xix |
| | Opening a Case with JTAC | xx |
| 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 | 8 |
| Part 2 | Configuring High Availability | |
| Chapter 2 | High Availability and Bidirectional Forwarding Detection (BFD) | 13 |
| | Understanding BFD for Static Routes for Faster Network Failure Detection | 13 |
| | Example: Configuring BFD for Static Routes for Faster Network Failure Detection | 18 |
| | Understanding BFD for BGP | 24 |
| | Example: Configuring BFD on Internal BGP Peer Sessions | 25 |
| | Understanding BFD for OSPF | 34 |
| | Example: Configuring BFD for OSPF | 36 |
| | Understanding BFD for IS-IS | 40 |
| | Example: Configuring BFD for IS-IS | 43 |
| | Understanding BFD for RIP | 48 |

| | | |
|------------------|---|-----------|
| | Example: Configuring BFD for RIP | 49 |
| | Understanding Independent Micro BFD Sessions for LAG | 55 |
| | Configuring Independent Micro BFD Sessions for LAG | 58 |
| | Example: Configuring Independent Micro BFD Sessions for LAG | 63 |
| | Configuring BFD for PIM | 73 |
| | Understanding Distributed BFD | 74 |
| | Understanding Static Route State When BFD is in Admin Down State | 77 |
| Chapter 3 | High Availability and Switching Control Board Redundancy | 79 |
| | Understanding Switching Control Board Redundancy | 79 |
| | Redundant CFEBs on the M10i Router | 80 |
| | Redundant FEBs on the M120 Router | 80 |
| | Redundant SSBs on the M20 Router | 82 |
| | Redundant SFMs on the M40e and M160 Routers | 83 |
| | Configuring CFEB Redundancy on the M10i Router | 83 |
| | Configuring FEB Redundancy on the M120 Router | 84 |
| | Example: Configuring FEB Redundancy on M120 Routers | 85 |
| | Configuring SFM Redundancy on M40e and M160 Routers | 86 |
| | Configuring SSB Redundancy on the M20 Router | 87 |
| Chapter 4 | High Availability and Routing Engine Redundancy | 89 |
| | Understanding Routing Engine Redundancy on Juniper Networks Routers | 89 |
| | Routing Engine Redundancy Overview | 89 |
| | Conditions That Trigger a Routing Engine Failover | 90 |
| | Default Routing Engine Redundancy Behavior | 91 |
| | Routing Engine Redundancy on a TX Matrix Router | 92 |
| | Routing Engine Redundancy on a TX Matrix Plus Router | 93 |
| | Situations That Require You to Halt Routing Engines | 94 |
| | Configuring Routing Engine Redundancy | 95 |
| | Modifying the Default Routing Engine Mastership | 95 |
| | Configuring Automatic Failover to the Backup Routing Engine | 96 |
| | Without Interruption to Packet Forwarding | 96 |
| | On Detection of a Hard Disk Error on the Master Routing Engine | 96 |
| | On Detection of a Loss of Keepalive Signal from the Master Routing Engine | 96 |
| | On Detection of the em0 Interface Failure on the Master Routing Engine | 98 |
| | When a Software Process Fails | 98 |
| | Manually Switching Routing Engine Mastership | 98 |
| | Verifying Routing Engine Redundancy Status | 98 |
| | Initial Routing Engine Configuration Example | 100 |
| | Copying a Configuration File from One Routing Engine to the Other | 101 |
| | Loading a Software Package from the Other Routing Engine | 102 |

| | | |
|------------------|--|------------|
| Chapter 5 | High Availability and Graceful Routing Engine Switchover (GRES) | 103 |
| | Understanding Graceful Routing Engine Switchover | 103 |
| | Graceful Routing Engine Switchover Concepts | 103 |
| | Effects of a Routing Engine Switchover | 108 |
| | Graceful Routing Engine Switchover System Requirements | 109 |
| | Graceful Routing Engine Switchover Platform Support | 110 |
| | Graceful Routing Engine Switchover Feature Support | 111 |
| | Graceful Routing Engine Switchover DPC Support | 112 |
| | Graceful Routing Engine Switchover and Subscriber Access | 112 |
| | Graceful Routing Engine Switchover PIC Support | 112 |
| | Requirements for Routers with a Backup Router Configuration | 113 |
| | Configuring Graceful Routing Engine Switchover | 114 |
| | Enabling Graceful Routing Engine Switchover | 114 |
| | Configuring Graceful Routing Engine Switchover with Graceful Restart | 114 |
| | Synchronizing the Routing Engine Configuration | 115 |
| | Verifying Graceful Routing Engine Switchover Operation | 116 |
| | Preventing Graceful Routing Engine Switchover in the Case of Slow Disks | 117 |
| | Resetting Local Statistics | 117 |
| Chapter 6 | High Availability and Nonstop Bridging | 119 |
| | Nonstop Bridging Concepts | 119 |
| | Nonstop Bridging System Requirements | 121 |
| | Platform Support | 121 |
| | Protocol Support | 122 |
| | Configuring Nonstop Bridging | 122 |
| | Enabling Nonstop Bridging | 122 |
| | Synchronizing the Routing Engine Configuration | 123 |
| | Verifying Nonstop Bridging Operation | 123 |
| Chapter 7 | High Availability and Nonstop Active Routing (NSR) | 125 |
| | Nonstop Active Routing Concepts | 125 |
| | Nonstop Active Routing System Requirements | 129 |
| | Nonstop Active Routing Platform and Switching Platform Support | 129 |
| | Nonstop Active Routing Protocol and Feature Support | 131 |
| | Nonstop Active Routing BFD Support | 134 |
| | Nonstop Active Routing BGP Support | 135 |
| | Nonstop Active Routing Layer 2 Circuit and VPLS Support | 136 |
| | Nonstop Active Routing PIM Support | 136 |
| | Nonstop Active Routing MSDP Support | 138 |
| | Nonstop Active Routing Support for RSVP-TE LSPs | 139 |
| | Configuring Nonstop Active Routing | 141 |
| | Enabling Nonstop Active Routing | 141 |
| | Synchronizing the Routing Engine Configuration | 142 |
| | Verifying Nonstop Active Routing Operation | 143 |
| | Preventing Automatic Reestablishment of BGP Peer Sessions After NSR Switchovers | 143 |
| | Example: Configuring Nonstop Active Routing | 144 |
| | Tracing Nonstop Active Routing Synchronization Events | 146 |
| | Resetting Local Statistics | 148 |

| | | |
|------------------|---|------------|
| Chapter 8 | High Availability and Graceful Restart | 149 |
| | Graceful Restart Concepts | 149 |
| | Graceful Restart System Requirements | 150 |
| | Enabling Graceful Restart | 151 |
| | Configuring Graceful Restart | 152 |
| | Graceful Restart for Aggregate and Static Routes | 177 |
| | Graceful Restart and Routing Protocols | 177 |
| | BGP | 178 |
| | IS-IS | 178 |
| | OSPF and OSPFv3 | 178 |
| | PIM Sparse Mode | 179 |
| | RIP and RIPng | 180 |
| | Configuring Routing Protocols Graceful Restart | 180 |
| | Enabling Graceful Restart | 180 |
| | Configuring Graceful Restart Options for BGP | 181 |
| | Configuring Graceful Restart Options for ES-IS | 182 |
| | Configuring Graceful Restart Options for IS-IS | 182 |
| | Configuring Graceful Restart Options for OSPF and OSPFv3 | 183 |
| | Configuring Graceful Restart Options for RIP and RIPng | 185 |
| | Configuring Graceful Restart Options for PIM Sparse Mode | 185 |
| | Tracking Graceful Restart Events | 186 |
| | Graceful Restart and MPLS-Related Protocols | 187 |
| | LDP | 187 |
| | RSVP | 187 |
| | CCC and TCC | 188 |
| | Configuring Graceful Restart for MPLS-Related Protocols | 188 |
| | Configuring Graceful Restart Globally | 188 |
| | Configuring Graceful Restart Options for RSVP, CCC, and TCC | 189 |
| | Configuring Graceful Restart Options for LDP | 189 |
| | Understanding Restart Signaling-Based Helper Mode Support for OSPF Graceful Restart | 190 |
| | Restart | 190 |
| | Example: Managing Helper Modes for OSPF Graceful Restart | 191 |
| | Tracing Restart Signaling-Based Helper Mode Events for OSPF Graceful Restart | 193 |
| | Restart | 193 |
| | Graceful Restart and Layer 2 and Layer 3 VPNs | 194 |
| | Configuring VPN Graceful Restart | 195 |
| | Configuring Graceful Restart Globally | 196 |
| | Configuring Graceful Restart for the Routing Instance | 196 |
| | Graceful Restart on Logical Systems | 197 |
| | Configuring Logical System Graceful Restart | 197 |
| | Enabling Graceful Restart Globally | 197 |
| | Configuring Graceful Restart for a Routing Instance | 198 |
| | Verifying Graceful Restart Operation | 198 |
| | Graceful Restart Operational Mode Commands | 199 |
| | Verifying BGP Graceful Restart | 199 |
| | Verifying IS-IS and OSPF Graceful Restart | 200 |
| | Verifying CCC and TCC Graceful Restart | 200 |

| | | |
|-------------------|---|------------|
| Chapter 9 | High Availability and Virtual Router Redundancy Protocol (VRRP) | 203 |
| | Understanding VRRP | 204 |
| | Junos OS Support for VRRPv3 | 206 |
| | Junos OS VRRP Support | 206 |
| | IPv6 VRRP Checksum Behavioral Differences | 207 |
| | VRRP Interoperability | 207 |
| | Upgrading from VRRPv2 to VRRPv3 | 208 |
| | Functionality of VRRPv3 Features | 210 |
| | VRRPv3 Authentication | 210 |
| | VRRPv3 Advertisement Intervals | 210 |
| | Unified ISSU for VRRPv3 | 210 |
| | VRRP failover-delay Overview | 211 |
| | When failover-delay Is Not Configured | 212 |
| | When failover-delay Is Configured | 213 |
| | Configuring Basic VRRP Support | 213 |
| | Configuring VRRP | 218 |
| | Configuring VRRP for IPv6 | 219 |
| | Configuring VRRP Authentication (IPv4 Only) | 220 |
| | Configuring the Advertisement Interval for the VRRP Master Router | 222 |
| | Modifying the Advertisement Interval in Seconds | 223 |
| | Modifying the Advertisement Interval in Milliseconds | 223 |
| | Configuring the Startup Period for VRRP Operations | 224 |
| | Configuring a Backup Router to Preempt the VRRP Master Router | 225 |
| | Modifying the Preemption Hold-Time Value for the VRRP Master Router | 225 |
| | Configuring the Asymmetric Hold Time for VRRP Routers | 226 |
| | Configuring Passive ARP Learning for Backup VRRP Routers | 226 |
| | Configuring VRRP Route Tracking | 227 |
| | Configuring a Logical Interface to Be Tracked for a VRRP Group | 229 |
| | Configuring a Route to Be Tracked for a VRRP Group | 231 |
| | Example: Configuring Multiple VRRP Owner Groups | 233 |
| | Configuring Inheritance for a VRRP Group | 239 |
| | Configuring an Interface to Accept All Packets Destined for the Virtual IP Address of a VRRP Group | 240 |
| | Configuring the Silent Period to Avoid Alarms Due to Delay in Receiving VRRP Advertisement Packets | 241 |
| | Enabling the Distributed Periodic Packet Management Process for VRRP | 242 |
| | Improving the Convergence Time for VRRP | 243 |
| | Configuring VRRP to Improve Convergence Time | 244 |
| | Tracing VRRP Operations | 246 |
| Chapter 10 | High Availability and Unified In-Service Software Upgrade (ISSU) | 249 |
| | Getting Started with Unified In-Service Software Upgrade | 249 |
| | Understanding the Unified ISSU Process | 250 |
| | Understanding the Unified ISSU Process on a Router | 250 |
| | Unified ISSU Process on a Router | 251 |
| | Understanding the Unified ISSU Process on the TX Matrix Router | 254 |
| | Unified ISSU Process on the TX Matrix Router | 255 |

| | |
|--|-----|
| Understanding the Unified ISSU Process on the TX Matrix Plus Router and on the TX Matrix Plus Router with 3D SIBs | 257 |
| Unified ISSU Process on the TX Matrix Plus Router and on the TX Matrix Plus Router with 3D SIBs | 257 |
| Unified ISSU System Requirements | 259 |
| General Unified ISSU Considerations for All Platforms | 260 |
| Unified ISSU Considerations for MX Series Routers | 261 |
| Unified ISSU Considerations for PTX Series Routers | 261 |
| Unified ISSU Considerations for M Series and T Series Routers | 262 |
| Unified ISSU Platform Support | 262 |
| Unified ISSU Protocol Support for M Series, MX Series, and T Series Routers and EX9200 Switches | 263 |
| Unified ISSU Feature Support | 264 |
| Unified ISSU PIC Support Considerations | 265 |
| PIC Considerations | 265 |
| SONET/SDH PICs | 266 |
| Fast Ethernet and Gigabit Ethernet PICs | 267 |
| Channelized PICs | 270 |
| Tunnel Services PICs | 271 |
| ATM PICs | 271 |
| Serial PICs | 272 |
| DS3, E1, E3, and T1 PICs | 272 |
| Enhanced IQ PICs | 273 |
| Enhanced IQ2 Ethernet Services Engine (ESE) PIC | 273 |
| Unified ISSU FPC Support on T4000 Routers | 274 |
| Unified ISSU Support on MX Series 3D Universal Edge Routers | 274 |
| Best Practices for Performing a Unified ISSU | 278 |
| Example: Performing a Unified ISSU | 279 |
| Verifying a Unified ISSU | 305 |
| Troubleshooting Unified ISSU Problems | 306 |
| Managing and Tracing BFD Sessions During Unified ISSU Procedures | 306 |

Part 3

Chapter 11

Configuration Statements and Operational Commands

High Availability Configuration Statements and Operational

| | |
|--|------------|
| Commands | 311 |
| Configuration Statements: Bidirectional Forwarding Detection | 311 |
| authentication (LAG) | 312 |
| bfd-liveness-detection (LAG) | 313 |
| detection-time (LAG) | 315 |
| traceoptions (Protocols BFD) | 316 |
| transmit-interval (LAG) | 318 |
| Configuration Statements: Graceful Routing Engine Switchover | 318 |
| graceful-switchover | 319 |
| Configuration Statements: Graceful Restart | 319 |
| disable | 320 |
| graceful-restart (Enabling Globally) | 321 |
| graceful-restart (Multicast Snooping) | 322 |
| helper-disable (Multiple Protocols) | 323 |

| | |
|--|-----|
| helper-disable (OSPF) | 324 |
| maximum-helper-recovery-time | 325 |
| maximum-helper-restart-time (RSVP) | 326 |
| maximum-neighbor-reconnect-time | 327 |
| maximum-neighbor-recovery-time | 328 |
| no-strict-lsa-checking | 329 |
| notify-duration | 330 |
| not-on-disk-underperform | 331 |
| reconnect-time | 332 |
| recovery-time | 333 |
| restart-duration | 334 |
| restart-time (BGP Graceful Restart) | 335 |
| stale-routes-time | 336 |
| traceoptions (Protocols) | 337 |
| Configuration Statements: Nonstop Active Routing | 338 |
| nonstop-routing | 339 |
| switchover-on-routing-crash | 340 |
| synchronize | 341 |
| traceoptions | 343 |
| Configuration Statements: Nonstop Bridging | 345 |
| nonstop-bridging | 345 |
| Configuration Statements: Routing Engine and Switching Control Board | |
| Redundancy | 345 |
| cfcb | 346 |
| description (Chassis Redundancy) | 347 |
| failover (Chassis) | 347 |
| failover (System Process) | 348 |
| feb (Creating a Redundancy Group) | 349 |
| feb (Assigning a FEB to a Redundancy Group) | 350 |
| keepalive-time | 351 |
| no-auto-failover | 352 |
| on-disk-failure (Chassis Redundancy Failover) | 352 |
| on-loss-of-keepalives | 353 |
| redundancy | 354 |
| redundancy-group | 355 |
| routing-engine (Chassis Redundancy) | 356 |
| sfm (Chassis Redundancy) | 357 |
| ssb | 358 |
| Configuration Statements: Unified ISSU | 358 |
| no-issu-timer-negotiation | 359 |
| traceoptions (Protocols BFD) | 360 |
| Configuration Statements: VRRP | 361 |
| accept-data | 363 |
| advertise-interval | 364 |
| asymmetric-hold-time | 365 |
| authentication-key | 366 |
| authentication-type | 367 |
| bandwidth-threshold | 368 |
| delegate-processing (VRRP) | 369 |

| | |
|---|-----|
| fast-interval | 370 |
| global-advertisements-threshold | 371 |
| hold-time (VRRP) | 372 |
| inherit-advertisement-interval | 373 |
| inet6-advertise-interval | 374 |
| interface | 375 |
| preempt (VRRP) | 376 |
| priority (Protocols VRRP) | 377 |
| priority-cost (VRRP) | 378 |
| priority-hold-time | 379 |
| route (Interfaces) | 380 |
| skew-timer-disable | 381 |
| startup-silent-period | 382 |
| traceoptions (Protocols VRRP) | 383 |
| track (VRRP) | 385 |
| version-3 | 386 |
| virtual-address | 387 |
| virtual-inet6-address | 388 |
| virtual-link-local-address | 389 |
| vrrp-group | 390 |
| vrrp-inet6-group | 392 |
| vrrp-inherit-from | 393 |
| Operational Commands | 393 |
| clear vrrp | 394 |
| request chassis ssb master switch | 395 |
| request system software in-service-upgrade | 397 |
| request system software in-service-upgrade (MX Series 3D Universal Edge Routers and EX9200 Switches) | 410 |
| request system software validate in-service-upgrade | 427 |
| show chassis ssb | 430 |
| show nonstop-routing | 432 |
| show pfe ssb | 435 |
| show system switchover | 441 |
| show task replication | 448 |
| show vrrp | 450 |

List of Figures

| | | |
|-------------------|---|------------|
| Part 2 | Configuring High Availability | |
| Chapter 2 | High Availability and Bidirectional Forwarding Detection (BFD) | 13 |
| | Figure 1: Customer Routes Connected to a Service Provider | 19 |
| | Figure 2: Typical Network with IBGP Sessions | 27 |
| | Figure 3: Configuring BFD for IS-IS | 43 |
| | Figure 4: RIP BFD Network Topology | 52 |
| | Figure 5: Configuring an Independent Micro BFD Session for LAG | 64 |
| Chapter 5 | High Availability and Graceful Routing Engine Switchover (GRES) | 103 |
| | Figure 6: Preparing for a Graceful Routing Engine Switchover | 106 |
| | Figure 7: Graceful Routing Engine Switchover Process | 107 |
| Chapter 6 | High Availability and Nonstop Bridging | 119 |
| | Figure 8: Nonstop Bridging Switchover Preparation Process | 120 |
| | Figure 9: Nonstop Bridging During a Switchover | 121 |
| Chapter 7 | High Availability and Nonstop Active Routing (NSR) | 125 |
| | Figure 10: Nonstop Active Routing Switchover Preparation Process | 127 |
| | Figure 11: Nonstop Active Routing During a Switchover | 128 |
| Chapter 8 | High Availability and Graceful Restart | 149 |
| | Figure 12: Layer 3 VPN Graceful Restart Topology | 153 |
| Chapter 9 | High Availability and Virtual Router Redundancy Protocol (VRRP) | 203 |
| | Figure 13: Basic VRRP | 205 |
| Chapter 10 | High Availability and Unified In-Service Software Upgrade (ISSU) | 249 |
| | Figure 14: Device Status Before Starting a Unified ISSU | 252 |
| | Figure 15: Device Status After the Backup Routing Engine Is Upgraded | 252 |
| | Figure 16: Device Status After One Packet Forwarding Engine Downloads the New Software | 253 |
| | Figure 17: Device Status Before the Routing Engine Switchover | 253 |
| | Figure 18: Device Status After the Routing Engine Switchover | 254 |
| | Figure 19: Device Status After the Unified ISSU Is Complete | 254 |
| | Figure 20: Unified ISSU Example Topology | 280 |

List of Tables

| | | |
|-------------------|--|------------|
| | About the Documentation | xv |
| | Table 1: Notice Icons | xvii |
| | Table 2: Text and Syntax Conventions | xviii |
| Part 2 | Configuring High Availability | |
| Chapter 2 | High Availability and Bidirectional Forwarding Detection (BFD) | 13 |
| | Table 3: Configuring BFD for IS-IS | 41 |
| Chapter 4 | High Availability and Routing Engine Redundancy | 89 |
| | Table 4: Routing Engine Mastership Log | 98 |
| Chapter 5 | High Availability and Graceful Routing Engine Switchover (GRES) | 103 |
| | Table 5: Effects of a Routing Engine Switchover | 108 |
| | Table 6: Graceful Routing Engine Switchover Feature Support | 111 |
| Chapter 7 | High Availability and Nonstop Active Routing (NSR) | 125 |
| | Table 7: Nonstop Active Routing Platform Support | 129 |
| | Table 8: Nonstop Active Routing Protocol and Feature Support | 131 |
| Chapter 9 | High Availability and Virtual Router Redundancy Protocol (VRRP) | 203 |
| | Table 9: VRRPv2 to VRRPv3 Transition Steps and Events | 209 |
| | Table 10: Interface State and Priority Cost Usage | 231 |
| Chapter 10 | High Availability and Unified In-Service Software Upgrade (ISSU) | 249 |
| | Table 11: Locating the Information You Need to Work With ISSU | 249 |
| | Table 12: Unified ISSU Support for Dual Routing Engine Platforms | 262 |
| | Table 13: Unified ISSU PIC Support: SONET/SDH | 266 |
| | Table 14: Unified ISSU PIC Support: Fast Ethernet and Gigabit Ethernet | 268 |
| | Table 15: Unified ISSU PIC Support: Channelized | 270 |
| | Table 16: Unified ISSU PIC Support: Tunnel Services | 271 |
| | Table 17: Unified ISSU PIC Support: ATM | 271 |
| | Table 18: Unified ISSU Support: Enhanced IQ2 Ethernet Services Engine (ESE) PIC | 273 |
| | Table 19: Unified ISSU Support: MX Series Router MPCs | 275 |
| | Table 20: Unified ISSU Support: MX Series Router MICs | 276 |
| | Table 21: Routing Engine Status Before Upgrading | 285 |
| | Table 22: Routing Engine Status After Upgrading and Rebooting Both Routing Engines | 288 |
| | Table 23: Routing Engine Status After Upgrading, Rebooting, and Switching Mastership | 289 |
| | Table 24: Routing Engine Status Before Upgrading and Manually Rebooting the Backup Routing Engine | 292 |

| | |
|--|-----|
| Table 25: Routing Engine Status After Upgrading and Before Manually Rebooting the Backup Routing Engine | 294 |
| Table 26: Routing Engine Status After Upgrading, Manually Rebooting, and Before Switching Mastership | 296 |
| Table 27: Routing Engine Status After Upgrading, Manually Rebooting, and Switching Mastership | 297 |
| Table 28: Routing Engine Status Before Upgrading and Rebooting One Routing Engine | 298 |
| Table 29: Routing Engine Status After Upgrading One Routing Engine and Before Upgrading the Other Routing Engine | 301 |
| Table 30: Routing Engine Status After Upgrading, Manually Rebooting, and Switching Mastership | 304 |

Part 3

Chapter 11

Configuration Statements and Operational Commands

| | |
|--|------------|
| High Availability Configuration Statements and Operational Commands | 311 |
| Table 31: show chassis ssb Output Fields | 430 |
| Table 32: show nonstop-routing Output Fields | 432 |
| Table 33: show pfe ssb Output Fields | 435 |
| Table 34: show system switchover Output Fields | 443 |
| Table 35: show task replication Output Fields | 448 |
| Table 36: show vrrp Output Fields | 451 |

About the Documentation

- [Documentation and Release Notes on page xv](#)
- [Supported Platforms on page xv](#)
- [Using the Examples in This Manual on page xv](#)
- [Documentation Conventions on page xvii](#)
- [Documentation Feedback on page xix](#)
- [Requesting Technical Support on page xix](#)

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 <http://www.juniper.net/techpubs/>.

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 <http://www.juniper.net/books>.

Supported Platforms

For the features described in this document, the following platforms are supported:

- [M Series](#)
- [T Series](#)
- [MX Series](#)
- [TX Matrix](#)
- [PTX Series](#)

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.xsl;
    }
  }
}
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 xvii](#) 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 xviii](#) 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 metric>; |
| (pipe symbol) | Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity. | broadcast multicast (string1 string2 string3) |
| # (pound sign) | Indicates a comment specified on the same line as the configuration statement to which it applies. | rsvp { # Required for dynamic MPLS only |
| [] (square brackets) | Encloses a variable for which you can substitute one or more values. | community name members [community-ids] |
| 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, comments, and suggestions so that we can improve the documentation. You can provide feedback by using either of the following methods:

- Online feedback rating system—On any page of the Juniper Networks TechLibrary site at <http://www.juniper.net/techpubs/index.html>, simply click the stars to rate the content, and use the pop-up form to provide us with information about your experience. Alternately, you can use the online feedback form at <http://www.juniper.net/techpubs/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 J-Care or Partner Support Service 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 <http://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf>.
- Product warranties—For product warranty information, visit <http://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: <http://www.juniper.net/customers/support/>
- Search for known bugs: <http://www2.juniper.net/kb/>
- Find product documentation: <http://www.juniper.net/techpubs/>
- Find solutions and answer questions using our Knowledge Base: <http://kb.juniper.net/>
- Download the latest versions of software and review release notes:
<http://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications:
<http://kb.juniper.net/InfoCenter/>
- Join and participate in the Juniper Networks Community Forum:
<http://www.juniper.net/company/communities/>
- Open a case online in the CSC Case Management tool: <http://www.juniper.net/cm/>

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

Opening a Case with JTAC

You can open a case with JTAC on the Web or by telephone.

- Use the Case Management tool in the CSC at <http://www.juniper.net/cm/>.
- 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 <http://www.juniper.net/support/requesting-support.html>.

PART 1

Overview

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

CHAPTER 1

High Availability Overview

Supported Platforms [M Series](#), [MX Series](#), [PTX Series](#), [T Series](#)

- [Understanding High Availability Features on Juniper Networks Routers on page 3](#)
- [High Availability-Related Features in Junos OS on page 8](#)

Understanding High Availability Features on Juniper Networks Routers

Supported Platforms [M Series](#), [MX Series](#), [PTX Series](#), [T Series](#)

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 103](#) and [“Nonstop Active Routing Concepts” on page 125](#).



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 3D Universal Edge Router 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)

For more information, see [“Nonstop Bridging Concepts” on page 119](#).

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 129](#).

For more information about nonstop active routing, see [“Nonstop Active Routing Concepts” on page 125](#).

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 149](#).

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 149](#).

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 125](#).

Effects of a Routing Engine Switchover

[“Effects of a Routing Engine Switchover” on page 103](#) 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 204](#).

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 249.

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 3D Universal Edge Routers, 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 3D Universal Edge Routers 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 8](#)

High Availability-Related Features in Junos OS

Supported Platforms [M Series](#), [PTX Series](#), [T Series](#)

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](#)

**Related
Documentation**

-

PART 2

Configuring High Availability

- [High Availability and Bidirectional Forwarding Detection \(BFD\) on page 13](#)
- [High Availability and Switching Control Board Redundancy on page 79](#)
- [High Availability and Routing Engine Redundancy on page 89](#)
- [High Availability and Graceful Routing Engine Switchover \(GRES\) on page 103](#)
- [High Availability and Nonstop Bridging on page 119](#)
- [High Availability and Nonstop Active Routing \(NSR\) on page 125](#)
- [High Availability and Graceful Restart on page 149](#)
- [High Availability and Virtual Router Redundancy Protocol \(VRRP\) on page 203](#)
- [High Availability and Unified In-Service Software Upgrade \(ISSU\) on page 249](#)

CHAPTER 2

High Availability and Bidirectional Forwarding Detection (BFD)

- [Understanding BFD for Static Routes for Faster Network Failure Detection on page 13](#)
- [Example: Configuring BFD for Static Routes for Faster Network Failure Detection on page 18](#)
- [Understanding BFD for BGP on page 24](#)
- [Example: Configuring BFD on Internal BGP Peer Sessions on page 25](#)
- [Understanding BFD for OSPF on page 34](#)
- [Example: Configuring BFD for OSPF on page 36](#)
- [Understanding BFD for IS-IS on page 40](#)
- [Example: Configuring BFD for IS-IS on page 43](#)
- [Understanding BFD for RIP on page 48](#)
- [Example: Configuring BFD for RIP on page 49](#)
- [Understanding Independent Micro BFD Sessions for LAG on page 55](#)
- [Configuring Independent Micro BFD Sessions for LAG on page 58](#)
- [Example: Configuring Independent Micro BFD Sessions for LAG on page 63](#)
- [Configuring BFD for PIM on page 73](#)
- [Understanding Distributed BFD on page 74](#)
- [Understanding Static Route State When BFD is in Admin Down State on page 77](#)

Understanding BFD for Static Routes for Faster Network Failure Detection

Supported Platforms [ACX Series, EX Series, M Series, MX Series, PTX Series, QFabric System, QFX Series, SRX Series, T Series](#)

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 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.

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

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



NOTE: Starting with Junos OS Release 15.1X49-D70, 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.

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 given in the table below:

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

| Type of BFD session | Description |
|---|---|
| Inline BFD | BFD sessions running on the FPC hardware. |
| <p>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.</p> <p>NOTE: Starting in Junos OS Release 16.1R1, the inline BFD sessions are supported on integrated routing and bridging (IRB) interfaces.</p> | |

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.



NOTE: SRX Series devices do not support distributed BFD.

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. |
| 15.1F6 | 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

- [Example: Configuring BFD for Static Routes for Faster Network Failure Detection on page 18](#)
- [Example: Enabling BFD on Qualified Next Hops in Static Routes for Route Selection](#)

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

Supported Platforms [ACX Series, EX4600, M Series, MX Series, PTX Series, QFabric System, SRX Series, T Series](#)

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

- [Requirements on page 18](#)
- [Overview on page 18](#)
- [Configuration on page 19](#)
- [Verification on page 22](#)

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 19 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 172.16.1/24;
    }
  }
}
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 22](#)
- [Viewing Detailed BFD Events on page 23](#)

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: pcmd_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 13](#)

Understanding BFD for BGP

Supported Platforms [ACX Series](#), [EX4600](#), [M Series](#), [MX Series](#), [PTX Series](#), [QFabric System](#), [QFX Series](#), [T Series](#)

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.

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: 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.)

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 |
|---------|--|
| 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

- [Example: Configuring BFD on Internal BGP Peer Sessions on page 25](#)

Example: Configuring BFD on Internal BGP Peer Sessions

Supported Platforms [ACX Series](#), [EX4600](#), [M Series](#), [MX Series](#), [PTX Series](#), [QFabric System](#), [QFX Series](#), [T Series](#)

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 25](#)
- [Overview on page 25](#)
- [Configuration on page 27](#)
- [Verification on page 31](#)

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 ms for Routing Engine-based sessions and less than 10 ms 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 seconds (5*1000 seconds) 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 6 seconds 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 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.

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 27](#) 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 10.10.10.1/30

[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 31](#)
- [Verifying That BFD Sessions Are Up on page 31](#)
- [Viewing Detailed BFD Events on page 32](#)
- [Viewing Detailed BFD Events After Deactivating and Reactivating a Loopback Interface on page 33](#)

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
```

| Address | State | Interface | Detect Time | Transmit Interval | Multiplier |
|--------------|-------|-----------|-------------|-------------------|------------|
| 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 • *Understanding BFD Authentication for BGP*

Understanding BFD for OSPF

Supported Platforms ACX Series, EX4600, M Series, MX Series, NFX Series, PTX Series, QFabric System, QFX Series, T Series

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: 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 36](#)
- [bfd-liveness-detection](#)

Example: Configuring BFD for OSPF

Supported Platforms [ACX Series, EX4600, M Series, MX Series, NFX Series, PTX Series, QFX Series, T Series, vSRX](#)

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

- [Requirements on page 36](#)
- [Overview on page 37](#)
- [Configuration on page 38](#)
- [Verification on page 40](#)

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 300 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 300 ms for distributed BFD sessions.



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 | <p>The output displays information about the BFD sessions.</p> <ul style="list-style-type: none">• 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 | <ul style="list-style-type: none">• Understanding BFD for OSPF on page 34• Understanding BFD Authentication for OSPF• Example: Configuring BFD Authentication for OSPF |

Understanding BFD for IS-IS

Supported Platforms [ACX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [SRX Series](#), [T Series](#)

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 15.2, 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.

To detect failures in the network, the set of statements in [Table 3 on page 41](#) 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. |

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

| Statement | Description |
|--|--|
| minimum-receive-interval <i>milliseconds</i> | Specify only the minimum receive interval for failure detection. 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. |
| multiplier number | Specify the number of hello packets not received by the neighbor that causes the originating interface to be declared down. The default is 3, and you can configure a value from 1 through 225. |
| no-adaptation | Disable BFD adaptation. In Junos OS Release 9.0 and later, you can specify that the BFD sessions not adapt to changing network conditions. NOTE: We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network. |
| threshold | Specify the threshold for the following: <ul style="list-style-type: none"> Adaptation of the detection time 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. Transmit interval NOTE: The threshold value must be greater than the minimum transmit interval multiplied by the multiplier number. |
| transmit-interval minimum-interval | 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. |
| version | 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 43](#)
- [Understanding BFD Authentication for IS-IS](#)

Example: Configuring BFD for IS-IS

Supported Platforms EX4600, M Series, PTX Series, QFX Series, SRX Series, T Series

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 43](#)
- [Overview on page 43](#)
- [Configuration on page 43](#)
- [Verification on page 46](#)

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 43](#) 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 46](#)
- [Verifying That IS-IS Is Configured on page 47](#)
- [Verifying That BFD Is configured on page 48](#)

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 Time | Transmit Interval | Multiplier |
|----------|-------|-----------|-------------|-------------------|------------|
| 10.0.0.2 | Up | so-0/0/0 | 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 40](#)

Understanding BFD for RIP

Supported Platforms ACX Series, EX4600, M Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, T Series

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. |

Related Documentation

- [Example: Configuring BFD for RIP on page 49](#)

Example: Configuring BFD for RIP

Supported Platforms

ACX Series, EX Series, M Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, T Series

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

- [Requirements on page 50](#)
- [Overview on page 50](#)

- [Configuration on page 52](#)
- [Verification on page 54](#)

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 255,000 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 52 shows the topology used in this example.

Figure 4: RIP BFD Network Topology



"CLI Quick Configuration" on page 52 shows the configuration for all of the devices in Figure 4 on page 52. The section "Step-by-Step Procedure" on page 53 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 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 </pre> |
| Device R2 | <pre> 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 </pre> |
| Device R3 | <pre> 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 </pre> |


```

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 54](#)
- [Checking the BFD Trace File on page 55](#)

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.

Related Documentation

- [Understanding BFD for RIP on page 48](#)

Understanding Independent Micro BFD Sessions for LAG

Supported Platforms [M Series](#), [MX Series](#), [PTX Series](#), [QFX Series](#), [T Series](#)

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-SFPP (Type 4 FPC)
- 24x10GE (LAN/WAN) SFPP, 12x10GE (LAN/WAN) SFPP, 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.

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.

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 with the AE interface address of the remote destination as the neighbor address.



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 |
|---------|--|
| 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 312](#)
- [bfd-liveness-detection on page 313](#)
- [detection-time on page 315](#)
- [transmit-interval on page 318](#)
- [Configuring Independent Micro BFD Sessions for LAG on page 58](#)
- [Example: Configuring Independent Micro BFD Sessions for LAG on page 63](#)

Configuring Independent Micro BFD Sessions for LAG

Supported Platforms [M Series](#), [MX Series](#), [PTX Series](#), [QFX Series](#), [T Series](#)

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;  
  }  
}
```

```

        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 312](#)
- [bfd-liveness-detection on page 313](#)
- [detection-time on page 315](#)
- [Example: Configuring Independent Micro BFD Sessions for LAG on page 63](#)
- [Understanding Independent Micro BFD Sessions for LAG on page 55](#)

Example: Configuring Independent Micro BFD Sessions for LAG

Supported Platforms [M Series](#), [MX Series](#), [PTX Series](#), [T Series](#)

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

- [Requirements on page 63](#)
- [Overview on page 64](#)
- [Configuration on page 64](#)
- [Verification on page 70](#)

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-SFPP (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 64 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::1:1/126 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::1:1/126
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::1/126;
    }
  }
}
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 70](#)
- [Viewing Detailed BFD Events on page 72](#)

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 | 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 Time | Transmit Interval | Multiplier |
|-----------------------|-------|-----------|-------------|-------------------|------------|
| 201:DB8:251::bb:bb:13 | | UP | xe-4/1/0 | 9.000 | 3.000 |

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 RtblIdx (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 312](#)
 - [bfd-liveness-detection on page 313](#)
 - [detection-time on page 315](#)
 - [Configuring Independent Micro BFD Sessions for LAG on page 58](#)
 - [Understanding Independent Micro BFD Sessions for LAG on page 55](#)

Configuring BFD for PIM

Supported Platforms EX4600, M Series, MX Series, NFX Series, OCX1100, PTX Series, QFX Series, SRX Series, T Series

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](#)

Understanding Distributed BFD

Supported Platforms [M Series](#), [MX Series](#), [PTX Series](#), [T Series](#)

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.

- **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 fate 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.

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
- VCCV BFD (Layer 2 circuit, Layer 3 VPN, and VPLS) (MPLS)



NOTE: Starting in 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, distribution of adjacency entry, that is, the IP addresses of adjacent routers, and transmit entry, or the IP address of the transmitting router, for a BFD session is asymmetric. This is because an adjacency entry that requires rules may or may 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 |
|---------|--|
| 13.3R5 | Starting in 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, distribution of adjacency entry, that is, the IP addresses of adjacent routers, and transmit entry, or the IP address of the transmitting router, for a BFD session is asymmetric. |

Related Documentation

- [show bfd session](#)
- [Understanding BFD for RIP on page 48](#)
- [Understanding BFD for Static Routes for Faster Network Failure Detection on page 13](#)
- [Understanding BFD for BGP on page 24](#)

- [Understanding BFD for IS-IS on page 40](#)
- [Understanding BFD for OSPF on page 34](#)
- [Understanding EBGP Multihop](#)

Understanding Static Route State When BFD is in Admin Down State

Supported Platforms [ACX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [SRX Series](#), [T Series](#)

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 13](#)
- [Example: Configuring BFD for Static Routes for Faster Network Failure Detection on page 18](#)

CHAPTER 3

High Availability and Switching Control Board Redundancy

Supported Platforms [M Series](#)

- [Understanding Switching Control Board Redundancy on page 79](#)
- [Configuring CFEB Redundancy on the M10i Router on page 83](#)
- [Configuring FEB Redundancy on the M120 Router on page 84](#)
- [Example: Configuring FEB Redundancy on M120 Routers on page 85](#)
- [Configuring SFM Redundancy on M40e and M160 Routers on page 86](#)
- [Configuring SSB Redundancy on the M20 Router on page 87](#)

Understanding Switching Control Board Redundancy

Supported Platforms [M Series](#)

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 80](#)
- [Redundant FEBs on the M120 Router on page 80](#)
- [Redundant SSBs on the M20 Router on page 82](#)
- [Redundant SFMs on the M40e and M160 Routers on page 83](#)

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 89](#)
- [Configuring CFEB Redundancy on the M10i Router on page 83](#)
- [Configuring FEB Redundancy on the M120 Router on page 84](#)
- [Configuring SFM Redundancy on M40e and M160 Routers on page 86](#)
- [Configuring SSB Redundancy on the M20 Router on page 87](#)
- *show chassis redundancy feb*
- *request chassis cb*

Configuring CFEB Redundancy on the M10i Router

Supported Platforms M10i

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 79](#)

Configuring FEB Redundancy on the M120 Router

Supported Platforms **M120**

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 79](#)
- [Example: Configuring FEB Redundancy on M120 Routers on page 85](#)

Example: Configuring FEB Redundancy on M120 Routers

Supported Platforms **M120**

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 79](#)
- [Configuring FEB Redundancy on the M120 Router on page 84](#)

Configuring SFM Redundancy on M40e and M160 Routers

Supported Platforms **M40e**

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 79](#)

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 79](#)

Related Documentation

- [Understanding High Availability Features on Juniper Networks Routers on page 3](#)
- [Understanding Routing Engine Redundancy on Juniper Networks Routers on page 89](#)
- *show chassis redundancy feb*
- *request chassis cb*

CHAPTER 4

High Availability and Routing Engine Redundancy

Supported Platforms [M Series, PTX Series, T Series](#)

- [Understanding Routing Engine Redundancy on Juniper Networks Routers on page 89](#)
- [Configuring Routing Engine Redundancy on page 95](#)
- [Initial Routing Engine Configuration Example on page 100](#)
- [Copying a Configuration File from One Routing Engine to the Other on page 101](#)
- [Loading a Software Package from the Other Routing Engine on page 102](#)

Understanding Routing Engine Redundancy on Juniper Networks Routers

Supported Platforms [M Series, PTX Series, T Series](#)

This topic contains the following sections:

- [Routing Engine Redundancy Overview on page 89](#)
- [Conditions That Trigger a Routing Engine Failover on page 90](#)
- [Default Routing Engine Redundancy Behavior on page 91](#)
- [Routing Engine Redundancy on a TX Matrix Router on page 92](#)
- [Routing Engine Redundancy on a TX Matrix Plus Router on page 93](#)
- [Situations That Require You to Halt Routing Engines on page 94](#)

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 103](#).
- 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 125](#).
- 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 149](#).

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 79](#)
- [Configuring Routing Engine Redundancy on page 95](#)

Configuring Routing Engine Redundancy

Supported Platforms [M Series](#), [PTX Series](#), [T Series](#)

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 95](#)
- [Configuring Automatic Failover to the Backup Routing Engine on page 96](#)
- [Manually Switching Routing Engine Mastership on page 98](#)
- [Verifying Routing Engine Redundancy Status on page 98](#)

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, you must modify the configuration and then activate it by issuing the **commit synchronize** command.

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 96](#)
- [On Detection of a Hard Disk Error on the Master Routing Engine on page 96](#)
- [On Detection of a Loss of Keepalive Signal from the Master Routing Engine on page 96](#)
- [On Detection of the em0 Interface Failure on the Master Routing Engine on page 98](#)
- [When a Software Process Fails on page 98](#)

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 114](#).

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 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 4 on page 98](#) lists the mastership log event codes and descriptions.

Table 4: 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. |

Table 4: Routing Engine Mastership Log (*continued*)

| Event Code | Description |
|------------|---|
| 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. |
| 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 89](#)
- [Understanding Switching Control Board Redundancy on page 79](#)

Initial Routing Engine Configuration Example

Supported Platforms [M Series](#), [PTX Series](#), [T Series](#)

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 *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 89](#)
- [Understanding Switching Control Board Redundancy on page 79](#)

Copying a Configuration File from One Routing Engine to the Other

Supported Platforms [M Series](#), [PTX Series](#), [T Series](#)

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 89](#)
- [Understanding Switching Control Board Redundancy on page 79](#)
- [Loading a Software Package from the Other Routing Engine on page 102](#)

Loading a Software Package from the Other Routing Engine

Supported Platforms [M Series, PTX Series, T Series](#)

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 89](#)
- [Understanding Switching Control Board Redundancy on page 79](#)
- [Copying a Configuration File from One Routing Engine to the Other on page 101](#)

Related Documentation

- [Understanding High Availability Features on Juniper Networks Routers on page 3](#)
- [Understanding Switching Control Board Redundancy on page 79](#)

CHAPTER 5

High Availability and Graceful Routing Engine Switchover (GRES)

- [Understanding Graceful Routing Engine Switchover on page 103](#)
- [Graceful Routing Engine Switchover System Requirements on page 109](#)
- [Requirements for Routers with a Backup Router Configuration on page 113](#)
- [Configuring Graceful Routing Engine Switchover on page 114](#)
- [Preventing Graceful Routing Engine Switchover in the Case of Slow Disks on page 117](#)
- [Resetting Local Statistics on page 117](#)

Understanding Graceful Routing Engine Switchover

Supported Platforms [EX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [QFX Series](#), [T Series](#)

This topic contains the following sections:

- [Graceful Routing Engine Switchover Concepts on page 103](#)
- [Effects of a Routing Engine Switchover on page 108](#)

Graceful Routing Engine Switchover Concepts

The graceful Routing Engine switchover (GRES) feature in Junos OS enables a routing platform 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

Any updates to the master Routing Engine are replicated to the backup Routing Engine as soon as they occur.



NOTE: Due to 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 149](#). For more information about nonstop active routing, see [“Nonstop Active Routing Concepts” on page 125](#).

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: takes 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 kept, 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 106 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

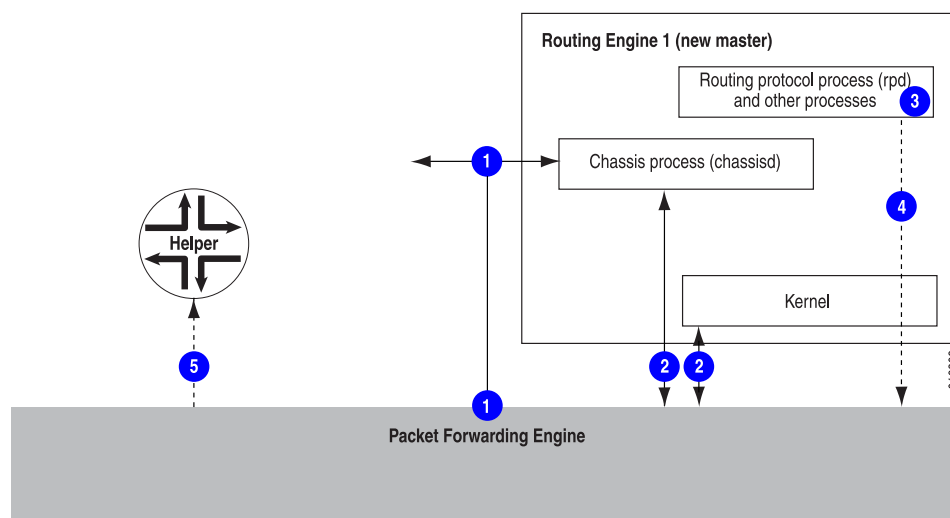
The switchover preparation process for GRES is as follows:

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 107 shows the effects of a switchover on the routing (or switching) platform.

Figure 7: Graceful Routing Engine Switchover Process



When a switchover occurs, the switchover process is as follows:

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: 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, we 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 5 on page 108 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 5: 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 nonstop active routing 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 5: 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. |

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 109](#)
- [Configuring Graceful Routing Engine Switchover on page 114](#)
- [Configuring Graceful Routing Engine Switchover in a Virtual Chassis \(CLI Procedure\)](#)
- [Configuring Graceful Routing Engine Switchover in a Virtual Chassis \(CLI Procedure\)](#)
- [Requirements for Routers with a Backup Router Configuration on page 113](#)
- [Example: Configuring IS-IS for GRES with Graceful Restart](#)
- [hold-time](#)

Graceful Routing Engine Switchover System Requirements

Supported Platforms [EX Series](#), [M Series](#), [MX240](#), [MX480](#), [MX960](#), [PTX Series](#), [QFX Series](#), [T Series](#)

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 110](#)
- [Graceful Routing Engine Switchover Feature Support on page 111](#)
- [Graceful Routing Engine Switchover DPC Support on page 112](#)
- [Graceful Routing Engine Switchover and Subscriber Access on page 112](#)
- [Graceful Routing Engine Switchover PIC Support on page 112](#)

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 6 on page 111](#).

Table 6: 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 |
| 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 3D Universal Edge Routers running the appropriate version of Junos OS as shown in “[Graceful Routing Engine Switchover Platform Support](#)” on page 110. 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.

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 103](#)
- [Configuring Graceful Routing Engine Switchover on page 114](#)
- [Configuring Graceful Routing Engine Switchover in a Virtual Chassis \(CLI Procedure\)](#)
- [Requirements for Routers with a Backup Router Configuration on page 113](#)

Requirements for Routers with a Backup Router Configuration

Supported Platforms [EX Series](#), [M Series](#), [PTX Series](#), [T Series](#)

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 103](#)
- [Graceful Routing Engine Switchover System Requirements on page 109](#)

Configuring Graceful Routing Engine Switchover

Supported Platforms EX Series, M Series, MX Series, PTX Series, QFX Series, T Series

This section contains the following topics:

- [Enabling Graceful Routing Engine Switchover on page 114](#)
- [Configuring Graceful Routing Engine Switchover with Graceful Restart on page 114](#)
- [Synchronizing the Routing Engine Configuration on page 115](#)
- [Verifying Graceful Routing Engine Switchover Operation on page 116](#)

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.

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 <code>CHASSISD_SNMP_TRAP7</code> statement at the [edit event-options policy <i>policy-name</i>] hierarchy level. |

Related Documentation

- [Understanding Graceful Routing Engine Switchover on page 103](#)
- [Graceful Routing Engine Switchover System Requirements on page 109](#)
- [Requirements for Routers with a Backup Router Configuration on page 113](#)
- [Resetting Local Statistics on page 117](#)
- [graceful-switchover on page 319](#)
- [graceful-switchover](#)
- [Example: Configuring IS-IS for GRES with Graceful Restart](#)

- *hold-time*

Preventing Graceful Routing Engine Switchover in the Case of Slow Disks

Supported Platforms [EX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [T Series](#)

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 331](#)
- [Understanding Graceful Routing Engine Switchover on page 103](#)

Resetting Local Statistics

Supported Platforms [M Series](#), [T Series](#)

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 103](#)
- [Configuring Graceful Routing Engine Switchover on page 114](#)

CHAPTER 6

High Availability and Nonstop Bridging

- [Nonstop Bridging Concepts on page 119](#)
- [Nonstop Bridging System Requirements on page 121](#)
- [Configuring Nonstop Bridging on page 122](#)

Nonstop Bridging Concepts


Supported Platforms [EX Series, MX Series, QFX Series](#)

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 103](#).

[Figure 8 on page 120](#) shows the system architecture of nonstop bridging and the process a routing (or switching) platform follows to prepare for a switchover.

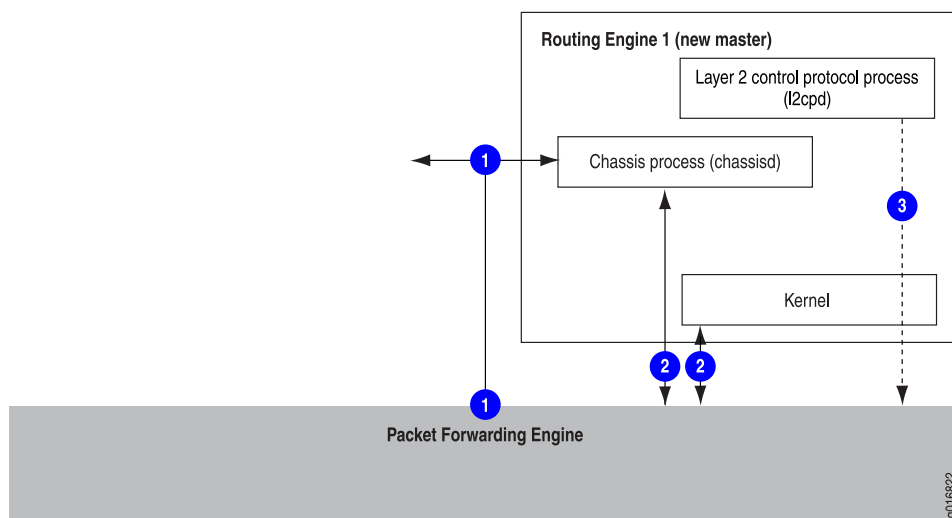
Figure 8: Nonstop Bridging Switchover Preparation Process

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 9 on page 121 shows the effects of a switchover on the routing platform.

Figure 9: 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 121](#)
- [Configuring Nonstop Bridging on page 122](#)
- [Configuring Nonstop Bridging on Switches \(CLI Procedure\)](#)

Nonstop Bridging System Requirements

Supported Platforms [EX Series, MX Series](#)

This topic contains the following sections:

- [Platform Support on page 121](#)
- [Protocol Support on page 122](#)

Platform Support

Nonstop bridging is supported on MX Series 3D Universal Edge Routers. 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 119](#)
- [Configuring Nonstop Bridging on page 122](#)
- *Configuring Nonstop Bridging on Switches (CLI Procedure)*

Configuring Nonstop Bridging

Supported Platforms [EX Series, MX Series](#)

This section includes the following topics:

- [Enabling Nonstop Bridging on page 122](#)
- [Synchronizing the Routing Engine Configuration on page 123](#)
- [Verifying Nonstop Bridging Operation on page 123](#)

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 119](#)
- [Nonstop Bridging System Requirements on page 121](#)
- [nonstop-bridging on page 345](#)
- [Configuring Nonstop Bridging on EX Series Switches \(CLI Procedure\)](#)

CHAPTER 7

High Availability and Nonstop Active Routing (NSR)

- [Nonstop Active Routing Concepts on page 125](#)
- [Nonstop Active Routing System Requirements on page 129](#)
- [Configuring Nonstop Active Routing on page 141](#)
- [Preventing Automatic Reestablishment of BGP Peer Sessions After NSR Switchovers on page 143](#)
- [Example: Configuring Nonstop Active Routing on page 144](#)
- [Tracing Nonstop Active Routing Synchronization Events on page 146](#)
- [Resetting Local Statistics on page 148](#)

Nonstop Active Routing Concepts

Supported Platforms [EX Series](#), [M Series](#), [PTX Series](#), [QFX Series](#), [T Series](#)

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 103](#).



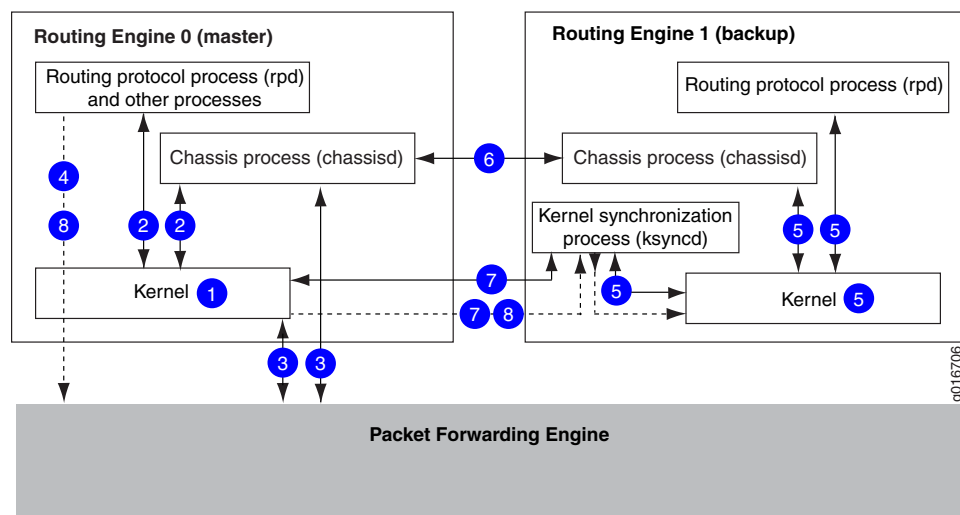
.....

NOTE: Starting with Junos OS Release 12.3, due to its synchronization requirements and logic, NSR/GRES performance is limited by the slowest Routing Engine in the system.

.....

Figure 10 on page 127 shows the system architecture of nonstop active routing and the process a routing (or switching) platform follows to prepare for a switchover.

Figure 10: Nonstop Active Routing Switchover Preparation Process

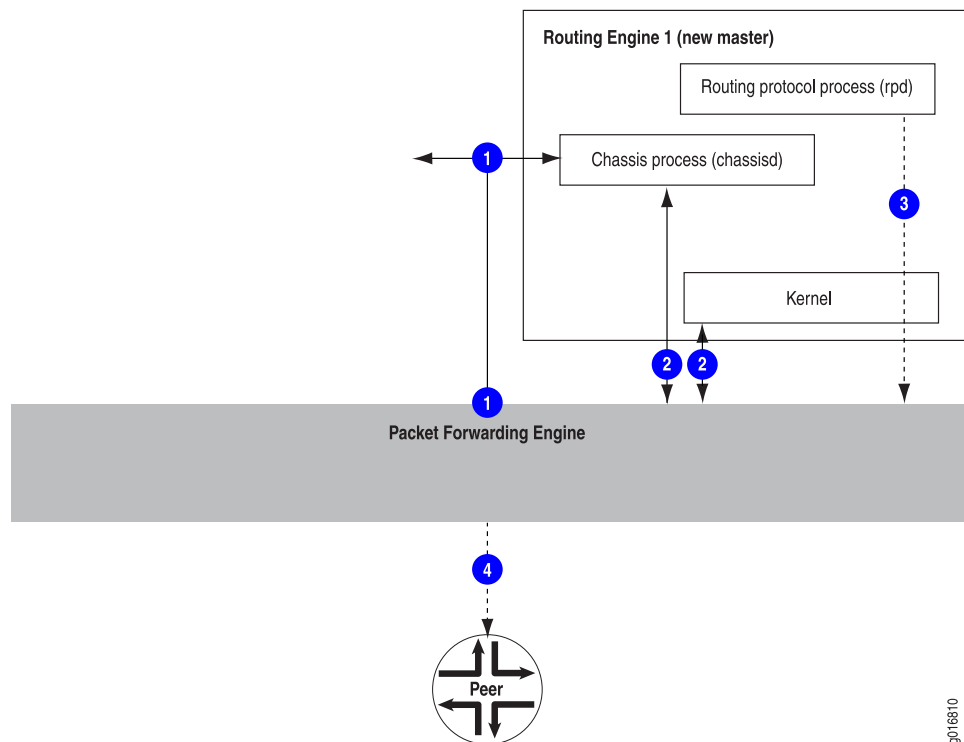


The switchover preparation process for NSR follows these steps:

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 graceful Routing Engine switchover and nonstop active routing 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 11 on page 128 shows the effects of a switchover on the routing platform.

Figure 11: Nonstop Active Routing 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 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, due to its synchronization requirements and logic, NSR/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 129](#)
- [Configuring Nonstop Active Routing on page 141](#)
- [Configuring Nonstop Active Routing on Switches](#)

Nonstop Active Routing System Requirements

Supported Platforms [EX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [QFX Series](#), [T Series](#)

This section contains the following topics:

- [Nonstop Active Routing Platform and Switching Platform Support on page 129](#)
- [Nonstop Active Routing Protocol and Feature Support on page 131](#)
- [Nonstop Active Routing BFD Support on page 134](#)
- [Nonstop Active Routing BGP Support on page 135](#)
- [Nonstop Active Routing Layer 2 Circuit and VPLS Support on page 136](#)
- [Nonstop Active Routing PIM Support on page 136](#)
- [Nonstop Active Routing MSDP Support on page 138](#)
- [Nonstop Active Routing Support for RSVP-TE LSPs on page 139](#)

Nonstop Active Routing Platform and Switching Platform Support

[Table 7 on page 129](#) lists the platforms that support nonstop active routing (NSR).

Table 7: 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 |

Table 7: Nonstop Active Routing Platform Support (*continued*)

| Platform | Junos OS Release |
|---|------------------|
| M320 router | 8.4 or later |
| MX Series routers | 9.0 or later |
| 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 | |
| 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 stitching) • Layer 3 VPNs • LDP • RSVP | |
| 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 stitching) • 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 |

Table 7: Nonstop Active Routing Platform Support (*continued*)

| Platform | Junos OS Release |
|--|--|
| 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 8 on page 131 lists the protocols that are supported by nonstop active routing.

Table 8: 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) For more information, see “Nonstop Active Routing BFD Support” on page 134. | 8.5 or later |
| BGP For more information, see “Nonstop Active Routing BGP Support” on page 135. | 8.4 or later |
| 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 |

Table 8: Nonstop Active Routing Protocol and Feature Support (*continued*)

| Protocol | Junos OS Release |
|---|--|
| LDP (PTX Series Packet Transport Routers only) | 12.3R4 or later |
| Nonstop active routing support for LDP includes: | (for LDP Point-to-Multipoint LSPs) 13.3R1 or later |
| <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 | (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). | |
| Layer 3 VPNs (see the first Note after this table for restrictions) | 9.2 or later |
| 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) | |
| Layer 3 VPNs (PTX Series Packet Transport Routers only) | 12.1R4 or later |
| Multicast Source Discovery Protocol (MSDP) | 12.1 or later |
| For more information, see "Nonstop Active Routing MSDP Support" on page 138 . | |
| OSPF/OSPFv3 | 8.4 or later |
| Protocol Independent Multicast (PIM) | (for IPv4) 9.3 or later |
| For more information, see "Nonstop Active Routing PIM Support" on page 136 . | (for IPv6) 10.4 or later |
| RIP and RIP next generation (RIPng) | 9.0 or later |

Table 8: Nonstop Active Routing Protocol and Feature Support (*continued*)

| Protocol | Junos OS Release |
|--|---|
| RSVP (PTX Series Packet Transport Routers only) | 12.1R4 or later |
| 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. | |
| RSVP-TE LSP | 9.5 or later |
| For more information, see “Nonstop Active Routing Support for RSVP-TE LSPs” on page 139 . | |
| 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 msp 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.

However, Junos OS nonstop active routing support for RSVP point-to-multipoint LSPs does not include support for dynamically created point-to-multipoint LSPs, such as VPLS.

Starting with Release 14.1, Junos OS extends nonstop active routing support to the next-generation multicast VPNs (MPVNs).

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**).

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, IP2MP LSPs used by VPLS and MVPN
- 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, IP2MP LSPs used by VPLS and MVPN |
| 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 14.1, Junos OS extends nonstop active routing support to the next-generation multicast VPNs (MPVNs). |
| 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 125](#)
- [Configuring Nonstop Active Routing on page 141](#)
- [Configuring Nonstop Active Routing on Switches](#)
- [Example: Configuring Nonstop Active Routing on Switches](#)

Configuring Nonstop Active Routing

Supported Platforms [EX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [T Series](#)

This section includes the following topics:

- [Enabling Nonstop Active Routing on page 141](#)
- [Synchronizing the Routing Engine Configuration on page 142](#)
- [Verifying Nonstop Active Routing Operation on page 143](#)

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]
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 can also issue the **show bgp replication** command.

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 125](#)
- [Nonstop Active Routing System Requirements on page 129](#)
- [Tracing Nonstop Active Routing Synchronization Events on page 146](#)
- [Resetting Local Statistics on page 148](#)
- [Example: Configuring Nonstop Active Routing on page 144](#)
- [nonstop-routing on page 339](#)

Preventing Automatic Reestablishment of BGP Peer Sessions After NSR Switchovers

Supported Platforms [M Series](#), [PTX Series](#), [T Series](#)

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

Supported Platforms EX Series, M Series, PTX Series, T Series

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 141](#)
- [Tracing Nonstop Active Routing Synchronization Events on page 146](#)

Tracing Nonstop Active Routing Synchronization Events

Supported Platforms EX Series, M Series, PTX Series, QFX Series, T Series

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 141](#)
- [Configuring Nonstop Active Routing on Switches](#)
- [Example: Configuring Nonstop Active Routing on Switches](#)
- [Example: Configuring Nonstop Active Routing on page 144](#)

Resetting Local Statistics

Supported Platforms [EX Series](#), [M Series](#), [PTX Series](#), [QFX Series](#), [T Series](#)

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 141](#)
- [Tracing Nonstop Active Routing Synchronization Events on page 146](#)

CHAPTER 8

High Availability and Graceful Restart

- [Graceful Restart Concepts on page 149](#)
- [Graceful Restart System Requirements on page 150](#)
- [Enabling Graceful Restart on page 151](#)
- [Configuring Graceful Restart on page 152](#)
- [Graceful Restart for Aggregate and Static Routes on page 177](#)
- [Graceful Restart and Routing Protocols on page 177](#)
- [Configuring Routing Protocols Graceful Restart on page 180](#)
- [Graceful Restart and MPLS-Related Protocols on page 187](#)
- [Configuring Graceful Restart for MPLS-Related Protocols on page 188](#)
- [Understanding Restart Signaling-Based Helper Mode Support for OSPF Graceful Restart on page 190](#)
- [Example: Managing Helper Modes for OSPF Graceful Restart on page 191](#)
- [Tracing Restart Signaling-Based Helper Mode Events for OSPF Graceful Restart on page 193](#)
- [Graceful Restart and Layer 2 and Layer 3 VPNs on page 194](#)
- [Configuring VPN Graceful Restart on page 195](#)
- [Graceful Restart on Logical Systems on page 197](#)
- [Configuring Logical System Graceful Restart on page 197](#)
- [Verifying Graceful Restart Operation on page 198](#)

Graceful Restart Concepts

Supported Platforms [EX Series](#), [M Series](#), [OCX1100](#), [PTX Series](#), [QFabric System](#), [QFX Series](#), [T Series](#)

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 150](#)
- [Graceful Restart for Aggregate and Static Routes on page 177](#)
- [Graceful Restart and Routing Protocols on page 177](#)
- [Graceful Restart and MPLS-Related Protocols on page 187](#)
- [Graceful Restart and Layer 2 and Layer 3 VPNs on page 194](#)
- [Graceful Restart on Logical Systems on page 197](#)
- [Configuring Graceful Restart on page 152](#)
- [*Configuring Graceful Restart for QFabric Systems*](#)

Graceful Restart System Requirements

Supported Platforms [EX Series](#), [M Series](#), [PTX Series](#), [T Series](#)

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 149](#)

Enabling Graceful Restart

Supported Platforms [EX Series](#), [M Series](#), [PTX Series](#), [T Series](#)

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 149](#)
- [Graceful Restart System Requirements on page 150](#)
- [Graceful Restart for Aggregate and Static Routes on page 177](#)
- [Configuring Graceful Restart on page 152](#)

Configuring Graceful Restart

Supported Platforms [M Series, PTX Series, T Series](#)

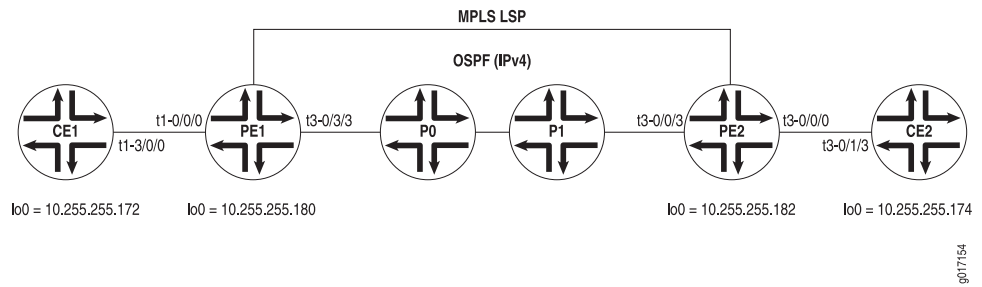
To enable graceful restart, include the **graceful-restart** statement at the **[edit routing-instance *instance-name* routing-options]** or **[edit routing-options]** hierarchy level. This enables graceful restart globally for all routing protocols. You can, optionally, modify or supplement the global settings at the individual protocol level.

For example:

```
protocols {
  bgp {
    group ext {
      graceful-restart {
        restart-time 400;
      }
    }
  }
}
routing-options {
  graceful-restart;
}
```

[Figure 12 on page 153](#) 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 12: 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 12vpn
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 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 all end-of-rib markers sent: inet-vpn-unicast 12vpn
Table bgp.13vpn.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.12vpn.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 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 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
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 |

| | | | |
|---------------------------------|-------|------------------------|-------|
| BGP-INET | vrf | l2circuit.0 | 0/0/0 |
| | | 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 | | |
| | | 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 151](#)
- [Configuring Routing Protocols Graceful Restart on page 180](#)
- [Configuring Graceful Restart for MPLS-Related Protocols on page 188](#)
- [Configuring VPN Graceful Restart on page 195](#)
- [Configuring Logical System Graceful Restart on page 197](#)
- [Verifying Graceful Restart Operation on page 198](#)

Graceful Restart for Aggregate and Static Routes

Supported Platforms [EX Series](#), [M Series](#), [PTX Series](#), [T Series](#)

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 149](#)
- [Graceful Restart System Requirements on page 150](#)
- [Enabling Graceful Restart on page 151](#)
- [Verifying Graceful Restart Operation on page 198](#)
- [Configuring Graceful Restart on page 152](#)

Graceful Restart and Routing Protocols

Supported Platforms [EX Series](#), [M Series](#), [PTX Series](#), [T Series](#)

This section covers the following topics:

- [BGP on page 178](#)
- [IS-IS on page 178](#)
- [OSPF and OSPFv3 on page 178](#)

- [PIM Sparse Mode on page 179](#)
- [RIP and RIPvng on page 180](#)

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 149](#)
- [Graceful Restart System Requirements on page 150](#)
- [Configuring Routing Protocols Graceful Restart on page 180](#)
- [Verifying Graceful Restart Operation on page 198](#)
- [Configuring Graceful Restart on page 152](#)
- [Example: Configuring IS-IS for GRES with Graceful Restart](#)

Configuring Routing Protocols Graceful Restart

Supported Platforms [EX Series](#), [M Series](#), [QFabric System](#), [QFX Series](#), [T Series](#)

This topic includes the following sections:

- [Enabling Graceful Restart on page 180](#)
- [Configuring Graceful Restart Options for BGP on page 181](#)
- [Configuring Graceful Restart Options for ES-IS on page 182](#)
- [Configuring Graceful Restart Options for IS-IS on page 182](#)
- [Configuring Graceful Restart Options for OSPF and OSPFv3 on page 183](#)
- [Configuring Graceful Restart Options for RIP and RIPng on page 185](#)
- [Configuring Graceful Restart Options for PIM Sparse Mode on page 185](#)
- [Tracking Graceful Restart Events on page 186](#)

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 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: If you configure Bidirectional Forwarding Detection (BFD) and graceful restart for IS-IS, graceful restart might not work as expected.



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 186](#).

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 | ospfv3) 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 | ospfv3) 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 | ospfv3) 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 reenable 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 186](#).



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.



NOTE: If you configure BFD and graceful restart for OSPF, graceful restart might not work as expected.

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 149](#)
- [Graceful Restart System Requirements on page 150](#)
- [Graceful Restart and Routing Protocols on page 177](#)

- [Verifying Graceful Restart Operation on page 198](#)
- [Configuring Graceful Restart on page 152](#)

Graceful Restart and MPLS-Related Protocols

Supported Platforms [EX Series](#), [M Series](#), [PTX Series](#), [QFX Series](#), [T Series](#)

This section contains the following topics:

- [LDP on page 187](#)
- [RSVP on page 187](#)
- [CCC and TCC on page 188](#)

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 149](#)
- [Graceful Restart System Requirements on page 150](#)
- [Configuring Graceful Restart for MPLS-Related Protocols on page 188](#)
- [Configuring Graceful Restart on page 152](#)

Configuring Graceful Restart for MPLS-Related Protocols

Supported Platforms [EX Series, M Series, PTX Series, T Series](#)

This section contains the following topics:

- [Configuring Graceful Restart Globally on page 188](#)
- [Configuring Graceful Restart Options for RSVP, CCC, and TCC on page 189](#)
- [Configuring Graceful Restart Options for LDP on page 189](#)

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.

Understanding Restart Signaling-Based Helper Mode Support for OSPF Graceful Restart

Supported Platforms [M Series](#), [T Series](#)

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 191](#)
- [Tracing Restart Signaling-Based Helper Mode Events for OSPF Graceful Restart on page 193](#)

Example: Managing Helper Modes for OSPF Graceful Restart

Supported Platforms [EX Series](#), [M Series](#), [T Series](#)

- [Requirements on page 191](#)
- [Overview on page 191](#)
- [Configuration on page 191](#)
- [Verification on page 193](#)

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 193](#)

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 190](#)
 - [Tracing Restart Signaling-Based Helper Mode Events for OSPF Graceful Restart on page 193](#)
 - [helper-disable \(OSPF\) on page 324](#)

Tracing Restart Signaling-Based Helper Mode Events for OSPF Graceful Restart

Supported Platforms [EX Series](#), [M Series](#), [T Series](#)

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 190](#)
- [Example: Managing Helper Modes for OSPF Graceful Restart on page 191](#)
- [helper-disable \(OSPF\) on page 324](#)

Graceful Restart and Layer 2 and Layer 3 VPNs

Supported Platforms [EX Series](#), [M Series](#), [PTX Series](#), [T Series](#)

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 149](#)
- [Graceful Restart System Requirements on page 150](#)
- [Configuring Logical System Graceful Restart on page 197](#)
- [Verifying Graceful Restart Operation on page 198](#)
- [Configuring Graceful Restart on page 152](#)

Configuring VPN Graceful Restart

Supported Platforms [EX Series](#), [M Series](#), [PTX Series](#), [T Series](#)

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 196](#)
- [Configuring Graceful Restart for the Routing Instance on page 196](#)

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 149](#)
- [Graceful Restart System Requirements on page 150](#)
- [Graceful Restart and Layer 2 and Layer 3 VPNs on page 194](#)
- [Verifying Graceful Restart Operation on page 198](#)
- [Configuring Graceful Restart on page 152](#)

Graceful Restart on Logical Systems

Supported Platforms [EX Series](#), [M Series](#), [PTX Series](#), [T Series](#)

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 149](#)
- [Graceful Restart System Requirements on page 150](#)
- [Configuring Logical System Graceful Restart on page 197](#)
- [Verifying Graceful Restart Operation on page 198](#)
- [Configuring Graceful Restart on page 152](#)

Configuring Logical System Graceful Restart

Supported Platforms [EX Series](#), [M Series](#), [PTX Series](#), [T Series](#)

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 197](#)
- [Configuring Graceful Restart for a Routing Instance on page 198](#)

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 149](#)
- [Graceful Restart System Requirements on page 150](#)
- [Graceful Restart on Logical Systems on page 197](#)
- [Verifying Graceful Restart Operation on page 198](#)
- [Configuring Graceful Restart on page 152](#)

Verifying Graceful Restart Operation

Supported Platforms EX Series, M Series, PTX Series, QFabric System, QFX Series, T Series

This topic contains the following sections:

- [Graceful Restart Operational Mode Commands on page 199](#)
- [Verifying BGP Graceful Restart on page 199](#)
- [Verifying IS-IS and OSPF Graceful Restart on page 200](#)
- [Verifying CCC and TCC Graceful Restart on page 200](#)

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 180](#)).

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

```

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 149](#)
 - *Configuring Graceful Restart for QFabric Systems*

CHAPTER 9

High Availability and Virtual Router Redundancy Protocol (VRRP)

- [Understanding VRRP on page 204](#)
- [Junos OS Support for VRRPv3 on page 206](#)
- [VRRP failover-delay Overview on page 211](#)
- [Configuring Basic VRRP Support on page 213](#)
- [Configuring VRRP on page 218](#)
- [Configuring VRRP for IPv6 on page 219](#)
- [Configuring VRRP Authentication \(IPv4 Only\) on page 220](#)
- [Configuring the Advertisement Interval for the VRRP Master Router on page 222](#)
- [Configuring the Startup Period for VRRP Operations on page 224](#)
- [Configuring a Backup Router to Preempt the VRRP Master Router on page 225](#)
- [Modifying the Preemption Hold-Time Value for the VRRP Master Router on page 225](#)
- [Configuring the Asymmetric Hold Time for VRRP Routers on page 226](#)
- [Configuring Passive ARP Learning for Backup VRRP Routers on page 226](#)
- [Configuring VRRP Route Tracking on page 227](#)
- [Configuring a Logical Interface to Be Tracked for a VRRP Group on page 229](#)
- [Configuring a Route to Be Tracked for a VRRP Group on page 231](#)
- [Example: Configuring Multiple VRRP Owner Groups on page 233](#)
- [Configuring Inheritance for a VRRP Group on page 239](#)
- [Configuring an Interface to Accept All Packets Destined for the Virtual IP Address of a VRRP Group on page 240](#)
- [Configuring the Silent Period to Avoid Alarms Due to Delay in Receiving VRRP Advertisement Packets on page 241](#)
- [Enabling the Distributed Periodic Packet Management Process for VRRP on page 242](#)
- [Improving the Convergence Time for VRRP on page 243](#)
- [Configuring VRRP to Improve Convergence Time on page 244](#)
- [Tracing VRRP Operations on page 246](#)

Understanding VRRP

Supported Platforms EX Series, M Series, MX Series, SRX100, SRX110, SRX210, SRX220, SRX240, SRX550, SRX650, T Series

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 fails, one of the backup routers becomes the new master router, 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 router can take over a failed default router within a few seconds. This is done with minimum VRRP traffic and without any interaction with the hosts.

Routers 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. The default interval is 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.




NOTE: To minimize network traffic, VRRP is designed in such a way that only the router that is acting as the master sends out VRRP advertisements at any given point in time. The backup routers 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.

[Figure 13 on page 205](#) 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 13: 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 Junos OS Release 14.2R1, 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.

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.

Release History Table

| Release | Description |
|---------|---|
| 14.2R1 | In Junos OS Release 14.2R1, 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. |

Related Documentation

- [Understanding High Availability Features on Juniper Networks Routers on page 3](#)
- [Junos OS Support for VRRPv3 on page 206](#)
- [Configuring Basic VRRP Support on page 213](#)
- [Configuring VRRP on page 218](#)

Junos OS Support for VRRPv3

Supported Platforms [EX Series, M320, MX Series, QFX Series, T Series](#)

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 206](#)
- [IPv6 VRRP Checksum Behavioral Differences on page 207](#)
- [VRRP Interoperability on page 207](#)
- [Upgrading from VRRPv2 to VRRPv3 on page 208](#)
- [Functionality of VRRPv3 Features on page 210](#)

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 207](#).

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 208](#) 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.

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 9 on page 209 illustrates the steps and events that take place during a VRRPv2 to VRRPv3 transition. In Table 9 on page 209, 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 9: VRRPv2 to VRRPv3 Transition Steps and Events

| For IPv4 | For IPv6 |
|--|---|
| <ol style="list-style-type: none"> 1. Upgrade Router R1 with Junos OS Release 12.2 or later. <ul style="list-style-type: none"> • 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. <ul style="list-style-type: none"> • 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. | <ol style="list-style-type: none"> 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 R1. <ul style="list-style-type: none"> • Router R1 becomes the master for both G1 and G2. 5. Enable VRRPv3 on Router R2. 6. Activate G1 and G2 on Router R2. <ul style="list-style-type: none"> • 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 210](#)
- [VRRPv3 Advertisement Intervals on page 210](#)
- [Unified ISSU for VRRPv3 on page 210](#)

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 204](#)
- [Configuring Basic VRRP Support on page 213](#)

VRRP failover-delay Overview

Supported Platforms [M Series, MX Series](#)

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.

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 212](#)
- [When failover-delay Is Configured on page 213](#)

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*](#)

Configuring Basic VRRP Support

Supported Platforms [EX Series](#), [M Series](#), [MX Series](#), [T Series](#)



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.

You configure VRRP by configuring a VRRP group.

To configure basic VRRP support, you configure VRRP groups on interfaces. An interface can be a member of one or more VRRP groups. Within a VRRP group, the master virtual router and the backup virtual router must be configured on different routing platforms.

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 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



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. If there are two or more backup routers with the same priority, the router that has the highest primary address becomes the master.

Configuring VRRP for IPv6 Groups



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 |
|---------|---|
| 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 229](#)
- [Configuring a Route to Be Tracked for a VRRP Group on page 231](#)
- [Junos OS Support for VRRPv3 on page 206](#)
- [Understanding VRRP on page 204](#)
- [Configuring the Startup Period for VRRP Operations on page 224](#)
- [Configuring VRRP Authentication \(IPv4 Only\) on page 220](#)

- [Configuring the Advertisement Interval for the VRRP Master Router on page 222](#)
- [Configuring VRRP on page 218](#)

Configuring VRRP

Supported Platforms [M Series](#), [MX Series](#), [T Series](#)

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 {
    gigether-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 204](#)
- [Configuring VRRP for IPv6 on page 219](#)
- [Configuring VRRP Route Tracking on page 227](#)

Configuring VRRP for IPv6

Supported Platforms [ACX Series, M Series, T Series](#)

Configure VRRP properties for IPv6 in one master (Router A) and one backup (Router B).

On Router A

```

[edit interfaces]
ge-1/0/0 {
    unit 0 {

```

```

family inet6 {
  address fe80::5:0:0:6/64;
  address fec0::5:0:0:6/64 {
    vrrp-inet6-group 3; # VRRP inet6 group number
    virtual-inet6-address fec0::5:0:0:7;
    virtual-link-local-address fe80::5:0:0:7;
    priority 200;
    preempt;
  }
}

```

```

[edit protocols]
router-advertisement {
  interface ge-1/0/0.0 {
    prefix fec0::/64;
    max-advertisement-interval 4;
  }
}

```

On Router B

```

[edit interfaces]
ge-1/0/0 {
  unit 0 {
    family inet6 {
      address fe80::5:0:0:8/64;
      address fec0::5:0:0:8/64 {
        vrrp-inet6-group 3; # VRRP inet6 group number
        virtual-inet6-address fec0::5:0:0:7;
        virtual-link-local-address fe80::5:0:0:7;
        priority 100;
        preempt;
      }
    }
  }
}

```

```

[edit protocols]
router-advertisement {
  interface ge-1/0/0.0 {
    prefix fec0::/64;
    max-advertisement-interval 4;
  }
}

```

**Related
Documentation**

- [Understanding VRRP on page 204](#)
- [Configuring VRRP on page 218](#)
- [Configuring VRRP Route Tracking on page 227](#)

Configuring VRRP Authentication (IPv4 Only)

Supported Platforms [EX Series](#), [M Series](#), [T Series](#)

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 204](#)
 - [Junos OS Support for VRRPv3 on page 206](#)
 - [Configuring Basic VRRP Support on page 213](#)
 - [Configuring VRRP on page 218](#)

Configuring the Advertisement Interval for the VRRP Master Router

Supported Platforms [EX Series](#), [M Series](#), [T Series](#)

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 223](#)
- [Modifying the Advertisement Interval in Milliseconds on page 223](#)

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 204](#)
- [Junos OS Support for VRRPv3 on page 206](#)
- [Configuring Basic VRRP Support on page 213](#)
- [Configuring a Backup Router to Preempt the VRRP Master Router on page 225](#)
- [Modifying the Preemption Hold-Time Value for the VRRP Master Router on page 225](#)
- [Configuring the Asymmetric Hold Time for VRRP Routers on page 226](#)
- [Configuring the Silent Period to Avoid Alarms Due to Delay in Receiving VRRP Advertisement Packets on page 241](#)
- [Configuring VRRP on page 218](#)

Configuring the Startup Period for VRRP Operations

Supported Platforms [EX Series](#), [M Series](#), [T Series](#)

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 204](#)
- [Configuring Basic VRRP Support on page 213](#)
- [Configuring VRRP Authentication \(IPv4 Only\) on page 220](#)
- [Configuring VRRP on page 218](#)

Configuring a Backup Router to Preempt the VRRP Master Router

Supported Platforms [EX Series, M Series, T Series](#)

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 204](#)
- [Configuring the Advertisement Interval for the VRRP Master Router on page 222](#)
- [Modifying the Preemption Hold-Time Value for the VRRP Master Router on page 225](#)
- [Configuring the Asymmetric Hold Time for VRRP Routers on page 226](#)
- [Configuring VRRP on page 218](#)

Modifying the Preemption Hold-Time Value for the VRRP Master Router

Supported Platforms [EX Series, M Series, T Series](#)

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 222](#)
- [Configuring a Backup Router to Preempt the VRRP Master Router on page 225](#)
- [Configuring the Asymmetric Hold Time for VRRP Routers on page 226](#)
- [Configuring VRRP on page 218](#)

Configuring the Asymmetric Hold Time for VRRP Routers

Supported Platforms [EX Series, M Series, T Series](#)

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.

**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 222](#)
- [Configuring a Backup Router to Preempt the VRRP Master Router on page 225](#)
- [Modifying the Preemption Hold-Time Value for the VRRP Master Router on page 225](#)
- [Configuring VRRP on page 218](#)

Configuring Passive ARP Learning for Backup VRRP Routers

Supported Platforms [EX Series, M Series, T Series](#)

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 204](#)

Configuring VRRP Route Tracking

Supported Platforms [M Series](#), [T Series](#)

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 204](#)
- [Configuring a Route to Be Tracked for a VRRP Group on page 231](#)
- [Configuring VRRP on page 218](#)
- [Configuring VRRP for IPv6 on page 219](#)

Configuring a Logical Interface to Be Tracked for a VRRP Group

Supported Platforms EX Series, M Series, T Series

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 226](#).

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 10 on page 231](#).

Table 10: Interface State and Priority Cost Usage

| Tracked Interface State | Priority Cost Usage |
|--|--|
| Down | priority-cost <i>priority</i> |
| 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 204](#)
- [Configuring a Route to Be Tracked for a VRRP Group on page 231](#)
- [Configuring VRRP on page 218](#)

Configuring a Route to Be Tracked for a VRRP Group

Supported Platforms [EX Series, M Series, T Series](#)

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 226](#).

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 204](#)
- [Configuring a Logical Interface to Be Tracked for a VRRP Group on page 229](#)
- [Configuring VRRP Route Tracking on page 227](#)

Example: Configuring Multiple VRRP Owner Groups

Supported Platforms [EX Series](#), [M Series](#), [MX Series](#), [T Series](#)

These examples show how to configure multiple virtual router redundancy protocol (VRRP) IPv4 and IPv6 owner groups.

- [Requirements on page 233](#)
- [Overview on page 233](#)
- [Configuration on page 233](#)
- [Verification on page 238](#)

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 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</pre> |
| Multiple IPv4 and IPv6 owner groups | <pre>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</pre> |

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 246](#)

- [Configuring Inheritance for a VRRP Group on page 239](#)

Configuring Inheritance for a VRRP Group

Supported Platforms ACX Series, EX Series, M Series, MX Series, QFabric System, QFX Series, SRX Series, T Series

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 204](#)

Configuring an Interface to Accept All Packets Destined for the Virtual IP Address of a VRRP Group

Supported Platforms [EX Series](#), [M Series](#), [T Series](#)

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 204](#)

- [Configuring VRRP on page 218](#)

Configuring the Silent Period to Avoid Alarms Due to Delay in Receiving VRRP Advertisement Packets

Supported Platforms [EX Series](#), [M Series](#), [T Series](#)

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 204](#)
 - [startup-silent-period on page 382](#)

Enabling the Distributed Periodic Packet Management Process for VRRP

Supported Platforms EX Series, M Series, T Series

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.

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 ppm 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 204](#)
- [delegate-processing \(VRRP\) on page 369](#)

Improving the Convergence Time for VRRP

Supported Platforms EX Series, M Series, MX Series

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 duplication 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. Detection 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 239](#)
- [Configuring VRRP to Improve Convergence Time on page 244](#)
- [delegate-processing on page 369](#)
- [global-advertisements-threshold on page 371](#)
- [skew-timer-disable on page 381](#)

Configuring VRRP to Improve Convergence Time

Supported Platforms EX Series, M Series, MX Series

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 218](#).

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 243](#)
- [Configuring Inheritance for a VRRP Group on page 239](#)
- [delegate-processing on page 369](#)
- [global-advertisements-threshold on page 371](#)
- [skew-timer-disable on page 381](#)

Tracing VRRP Operations

Supported Platforms [EX Series](#), [M Series](#), [T Series](#)

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 204](#)

CHAPTER 10

High Availability and Unified In-Service Software Upgrade (ISSU)

- [Getting Started with Unified In-Service Software Upgrade on page 249](#)
- [Understanding the Unified ISSU Process on page 250](#)
- [Unified ISSU System Requirements on page 259](#)
- [Best Practices for Performing a Unified ISSU on page 278](#)
- [Example: Performing a Unified ISSU on page 279](#)
- [Verifying a Unified ISSU on page 305](#)
- [Troubleshooting Unified ISSU Problems on page 306](#)
- [Managing and Tracing BFD Sessions During Unified ISSU Procedures on page 306](#)

Getting Started with Unified In-Service Software Upgrade

Supported Platforms [EX Series, M10i, M120, M320, MX104, MX2010, MX2020, MX240, MX480, MX960, PTX3000, PTX5000, T1600, T320, T4000, T640, TX Matrix, TX Matrix Plus](#)

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 11 on page 249](#).

Table 11: 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 259 |
| Perform a unified ISSU | "Example: Performing a Unified ISSU" on page 279 |
| Verify that the unified ISSU is successful | "Verifying a Unified ISSU" on page 305 |
| Understand how the unified ISSU process works | "Understanding the Unified ISSU Process" on page 250 |

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

Supported Platforms [EX9204, EX9208, EX9214, M10i, M120, M320, MX104, MX2010, MX2020, MX240, MX480, MX960, PTX3000, PTX5000, T1600, T320, T4000, T640, TX Matrix, TX Matrix Plus](#)

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 250](#)
- [Understanding the Unified ISSU Process on the TX Matrix Router on page 254](#)
- [Understanding the Unified ISSU Process on the TX Matrix Plus Router and on the TX Matrix Plus Router with 3D SIBs on page 257](#)

Understanding the Unified ISSU Process on a Router

Supported Platforms [EX Series, M10i, M120, M320, MX Series, PTX5000, T Series](#)

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 14 on page 252](#) through [Figure 19 on page 254](#) 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 14: 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 15: 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 16: 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 17: 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 18: 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 19: 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

Supported Platforms TX Matrix

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 255](#)

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

Supported Platforms TX Matrix

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 257](#)

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 249](#)
- [Best Practices for Performing a Unified ISSU on page 278](#)
- [Unified ISSU System Requirements on page 259](#)
- [Example: Performing a Unified ISSU on page 279](#)
- [request system software validate in-service-upgrade on page 427](#)

Unified ISSU System Requirements

Supported Platforms M10i, M120, M320, MX104, MX2010, MX2020, MX240, MX480, MX960, PTX3000, PTX5000, T1600, T320, T4000, T640, TX Matrix, TX Matrix Plus

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 (<http://pathfinder.juniper.net/feature-explorer/>).

This section contains the following topics:

- [General Unified ISSU Considerations for All Platforms on page 260](#)
- [Unified ISSU Considerations for MX Series Routers on page 261](#)
- [Unified ISSU Considerations for PTX Series Routers on page 261](#)
- [Unified ISSU Considerations for M Series and T Series Routers on page 262](#)
- [Unified ISSU Platform Support on page 262](#)
- [Unified ISSU Protocol Support for M Series, MX Series, and T Series Routers and EX9200 Switches on page 263](#)
- [Unified ISSU Feature Support on page 264](#)
- [Unified ISSU PIC Support Considerations on page 265](#)

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 http://www.juniper.net/techpubs/en_US/junos14.2/information-products/topic-collections/release-notes/14.2/junos-release-notes-14.2.pdf. For more information about Junos OS Release 15.1, see http://www.juniper.net/techpubs/en_US/junos15/information-products/topic-collections/release-notes/15/junos-release-notes-15.pdf.
- 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 *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 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 MPCs 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.

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 PTX 3000 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 M Series and T Series Routers

Unified ISSU has the following caveats for M Series and 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.

Unified ISSU Platform Support

Table 12 on page 262 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 (<http://pathfinder.juniper.net/feature-explorer/>).

Table 12: 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 |

Table 12: Unified ISSU Support for Dual Routing Engine Platforms (*continued*)

| Platform | Junos OS Release |
|-------------------------------------|------------------|
| M10i router | 9.5R1 |
| M120 router | 9.2R1 |
| M320 router | 9.0R1 |
| MX240 router | 9.3R1 |
| 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 |
| 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 login 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.



NOTE: 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.



NOTE: LACP is not supported on the PTX5000 and PTX 3000 routers during a unified ISSU.

- 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 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](#)

Unified ISSU PIC Support Considerations

The following sections list the PICs that are supported by unified ISSU.

- [PIC Considerations on page 265](#)
- [SONET/SDH PICs on page 266](#)
- [Fast Ethernet and Gigabit Ethernet PICs on page 267](#)
- [Channelized PICs on page 270](#)
- [Tunnel Services PICs on page 271](#)
- [ATM PICs on page 271](#)
- [Serial PICs on page 272](#)
- [DS3, E1, E3, and T1 PICs on page 272](#)
- [Enhanced IQ PICs on page 273](#)
- [Enhanced IQ2 Ethernet Services Engine \(ESE\) PIC on page 273](#)
- [Unified ISSU FPC Support on T4000 Routers on page 274](#)
- [Unified ISSU Support on MX Series 3D Universal Edge Routers on page 274](#)



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 brought 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 13 on page 266 lists the SONET/SDH PICs that are supported during a unified ISSU.

Table 13: 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) | |
| | 2 | PE-4OC3-SON-MM—(EOL) | M10i |
| | | PE-4OC3-SON-SMIR—(EOL) | |
| OC3c/STM1 with SFP | 2 | PE-2OC3-SON-MM—(EOL) | M10i |
| | | PE-2OC3-SON-SMIR—(EOL) | |
| 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 |
| | | PB-4OC3-1OC12-SON-SFP | |
| | 4 OC3 ports, 1 OC12 port | PB-4OC3-1OC12-SON2-SFP | M10i |
| | | PE-4OC3-1OC12-SON-SFP | |

Table 13: Unified ISSU PIC Support: SONET/SDH (*continued*)

| PIC Type | Number of Ports | Model Number | Device |
|------------------|-----------------|-------------------------|---|
| OC12c/STM4 | 1 | PE-1OC12-SON-SFP | M10i |
| | | PE-1OC12-SON-MM—(EOL) | |
| | | PE-1OC12-SON-SMIR—(EOL) | |
| | 4 | PB-1OC12-SON-MM—(EOL) | M120, M320, T320, T640, T1600, T4000, TX Matrix, TX Matrix Plus with 3D SIBs |
| | | PB-1OC12-SON-SMIR—(EOL) | |
| | | PB-4OC12-SON-MM | |
| OC12c/STM4, SFP | 1 | 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 | M120, M320, MX Series, T320, T640, T1600, TX Matrix, T4000, TX Matrix Plus, TX Matrix Plus with 3D SIBs |
| | | PB-1OC48-SON-B-SFP | |
| | | | |
| OC48c/STM16, SFP | 4 | PC-4OC48-SON-SFP | |
| | | | |
| OC192/STM64 | 1 | PC-1OC192-SON-VSR | MX Series routers |
| OC192/STM64, XFP | 1 | PC-1OC192-SON-LR | M320, T320, T640, T1600, T4000, TX Matrix Plus with 3D SIBs |
| | | PC-1OC192-SON-SR2 | |
| | | PC-1OC192-VSR | |
| OC192/STM64, XFP | 4 | PD-4OC192-SON-XFP | M120, T640, T1600, T4000, TX Matrix Plus with 3D SIBs |
| | | | |
| OC192/STM64, XFP | 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 14 on page 268 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 14: 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 | |
| | 40 | EX9200-40F | EX9200 |
| 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 | |

Table 14: Unified ISSU PIC Support: Fast Ethernet and Gigabit Ethernet (*continued*)

| PIC Type | Number of Ports | Model Number | Device |
|---|-----------------|--|--|
| 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 | |
| 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 | 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 | PD-5-10XGE-SFPP | T640, T1600, T4000, TX Matrix Plus, TX Matrix Plus with 3D SIBs |
| | 24 | P1-PTX-24-10GE-SFPP | PTX5000 |
| 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 | 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 |
| 40-Gigabit Ethernet, QFSP+ | 4 | EX9200-4QS | EX9200 |
| 10-Gigabit Ethernet, 40-Gigabit Ethernet, QSFPP | 48/12 | P2-10G-40G-QSFPP | PTX5000 |

Table 14: 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 15 on page 270 lists the channelized PICs that are supported during a unified ISSU.

Table 15: 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 15: 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 16 on page 271 lists the Tunnel Services PICs that are supported during a unified ISSU.

Table 16: 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 17 on page 271 lists the ATM PICs that are supported during a unified ISSU. The table includes support on Enhanced III FPCs.

Table 17: 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 17: 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 18 on page 273](#).

Table 18: 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 18: 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 274](#)
- [Unified ISSU MIC and MPC Support on MX Series Routers on page 275](#)
- [Unified ISSU Limitations on MX Series Routers on page 277](#)

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 http://www.juniper.net/techpubs/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 19 on page 275](#) and [Table 20 on page 276](#). 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 19: 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 19: 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 |

Table 20: 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 |
| 100-Gigabit Ethernet MPC with QSFP | 2 | MIC3-3D-2X40GE-QSFP | MX Series routers |
| 100-Gigabit Ethernet MPC with SFPP | 10 | MIC3-3D-10XGE-SFPP | MX Series routers |
| 100-Gigabit Ethernet MPC with CXP | 1 | MIC3-3D-1X100GE-CXP | MX Series routers |
| 100-Gigabit Ethernet MPC 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 |

Table 20: Unified ISSU Support: MX Series Router MICs (*continued*)

| MIC Type | Number of Ports | Model Number | Platform |
|--|-----------------|-----------------------|-------------------|
| 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 |



NOTE: Note that unified ISSU is supported only by the MICs listed in Table 20 on page 276.



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 MPCEs 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 |
|---------|---|
| 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. |
| 13.2 | Starting with Junos OS Release 13.2 |

Related Documentation

- [Getting Started with Unified In-Service Software Upgrade on page 249](#)
- [Best Practices for Performing a Unified ISSU on page 278](#)
- [Understanding the Unified ISSU Process on page 250](#)
- [Example: Performing a Unified ISSU on page 279](#)
- [request system software validate on \(Junos OS with Upgraded FreeBSD\)](#)

Best Practices for Performing a Unified ISSU

Supported Platforms

EX Series, M10i, M120, M320, MX104, MX2010, MX2020, MX240, MX480, MX960, PTX3000, PTX5000, T1600, T320, T4000, T640, TX Matrix, TX Matrix Plus

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 259](#) to anticipate any special circumstances that might affect your upgrade.

Related Documentation

- [Example: Performing a Unified ISSU on page 279](#)
- [Verifying a Unified ISSU on page 305](#)

- [Troubleshooting Unified ISSU Problems on page 306](#)

Example: Performing a Unified ISSU

Supported Platforms EX Series, M10i, M120, M320, MX104, MX2010, MX2020, MX240, MX480, MX960, PTX3000, PTX5000, T1600, T320, T4000, T640, TX Matrix, TX Matrix Plus

This example shows how to perform a unified in-service software upgrade (ISSU).

- [Requirements on page 279](#)
- [Overview on page 280](#)
- [Configuration on page 280](#)

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 259](#) 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 <http://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

<http://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 20 on page 280 shows the topology used in this example.

Figure 20: 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 281](#)
- [Verifying the Software Versions and Backing Up the Device Software on page 282](#)
- [Adjusting Timers and Changing Feature-Specific Configuration on page 283](#)
- [Upgrading and Rebooting Both Routing Engines Automatically on page 285](#)
- [Restoring Feature-Specific Configuration on page 290](#)
- [Upgrading Both Routing Engines and Rebooting the New Backup Routing Engine Manually on page 291](#)
- [Upgrading and Rebooting Only One Routing Engine on page 298](#)

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 21 on page 285 shows the Routing Engine status prior to starting the unified ISSU.

Table 21: 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 <http://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 22 on page 288](#) shows the Routing Engine status after the unified ISSU.

Table 22: 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 23 on page 289](#) shows the Routing Engine status after Step 5 is completed.

Table 23: 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 290](#).
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 281.
2. Perform the steps in “[Verifying the Software Versions and Backing Up the Device Software](#)” on page 282.
3. Perform the steps in “[Adjusting Timers and Changing Feature-Specific Configuration](#)” on page 283.
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 <http://kb.juniper.net/InfoCenter/index?page=content&id=KB17665>.

Table 24 on page 292 shows the Routing Engine status prior to starting the unified ISSU.

Table 24: Routing Engine Status Before Upgrading and Manually Rebooting the Backup Routing Engine

| RE0 | RE1 |
|--------------------------------|--------------------------------|
| Master | Backup |
| 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 25 on page 294 shows the Routing Engine status after the unified ISSU and before manually rebooting the backup Routing Engine.

Table 25: 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 |

6. 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 26 on page 296](#) shows the Routing Engine status after the unified ISSU, after rebooting the backup Routing Engine, but before switching mastership.

Table 26: 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 27 on page 297 shows the Routing Engine status after the unified ISSU, after rebooting the backup Routing Engine, and after switching mastership.

Table 27: 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 290.
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 28 on page 298 shows the Routing Engine status prior to starting the unified ISSU.

Table 28: 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 281.
2. Perform the steps in “[Verifying the Software Versions and Backing Up the Device Software](#)” on page 282.
3. Perform the applicable steps in “[Adjusting Timers and Changing Feature-Specific Configuration](#)” on page 283.
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 <http://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/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
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 29 on page 301](#) shows the Routing Engine status after the unified ISSU upgrades the master Routing Engine but before the backup Routing Engine is upgraded.

Table 29: 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 |

6. 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. 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/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
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 30 on page 304](#) shows the Routing Engine status after the unified ISSU, after rebooting the backup Routing Engine, and after switching mastership.

Table 30: 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 |

14. Enable GRES and NSR again by performing the steps in [“Verifying Dual Routing Engines and Enabling GRES and NSR” on page 281](#).
15. Perform the applicable steps in [“Restoring Feature-Specific Configuration” on page 290](#).
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 249](#)
- [Understanding the Unified ISSU Process on page 250](#)
- [Unified ISSU System Requirements on page 259](#)
- [Best Practices for Performing a Unified ISSU on page 278](#)
- [Verifying a Unified ISSU on page 305](#)
- [Troubleshooting Unified ISSU Problems on page 306](#)
- [Managing and Tracing BFD Sessions During Unified ISSU Procedures on page 306](#)

Verifying a Unified ISSU

Supported Platforms EX Series, M10i, M120, M320, MX104, MX2010, MX2020, MX240, MX480, MX960, PTX3000, PTX5000, T1600, T320, T4000, T640, TX Matrix, TX Matrix Plus

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 279](#)
 - [Troubleshooting Unified ISSU Problems on page 306](#)
 - [Understanding the Unified ISSU Process on page 250](#)
 - [Unified ISSU System Requirements on page 259](#)
 - [Managing and Tracing BFD Sessions During Unified ISSU Procedures on page 306](#)

Troubleshooting Unified ISSU Problems

Supported Platforms EX9204, EX9208, EX9214, M10i, M120, M320, MX104, MX2010, MX2020, MX240, MX480, MX960, PTX3000, PTX5000, T1600, T320, T4000, T640, TX Matrix, TX Matrix Plus

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 system software abort in-service-upgrade (ICU)* for more information.

- Related Documentation**
- [Understanding the Unified ISSU Process on page 250](#)
 - [Unified ISSU System Requirements on page 259](#)
 - [Best Practices for Performing a Unified ISSU on page 278](#)
 - [Example: Performing a Unified ISSU on page 279](#)
 - [Verifying a Unified ISSU on page 305](#)
 - [Managing and Tracing BFD Sessions During Unified ISSU Procedures on page 306](#)

Managing and Tracing BFD Sessions During Unified ISSU Procedures

Supported Platforms EX Series, M Series, MX104, MX2010, MX2020, MX240, MX480, PTX3000, PTX5000, T1600, T320, T320, T4000, T640, TX Matrix, TX Matrix Plus

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 249](#)
- [Understanding the Unified ISSU Process on page 250](#)
- [Unified ISSU System Requirements on page 259](#)
- [Best Practices for Performing a Unified ISSU on page 278](#)
- [Example: Performing a Unified ISSU on page 279](#)
- [Verifying a Unified ISSU on page 305](#)
- [Troubleshooting Unified ISSU Problems on page 306](#)

PART 3

Configuration Statements and Operational Commands

- [High Availability Configuration Statements and Operational Commands on page 311](#)

CHAPTER 11

High Availability Configuration Statements and Operational Commands

- Configuration Statements: Bidirectional Forwarding Detection on page 311
- Configuration Statements: Graceful Routing Engine Switchover on page 318
- Configuration Statements: Graceful Restart on page 319
- Configuration Statements: Nonstop Active Routing on page 338
- Configuration Statements: Nonstop Bridging on page 345
- Configuration Statements: Routing Engine and Switching Control Board Redundancy on page 345
- Configuration Statements: Unified ISSU on page 358
- Configuration Statements: VRRP on page 361
- Operational Commands on page 393

Configuration Statements: Bidirectional Forwarding Detection

- authentication (LAG) on page 312
- bfd-liveness-detection (LAG) on page 313
- detection-time (LAG) on page 315
- traceoptions (Protocols BFD) on page 316
- transmit-interval (LAG) on page 318

authentication (LAG)

Supported Platforms [M Series](#), [MX Series](#), [PTX Series](#), [T Series](#)

Syntax

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

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 **algorithm** *algorithm-name*—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

key-chain *key-chain-name*—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.

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.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [bfd-liveness-detection on page 313](#)
- [detection-time on page 315](#)
- [transmit-interval on page 318](#)
- [Configuring Independent Micro BFD Sessions for LAG on page 58](#)
- [Example: Configuring Independent Micro BFD Sessions for LAG on page 63](#)
- [Understanding Independent Micro BFD Sessions for LAG on page 55](#)

bfd-liveness-detection (LAG)

Supported Platforms M Series, MX Series, PTX Series, QFX Series, T Series

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).

Default: automatic

The remaining statements are explained 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">• authentication on page 312• detection-time on page 315• transmit-interval on page 318• Configuring Independent Micro BFD Sessions for LAG on page 58• Example: Configuring Independent Micro BFD Sessions for LAG on page 63• Understanding Independent Micro BFD Sessions for LAG on page 55 |
|------------------------------|--|

detection-time (LAG)

| | |
|--------------------------|--|
| Supported Platforms | M Series , MX Series , PTX Series , T Series |
| Syntax | <pre>detection-time { 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 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 312• bfd-liveness-detection on page 313• transmit-interval on page 318• Configuring Independent Micro BFD Sessions for LAG on page 58• Example: Configuring Independent Micro BFD Sessions for LAG on page 63• Understanding Independent Micro BFD Sessions for LAG on page 55 |

traceoptions (Protocols BFD)

| | |
|---------------------|---|
| Supported Platforms | ACX Series, M Series, MX Series, SRX240, SRX3400, T1600, T4000, T640 |
| 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 306 |
|------------------------------|--|

transmit-interval (LAG)

| | |
|--------------------------|---|
| Supported Platforms | M Series, MX Series, PTX Series, T Series |
| 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 312• bfd-liveness-detection on page 313• detection-time on page 315• Configuring Independent Micro BFD Sessions for LAG on page 58• Example: Configuring Independent Micro BFD Sessions for LAG on page 63• Understanding Independent Micro BFD Sessions for LAG on page 55 |

Configuration Statements: Graceful Routing Engine Switchover

- [graceful-switchover on page 319](#)

graceful-switchover

| | |
|---------------------------------|---|
| Supported Platforms | EX Series, M Series, MX Series, PTX Series, T Series |
| 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. |
| 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 Routing Engine Switchover on page 114 |

Configuration Statements: Graceful Restart

- [disable on page 320](#)
- [graceful-restart \(Enabling Globally\) on page 321](#)
- [graceful-restart \(Multicast Snooping\) on page 322](#)
- [helper-disable \(Multiple Protocols\) on page 323](#)
- [helper-disable \(OSPF\) on page 324](#)
- [maximum-helper-recovery-time on page 325](#)
- [maximum-helper-restart-time \(RSVP\) on page 326](#)
- [maximum-neighbor-reconnect-time on page 327](#)
- [maximum-neighbor-recovery-time on page 328](#)
- [no-strict-lsa-checking on page 329](#)
- [notify-duration on page 330](#)
- [not-on-disk-underperform on page 331](#)
- [reconnect-time on page 332](#)
- [recovery-time on page 333](#)
- [restart-duration on page 334](#)
- [restart-time \(BGP Graceful Restart\) on page 335](#)
- [stale-routes-time on page 336](#)
- [traceoptions \(Protocols\) on page 337](#)

disable

| | |
|---------------------------------|---|
| Supported Platforms | ACX Series, EX Series, M Series, OCX1100, PTX Series, QFabric System, QFX Series, SRX Series, T Series |
| 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 esis 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 151 • Configuring Routing Protocols Graceful Restart on page 180 • Configuring Graceful Restart for MPLS-Related Protocols on page 188 • Configuring VPN Graceful Restart on page 195 • Configuring Logical System Graceful Restart on page 197 • Graceful Restart Configuration Statements • Configuring Graceful Restart for QFabric Systems |

graceful-restart (Enabling Globally)

Supported Platforms ACX Series, EX Series, M Series, MX Series, OCX1100, PTX Series, QFabric System, QFX Series, SRX210, SRX3400, T1600, T640

Syntax `graceful-restart {
 disable;
 helper-disable;
 maximum-helper-recovery-time seconds;
 maximum-helper-restart-time seconds;
 notify-duration seconds;
 recovery-time seconds;
 restart-duration seconds;
 stale-routes-time seconds;
 }`

Hierarchy Level [edit logical-systems *logical-system-name* routing-options],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
 [edit routing-options],
 [edit routing-instances *routing-instance-name* routing-options]

Release Information Statement introduced before Junos OS Release 7.4.
 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 Configure graceful restart globally to enable the feature. You cannot enable graceful restart for specific protocols unless graceful restart is also enabled globally. You can, optionally, modify the global settings at the individual protocol level.



NOTE:

- 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.

Default Graceful restart is disabled by default.

Options The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
 routing-control—To add this statement to the configuration.

- Related Documentation**
- [Enabling Graceful Restart on page 151](#)
 - [Configuring Routing Protocols Graceful Restart on page 180](#)
 - [Configuring Graceful Restart for MPLS-Related Protocols on page 188](#)
 - [Configuring VPN Graceful Restart on page 195](#)
 - [Configuring Logical System Graceful Restart on page 197](#)
 - [Configuring Graceful Restart for QFabric Systems](#)

graceful-restart (Multicast Snooping)

Supported Platforms [MX Series, PTX Series](#)

Syntax `graceful-restart {
 disable;
 restart-duration seconds;
}`

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**
- [Example: Configuring Multicast Snooping](#)
 - [query-response-interval \(Bridge Domains\)](#)

helper-disable (Multiple Protocols)

| | |
|---------------------------------|--|
| Supported Platforms | ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, SRX210, SRX3400, T640 |
| 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 180 • Configuring Graceful Restart for MPLS-Related Protocols on page 188 |

helper-disable (OSPF)

| | |
|--|---|
| Supported Platforms | EX4600, M Series, MX Series, OCX1100, QFabric System, QFX Series, SRX210, SRX3400, T1600, T640 |
| 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>  <p>NOTE: Restart signaling-based helper mode is not supported for OSPFv3 configurations.</p> </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 180 • Configuring Graceful Restart for MPLS-Related Protocols on page 188 |

maximum-helper-recovery-time

| | |
|---------------------------------|---|
| Supported Platforms | ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, SRX Series, T Series |
| Syntax | maximum-helper-recovery-time <i>seconds</i> ; |
| 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 188 • maximum-helper-restart-time (RSVP) on page 326 |

maximum-helper-restart-time (RSVP)

| | |
|---------------------------------|--|
| Supported Platforms | ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, SRX210, T1600, T640 |
| Syntax | maximum-helper-restart-time <i>seconds</i> ; |
| 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 | seconds —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 188• maximum-helper-recovery-time on page 325 |

maximum-neighbor-reconnect-time

| | |
|--------------------------|--|
| Supported Platforms | ACX Series , EX Series , M Series , MX Series , PTX Series , SRX Series , T Series |
| 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 188 |

maximum-neighbor-recovery-time

| | |
|--------------------------|---|
| Supported Platforms | ACX Series , EX Series , M Series , MX Series , PTX Series , QFX Series , SRX210 , SRX3400 , T1600 , T640 |
| Syntax | maximum-neighbor-recovery-time <i>seconds</i> ; |
| 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 | seconds —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 LDP Graceful Restart• Configuring Graceful Restart Options for LDP on page 188• no-strict-lsa-checking on page 329• recovery-time on page 333 |

no-strict-lsa-checking

| | |
|---------------------------------|---|
| Supported Platforms | ACX Series, EX Series, M Series, MX Series, OCX1100, PTX Series, QFabric System, QFX Series, SRX Series, T Series |
| 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 180• <i>Configuring Graceful Restart for QFabric Systems</i>• maximum-neighbor-recovery-time on page 328• recovery-time on page 333 |

notify-duration

| | |
|--------------------------|---|
| Supported Platforms | ACX Series, EX Series, M Series, MX Series, OCX1100, PTX Series, QFabric System, QFX Series, SRX Series, T Series |
| Syntax | notify-duration <i>seconds</i> ; |
| Hierarchy Level | [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] |
| Release Information | 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. |
| Description | Specify the length of time the router or switch notifies helper OSPF routers that it has completed graceful restart. |
| Options | seconds —Length of time in the router notifies helper OSPF routers that it has completed graceful restart. Range: 1 through 3600 Default: 30 |
| 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 180• Configuring Graceful Restart for QFabric Systems• restart-duration on page 334 |

not-on-disk-underperform

Supported Platforms [EX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [T Series](#)

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 117](#)

reconnect-time

| | |
|---------------------------------|--|
| Supported Platforms | ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, SRX Series, T Series |
| Syntax | reconnect-time <i>seconds</i> ; |
| 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 188 |

recovery-time

| | |
|---------------------------------|--|
| Supported Platforms | EX Series, M Series, PTX Series, T Series |
| Syntax | recovery-time <i>seconds</i> ; |
| 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 | <i>seconds</i> —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 188 • maximum-neighbor-recovery-time on page 328 • no-strict-lsa-checking on page 329 |

restart-duration

| | |
|----------------------------|---|
| Supported Platforms | ACX Series, EX Series, M Series, MX Series, OCX1100, PTX Series, QFabric System, SRX Series, T Series |
| Syntax | restart-duration <i>seconds</i> ; |
| Hierarchy Level | <p>[edit logical-systems <i>logical-system-name</i> protocols (isis ospf ospf3 pim) <i>graceful-restart</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3 pim) <i>graceful-restart</i>],</p> <p>[edit protocols (esis isis ospf ospf3 pim) <i>graceful-restart</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3 pim) <i>graceful-restart</i>],</p> <p>[edit routing-options <i>graceful-restart</i>]</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 <i>graceful-restart</i>] (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 <i>graceful-restart</i>] (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

- [Enabling Graceful Restart on page 151](#)
- [Configuring Graceful Restart for MPLS-Related Protocols on page 188](#)
- [Configuring VPN Graceful Restart on page 195](#)
- [Configuring Graceful Restart for VPNs](#)
- [Configuring Logical System Graceful Restart on page 197](#)

restart-time (BGP Graceful Restart)

Supported Platforms ACX Series, EX Series, M Series, MX Series, OCX1100, OCX1100, PTX Series, QFabric System, QFX Series, SRX Series, T Series

Syntax restart-time *seconds*;

Hierarchy Level [edit protocols (bgp | rip | ripng) [graceful-restart](#)],
[edit logical-systems *logical-system-name* protocols (bgp | rip | ripng) [graceful-restart](#) ([Enabling Globally](#))],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp [graceful-restart](#)],
[edit routing-instances *routing-instance-name* 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 Configure the duration of the BGP, RIP, or next-generation RIP (RIPng) graceful restart period.

Options *seconds*—Length of time for the graceful restart period.

Range: 1 through 600 seconds

Default: Varies by protocol:

- BGP—120 seconds
- RIP and RIPng—60 seconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [Configuring Graceful Restart Options for BGP on page 180](#)
- [Configuring Graceful Restart Options for RIP and RIPng on page 180](#)
- [Configuring Graceful Restart for QFabric Systems](#)
- [stale-routes-time on page 336](#)

stale-routes-time

| | |
|--------------------------|---|
| Supported Platforms | ACX Series, EX Series, M Series, MX Series, OCX1100, OCX1100, PTX Series, QFabric System, QFX Series, SRX Series, T Series |
| Syntax | stale-routes-time <i>seconds</i> ; |
| 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 180• Configuring Graceful Restart for QFabric Systems• restart-time (BGP Graceful Restart) on page 335 |

traceoptions (Protocols)

Supported Platforms EX Series, M Series, PTX Series, T Series

Syntax

```
traceoptions {
    file name <size size> <files number> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
```

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 Define tracing operations that graceful restart functionality in the router or switch.

To specify more than one tracing operation, include multiple **flag** statements.

Default If you do not include this statement, no global tracing operations are performed.

Options **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 *name*—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.

Range: 2 through 1000 files

Default: 2 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. The nonstop active routing tracing option is:

- **graceful-restart**—Tracing operations for nonstop active routing

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">• Tracking Graceful Restart Events on page 180 |
|------------------------------|--|

Configuration Statements: Nonstop Active Routing

- [nonstop-routing on page 339](#)
- [switchover-on-routing-crash on page 340](#)
- [synchronize on page 341](#)
- [traceoptions on page 343](#)

nonstop-routing

Supported Platforms ACX Series, EX Series, M Series, MX Series, PTX Series, SRX210, SRX3400, T1600, T640

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 141](#)

switchover-on-routing-crash

| | |
|---------------------------------|---|
| Supported Platforms | EX4300, M Series, MX Series, PTX Series, QFX Series, T4000, TX Matrix |
| 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 141 |

synchronize

| | |
|----------------------------|--|
| Supported Platforms | EX Series, M Series, PTX Series, T Series |
| 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 141](#)
- *Configuring Multiple Routing Engines to Synchronize Committed Configurations Automatically*

traceoptions

| | |
|----------------------------|--|
| Supported Platforms | ACX Series, EX Series, M Series, MX Series, QFX Series, SRX210, T1600, T640, vSRX |
| 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> |
| Default | If you do not include this statement, no global tracing operations are performed. |
| Options | <p>Values:</p> <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>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. We recommend that you place global routing protocol tracing output in the file routing-log.</p> |

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 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

- *Example: Tracing Global Routing Protocol Operations*

Configuration Statements: Nonstop Bridging

- [nonstop-bridging on page 345](#)

nonstop-bridging

Supported Platforms [EX Series](#), [MX Series](#), [QFX Series](#)

Syntax nonstop-bridging;

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

- [Synchronizing the Routing Engine Configuration on page 141](#)
- [Configuring Nonstop Bridging on page 122](#)
- 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)*
- For information about configuring NSB on switches that support ELS, see *Configuring Nonstop Bridging on Switches (CLI Procedure)*

Configuration Statements: Routing Engine and Switching Control Board Redundancy

- [cfeb on page 346](#)
- [description \(Chassis Redundancy\) on page 347](#)
- [failover \(Chassis\) on page 347](#)
- [failover \(System Process\) on page 348](#)
- [feb \(Creating a Redundancy Group\) on page 349](#)
- [feb \(Assigning a FEB to a Redundancy Group\) on page 350](#)
- [keepalive-time on page 351](#)

- [no-auto-failover](#) on page 352
- [on-disk-failure](#) (Chassis Redundancy Failover) on page 352
- [on-loss-of-keepalives](#) on page 353
- [redundancy](#) on page 354
- [redundancy-group](#) on page 355
- [routing-engine](#) (Chassis Redundancy) on page 356
- [sfm](#) (Chassis Redundancy) on page 357
- [ssb](#) on page 358

cfeb

| | |
|---------------------------------|---|
| Supported Platforms | M10i |
| Syntax | <code>cfeb slot-number</code> (always preferred); |
| 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 | slot-number —Specify which slot is the master and which is the backup. always —Define this CFEB as the sole device. preferred —Define this CFEB as the preferred device of at least two. |
| 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 CFEB Redundancy on the M10i Router on page 83 |

description (Chassis Redundancy)

| | |
|---------------------------------|--|
| Supported Platforms | M120 |
| Syntax | <code>description <i>description</i>;</code> |
| 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 | <i>description</i> —Provide a description for the FEB redundancy 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 FEB Redundancy on the M120 Router on page 84 |

failover (Chassis)

| | |
|---------------------------------|---|
| Supported Platforms | M Series, MX Series, PTX Series, T Series |
| 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 | Specify conditions on the master Routing Engine that cause the backup router to take mastership. The remaining statements are explained 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"> • On Detection of a Hard Disk Error on the Master Routing Engine on page 95 |

failover (System Process)

| | |
|--------------------------|--|
| Supported Platforms | M Series, MX Series, PTX Series, T Series |
| 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 95• <i>processes</i> |

feb (Creating a Redundancy Group)

Supported Platforms [M120](#)

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 Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring FEB Redundancy on the M120 Router on page 84](#)

feb (Assigning a FEB to a Redundancy Group)

| | |
|---------------------------------|---|
| Supported Platforms | M120 |
| Syntax | feb <i>slot-number</i> (backup primary); |
| Hierarchy Level | [edit chassis redundancy feb redundancy-group <i>group-name</i>] |
| 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 84 |

keepalive-time

| | |
|---------------------------------|---|
| Supported Platforms | M Series, PTX Series, T Series |
| 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 | <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 95 • failover (Chassis) on page 347 • on-loss-of-keepalives on page 353 |

no-auto-failover

| | |
|---------------------------------|--|
| Supported Platforms | M120 |
| Syntax | no-auto-failover; |
| Hierarchy Level | [edit chassis redundancy feb redundancy-group group-name] |
| 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 84 |

on-disk-failure (Chassis Redundancy Failover)

| | |
|---------------------------------|---|
| Supported Platforms | M Series, PTX Series, T Series |
| 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 95 |

on-loss-of-keepalives

| | |
|--------------------------|---|
| Supported Platforms | M Series, PTX Series, T Series |
| 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 95• keepalive-time on page 351 |

redundancy

Supported Platforms [M Series](#), [MX Series](#), [PTX Series](#), [T Series](#)

Syntax

```
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);  
}
```

Hierarchy Level [edit chassis]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure redundancy options.

Options The statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

- Related Documentation**
- [Configuring Routing Engine Redundancy on page 95](#)
 - [Configuring CFEB Redundancy on the M10i Router on page 83](#)
 - [Configuring FEB Redundancy on the M120 Router on page 84](#)
 - [Configuring SFM Redundancy on M40e and M160 Routers on page 86](#)
 - [Configuring SSB Redundancy on the M20 Router on page 87](#)

redundancy-group

Supported Platforms [SRX210, SRX3400](#)

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 84](#)

routing-engine (Chassis Redundancy)

| | |
|---------------------------------|---|
| Supported Platforms | M Series, PTX Series, T Series |
| Syntax | routing-engine <i>slot-number</i> (backup disabled master); |
| 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 | 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 95 |

sfm (Chassis Redundancy)

| | |
|---------------------------------|--|
| Supported Platforms | M40e |
| Syntax | <code>sfm slot-number</code> (always preferred); |
| 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 86 |

ssb

| | |
|---------------------------------|--|
| Syntax | <code>ssb slot-number</code> (<code>always</code> <code>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 | <code>slot-number</code> —Specify which slot is the master and which is the backup. <code>always</code> —Define this SSB as the sole device. <code>preferred</code> —Define this SSB as the preferred device of at least two. |
| 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 SSB Redundancy on the M20 Router on page 87 |

Configuration Statements: Unified ISSU

- [no-issu-timer-negotiation on page 359](#)
- [traceoptions \(Protocols BFD\) on page 360](#)

no-issu-timer-negotiation

| | |
|----------------------------|--|
| Supported Platforms | ACX Series, EX Series, M Series, MX Series, PTX5000, SRX Series, T Series |
| Syntax | no-issu-timer-negotiation; |
| Hierarchy Level | [edit protocols bfd], [edit logical-systems <i>logical-system-name</i> protocols bfd], [edit routing-instances <i>routing-instance-name</i> 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 | <ul style="list-style-type: none"> • Managing and Tracing BFD Sessions During Unified ISSU Procedures on page 306 • <i>Junos OS Routing Protocols Library</i> |

traceoptions (Protocols BFD)

| | |
|---------------------|---|
| Supported Platforms | ACX Series, M Series, MX Series, SRX240, SRX3400, T1600, T4000, T640 |
| 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 306 |
|------------------------------|--|

Configuration Statements: VRRP


- [accept-data on page 363](#)
- [advertise-interval on page 364](#)
- [asymmetric-hold-time on page 365](#)
- [authentication-key on page 366](#)
- [authentication-type on page 367](#)
- [bandwidth-threshold on page 368](#)
- [delegate-processing \(VRRP\) on page 369](#)
- [fast-interval on page 370](#)
- [global-advertisements-threshold on page 371](#)
- [hold-time \(VRRP\) on page 372](#)

- [inherit-advertisement-interval on page 373](#)
- [inet6-advertise-interval on page 374](#)
- [interface on page 375](#)
- [preempt \(VRRP\) on page 376](#)
- [priority \(Protocols VRRP\) on page 377](#)
- [priority-cost \(VRRP\) on page 378](#)
- [priority-hold-time on page 379](#)
- [route \(Interfaces\) on page 380](#)
- [skew-timer-disable on page 381](#)
- [startup-silent-period on page 382](#)
- [traceoptions \(Protocols VRRP\) on page 383](#)
- [track \(VRRP\) on page 385](#)
- [version-3 on page 386](#)
- [virtual-address on page 387](#)
- [virtual-inet6-address on page 388](#)
- [virtual-link-local-address on page 389](#)
- [vrrp-group on page 390](#)
- [vrrp-inet6-group on page 392](#)
- [vrrp-inherit-from on page 393](#)

accept-data

| | |
|--|--|
| Supported Platforms | EX Series, M Series, MX Series, NFX Series, OCX1100, QFabric System, QFX Series, SRX Series, T Series |
| Syntax | (accept-data no-accept-data); |
| 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>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.</p> <ul style="list-style-type: none"> • 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. |
| <div>  NOTE: <ul style="list-style-type: none"> • 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, <i>Virtual Router Redundancy Protocol (VRRP)</i>). </div> | |
| 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 an Interface to Accept All Packets Destined for the Virtual IP Address of a VRRP Group on page 240 |

advertise-interval

| | |
|--|--|
| Supported Platforms | EX Series, M Series, NFX Series, OCX1100, QFabric System, QFX Series, T Series |
| Syntax | advertise-interval <i>seconds</i> ; |
| Hierarchy Level | [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 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>] |
| 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. |
| <div>  <p>NOTE: When VRRPv3 is enabled, the <code>advertise-interval</code> statement cannot be used to configure advertisement intervals. Instead, use the <code>fast-interval</code> statement to configure advertisement intervals.</p> </div> | |
| Options | <i>seconds</i> —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 | <ul style="list-style-type: none"> • Configuring the Advertisement Interval for the VRRP Master Router on page 222 • fast-interval on page 370 • inet6-advertise-interval on page 374 • version-3 on page 386 |

asymmetric-hold-time

| | |
|---------------------------------|--|
| Supported Platforms | EX Series, M Series, MX Series, SRX Series, T Series |
| 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 route goes down. However, when the route comes back online, the backup router that is acting as the master waits for the priority hold time to expire before switching the mastership back to the original master VRRP router. |
| 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 226 |

authentication-key

| | |
|----------------------------|--|
| Supported Platforms | EX Series, M Series, MX Series, NFX Series, OCX1100, QFX Series, SRX Series, T Series |
| Syntax | authentication-key <i>key</i> ; |
| Hierarchy Level | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrp-group <i>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> vrp-group <i>group-id</i>] |
| 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 a Virtual Router Redundancy Protocol (VRRP) IPv4 authentication key. You also must specify a VRRP authentication scheme by including the authentication-type statement. All routers in the VRRP group must use the same authentication scheme and password. |



NOTE: When VRRPv3 is enabled, the **authentication-type** and **authentication-key** statements cannot be configured for any VRRP groups.

| | |
|---------------------------------|--|
| Options | 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 (" "). |
| 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 Authentication (IPv4 Only) on page 220 • Configuring VRRP Authentication (IPv4 Only) • authentication-type on page 367 • version-3 on page 386 |

authentication-type

Supported Platforms EX Series, M Series, MX Series, NFX Series, OCX1100, QFabric System, QFX Series, SRX Series, T Series

Syntax authentication-type *authentication*;

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 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.

All routers in the VRRP group must use the same authentication scheme and password.



NOTE: When VRRPv3 is enabled, the **authentication-type** and **authentication-key** statements cannot be configured for any VRRP groups.

Options *authentication*—Authentication scheme:

- **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.

Default: none (no authentication is performed).

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring VRRP Authentication \(IPv4 Only\) on page 220](#)
- [Configuring VRRP Authentication \(IPv4 Only\)](#)
- [authentication-key on page 366](#)
- [version-3 on page 386](#)


bandwidth-threshold

| | |
|---------------------------------|--|
| Supported Platforms | EX Series, M Series, MX Series, NFX Series, OCX1100, QFabric System, QFX Series, SRX Series, T Series |
| 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> <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 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 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 229 Configuring a Logical Interface to Be Tracked |

delegate-processing (VRRP)

| | |
|---------------------------------|---|
| Supported Platforms | EX Series, M Series, T Series |
| Syntax | <pre>delegate-processing { ae-irb; }</pre> |
| Hierarchy Level | [edit protocols vrrp] |
| Release Information | <p>Statement introduced in Junos OS Release 9.6.</p> <p>ae-irb option introduced in Junos OS Release 15.1.</p> |
| 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 ppmd 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. Using the ae-irb option requires use of the enhanced-ip mode.</p> |
| 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 the Distributed Periodic Packet Management Process for VRRP on page 242 |

fast-interval

| | |
|---|--|
| Supported Platforms | EX Series, M Series, MX Series, NFX Series, OCX1100, QFX Series, SRX Series, T Series |
| Syntax | <code>fast-interval milliseconds;</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 222 • Configuring the Advertisement Interval for the VRRP Master • advertise-interval on page 364 • advertise-interval on page 364 • inet6-advertise-interval on page 374 • version-3 on page 386 |

global-advertisements-threshold

Supported Platforms [EX Series](#), [M Series](#), [MX Series](#)

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 243](#)
 - [Configuring VRRP to Improve Convergence Time on page 244](#)

hold-time (VRRP)

| | |
|--------------------------|---|
| Supported Platforms | EX Series, M Series, NFX Series, QFX Series, T Series |
| Syntax | hold-time <i>seconds</i> ; |
| Hierarchy Level | [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>], [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>], [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>], [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>] |
| 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 225• Configuring VRRP Preemption and Hold Time |

inherit-advertisement-interval

Supported Platforms [M Series](#), [MX Series](#), [T Series](#)

Syntax inherit-advertisement-interval *seconds*;

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 *seconds*—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 •

inet6-advertise-interval

| | |
|---------------------|--|
| Supported Platforms | EX Series, M Series, MX Series, PTX Series, SRX Series, T Series |
| 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 group-id], [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] |
| 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 | <i>milliseconds</i> —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 | <ul style="list-style-type: none">• Configuring the Advertisement Interval for the VRRP Master Router on page 222• advertise-interval on page 364• fast-interval on page 370• version-3 on page 386 |

interface

| | |
|---------------------------------|--|
| Supported Platforms | M Series, NFX Series, QFX Series, T Series |
| Syntax | <pre>interface <i>interface-name</i> { bandwidth-threshold <i>bits-per-second</i> <i>priority-cost</i> <i>priority</i>; priority-cost <i>priority</i>; }</pre> |
| 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 before Junos OS Release 7.4.</p> <p>bandwidth-threshold statement added in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p> |
| Description | Enable logical interface tracking for a Virtual Router Redundancy Protocol (VRRP) group. |
| Options | <p><i>interface-name</i>—Interface to be tracked for this VRRP group.</p> <p>Range: 1 through 10 interfaces</p> <p>The remaining statements are described separately.</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 229 • Configuring a Logical Interface to Be Tracked • Junos OS Services Interfaces Library for Routing Devices |

preempt (VRRP)

Supported Platforms EX9200, M Series, NFX Series, OCX1100, QFabric System, QFX Series, T Series

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.

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 225](#)
- [Configuring VRRP Preemption and Hold Time](#)


priority (Protocols VRRP)

| | |
|---------------------------------|--|
| Supported Platforms | EX Series, M Series, NFX Series, OCX1100, QFabric System, QFX Series, T Series |
| 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> |
| Description | Configure a Virtual Router Redundancy Protocol (VRRP) router's priority for becoming the master default router. The router with the highest priority within the group becomes the master. |
| Options | <p>priority—Router's priority for being elected to be the master router in the VRRP group. A larger value indicates a higher priority for being elected.</p> <p>Range: 1 through 255</p> <p>Default: 100. If two or more routers have the highest priority in the VRRP group, the router 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 213 • Configuring Basic VRRP Support for QFX |

priority-cost (VRRP)

| | |
|---------------------------------|---|
| Supported Platforms | ACX2000 and ACX2100, EX Series, M120, MX240, NFX Series, OCX1100, QFabric System, QFX Series, SRX210, SRX3400, T1600, T640 |
| Syntax | <code>priority-cost <i>priority</i>;</code> |
| Hierarchy Level | <code>[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>],</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> track interface <i>interface-name</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> vrrp-group <i>group-id</i> track interface <i>interface-name</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> vrrp-inet6-group <i>group-id</i> track interface <i>interface-name</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 12.2 for ACX2000 Universal Access Routers. |
| 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 | <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. Range: 1 through 254 |
| 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 229• Configuring a Logical Interface to Be Tracked |

priority-hold-time

| | |
|--|--|
| Supported Platforms | EX Series, M Series, MX Series, NFX Series, OCX1100, QFabric System, QFX Series, SRX Series, T Series |
| 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> |
| <div>  <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> </div> | |
| Options | <p>seconds—Minimum length of time that must elapse between dynamic priority changes.</p> <p>Range: 0 through 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 229 • Configuring a Logical Interface to Be Tracked |

route (Interfaces)

| | |
|--------------------------|---|
| Supported Platforms | EX Series, M120, MX240, NFX Series, OCX1100, QFabric System, QFX Series, SRX210, SRX3400, T1600, T640 |
| Syntax | route <i>prefix</i> routing-instance <i>instance-name</i> priority-cost <i>priority</i> ; |
| Hierarchy Level | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrrp-group group-id track], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group group-id 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 group-id 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 group-id track] |
| Release Information | Statement introduced in Junos OS Release 9.0. Statement introduced in Junos OS 11.3 for QFX Series. Statement introduced in Junos OS 12.1 for EX Series switches. |
| Description | Enable route tracking for a Virtual Router Redundancy Protocol (VRRP) group. |
| Options | <i>prefix</i> —Route to be tracked for this VRRP group. <i>priority-cost priority</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. <i>routing-instance instance-name</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 . |
| 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 Route to Be Tracked for a VRRP Group on page 231• Configuring a Route to Be Tracked |

skew-timer-disable

Supported Platforms [M Series, MX Series](#)

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 243](#)
- [Configuring VRRP to Improve Convergence Time on page 244](#)

startup-silent-period

| | |
|---------------------------------|---|
| Supported Platforms | EX Series, M Series, MX Series, NFX Series, OCX1100, QFabric System, QFX Series, SRX Series, T Series |
| Syntax | startup-silent-period <i>seconds</i> ; |
| 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 224• <i>Configuring the Startup Period for VRRP Operations</i> |

traceoptions (Protocols VRRP)

| | |
|----------------------------|---|
| Supported Platforms | M120, MX240, SRX210, SRX3400, T1600, T640 |
| 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 246 |
|------------------------------|---|

track (VRRP)

| | |
|---------------------------------|--|
| Supported Platforms | EX Series, M Series, NFX Series, OCX1100, QFabric System, QFX Series, T Series |
| 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 | <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>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 229 • Configuring a Route to Be Tracked for a VRRP Group on page 231 • <i>Configuring a Logical Interface to Be Tracked</i> • <i>Configuring a Route to Be Tracked</i> |

version-3

Supported Platforms [EX Series, M320, MX Series](#)

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 206](#)


virtual-address

| | |
|---------------------------------|--|
| Supported Platforms | EX Series, M Series, MX Series, NFX Series, OCX1100, QFabric System, QFX Series, SRX Series, T Series |
| Syntax | virtual-address [<i>addresses</i>]; |
| Hierarchy Level | [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 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>] |
| 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 addresses of the virtual routers in a Virtual Router Redundancy Protocol (VRRP) IPv4 or 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 | 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 213 • Configuring Basic VRRP Support for QFX |

virtual-inet6-address

| | |
|---------------------------------|--|
| Supported Platforms | EX Series, M Series, MX Series, QFX Series, SRX Series, T Series |
| Syntax | virtual-inet6-address [<i>addresses</i>]; |
| Hierarchy Level | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group group-id], [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] |
| 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 | 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 213 |

virtual-link-local-address

| | |
|--|--|
| Supported Platforms | EX Series, M Series, MX Series, QFX Series, SRX Series, T Series |
| Syntax | virtual-link-local-address <i>ipv6-address</i> ; |
| Hierarchy Level | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group group-id], [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] |
| 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. |
| <div style="display: flex; align-items: center;">  <div> <p>NOTE: You do <i>not</i> 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.</p> </div> </div> | |
| Options | <i>ipv6-address</i> —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 | <ul style="list-style-type: none"> • Configuring Basic VRRP Support on page 213 • Junos OS Support for VRRPv3 on page 206 |

vrrp-group

| | |
|---------------------|---|
| Supported Platforms | EX Series, M Series, MX Series, NFX Series, OCX1100, QFabric System, QFX Series, SRX Series, T Series |
| Syntax | <pre> vrrp-group group-id { (accept-data no-accept-data); advertise-interval seconds; global-advertisements-threshold number; authentication-key key; authentication-type authentication; fast-interval milliseconds; (preempt no-preempt) { hold-time seconds; } priority number; track { interface interface-name { bandwidth-threshold bits-per-second priority-cost priority; priority-cost priority; } priority-hold-time seconds; route prefix/prefix-length routing-instance instance-name priority-cost priority; } virtual-address [addresses]; vrrp-inherit-from vrrp-group; } </pre> |
| Hierarchy Level | <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet 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 inet address <i>address</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 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> |



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:00:53:00 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.

Range: 0 through 255

The remaining 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 | <ul style="list-style-type: none">• Configuring Basic VRRP Support on page 213• Configuring VRRP on page 218• <i>Configuring Basic VRRP Support for QFX</i>• <i>Example: Configuring VRRP for Load Sharing</i>• vrrp-inet6-group on page 392• nonstop-routing on page 339 |
|------------------------------|--|

vrrp-inet6-group

Supported Platforms EX Series, M Series, MX Series, QFX Series, SRX Series, T Series

Syntax

```
vrrp-inet6-group group-id {
  (accept-data | no-accept-data);
  advertisements-threshold number;
  fast-interval milliseconds;
  inet6-advertise-interval seconds;
  (preempt | no-preempt) {
    hold-time seconds;
  }
  priority number;
  track {
    interface interface-name {
      bandwidth-threshold bits-per-second priority-cost priority;
      priority-cost priority;
    }
    priority-hold-time seconds;
    route prefix/prefix-length routing-instance instance-name priority-cost priority;
  }
  virtual-inet6-address [ addresses ];
  virtual-link-local-address ipv6-address;
  vrrp-inherit-from vrrp-group;
}
```

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family inet6 address *address*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*
family inet6 address *address*]

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.

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 213](#)

vrrp-inherit-from

Supported Platforms [M Series](#), [MX Series](#), [T Series](#)

Syntax

```
vrrp-inherit-from {
    active-group group-index;
    active-interface active-interface-name;
}
```

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family inet6 [vrrp-inet6-group group-id](#)]
[edit interfaces *interface-name* unit *logical-unit-number* family inet [vrrp-group group-id](#)]

Release Information Statement introduced before Junos OS Release 7.4.

Description VRRP group to follow for the vrrp-group or vrrp-inet6-group.

Options *group-index*—Identifier for VRRP active group.
Range: 0 through 255

active-interface-name—Interface name of VRRP active group.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Understanding VRRP on page 204](#)

Operational Commands

- [clear vrrp](#)
- [request chassis ssb master switch](#)
- [request system software in-service-upgrade](#)
- [request system software in-service-upgrade \(MX Series 3D Universal Edge Routers and EX9200 Switches\)](#)
- [request system software validate in-service-upgrade](#)
- [show chassis ssb](#)
- [show nonstop-routing](#)
- [show pfe ssb](#)
- [show system switchover](#)
- [show task replication](#)
- [show vrrp](#)

clear vrrp

| | |
|--------------------------|--|
| Supported Platforms | EX Series , M Series , MX Series , PTX Series , T Series |
| Syntax | clear vrrp (all interface <i>interface-name</i>) |
| 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 450 |
| List of Sample Output | clear vrrp all on page 394 |
| 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 430 |
| List of Sample Output | request chassis ssb master switch on page 395 request chassis ssb master switch no-confirm on page 395 |
| 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 system software in-service-upgrade

| | |
|----------------------------------|---|
| Supported Platforms | ACX5000, EX4600, M10i, M120, M320, MX Series, PTX5000, QFX Series, T Series |
| Syntax | request system software in-service-upgrade <i>package-name</i> <no-old-master-upgrade> <reboot> |
| Syntax (QFX5100 Switches) | request system software in-service-upgrade <i>package-name</i> |
| 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. In addition, graceful Routing Engine switchover (GRES) and nonstop active routing (NSR) must be enabled. On QFX5100 switches, nonstop bridging (NSB) must be enabled if you are using the Layer 2 Control Protocol process (l2cpd) to transmit Layer 2 spanning tree protocols in a Layer 2 bridge environment.</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"> <i>ftp</i>—File Transfer Protocol <i>http</i>—Hypertext Transfer Protocol <i>scp</i>—Secure copy (available only for Canada and U.S. version) <p><i>no-old-master-upgrade</i>—(Optional) When the <i>no-old-master-upgrade</i> 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 <i>no-old-master-upgrade</i> option, the system will automatically upgrade the former master Routing Engine.</p> <p><i>reboot</i>—(Optional) When the <i>reboot</i> option is included, the former master (new backup) Routing Engine is automatically rebooted after being upgraded to the new software.</p> |

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: The reboot option is not available on the QFX5100 switch.

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 259](#).
- Unsupported PICs are restarted during a unified ISSU on certain routing devices. For information about supported PICs, see the *Junos OS High Availability Library for Routing Devices*.
- Unsupported protocols will experience packet loss during a unified ISSU. For information about supported protocols, see the *Junos OS High Availability Library for Routing Devices*.
- During a unified ISSU, you cannot bring any PICs online or offline on certain routing devices.

For more information, see the *Junos OS High Availability Library for Routing Devices*.

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 249](#)
- [Example: Performing a Unified ISSU on page 279](#)

List of Sample Output

[request system software-in-service upgrade reboot on page 398](#)
[request system software-in-service upgrade reboot \(TX Matrix Plus Router\) on page 400](#)
[request system software-in-service upgrade \(QFX5100 Switch\) on page 408](#)

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 3D Universal Edge Routers and EX9200 Switches)

Supported Platforms [EX Series](#), [MX Series](#)

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 3D Universal Edge Routers 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 *Junos OS High Availability Library for Routing Devices*. For information about supported EX9200 line cards, see [“Unified ISSU System Requirements” on page 259](#).
- Unsupported protocols will experience packet loss during a unified ISSU. For information about supported protocols, see the *Junos OS High Availability Library for Routing Devices* or, for EX9200, see [“Unified ISSU System Requirements” on page 259](#).
- During a unified ISSU, you cannot bring any PICs online or offline.

For more information, see the *Junos OS High Availability Library for Routing Devices* or the *High Availability Feature Guide for EX9200 Switches*.

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 411](#)
[request system software in-service-upgrade \(MX Series Virtual Chassis\) on page 422](#)

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-rel:
```

```
-----
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 validate in-service-upgrade

| | |
|--------------------------|---|
| Supported Platforms | EX Series, M10i, M120, M320, MX Series, PTX5000, T Series |
| Syntax | request system software validate in-service-upgrade <i>package-name</i> |
| 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>package-name—Location from which the software package or bundle is to be installed. For example:</p> <ul style="list-style-type: none"> /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: <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 259.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> request system software in-service-upgrade on page 397 request system software abort show chassis in-service-upgrade Getting Started with Unified In-Service Software Upgrade on page 249 Example: Performing a Unified ISSU on page 279 |
| List of Sample Output | request system software validate in-service-upgrade on page 428 |
| 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 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 | <ul style="list-style-type: none"> request chassis ssb master switch on page 395 |
| List of Sample Output | show chassis ssb on page 431 |
| Output Fields | Table 31 on page 430 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 31: 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. |
| Buffer utilization | Percentage of buffer space being used by the SSB's processor. |
| DRAM | Total DRAM available to the SSB's processor. |

Table 31: show chassis ssb Output Fields (*continued*)

| Field Name | Field Description |
|-------------------|---|
| 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 339 |
| List of Sample Output | show nonstop-routing (MX Series Router) on page 433 show nonstop-routing (MX Series Router) on page 434 |
| Output Fields | Table 32 on page 432 describes the output fields for the show nonstop-routing command. Output fields are listed in the approximate order in which they appear. |

Table 32: 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. |
| Automerger | <p>Status of the automerger.</p> <ul style="list-style-type: none"> • Active—Automerger of sockets by the kernel after switchover is active. • Inactive—Automerger of sockets by the kernel after switchover is inactive. |

Table 32: show nonstop-routing Output Fields (*continued*)

| Field Name | Field Description |
|-----------------------|---|
| Batching | Status of Batching. <ul style="list-style-type: none"> • Yes—Automerger of sockets by the kernel after a switchover. • No—Automerger of sockets by the kernel after switchover is inactive. |
| Batch count | Number of sockets merged per batch. |
| Batch count adjust | Speed at which the batch count is adjusted. <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 automerger operation. |
| Batch interval adjust | Speed at which the batch interval is adjusted. <ul style="list-style-type: none"> • Exp—Time interval between automerger of batches is increased exponentially. • None—Time interval between automerger of batches is not adjusted. |
| Automerger State | State of the automerger <ul style="list-style-type: none"> • Ready—Ready to automerger socket pairs from secondary to primary routing engine • InProgress—Kernel is performing automerger after switchover • Switchover Completed—Sessions merged after switchover |
| Sessions Processed | Count of sessions that are automerger. |

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
Automerger : Active
Batching: No
Batch count: 200
Batch count adjust: Exponential
Batch interval: 20 msec
Batch interval adjust: None
Automerger 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 437 |
| Output Fields | Table 33 on page 435 lists the output fields for the show pfe ssb command. Output fields are listed in the approximate order in which they appear. |

Table 33: 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. |
| Nextthop | Peer message type Nextthop. |
| COS | Peer message type COS. |
| Route | Peer message type Route. |
| SW Firewall | Peer message type SW Firewall. |
| HW Firewall | Peer message type HW Firewall. |

Table 33: show pfe ssb Output Fields (*continued*)

| Field Name | Field Description |
|--------------------------------|--|
| 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 | <p>Information about Packet Forwarding Engine listener statistics:</p> <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 33: 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

| | |
|---------------------------------------|---|
| Supported Platforms | EX Series, M Series, MX Series, PTX Series, QFX Series, T Series |
| List of Syntax | Syntax on page 441 Syntax (TX Matrix Router) on page 441 Syntax (TX Matrix Plus Router) on page 441 Syntax (MX Series Router) on page 441 |
| Syntax | show system switchover |
| Syntax (TX Matrix Router) | show system switchover <all-chassis all-lcc lcc <i>number</i> scc> |
| Syntax (TX Matrix Plus Router) | show system switchover <all-chassis all-lcc lcc <i>number</i> sfc <i>number</i> > |
| Syntax (MX Series Router) | show system switchover <all-members> <local> <member <i>member-id</i> > |
| 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 444](#)
[show system switchover \(Backup Routing Engine - Not Ready\) on page 444](#)
[show system switchover \(MX Virtual Chassis\) on page 444](#)
[show system switchover \(MX Virtual Chassis\) on page 445](#)
[show system switchover \(Routing Matrix and Routing Matrix Plus\) - Master Ready on page 445](#)
[show system switchover \(Routing Matrix and Routing Matrix Plus\) - Master Not Ready on page 445](#)
[show system switchover \(Routing Matrix and Routing Matrix Plus\) - Backup Ready on page 445](#)
[show system switchover \(Routing Matrix and Routing Matrix Plus\) - Backup Not Ready on page 446](#)
[show system switchover all-lcc \(Routing Matrix and Routing Matrix Plus\) on page 446](#)

Output Fields [Table 34 on page 443](#) describes the output fields for the **show system switchover** command. Output fields are listed in the approximate order in which they appear.

Table 34: 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 34: 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
```

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 StateSwitchover Status: Not Ready
```

show system switchover (MX Virtual Chassis)

```
{master:member1-re1}

user@host> show system switchover
member0:
-----
Graceful switchover: On
Configuration database: Ready
Kernel database: Ready
Switchover Status: Ready

member1:
```



```
-----
Command is not applicable on this member of the virtual-chassis
```

show system switchover (MX Virtual Chassis)

```
{master:member1-re1}
user@host> show system switchover
member0:
-----
Graceful switchover: On
Configuration database: Ready
Kernel database: Ready
Switchover Ready
```

```
member1:
-----
Command is not applicable on this member of the virtual-chassis
```

show system switchover (Routing Matrix and Routing Matrix Plus) - Master Ready

```
user@host> show system switchover
lcc0-re1:
-----
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 system switchover (Routing Matrix and Routing Matrix Plus) - Master Not Ready

```
user@host> show system switchover
lcc0-re1:
-----
Multichassis replication: On
Configuration database: Ready
Kernel database: Ready
Peer state: Steady State
Switchover Status: Ready
```

```
lcc2-re1:
-----
Multichassis replication: On
Configuration database: Ready
Kernel database: Ready
Peer state: Steady State
Switchover Status: Not Ready
```

show system switchover (Routing Matrix and Routing Matrix Plus) - Backup Ready

```
user@host> show system switchover
```

scc-re0:

Graceful switchover: On
Configuration database: Ready
Kernel database: Ready
Switchover Status: Ready

lcc0-re0:

Graceful switchover: On
Configuration database: Ready
Kernel database: Ready
Switchover Status: Ready

lcc2-re1:

Graceful switchover: On
Configuration database: Ready
Kernel database: Ready
Switchover Status: Ready

show system switchover (Routing Matrix and Routing Matrix Plus) - Backup Not Ready

```
user@host> show system switchover
```

scc-re0:

Graceful switchover: On
Configuration database: Ready
Kernel database: Ready
Switchover Status: Not Ready

lcc0-re0:

Graceful switchover: On
Configuration database: Ready
Kernel database: Ready
Switchover Status: Ready

lcc2-re1:

Graceful switchover: On
Configuration database: Ready
Kernel database: Ready
Switchover Status: Ready

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

Supported Platforms [EX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [QFX Series](#), [T Series](#)

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.

Options This command has no options.

Required Privilege Level view

List of Sample Output [show task replication \(Issued on the Master Routing Engine\) on page 448](#)
[show task replication \(Issued on the Backup Routing Engine\) on page 449](#)

Output Fields [Table 35 on page 448](#) lists the output fields for the `show task replication` command. Output fields are listed in the approximate order in which they appear.

Table 35: 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 . |

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 vrrp

Supported Platforms [EX Series](#), [M Series](#), [MX Series](#), [T Series](#)

Syntax `show vrrp`
`<brief | detail | extensive | summary>`
`<interface interface-name <group number>>`
`<logical-system logical-system-name >`
`<nsr>`

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 `none`—(Same as `brief`) Display brief status information about all VRRP interfaces.

`brief | detail | extensive | summary`—(Optional) Display the specified level of output.

`interface interface-name <group number>`—(Optional) Display information and status about the specified VRRP interface and, optionally, the group number.

`logical-system logical-system-name`—(Optional) Perform this operation on a particular logical system.

`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.

Required Privilege Level view

Related Documentation

- [show vrrp track](#)
- [clear vrrp on page 394](#)

List of Sample Output [show vrrp on page 456](#)
[show vrrp brief on page 456](#)
[show vrrp detail \(IPv6\) on page 456](#)
[show vrrp detail \(Route Track\) on page 457](#)
[show vrrp detail \(Route Track\) on page 457](#)
[show vrrp extensive on page 457](#)
[show vrrp interface on page 458](#)
[show vrrp nsr on page 459](#)
[show vrrp summary on page 460](#)

Output Fields [Table 36 on page 451](#) lists the output fields for the `show vrrp` command. Output fields are listed in the approximate order in which they appear

Table 36: 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 36: 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 36: 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 36: 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 36: 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 36: 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 456](#).

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::13:1:1:1/120

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 | lcl | 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 1cl yes yes 50.0.0.1
vip 50.0.0.3
ge-1/0/1.0 up 1 master Active 1cl yes yes 1000::1
vip
fe80::200:5eff:fe00:1
vip 1000::3
ge-1/0/1.0 up 2 master Active 1cl yes yes 2000::1
vip
fe80::200:5eff:fe00:2
vip 2000::3
ge-1/0/1.0 up 3 master Active 1cl yes yes 3000::1
vip
fe80::200:5eff:fe00:3
vip 3000::3
ge-1/0/1.0 up 4 master Active 1cl yes yes 4000::1
vip
fe80::200:5eff:fe00:4
vip 4000::3
ge-1/0/1.0 up 5 master Active 1cl yes yes 5000::1
vip
fe80::200:5eff:fe00:5
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   1cl   10.57.0.2
vip           10.57.0.100

```