



Junos[®] OS for EX Series Ethernet Switches

BGP for EX9200 Switches

Release
12.3



Published: 2013-04-03

Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

This product includes the Envoy SNMP Engine, developed by Epilogue Technology, an Integrated Systems Company. Copyright © 1986-1997, Epilogue Technology Corporation. All rights reserved. This program and its documentation were developed at private expense, and no part of them is in the public domain.

This product includes memory allocation software developed by Mark Moraes, copyright © 1988, 1989, 1993, University of Toronto.

This product includes FreeBSD software developed by the University of California, Berkeley, and its contributors. All of the documentation and software included in the 4.4BSD and 4.4BSD-Lite Releases is copyrighted by the Regents of the University of California. Copyright © 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994. The Regents of the University of California. All rights reserved.

GateD software copyright © 1995, the Regents of the University. All rights reserved. Gate Daemon was originated and developed through release 3.0 by Cornell University and its collaborators. Gated is based on Kirton's EGP, UC Berkeley's routing daemon (routed), and DCN's HELLO routing protocol. Development of Gated has been supported in part by the National Science Foundation. Portions of the GateD software copyright © 1988, Regents of the University of California. All rights reserved. Portions of the GateD software copyright © 1991, D. L. S. Associates.

This product includes software developed by Maker Communications, Inc., copyright © 1996, 1997, Maker Communications, Inc.

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.

Products made or sold by Juniper Networks or components thereof might be covered by one or more of the following patents that are owned by or licensed to Juniper Networks: U.S. Patent Nos. 5,473,599, 5,905,725, 5,909,440, 6,192,051, 6,333,650, 6,359,479, 6,406,312, 6,429,706, 6,459,579, 6,493,347, 6,538,518, 6,538,899, 6,552,918, 6,567,902, 6,578,186, and 6,590,785.

Junos® OS for EX Series Ethernet Switches BGP for EX9200 Switches

Release 12.3

Copyright © 2013, 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	vii
	Documentation and Release Notes	vii
	Supported Platforms	vii
	Using the Examples in This Manual	vii
	Merging a Full Example	viii
	Merging a Snippet	viii
	Documentation Conventions	ix
	Documentation Feedback	xi
	Requesting Technical Support	xi
	Self-Help Online Tools and Resources	xi
	Opening a Case with JTAC	xii
Part 1	Overview	
Chapter 1	Feature Support	3
	EX Series Switch Software Features Overview	3
Part 2	Configuration	
Chapter 2	Configuration Statements	49
	[edit protocols bgp] Hierarchy Level	49
	Common BGP Family Options	49
	Complete [edit protocols bgp] Hierarchy	50
	accept-remote-nexthop	55
	advertise-external	56
	advertise-inactive	57
	advertise-peer-as	58
	aggregate-label	59
	allow	60
	as-override	61
	authentication-algorithm	62
	authentication-key (Protocols BGP)	63
	authentication-key-chain (Protocols BGP)	64
	autonomous-system	65
	bfd-liveness-detection (Protocols BGP)	68
	bgp	72
	bgp-orf-cisco-mode	73
	cluster	75
	damping (Protocols BGP)	77
	description (Protocols BGP)	79
	disable (Protocols BGP)	80

explicit-null (Protocols BGP)	81
export (Protocols BGP)	82
family (Protocols BGP)	83
graceful-restart (Protocols BGP)	87
group (Protocols BGP)	89
hold-time (Protocols BGP)	92
idle-after-switch-over	94
import (Protocols BGP)	95
include-mp-next-hop	96
keep	97
labeled-unicast (Protocols BGP)	98
local-address (Protocols BGP)	100
local-as	102
local-interface (IPv6)	104
local-preference	105
log-updown (Protocols BGP)	106
metric-out (Protocols BGP)	107
mtu-discovery	109
multihop	111
multipath (Protocols BGP)	113
neighbor (Protocols BGP)	114
no-aggregator-id	117
no-client-reflect	118
no-validate	119
out-delay	121
outbound-route-filter	122
passive (Protocols BGP)	123
peer-as (Protocols BGP)	124
preference (Protocols BGP)	126
prefix-limit	127
remove-private	129
rib-group (Protocols BGP)	131
tcp-mss (Protocols BGP)	132
traceoptions (Protocols BGP)	133
type (Protocols BGP)	136

Part 3

Chapter 3

Administration

Operational Commands	139
clear bgp damping	140
clear bgp neighbor	141
clear bgp table	143
show bgp bmp	145
show bgp group	146
show bgp neighbor	153
show bgp summary	167
show policy damping	172

List of Tables

	About the Documentation	vii
	Table 1: Notice Icons	ix
	Table 2: Text and Syntax Conventions	ix
Part 1	Overview	
Chapter 1	Feature Support	3
	Table 3: First Junos OS Release for Each EX Series Switch	4
	Table 4: Access Control Features on Switches by Junos OS Release	5
	Table 5: Administration Features on Switches by Junos OS Release	6
	Table 6: Class-of-Service (CoS) Features on Switches by Junos OS Release	7
	Table 7: Class-of-Service (CoS) Features on EX9200 Switches by Junos OS Release	8
	Table 8: Converged Networks (LAN and SAN) Features on Switches by Junos OS Release	9
	Table 9: Device Security Features on Switches by Junos OS Release	10
	Table 10: High Availability and Resiliency Features on Switches by Junos OS Release	11
	Table 11: High Availability and Resiliency Features on EX9200 Switches by Junos OS Release	13
	Table 12: Interfaces Features on Switches by Junos OS Release	14
	Table 13: Interfaces Features on EX9200 Switches by Junos OS Release	15
	Table 14: IP Address Management Features on Switches by Junos OS Release	16
	Table 15: IP Address Management Features on EX9200 Switches by Junos OS Release	16
	Table 16: IPv6 Features on Switches by Junos OS Release	17
	Table 17: Layer 2 Network Protocols Features on Switches by Junos OS Release	18
	Table 18: Layer 2 Networking Features on EX9200 Switches by Junos OS Release	20
	Table 19: Layer 3 Protocols Features on Switches by Junos OS Release	21
	Table 20: Layer 3 Protocols Features on EX9200 Switches by Junos OS Release	23
	Table 21: Logical Systems Features on EX9200 Switches by Junos OS Release	26
	Table 22: MPLS Features on Switches by Junos OS Release	26
	Table 23: MPLS Features on EX9200 Switches by Junos OS Release	28
	Table 24: Multicast Features on Switches by Junos OS Release	30
	Table 25: Multicast Features on EX9200 Switches by Junos OS Release	31

Table 26: Network Management and Monitoring Features on Switches by Junos OS Release	33
Table 27: Network Management and Monitoring Features on EX9200 Switches by Junos OS Release	34
Table 28: Power over Ethernet (PoE) Features on Switches by Junos OS Release	36
Table 29: Port Security Features on Switches by Junos OS Release	37
Table 30: Routing Policy and Packet Filtering Features on Switches by Junos OS Release	38
Table 31: Routing Policy and Firewall Filters on EX9200 Switches by Junos OS Release	39
Table 32: Spanning-Tree Protocols Features on Switches by Junos OS Release	40
Table 33: System Management Features on Switches by Junos OS Release	41
Table 34: User Interface and Configuration Features on EX9200 Switches by Junos OS Release	42
Table 35: VPN Features on EX9200 Switches by Junos OS Release	43

Part 3

Chapter 3

Administration

Operational Commands	139
Table 36: show bgp bmp Output Fields	145
Table 37: show bgp group Output Fields	147
Table 38: show bgp neighbor Output Fields	154
Table 39: show bgp summary Output Fields	167
Table 40: show policy damping Output Fields	172

About the Documentation

- Documentation and Release Notes on page vii
- Supported Platforms on page vii
- Using the Examples in This Manual on page vii
- Documentation Conventions on page ix
- Documentation Feedback on page xi
- Requesting Technical Support on page xi

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:

- EX 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.xml;
    }
  }
}
interfaces {
  fxp0 {
    disable;
    unit 0 {
      family inet {
        address 10.0.0.1/24;
      }
    }
  }
}
```

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

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

Merging a Snippet

To merge a snippet, follow these steps:

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

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

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

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


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

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

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

For more information about the **load** command, see the CLI User Guide.

Documentation Conventions

Table 1 on page ix 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.

Table 2 on page ix 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

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

Convention	Description	Examples
<i>Italic text like this</i>	<ul style="list-style-type: none"> Introduces or emphasizes important new terms. Identifies book 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 System Basics Configuration 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)	Enclose 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 <i>(string1 string2 string3)</i>
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	rsvp { # Required for dynamic MPLS only
[] (square brackets)	Enclose a variable for which you can substitute one or more values.	community name members [community-ids]
Indentation and braces ({ })	Identify 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.	
J-Web GUI Conventions		
Bold text like this	Represents J-Web 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 J-Web 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 send your comments to techpubs-comments@juniper.net, or fill out the documentation feedback form at <https://www.juniper.net/cgi-bin/docbugreport/>. If you are using e-mail, be sure to include the following information with your comments:

- Document or topic name
- URL or page number
- Software release 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 JNASC 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: <https://www.juniper.net/alerts/>

- 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

- [Feature Support on page 3](#)

CHAPTER 1

Feature Support

- [EX Series Switch Software Features Overview on page 3](#)

EX Series Switch Software Features Overview

This topic lists the Juniper Networks EX Series Ethernet Switches software features, the Juniper Networks Junos operating system (Junos OS) release in which they were introduced, and the first Junos OS release for each switch.



NOTE: For Virtual Chassis features, see [EX Series Virtual Chassis Software Features Overview](#).

- [Table 3 on page 4](#)—First Junos OS Release for Each EX Series Switch
- [Table 4 on page 5](#)—Access Control Features
- [Table 5 on page 6](#)—Administration Features
- [Table 6 on page 7](#)—Class-of-Service (CoS) Features
- [Table 7 on page 8](#)—Class-of-Service (CoS) Features for EX9200 Switches
- [Table 8 on page 9](#)—Converged Networks (LAN and SAN) Features
- [Table 9 on page 10](#)—Device Security Features
- [Table 10 on page 11](#)—High Availability and Resiliency Features
- [Table 11 on page 13](#)—High Availability and Resiliency Features on EX9200 Switches
- [Table 12 on page 14](#)—Interfaces Features
- [Table 13 on page 15](#)—Interfaces Features on EX9200 Switches
- [Table 14 on page 16](#)—IP Address Management Features
- [Table 15 on page 16](#)—IP Address Management Features on EX9200 Switches
- [Table 16 on page 17](#)—IPv6 Features
- [Table 17 on page 18](#)—Layer 2 Network Protocols Features
- [Table 18 on page 20](#)—Layer 2 Network Protocols Features on EX9200 Switches
- [Table 19 on page 21](#)—Layer 3 Protocols Features

- [Table 19 on page 21](#)—Layer 3 Protocols Features for EX9200 Switches
- [Table 21 on page 26](#)—Logical Systems Features on EX9200 Switches
- [Table 22 on page 26](#)—MPLS Features
- [Table 23 on page 28](#)—MPLS Features on EX9200 Switches
- [Table 24 on page 30](#)—Multicast Features
- [Table 26 on page 33](#)—Network Management and Monitoring Features
- [Table 27 on page 34](#)—Network Management and Monitoring Features on EX9200 Switches
- [Table 28 on page 36](#)—Power over Ethernet (PoE) Features
- [Table 29 on page 37](#)—Port Security Features
- [Table 30 on page 38](#)—Routing Policy and Packet Filtering Features
- [Table 31 on page 39](#)—Routing Policy and Packet Filtering Features
- [Table 32 on page 40](#)—Spanning-Tree Protocols Features
- [Table 33 on page 41](#)—System Management Features
- [Table 34 on page 42](#)—User Interface and Configuration Features on EX9200 Switches
- [Table 35 on page 43](#)—VPN Features on EX9200 Switches

The Junos OS release for software features on a switch cannot be earlier than the first Junos OS release for that switch.

Table 3: First Junos OS Release for Each EX Series Switch

Switch	Junos OS Release
EX2200 switches*	Junos OS Release 10.1R1 *EX2200-C models: Junos OS Release 11.3R1
EX3200 switches	Junos OS Release 9.0R1
EX3300 switches	Junos OS Release 11.3R1
EX4200 switches	Junos OS Release 9.0R1
EX4500 switches*	Junos OS Release 10.2R1* *EX4500-C models: Junos OS Release 10.3R2
EX4550 switches	Junos OS Release 12.2R1
EX6200 switch	Junos OS Release 11.3R2
EX8208 switches	Junos OS Release 9.4R1
EX8216 switches	Junos OS Release 9.5R1

Table 3: First Junos OS Release for Each EX Series Switch (*continued*)

Switch	Junos OS Release
EX9200 switches	Junos OS Release 12.3R2

Table 4: Access Control Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4500	EX4550	EX6200	EX8200	EX9200
802.1X authentication (port-based, multiple supplicant)	10.1R1	9.0R2	11.3R1	12.1R1	12.2R1	11.3R2	10.2R1	Not supported
802.1X authentication with authentication bypass	10.1R1	9.0R2	11.3R1	12.1R1	12.2R1	11.3R2	10.2R1	Not supported
802.1X authentication with VLAN assignment, VoIP VLAN support	10.1R1	9.0R1	11.3R1	12.1R1	12.2R1	11.3R2	10.3R1	Not supported
802.1X user-based dynamic firewall filters	10.1R1	9.0R2	11.3R1	12.1R1	12.2R1	11.3R2	10.3R1	Not supported
802.1X user-based dynamic firewall filters on multiple-suppliant ports	10.1R1	9.5R2	11.3R1	12.1R1	12.2R1	11.3R2	10.3R1	Not supported
802.1X per-user statistics	10.1R1	9.2R1	11.3R1	12.1R1	12.2R1	11.3R1	10.3R1	Not supported
Authentication fallback	11.3R1	10.3R1	12.3R1	12.1R1	12.2R1	11.3R2	Not supported	Not supported
Captive portal authentication for Layer 3 interfaces	11.3R1	10.1R1	Not supported	12.1R1	12.2R1	Not supported	Not supported	Not supported
Captive portal authentication for Layer 2 interfaces	11.3R1	10.3R1	12.3R1	12.1R1	12.2R1	11.3R2	Not supported	Not supported
Energy Efficient Ethernet (EEE)	12.2R1 (EX2200 -C models only)	Not supported	12.2R1	Not supported	Not supported	12.2R1	12.2R1	Not supported

Table 4: Access Control Features on Switches by Junos OS Release (*continued*)

Feature	EX2200	EX3200, EX4200	EX3300	EX4500	EX4550	EX6200	EX8200	EX9200
Infranet Controller (IC) as an external captive-portal server	12.2R1	12.2R1	12.2R1	12.2R1	12.2R1	Not supported	Not supported	Not supported
MAC RADIUS authentication	10.1R1	9.3R2	11.3R1	10.2R1	12.2R1	11.3R2	10.3R1	Not supported
NetBIOS snooping	11.3R5	11.1R1	11.3R5	Not supported	Not supported	11.3R5	11.1R1	Not supported
Server fail fallback	10.1R1	9.3R2	11.3R1	10.2R1	12.2R1	Not supported	10.2R1	Not supported
TACACS+	10.1R1	9.0R2	11.3R1	10.2R1	12.2R1	11.3R2	9.4R1	12.3R2

Table 5: Administration Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4500	EX4550	EX6200	EX8200	EX9200
System logging (syslog) over IPv4	10.1R1	9.0R2	11.3R1	10.2R1	12.2R1	Not supported	9.4R1	12.3R2
System logging (syslog) over IPv6	10.3R1	9.3R2	11.3R1	10.4R1	12.2R1	11.3R2	10.1R1	12.3R2
System snapshot	Not supported	10.0R1	Not supported	10.2R1	12.2R1	Not supported	10.0R1	12.3R2

Table 6: Class-of-Service (CoS) Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4500	EX4550	EX6200	EX8200	EX9200
Class of service (CoS)—Class-based queuing with prioritization, Layer 2 and Layer 3 classification, rewrite, and queuing; strict priority queuing on egress	10.1R1	9.0R2	11.3R1	10.2R1	12.2R1	11.3R2	9.4R1	See Table 7 on page 8 for a list of EX9200 CoS features
CoS—DSCP, IEEE 802.1p, and IP precedence packet rewrites on routed VLAN interfaces (RVIs)	Not supported	9.5R1	11.3R1	10.2R1	12.2R1	11.3R2	9.4R1	
CoS—Interface-specific classifiers on routed VLAN interfaces (RVIs)	Not supported	9.4R1	11.3R1	11.3R1	12.2R1	Not supported	10.2R1	
CoS—Multidestination	—	—	—	—	—	Not supported	9.5R1	
CoS—Per-interface classification	Not supported	9.3R1	11.3R1	10.2R1	12.2R1	11.3R2	10.2R1	
CoS support on link aggregation groups (LAGs)	10.1R1	9.2R1	11.3R1	10.2R1	12.2R1	Not supported	9.4R1	
CoS support on routed VLAN interfaces (RVIs)	10.3R1	9.4R1	11.3R1	10.4R1	12.2R1	Not supported	9.4R1	
DSCP setting on ingress interface	Not supported	Not supported	Not supported	Not supported	Not supported	Not supported	12.3R1	
Flexible CoS—outer 802.1p marking	Not supported	9.6R1	12.3R1	12.1R1	12.2R1	Not supported	Not supported	
Interface-specific CoS rewrite rules	10.3R1	9.4R1	Not supported	11.2R1	12.2R1	Not supported	10.2R1	
Junos EZQoS for CoS	10.1R1	9.3R2	11.3R1	10.2R1	12.2R1	11.3R2	9.4R1	
Port shaping and queue shaping	10.1R1	9.3R2	11.3R1	10.2R1	12.2R1	11.3R2	10.1R1	

Table 6: Class-of-Service (CoS) Features on Switches by Junos OS Release (*continued*)

Feature	EX2200	EX3200, EX4200	EX3300	EX4500	EX4550	EX6200	EX8200	EX9200
Re-marking of bridged packets	11.2R1	9.4R1	Not supported	10.2R1	12.2R1	11.3R2	10.2R1	
Shaped-deficit weighted round-robin (SDWRR)	10.1R1	9.0R2	11.3R1	10.2R1	12.2R1	11.3R2	9.4R1	
Single-rate two-color marking	10.1R1	9.0R2	11.3R1	10.2R1	12.2R1	11.3R2	9.4R1	

Table 7: Class-of-Service (CoS) Features on EX9200 Switches by Junos OS Release

Feature	Junos OS Release
Assigning forwarding class and DSCP value for Routing Engine generated traffic	12.3R2
BA classification for VPLS based on IEEE 802.1p bits	12.3R2
Classification—Associate packets with CoS servicing levels. Types of classification: <ul style="list-style-type: none"> • Behavior aggregate (BA)—Operates on packets as they enter the switch • Multifield classification—Examines multiple fields in packets. • Fixed classification—Associate a forwarding class with a packet regardless of packet contents. 	12.3R2
Classification and DSCP marking of distributed protocol handler traffic	12.3R2
Classification of control-plane traffic	12.3R2
CoS classification and rewrite for IRB and Layer 2 interfaces and for other Layer 3 interfaces. Port-level queuing, scheduling, and shaping are supported.	12.3R2
Egress GRE classification based on DSCP	12.3R2
IEEE 802.1p inheritance push and swap from transparent tags configuration	12.3R2
Elevated packet drops during oversubscription	12.3R2
Layer 2 policers for the ingress and egress interfaces. Policer types: <ul style="list-style-type: none"> • Single-rate two-color • Single-rate three-color (color-blind and color-aware) • Two-rate three-color (color-blind and color-aware) 	12.3R2
Independent values for DSCP and EXP bits	12.3R2
Apply CoS schedulers on ingress interfaces	12.3R2

Table 7: Class-of-Service (CoS) Features on EX9200 Switches by Junos OS Release (*continued*)

Feature	Junos OS Release
Ingress DSCP bits for multicast traffic over Layer 3 VPNs	12.3R2
Layer 2 traffic policing	12.3R2
Policer support for aggregated Ethernet bundles (link aggregation groups, or LAGs)	12.3R2
Queuing support for logical tunnel interfaces	12.3R2
Rate-limit and excess rate or excess priority option	12.3R2
Re-marking of MVPN GRE encapsulation DCSP at ASBR	12.3R2
Scheduling	12.3R2
Set IPv6 DSCP and MPLS EXP independently	
Set IPv6 DiffServ code point (DSCP) and MPLS EXP independently	12.3R2
Support for Layer 2 policers at the VLAN level	12.3R2
Support for applying a transmit rate limit to logical interfaces on Type 1, 2, or 3 Multiservices PICs	12.3R2
Support for configuring ToS rewrite rules so that DCSP bits of GRE packets are consistent with service providers' CoS policy	12.3R2
Support for copying the TOS bits to the outer IP header on GRE tunnel traffic sent by the Routing Engine	12.3R2
Support for setting the forwarding class and DSCP value for traffic generated by the Routing Engine	12.3R2
Unified command to display all CoS statistics	12.3R2

Table 8: Converged Networks (LAN and SAN) Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4500	EX4550	EX6200	EX8200	EX9200
Data Center Bridging Capability Exchange protocol (DCBX)	Not supported	Not supported	Not supported	11.3R1	12.2R1	Not supported	Not supported	Not supported
DCBX application protocol TLV exchange	Not supported	Not supported	Not supported	12.1R1	12.2R1	Not supported	Not supported	Not supported

NOTE: The EX4500 switch models that support Fibre Channel over Ethernet features must be Converged Enhanced Ethernet (CEE) capable. The CEE-capable EX4500 switch models have a “-C” in the hardware model number. See EX4500 Switch Models.

Table 8: Converged Networks (LAN and SAN) Features on Switches by Junos OS Release (*continued*)

Feature	EX2200	EX3200, EX4200	EX3300	EX4500	EX4550	EX6200	EX8200	EX9200
FIP snooping	Not supported	Not supported	Not supported	10.4R1	12.2R1	Not supported	Not supported	Not supported
Priority-based flow control (PFC)	Not supported	Not supported	Not supported	10.4R1	12.2R1	Not supported	Not supported	Not supported

Table 9: Device Security Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4500	EX4550	EX6200	EX8200	EX9200
Automatic recovery for port error disable conditions	10.1R1	9.6R1	11.3R1	10.2R1	12.2R1	11.3R2	10.0R1	Not supported
Storm control (broadcast and unicast)	10.1R1	9.1R1	11.3R1	10.2R1	12.2R1	11.3R2	9.4R1	Not supported
Storm control (multicast)	10.3R2	10.3R2	Not supported	10.3R2	12.2R1	11.3R2	10.3R2	Not supported
Unknown Layer 2 unicast forwarding	10.1R1	9.3R2	11.3R1	10.2R1	12.2R1	Not supported	10.0R1	12.3R2

Table 10: High Availability and Resiliency Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4500	EX4550	EX6200	EX8200	EX9200
NOTE: For complete lists of Virtual Chassis features, see EX Series Virtual Chassis Software Features Overview.								See Table 11 on page 13 for a list of EX9200 HA features.
Graceful protocol restart for BGP	–	9.0R2	Not supported	–	–	11.3R2	9.4R1	
Graceful protocol restart for IS-IS	–	9.3R2	Not supported	–	–	11.3R2	9.4R1	
Graceful protocol restart for OSPF	–	9.0R2	Not supported	–	–	11.3R2	9.4R1	
Graceful protocol restart for RSVP and LDP	Not supported	Not supported	Not supported	Not supported	Not supported	Not supported	12.3R1	
GRES for ARP entries, forwarding database, and Layer 3 protocols	–	9.2R1 (Virtual Chassis only)	11.3R1	11.2R1 (Virtual Chassis only)	12.2R1 (Virtual Chassis only)	11.3R2	9.4R1	
GRES for IGMP snooping	–	11.3R1 (Virtual Chassis only)	12.1R1 (Virtual Chassis only)	11.4R1 (Virtual Chassis only)	12.2R1 (Virtual Chassis only)	Not supported	11.3R1	
GRES for LACP	Not supported	11.3R1	Not supported	11.3R1	12.2R1	Not supported	11.3R1	
GRES for Layer 2 and Layer 3 VPN LSPs	Not supported	Not supported	Not supported	Not supported	Not supported	Not supported	12.3R1	
GRES for port security (DHCP snooping, DAI, and IP source guard)	–	9.2R1 (Virtual Chassis only)	Not supported	–	–	11.3R2	9.6R1	
LACP support for dual-homing applications in data centers	10.1R1	10.0R1	Not supported	10.2R1	12.2R1	Not supported	10.0R1	
Link Aggregation Control Protocol (LACP)	10.1R1	9.0R2	11.3R1	10.2R1	12.2R1	11.3R2	9.4R1	
Link aggregation groups (LAGs)	10.1R1	9.0R2	11.3R1	10.2R1	12.2R1	11.3R2	9.4R1	
	–					11.3R2	11.1R1	

Table 10: High Availability and Resiliency Features on Switches by Junos OS Release (*continued*)

Feature	EX2200	EX3200, EX4200	EX3300	EX4500	EX4550	EX6200	EX8200	EX9200
Nonstop active routing (NSR) for BGP, IS-IS, IGMP with BFD, and RIP		11.1R1 (Virtual Chassis only)	12.1R1 (Virtual Chassis only)	11.3R1 (Virtual Chassis only)	12.2R1 (Virtual Chassis only)			
Nonstop active routing (NSR) for IPv6 IS-IS, RIPng, and OSPFv3 with BFD	–	12.2R1 (Virtual Chassis only)	12.2R1 (Virtual Chassis only)	12.2R1 (Virtual Chassis only)	12.2R1 (Virtual Chassis only)	Not supported	11.2R1	
Nonstop active routing (NSR) for OSPFv2	–	11.1R1	12.1R1 (Virtual Chassis only)	11.2R1 (Virtual Chassis only)	12.2R1 (Virtual Chassis only)	11.3R2	10.4R1	
Nonstop active routing (NSR) for Protocol Independent Multicast (PIM)	Not supported	12.2R1 (Virtual Chassis only)	12.2R1 (Virtual Chassis only)	12.2R1 (Virtual Chassis only)	12.2R1 (Virtual Chassis only)	Not supported	11.4R2	
Nonstop active routing (NSR) for RSVP and LDP	Not supported	Not supported	Not supported	Not supported	Not supported	Not supported	12.3R1	
Nonstop bridging (NSB) for LAGs and LACP	–	11.4R1 (Virtual Chassis only)	12.2R1 (Virtual Chassis only)	11.4R1 (Virtual Chassis only)	12.2R1 (Virtual Chassis only)	12.1R1	11.3R1	
Nonstop bridging (NSB) for LLDP and LLDP-MED	–	11.3R1 (Virtual Chassis only)	12.2R1 (Virtual Chassis only)	Not supported	Not supported	Not supported	11.3R1	
Nonstop bridging (NSB) for spanning-tree protocols	–	11.3R1 (Virtual Chassis only)	12.2R1 (Virtual Chassis only)	12.1R1 (Virtual Chassis only)	12.2R1 (Virtual Chassis only)	12.1R1	11.3R1	
Nonstop software upgrade (NSSU)	–	12.1R1 (Virtual Chassis only)	12.2R1 (Virtual Chassis only)	12.1R1 (Virtual Chassis only)	12.1R1 (Virtual Chassis only)	12.2R1	10.4R1	
Power budget management	–	–	Not supported	–	–	11.3R2	10.2R1	
	Not supported	Not supported	Not supported	Not supported	Not supported	Not supported	11.4R1	

Table 10: High Availability and Resiliency Features on Switches by Junos OS Release (*continued*)

Feature	EX2200	EX3200, EX4200	EX3300	EX4500	EX4550	EX6200	EX8200	EX9200
Virtual Router—Network Time Protocol (NTP), system logging, Simple Network Management Protocol (SNMP), RADIUS, and TACACS support in a virtual router								
Virtual Router Redundancy Protocol (VRRP)	12.3R1	9.0R2	12.1R1	10.2R1	12.2R1	11.3R2	9.4R1	
Virtual Router Redundancy Protocol (VRRP)—Support for multiple VRRP owners per physical interface	12.3R1	12.3R1	12.3R1	12.3R1	12.3R1	Not supported	12.3R1	
Virtual Router Redundancy Protocol (VRRP) for IPv6 (except authentication type and authentication key)	Not supported	10.2R1	12.3R1	11.2R1	12.2R1	12.1R1	10.1R1	

Table 11: High Availability and Resiliency Features on EX9200 Switches by Junos OS Release

Feature	Junos OS Release
Graceful Routing Engine switchover (GRES)	12.3R2
Nonstop active routing (NSR) support for Protocol Independent Multicast (PIM) for IPv4 and IPv6	12.3R2
Nonstop active routing (NSR) support for VPLS and for LDP-based VPLS	12.3R2
Nonstop active routing (NSR) support for LDP OAM features	12.3R2
Nonstop active routing (NSR) support for Layer 2 VPNs	12.3R2
Virtual Router Redundancy Protocol version 3 (VRRPv3)	12.3R2

Table 12: Interfaces Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4500	EX4550	EX6200	EX8200	EX9200
Digital optical monitoring (DOM)	Not supported	10.0R1	11.3R1	Not supported	Not supported	Not supported	10.0R1	See Table 13 on page 15 for a list of EX9200 interfaces features.
Interface ranges	10.1R1	10.0R1	11.3R1	10.2R	12.2R1	11.3R2	10.1R1	
IPv4 over generic routing encapsulation (GRE) tunnels—encapsulation support	Not supported	12.1R1	Not supported	Not supported	Not supported	Not supported	12.1R1	
IPv4 over generic routing encapsulation (GRE) tunnels—decapsulation support	Not supported	12.1R1	Not supported	Not supported	Not supported	Not supported	12.3R1	
IPv6 over generic routing encapsulation (GRE) tunnels using IPv4 transport—encapsulation support	Not supported	12.1R1	Not supported	Not supported	Not supported	Not supported	12.1R1	
IPv6 over generic routing encapsulation (GRE) tunnels using IPv4 transport—decapsulation support	Not supported	12.1R1	Not supported	Not supported	Not supported	Not supported	Not supported	
IP directed broadcast	Not supported	9.4R1	12.3R1	10.2R1	12.2R1	11.3R2	9.4R1	
Time domain reflectometry (TDR)	10.1R1	9.0R2	11.3R1	Not supported	Not supported	11.3R2	9.4R1	
Unicast reverse-path forwarding (RPF)	Not supported	9.3R2	12.3R1	11.2R1	12.2R1	11.3R2	10.1R1	
VLAN-tagged Layer 3 subinterfaces	Not supported	9.2R1	Not supported	11.2R1	12.2R1	11.3R2	9.4R1	

Table 13: Interfaces Features on EX9200 Switches by Junos OS Release

Feature	Junos OS Release
ICMP redirect	12.3R2
Clear MAC address information	12.3R2
IPv6 subnet support on loopback interfaces	12.3R2
IPv6 support for unnumbered Ethernet interfaces	12.3R2
Multichassis link aggregation (MC-LAG)	12.3R2
Nonstop active routing (NSR) support for Bidirectional Forwarding Detection (BFD)	12.3R2
Protection against distributed denial-of-service (DDOS) attacks	12.3R2
Software support for IPv4 to IPv6 transition	12.3R2
Static mapping for port forwarding	12.3R2
Support for active monitoring on logical systems	12.3R2
Support for VRF in Routing Engine-based sampling	12.3R2
Support for integrated routing and bridging (IRB) MAC synchronization in MC-LAG for aggregated Ethernet	12.3R2
Targeted broadcast for virtual routing and forwarding (VRF) instances	12.3R2
Trunk interface enhancements: <ul style="list-style-type: none"> • Configure a single logical trunk interface to support a list of VLANs or to accept packets with no VLAN tag. • Configure multiple logical trunk interfaces on a single physical interface. 	12.3R2
Unicast reverse-path forwarding (RPF) loose mode, with ability to discard packets with source addresses pointing to the discard interface	12.3R2
Unnumbered Ethernet—Configure IPv4 processing on interfaces without assigning explicit IP addresses to the interfaces.	12.3R2
VLAN rewrite operations on incoming and outgoing frames	12.3R2

Table 14: IP Address Management Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4500	EX4550	EX6200	EX8200	EX9200
DHCP server and relay with option 82 for Layer 2 VLANs	10.1R1	9.3R2	11.3R1	10.2R1	12.2R1	11.3R2	9.4R1	See Table 15 on page 16 for a list of EX9200 IP address management features.
DHCP server and relay with option 82 for Layer 3 interfaces	10.1R1	9.0R1	11.3R1	10.2R1	12.2R1	11.3R2	9.4R1	
DNS for IPv6	Not supported	9.3R2	Not supported	Not supported	Not supported	Not supported	Not supported	
Local DHCP server	10.1R1	9.3R2	11.3R1	12.1R1	12.2R1	11.3R2	9.4R1	
Virtual router aware DHCP (VR-aware DHCP)	Not supported	12.1R1	Not supported	Not supported	Not supported	Not supported	12.1R1	
Virtual router aware DHCPv6 (VR-aware DHCPv6)	Not supported	Not supported	Not supported	12.3R1	12.3R1	12.3R1	Not supported	
Static addresses	10.1R1	9.0R2	11.3R1	10.2R1	12.2R1	Not supported	9.4R1	

Table 15: IP Address Management Features on EX9200 Switches by Junos OS Release

Feature	Junos OS Release
DHCP server and relay	12.3R2
DHCPv6 local server	12.3R2
DHCPv6 support	12.3R2
Distinguishing DHCP duplicate clients by subinterface	12.3R2
Dynamic reconfiguration of extended DHCP and DHCPv6 local server clients	12.3R2
Dynamic IPv6 filters	12.3R2
Expression support for dynamic profiles	12.3R2
Extended DHCP relay proxy	12.3R2
Optional disabling of automatic ARP table population	12.3R2

Table 15: IP Address Management Features on EX9200 Switches by Junos OS Release (*continued*)

Feature	Junos OS Release
IPv6 address assignment pools	12.3R2
Overriding DHCP settings on specific interfaces	12.3R2
Per-interface DHCP tracing operations	12.3R2
S-VLAN-based shaping for dynamic profiles	12.3R2
Sending a DHCP relay and relay proxy release message	12.3R2
Specifying the DHCP source address used for IP packets	12.3R2
Support for MAC address validation	12.3R2
Support for address pool threshold traps	12.3R2
Address assignment pools	12.3R2
Per-interface DHCP lease limits	12.3R2

Table 16: IPv6 Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4500	EX4550	EX6200	EX8200	EX9200
NOTE: A separate software license is required for IPv6. See Understanding Software Licenses for EX Series Switches.								
Application identification (APPID) for IPv6 packets	Not supported	Not supported	Not supported	Not supported	Not supported	Not supported	Not supported	12.3R2
BFD for IPv6	Not supported	12.3R1	Not supported	12.3R1	12.3R1	Not supported	12.3R1	12.3R2 (also for static routes)
BGP for IPv6	Not supported	9.4R1	12.3R1	11.1R1	12.2R1	12.1R1	10.1R1	12.3R2
IPv6 CoS (multifield classification and rewrite)	Not supported	10.2R1	12.3R1	12.1R1	12.2R1	12.1R1	10.4R1	12.3R2
IPv6 management	10.3R1	9.3R2	11.3R1	10.4R1	12.2R1	Not supported	10.1R1	12.3R2

Table 16: IPv6 Features on Switches by Junos OS Release (*continued*)

Feature	EX2200	EX3200, EX4200	EX3300	EX4500	EX4550	EX6200	EX8200	EX9200
IPv6 multicast protocols (PIM, MLDv1/v2)	Not supported	10.1R1	12.3R1	11.2R1	12.2R1	12.1R1	10.2R1	12.3R2
IPv6 neighbor redirect compliance with RFC 4861	12.3R1	12.3R1	12.3R1	12.3R1	12.3R1	12.3R1	12.3R1	12.3R2
IPv6 path MTU discovery	10.3R1	9.3R1	12.3R1	10.4R1	12.2R1	Not supported	10.2R1	12.3R2
IS-IS for IPv6	Not supported	9.4R1	Not supported	11.2R1	12.2R1	12.1R1	10.1R1	12.3R2
MBGP for IPv6	Not supported	9.3R1	12.3R1	Not supported	Not supported	12.1R1	10.1R1	12.3R2
OSPFv3	Not supported	9.3R1	12.3R1	11.1R1	12.2R1	12.1R1	10.1R1	12.3R2
RFC 4291 Compliance	12.3R1	12.3R1	12.3R1	12.3R1	12.3R1	12.3R1	12.3R1	12.3R2
RIPng	Not supported	9.3R1	12.3R1	11.1R1	12.2R1	12.1R1	10.1R1	12.3R2
VRRPv3 (RFC 5798 Compliance, ability to send SNMP traps)	Not supported	12.3R1	Not supported	12.3R1	12.3R1	Not supported	12.3R1	12.3R2

Table 17: Layer 2 Network Protocols Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4500	EX4550	EX6200	EX8200	EX9200
802.1Q VLAN tagging	10.1R1	9.0R2	11.3R1	10.2R1	12.2R1	11.3R2	9.4R1	12.3R2
Edge virtual bridging (EVB) support with virtual Ethernet port aggregator (VEPA)	Not supported	Not supported	Not supported	12.1R1	12.2R1	Not supported	12.1R1	Not supported
Ethernet ring protection switching (ERPS, G.8032/Y.1344)	12.1R1	12.1R1	12.3R1	12.3R1	12.3R1	Not supported	12.3R1	12.3R2

Table 17: Layer 2 Network Protocols Features on Switches by Junos OS Release (*continued*)

Feature	EX2200	EX3200, EX4200	EX3300	EX4500	EX4550	EX6200	EX8200	EX9200
Layer 2 protocol tunneling (L2PT)	11.1R1	10.0R1	12.3R1	11.2R1	12.2R1	12.1R1	Not supported	12.3R2
Link Layer Discovery Protocol (LLDP)	10.1R1	9.0R2	11.3R1	10.2R1	12.2R1	11.3R2	9.4R1	12.3R2
Link Layer Discovery Protocol—Media Endpoint Discovery (LLDP-MED) with voice over IP (VoIP) integration	10.1R1	9.0R2	11.3R1	Not supported	Not supported	Not supported	Not supported	Not supported
MAC-based VLANs	10.1R1	9.0R2	11.3R1	10.2R1	12.2R1	11.3R2	9.4R1	12.3R2
Multiple VLAN Registration Protocol (MVRP, IEEE 802.1ak)	11.3R1	10.0R1	12.3R1	11.2R1	12.2R1	12.1R1	10.0R1	12.3R2
Private VLANs (PVLANS)	11.1R1	9.3R2	12.1R1	11.2R1	12.2R1	11.3R2	10.1R1	Not supported
Private VLANs (PVLANS) support across switches	11.1R1	10.4R1	12.1R1	11.2R1	12.2R1	11.3R2	11.2R1	Not supported
Proxy ARP—Restricted	10.1R1	10.0R1	11.3R1	10.2R1	12.2R1	Not supported	10.0R1	12.3R2
Proxy ARP—Unrestricted	10.1R1	9.6R1	11.3R1	10.2R1	12.2R1	12.1R1	10.0R1	12.3R2
Proxy ARP per VLAN	10.1R1	10.1R1	Not supported	10.2R1	12.2R1	Not supported	10.1R1	
Q-in-Q tunneling	11.1R1	9.3R2	11.4R1	11.2R1	12.2R1	12.1R1	11.1R1	12.3R2
Q-in-Q VLAN extended support for multiple S-VLANs per access interface, firewall-filter-based VLAN assignment, and routed VLAN interfaces (RVIs)	Not supported	9.6R1	12.3R1	11.2R1	12.2R1	12.1R1	11.1R1	12.3R2

Table 17: Layer 2 Network Protocols Features on Switches by Junos OS Release (*continued*)

Feature	EX2200	EX3200, EX4200	EX3300	EX4500	EX4550	EX6200	EX8200	EX9200
Redundant trunk groups	10.1R1	9.0R2	11.3R1	11.2R1	12.2R1	11.3R2	9.4R1	Not supported
Routed VLAN interfaces (RVIs)—Also known as integrated routing and bridging (IRB) interfaces	10.1R1	9.0R2	11.3R1	10.2R1	12.2R1	11.3R2	9.4R1	12.3R2
VLAN ID translation	11.1R1	10.0R1	Not supported	11.2R1	12.2R1	Not supported	11.1R1	
VLAN ranges	10.1R1	9.2R1	11.3R1	10.2R1	12.2R1	11.3R2	9.4R1	12.3R2

Table 18: Layer 2 Networking Features on EX9200 Switches by Junos OS Release

Feature	Junos OS Release
VLANs and virtual switching	12.3R2
DHCP support for integrated routing and bridging (IRB)	12.3R2
MC-LAG support for IGMP snooping in IRB	12.3R2
Hash-key load-balancing support for Layer 3 and Layer 4 fields	12.3R2
IP multicast over Layer 2 trunk port support	12.3R2
Integrated routing and bridging (IRB)	12.3R2
Layer 2 Ethernet OAM: <ul style="list-style-type: none"> Distributed periodic packet management process (ppmd) for improved scaling Graceful Routing Engine switchover (GRES) Remote defect indication (RDI) Configuration of action profiles 	12.3R2
Layer 2 address learning in logical systems	12.3R2
Layer 2 forwarding support for bridging and VPLS	12.3R2
Layer 2 policer statistics MIB	12.3R2
Firewall filter match conditions for Layer 2 bridging and VPLS	12.3R2
Next-hop groups using either IP addresses or Layer 2 addresses for the next hop	12.3R2

Table 18: Layer 2 Networking Features on EX9200 Switches by Junos OS Release (*continued*)

Feature	Junos OS Release
Unicast reverse-path forwarding (RPF) loose mode, with ability to discard packets with source addresses pointing to the discard interface	12.3R2
Spanning-tree protocols support for Layer 2 bridging and VPLS	12.3R2
VLAN rewrite operations on incoming and outgoing frames	12.3R2
STP root guard (root protection)	12.3R2
Support for Layer 2 and Layer 2.5 features: <ul style="list-style-type: none"> Extensive set of Layer 2 label-manipulation capabilities, Q-in-Q support MC-LAG active / standby and MC-LAG active / active xSTP protocol support Integrated Routing and Bridging (IRB) interface support IGMP snooping for multichassis link aggregation group (MC-LAG) interfaces Configurable label block sizes for VPLS Connectivity fault management process flooding to interfaces based on mesh groups Layer 2 address learning in logical systems Virtual switch support, providing virtual Layer 2 switch instances with separate Layer 2 learning domains, isolated 4K VLAN ID spaces, and STP instances Ethernet Ring Protocol (ERP) for multiple ring instances on the same physical ring Transit and bypass static label-switched paths (LSPs) Layer 2 Gigabit Ethernet logical interface policing Static LSP statistics Multiple VLAN Registration Protocol (MVRP)-IEEE 802.1ak-2007 	12.3R2
VPLS root protection topology change-action control	12.3R2

Table 19: Layer 3 Protocols Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4500	EX4550	EX6200	EX8200	EX9200
Bidirectional Forwarding Detection (BFD)	11.3R1	9.0R2	12.1R1	10.2R1	12.2R1	12.1R1	9.4R1	See Table 20 on page 23 for a list of EX9200 Layer 3 protocols features.
Border Gateway Protocol (BGP)	Not supported	9.0R2	12.1R1	11.1R1	12.2R1	11.3R2	9.4R1	

Table 19: Layer 3 Protocols Features on Switches by Junos OS Release (*continued*)

Feature	EX2200	EX3200, EX4200	EX3300	EX4500	EX4550	EX6200	EX8200	EX9200
Multiprotocol Border Gateway Protocol (MBGP)	Not supported	9.3R1	12.3R1	11.2R1	12.2R1	12.1R1	9.4R1	
A separate software license is required for BGP and MBGP. See Understanding Software Licenses for EX Series Switches.								
Distributed periodic packet management (PPM) with BFD	Not supported	10.4R1	Not supported	Not supported	Not supported	12.1R1	10.4R1	
Distributed periodic packet management (PPM) with LACP	Not supported	10.2R1	Not supported	11.1R1	12.2R1	11.3R2	10.2R1	
Filter-based forwarding	Not supported	9.4R1	12.3R1	11.2R1	12.2R1	11.3R2	9.6R1	
Filter-based forwarding over IPv6	Not supported	10.1R1	Not supported	Not supported	Not supported	Not supported	10.3R1	
Intermediate System-to-Intermediate System (IS-IS)	Not supported	9.0R2	Not supported	11.1R1	12.2R1	11.3R2	9.4R1	
A separate software license is required for IS-IS. See Understanding Software Licenses for EX Series Switches.								
IPv6 Layer 3 multicast protocols	Not supported	10.1R1	Not supported	Not supported	Not supported	Not supported	10.2R1	
Jumbo frames on routed VLAN interfaces (RVIs)	Not supported	9.4R1	11.3R1	10.2R1	12.2R1	Not supported	9.4R1	
OSPF Multitopology Routing (MT-OSPF)	Not supported	9.5R1	Not supported	Not supported	Not supported	Not supported	Not supported	
See the <i>Junos OS Routing Protocols Configuration Guide</i> .								
OSPFv2	11.1R1	9.0R2	11.4R1	10.2R1	12.2R1	11.3R2	9.4R1	
OSPFv3 IPsec support	Not supported	10.3R1	Not supported	Not supported	Not supported	Not supported	Not supported	

Table 19: Layer 3 Protocols Features on Switches by Junos OS Release (*continued*)

Feature	EX2200	EX3200, EX4200	EX3300	EX4500	EX4550	EX6200	EX8200	EX9200
Routing Information Protocol version 1 (RIPv1) and RIPv2	10.1R1	9.0R2	11.3R1	10.2R1	12.2R1	11.3R2	9.4R1	
Static routes	10.1R1	9.0R2	11.3R1	10.2R1	12.2R1	11.3R2	9.4R1	
Virtual routing and forwarding (VRF) with IPv4—Virtual routing instances	12.3R1	9.2R1	12.3R1	11.1R1	12.2R1	11.3R2	9.6R1	
VRF with IPv4—Virtual routing instances for PIM and IGMP	Not supported	10.0R1	Not supported	11.1R1	12.2R1	11.3R2	10.0R1	
VRF with IPv4—Virtual routing instances for IGMP snooping	Not supported	11.4R1	Not supported	12.1R1	12.2R1	Not supported	11.3R1	
VRF with IPv6—Virtual routing instances for multicast traffic	Not supported	10.1R1	Not supported	Not supported	Not supported	Not supported	10.1R1	
VRF with IPv6—Virtual routing instances for unicast traffic	Not supported	10.1R1	12.3R1	Not supported	Not supported	Not supported	10.1R1	

Table 20: Layer 3 Protocols Features on EX9200 Switches by Junos OS Release

Feature	Junos OS Release
Accumulated IGP attribute for BGP	12.3R2
Advertisement of the best external BGP path to internal peers	12.3R2
Alias support for local autonomous system numbers for BGP	12.3R2
BFD liveness detection	12.3R2
BFD protocol support for OSPFv3	12.3R2
BGP remote next-hop support for single-hop EBGP peers	12.3R2
BGP support for 4-byte autonomous system numbers	12.3R2

Table 20: Layer 3 Protocols Features on EX9200 Switches by Junos OS Release (*continued*)

Feature	Junos OS Release
BGP support for MDT-SAFI updates without a route target	12.3R2
Behavior change for BGP-independent autonomous system (AS) domains	12.3R2
Bidirectional Forwarding Detection (BFD) hold-down timer	12.3R2
Distributed periodic packet management support for aggregate interfaces	12.3R2
Egress filtering PIMv4/v6 join messages	12.3R2
For internal BGP (IBGP), advertise multiple paths to a destination	12.3R2
Frequent BGP keepalive messages and short BGP hold time	12.3R2
Hitless authentication key rollover for IS-IS	12.3R2
Hub-and-spoke support for multiprotocol BGP-based multicast VPNs with PIM-SSM GRE S-PMSI transport	12.3R2
IPv4 subnet support on loopback interfaces	12.3R2
IS-IS hold-down timer for subsequent SPF calculations	12.3R2
Keepalive support for GRE interfaces	12.3R2
Multitopology routing (MTR)	12.3R2
Nonstop active routing (NSR) support for the Routing Information Protocol (RIP) and RIP next generation (RIPng)	12.3R2
Nonstop active routing (NSR) support	12.3R2
OSPF graceful restart enhancement	12.3R2
OSPF hold-down timer for subsequent SPF calculations	12.3R2
Only the system log notes failure to add routes	12.3R2
Origin validation for BGP	12.3R2
PIM join suppression support	12.3R2
Priority assignment for prefixes in OSPF import policies	12.3R2
Reduction in flooding of self-originated OSPF LSAs	12.3R2
Support for BFD over multihop static routes	12.3R2

Table 20: Layer 3 Protocols Features on EX9200 Switches by Junos OS Release (*continued*)

Feature	Junos OS Release
Support for BFD on logical switches	12.3R2
Support for IPSec authentication for OSPFv2	12.3R2
Support for OSPF database protection for OSPF and OSPFv3	12.3R2
Support for OSPF export and import policies for network-summary LSAs	12.3R2
Support for alternate loop-free routes for IS-IS and OSPF	12.3R2
Support for clearing the VPN tag	12.3R2
Support for disabling the attribute set messages on independent AS domains for BGP loop detection	12.3R2
Support for disabling traps for passive OSPFv2 interfaces	12.3R2
Support for display of flood next-hop branch overflow condition	12.3R2
Support for dropping and ignoring path attributes during BGP neighbor updates	12.3R2
Support for the algorithm that determines the single best path to skip the step that evaluates an AS path	12.3R2
Support for limiting the number of prefixes accepted from a BGP peer	12.3R2
Support for multiarea adjacency in OSPFv2	12.3R2
Support for multiple address families in OSPFv3	12.3R2
Support for route leaking when the switch is in overload mode	12.3R2
Support for route-filter-based BGP outbound route filtering	12.3R2
Support for the BGP Monitoring Protocol	12.3R2
Support to hold down BGP peering sessions after a nonstop active routing (NSR) switchover Timer to delay MED updates for routes advertised by BGP groups or peers configured with the metric-out igp statement Virtual Router Redundancy Protocol (VRRP)	12.3R2
Timer to delay MED updates for routes advertised by BGP groups or peers configured with the metric-out igp statement	12.3R2
Virtual Router Redundancy Protocol (VRRP)	12.3R2

Table 21: Logical Systems Features on EX9200 Switches by Junos OS Release

Feature	Junos OS Release
A separate software license is required for logical systems. See Understanding Software Licenses for EX Series Switches.	
Layer 2 address learning in logical systems	12.3R2
	12.3R2

Table 22: MPLS Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4500	EX4550	EX6200	EX8200	EX9200
A separate software license is required for MPLS. See Understanding Software Licenses for EX Series Switches.								
Aggregated Ethernet interfaces (LAGs) on circuit cross-connects (CCCs)	Not supported	Not supported	Not supported	12.2R1	12.2R1	Not supported	11.1R1	See Table 23 on page 28 for a list of EX9200 MPLS features
BFD for an LDP-based LSP	Not supported	Not supported	Not supported	Not supported	Not supported	Not supported	12.2R1	
BFD for an RSVP-based LSP	Not supported	Not supported	Not supported	Not supported	Not supported	Not supported	12.2R1	
CCC between 2 interfaces in the same switch	Not supported	Not supported	Not supported	12.2R1	12.2R1	Not supported	11.1R1	
Interior gateway protocol (IGP) IS-IS and OSPF shortcuts	Not supported	Not supported	Not supported	12.2R1	12.2R1	Not supported	11.1R1	
IP over MPLS	Not supported	10.1R1	Not supported	12.2R1	12.2R1	Not supported	11.1R1	
IPv6 over MPLS label-switched paths (LSPs)	Not supported	Not supported	Not supported	12.2R1	12.2R1	Not supported	12.1R1	
LDP-based MPLS	Not supported	Not supported	Not supported	12.2R1	12.2R1	Not supported	11.1R1	
LDP tunneling (LDP over RSVP)	Not supported	Not supported	Not supported	12.2R1	12.2R1	Not supported	11.1R1	

Table 22: MPLS Features on Switches by Junos OS Release (*continued*)

Feature	EX2200	EX3200, EX4200	EX3300	EX4500	EX4550	EX6200	EX8200	EX9200
MPLS-based circuit cross-connects (CCC)	Not supported	9.5R1	Not supported	12.2R1	12.2R1	Not supported	11.1R1	
MPLS label-switched router (LSR) support	Not supported	Not supported	Not supported	12.2R1	12.2R1	Not supported	11.1R1	
MPLS Layer 2 CCC on Ethernet-encapsulated interfaces (RFC 6624)	Not supported	9.5R1	Not supported	12.2R1	12.2R1	Not supported	11.1R1	
MPLS Layer 2 CCC on VLAN-encapsulated interfaces (RFC 4905)	Not supported	Not supported	Not supported	12.2R1	12.2R1	Not supported	11.1R1	
MPLS Layer 2 VLAN CCC on Ethernet-encapsulated interfaces (RFC 6624)	Not supported	9.5R1	Not supported	12.2R1	12.2R1	Not supported	11.3R1	
MPLS Layer 2 VLAN CCC on VLAN-encapsulated interfaces (RFC 4905)	Not supported	Not supported	Not supported	12.2R1	12.2R1	Not supported	11.3R1	
MPLS Layer 2 VPN over CCC	Not supported	Not supported	Not supported	12.2R1	12.2R1	Not supported	11.1R1	
MPLS Layer 2 VPN over VLAN CCC	Not supported	Not supported	Not supported	12.2R1	12.2R1	Not supported	11.3R1	
MPLS OAM-LSP ping	Not supported	Not supported	Not supported	Not supported	Not supported	Not supported	11.1R1	
MPLS over untagged Layer 3 interfaces	Not supported	Not supported	Not supported	12.2R1	12.2R1	Not supported	11.1R1	
MPLS with class of service (CoS)	Not supported	Not supported	Not supported	Not supported	Not supported	Not supported	12.1R1	
MPLS Layer 3 VPNs	Not supported	Not supported	Not supported	12.2R1	12.2R1	Not supported	11.1R1	

Table 22: MPLS Features on Switches by Junos OS Release (*continued*)

Feature	EX2200	EX3200, EX4200	EX3300	EX4500	EX4550	EX6200	EX8200	EX9200
MPLS with RSVP-based label-switched paths (LSPs)	Not supported	9.5R1	Not supported	12.2R1	12.2R1	Not supported	11.1R1	
Layer 3 subinterfaces as MPLS core interfaces	Not supported	Not supported	Not supported	12.2R1 See “Note” on page 28	12.2R1 See “Note” on page 28	Not supported	12.1R1	
Routed VLAN interfaces (RVIs) as MPLS core interfaces	Not supported	Not supported	Not supported	Not supported	Not supported	Not supported	12.1R1	
Path maximum transmission unit (MTU) and unicast reverse-path forwarding (RPF) checks for VPNs	Not supported	Not supported	Not supported	12.2R1	12.2R1	Not supported	11.1R1	
Resource Reservation Protocol—traffic engineering (RSVP-TE)	Not supported	Not supported	Not supported	12.2R1	12.2R1	Not supported	11.1R1	
Standby secondary path protection	Not supported	12.1R1	Not supported	12.2R1	12.2R1	Not supported	11.1R1	
Static label-switched paths (LSPs)	Not supported	Not supported	Not supported	12.2R1	12.2R1	Not supported	12.1R1	



NOTE: For EX4500 and EX4550 switches to support Layer 3 subinterfaces as MPLS core interfaces, the peer switch that the Layer 3 subinterfaces connect to, must be an EX8200 switch.

Table 23: MPLS Features on EX9200 Switches by Junos OS Release

Feature	Junos OS Release
A separate software license is required for MPLS. See Understanding Software Licenses for EX Series Switches.	
Bypass static LSPs	12.3R2

Table 23: MPLS Features on EX9200 Switches by Junos OS Release (*continued*)

Feature	Junos OS Release
LDP LSP action based on a BFD failure event	12.3R2
LDP downstream on demand	12.3R2
LDP, BGP, and VPLS interworking	12.3R2
P2MP LSP traceroute	12.3R2
Static LSP: <ul style="list-style-type: none"> • Revert timer • Statistics • Traceoptions • At the ingress switch • At the transit switch 	12.3R2
Statistics for P2MP LSPs	12.3R2

Table 24: Multicast Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4500	EX4550	EX6200	EX8200	EX9200
IGMP (Internet Group Management Protocol) version 1 (IGMPv1) and IGMPv2	11.1R1	9.0R2	12.1R1	10.2R1	12.2R1	11.3R2	9.4R1	See Table 25 on page 31 for a list of EX9200 multicast features
IGMP filtering	11.3R1	9.5R1	12.3R1	11.3R1	11.3R1	11.3R1	9.5R1	
IGMP snooping with routed VLAN interfaces (RVIs)	10.1R1	9.2R1	12.1R1	10.2R1	12.2R1	Not supported	9.4R1	
IGMPv3	11.1R1	9.3R2	12.1R1	10.2R1	12.2R1	11.3R2	9.6R1	
IGMPv1 and IGMPv2 snooping	10.1R1	9.1R1	11.3R1	10.2R1	12.2R1	11.3R2	9.4R1	
IGMPv3 snooping	10.1R1	9.6R1	11.3R1	10.2R1	12.2R1	11.3R2	9.6R1	
Multicast Listener Discovery version 1 and 2 (MLDv1 and MLDv2)	Not supported	10.1R1	Not supported	11.2R1	12.2R1	12.1R1	10.2R1	
Multicast Listener Discovery version 1 (MLDv1) snooping (MLDv1 snooping)	12.1R1	12.1R1	12.1R1	12.1R1	12.2R1	12.1R1	12.1R1	
Multicast Listener Discovery version 2 (MLDv2) snooping (MLDv2 snooping)	12.1R1	12.1R1	12.1R1	12.1R1	12.2R1	12.1R1	12.1R1	
Multicast Source Discovery Protocol (MSDP)	Not supported	9.4R1	12.3R1	10.2R1	12.2R1	12.1R1	9.4R1	
See the Junos OS Multicast Protocols Configuration Guide .								
Multicast VLAN registration (MVR)	11.3R1	9.6R1	12.1R1	Not supported	Not supported	Not supported	Not supported	
Protocol Independent Multicast dense mode (PIM DM)	11.1R1	9.2R1	12.1R1	11.2R1	12.2R1	11.3R2	9.4R1	
See the Junos OS Multicast Protocols Configuration Guide .								

Table 24: Multicast Features on Switches by Junos OS Release (*continued*)

Feature	EX2200	EX3200, EX4200	EX3300	EX4500	EX4550	EX6200	EX8200	EX9200
Protocol Independent Multicast sparse mode (PIM SM)	11.1R1	9.0R2	12.1R1	10.2R1	12.2R1	11.3R2	9.4R1	
See the Junos OS Multicast Protocols Configuration Guide .								
Protocol Independent Multicast source-specific multicast (PIM SSM)	11.1R1	9.3R1	12.1R1	10.2R1	12.2R1	11.3R2	9.4R1	
See the Junos OS Multicast Protocols Configuration Guide .								
Single-source multicast	Not supported	9.0R2	Not supported	Not supported	Not supported	Not supported	9.4R1	

Table 25: Multicast Features on EX9200 Switches by Junos OS Release

Feature	Junos OS Release
BFD for PIM—IPv6	12.3R2
BFD support for ECMP LSPs signaled using LDP	12.3R2
Bidirectional PIM (RFC 5015)	12.3R2
Control of PIM resources for multicast VPNs	12.3R2
Disable PIM for IPv6 only	12.3R2
Dynamic reuse of data multicast distribution tree (MDT) group addresses	12.3R2
Flexible configuration for IGMP or MLD static-join	12.3R2
IGMPv3 and MLDv2 full support	12.3R2
IGMP and MLD enhancements— <ul style="list-style-type: none"> • immediate-leave (IGMP and MLD) • promiscuous-mode (IGMP only) 	12.3R2
IGMP and PIM support for unnumbered interfaces	12.3R2
IGMP join and leave recording for system or for specific interfaces	12.3R2

Table 25: Multicast Features on EX9200 Switches by Junos OS Release (*continued*)

Feature	Junos OS Release
IGMP and MLD source or group access lists and MLD join and leave recording	12.3R2
IGMP and MLD support for dynamic interfaces	12.3R2
Independently configurable loopback addresses for VRF VPNs	12.3R2
Internet multicast using ingress replication provider tunnels	12.3R2
Software support for configuring accept any-source multicast (ASM) join messages (*G) for group addresses	12.3R2
Software support for configuring a provider network to operate in source-specific multicast (SSM) mode	12.3R2
LDP signaling for point-to-multipoint LSPs in next-generation MBGP multicast VPNs	12.3R2
Load-balancing PIM join messages on multicast VPNs	12.3R2
Multicast flow maps	12.3R2
Nonstop active routing (NSR) PIM for Draft-Rosen VPNs	12.3R2
PIM automatic make-before-break (MBB) join load balancing	12.3R2
PIM join load balancing	12.3R2
Source-specific multicast (SSM)-map definition for different groups to different sources	12.3R2
Support for filtering unwanted PIM neighbors	12.3R2
Support for multicast output interface (OIF) mapping	12.3R2
Translation of PIM join/prune messages to IGMP or MLD report/leave messages	12.3R2
Turn off spanning-tree interface state in multicast snooping	12.3R2

Table 26: Network Management and Monitoring Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4500	EX4550	EX6200	EX8200	EX9200
802.1ag Ethernet OAM connectivity fault management (CFM)	11.2R1	10.2R1	12.3R1	12.2R1	12.2R1	Not supported	11.4R1	See Table 27 on page 34 for a list of EX9200 network management and monitoring features.
Ethernet frame delay measurement (ETH-DM, Y.1731)	Not supported	11.4R1 (EX4200 only)	Not supported	11.4R1	12.2R1	Not supported	11.4R1	
Ethernet OAM link fault management (LFM—also known as Ethernet in the First Mile, EFM)	11.1R1	9.4R1	12.2R1	12.2R1	12.2R1	Not supported	10.0R1	
Port mirroring	10.1R1	9.0R2	11.3R1	10.2R1	12.2R1	11.3R2	9.4R1	
Port mirroring enhancements <ul style="list-style-type: none"> Layer 3 interface support Multiple VLAN support 	Not supported	9.5R1	Not supported	Not supported	Not supported	Not supported	9.5R1	
Port mirroring enhancements <ul style="list-style-type: none"> For remote port mirroring, ingress and egress options on VLAN member interfaces on the intermediate (transit) switch to avoid flooding mirrored traffic to those interfaces 	Not supported	10.0R1	Not supported	Not supported	Not supported	Not supported	Not supported	
Port mirroring support for multiple analyzers per session	Not supported	Not supported	Not supported	11.2R1	12.2R1	Not supported	Not supported	
Real-time performance monitoring (RPM)	10.1R1	9.3R2	12.2R1	10.2R1	12.2R1	12.1R1	10.1R1	
	Not supported	10.3R1	12.2R1	10.2R1	12.2R1	12.1R1	10.3R1	

Table 26: Network Management and Monitoring Features on Switches by Junos OS Release (*continued*)

Feature	EX2200	EX3200, EX4200	EX3300	EX4500	EX4550	EX6200	EX8200	EX9200
Real-time performance monitoring (RPM)—hardware timestamps with routed VLAN interfaces (RVIs)								
Real-time performance monitoring (RPM)—client and server on same interface	10.3R1	10.3R1	12.2R1	11.1R1	12.2R1	Not supported	10.3R1	
Routing Engine Software Development Kit (SDK)	Not supported	12.2R1	12.2R1 (EX4200 only)	12.2R1	12.2R1	Not supported	12.2R1	
RMON	10.1R1	9.0R2	11.3R1	10.2R1	12.2R1	11.3R2	9.4R1	
sFlow monitoring technology	11.1R1	9.3R2	12.1R1	11.2R1	12.2R1	12.1R1	10.0R1	
sFlow monitoring technology—Persistent IP addresses for agent IDs and use in datagrams	11.1R1	10.2R1	Not supported	Not supported	Not supported	12.1R1	10.2R1	
Simple Network Management Protocol version 1 (SNMPv1), SNMPv2, and SNMPv3	10.1R1	9.0R2	11.3R1	10.2R1	12.2R1	11.3R2	9.4R1	
Uplink failure detection	11.1R1	11.1R1	12.1R2	11.1R1	12.2R1	Not supported	12.1R1	

Table 27: Network Management and Monitoring Features on EX9200 Switches by Junos OS Release

Feature	Junos OS Release
Junos OS XML API and scripting—NETCONF Java toolkit for rapid development of Java applications to manage devices running Junos OS	12.3R2

Table 27: Network Management and Monitoring Features on EX9200 Switches by Junos OS Release (*continued*)

Feature	Junos OS Release
Junos OS XML API and scripting—NETCONF Perl client installation—Supports loading prerequisites from Comprehensive Perl Archive Network (CPAN) global repository	12.3R2
Junos OS XML API and scripting—NETCONF tracing operations	12.3R2
Junos OS XML API and scripting: <ul style="list-style-type: none"> • Dedicated directory for user script library • Global variable provided to Junos OS automation scripts • References to a correlating event in a policy action • Trigger a policy based on the event count • Unique filenames for uploaded files • Upload files created by event scripts • XML schemata for Junos OS XML operational tag elements • jcs:open() extension function support for routing instances 	12.3R2
Configuration options to filter out interfaces from SNMP Get and GetNext operations	12.3R2
Enhanced SNMP support for logical switches and routing instances	12.3R2
Generating SNMP traps when MAC address table is full	12.3R2
Junos OS MIB support for VPLS	12.3R2
MIB support for VRF route entries	12.3R2
Proxy SNMP agent	12.3R2
SNMP MIB support for OSPFv3	12.3R2
SNMP poll and trap support for DHCP leases	12.3R2
SNMP support for the DHCP bindings table	12.3R2
SNMP support for the authd daemon and for radius-acc-server-mib and radius-auth-server-mib	12.3R2
SNMP support for spanning-tree protocols	12.3R2
Support for Internet draft draft-ietf-bfd-mib-02.txt—MIB for BFD liveness detection	12.3R2
Support for MIB objects in accounting profiles	12.3R2
Support for an enterprise-specific event MIB (mib-jnx-event.txt)	12.3R2
Support for sending traps over routing instances	12.3R2

Table 27: Network Management and Monitoring Features on EX9200 Switches by Junos OS Release (*continued*)

Feature	Junos OS Release
Support for adding lists of clients to the SNMP community	12.3R2
Support for the enterprise-specific Packet Forwarding Engine MIB (mib-jnx-pfe.txt)	12.3R2
Support for the pimNeighborLoss trap	12.3R2
Support for trap spoofing	12.3R2
IEEE 802.3ah link fault management (LFM) for Ethernet OAM (also known as Ethernet in the First Mile, or EFM)	12.3R2
Port mirroring of Layer 2 VLAN and VPLS traffic	12.3R2
Fast update filters for dynamic profiles	12.3R2
Flow aggregation to multiple collectors	12.3R2
IPv6 flow aggregation templates	12.3R2
Inline flow monitoring	12.3R2

Table 28: Power over Ethernet (PoE) Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4500	EX4550	EX6200	EX8200	EX9200
Link Layer Discovery Protocol (LLDP) with granular Power over Ethernet (PoE) management	12.2R1	12.2R1 (EX4200-24PX and EX4200-48PX switch models only)	12.2R1	Not supported	Not supported	12.2R1	12.2R1	Not supported
NOTE: EX4200 switches must be running PoE controller software firmware version 4.04 or later to support the Link Layer Discovery Protocol (LLDP) with granular Power over Ethernet (PoE) management feature. See show chassis firmware detail and request system firmware upgrade poe to check or upgrade this firmware.								Not supported
Power over Ethernet (PoE)	10.1R1	9.0R2	11.3R1	—	—	11.3R2	11.2R1	Not supported
Power over Ethernet Plus (PoE+)	10.3R1	Not supported	11.3R1	—	—	11.3R2	11.2R1	Not supported

Table 28: Power over Ethernet (PoE) Features on Switches by Junos OS Release (*continued*)

Feature	EX2200	EX3200, EX4200	EX3300	EX4500	EX4550	EX6200	EX8200	EX9200
Power over Ethernet (PoE) power management mode	10.1R1	9.3R2	11.3R1	—	—	11.3R2	11.2R1	Not supported

Table 29: Port Security Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4500	EX4550	EX6200	EX8200	EX9200
Automatic recovery for port error disable conditions	10.1R1	9.6R1	11.3R1	10.2R1	12.2R1	11.3R2	10.0R1	Not supported
DHCP option 82	10.1R1	9.3R2	11.3R1	10.2R1	12.2R1	Not supported	9.4R1	12.3R2
DHCP snooping	10.1R1	9.0R2	11.3R1	12.1R1	12.2R1	11.3R2	10.3R1	Not supported
Dynamic ARP inspection (DAI)	10.1R1	9.0R2	11.3R1	12.1R1	12.2R1	11.3R2	10.3R1	Not supported
IP source guard	10.1R1	9.2R1	11.3R1	12.1R1	12.2R1	11.3R2	10.3R1	Not supported
Layer 3 virtual private network (VPN) for IPv4 (RFC 2547 and 4364)	Not supported	Not supported	Not supported	12.2R1	12.2R1	Not supported	11.1R1	12.3R2
Layer 3 virtual private network (VPN) for IPv6 through IPv4 MPLS	Not supported	Not supported	Not supported	12.2R1	12.2R1	Not supported	11.1R1	12.3R2
MAC limiting	10.1R1	9.0R2	11.3R1	10.2R1	12.2R1	11.3R2	9.4R1	12.3R2
MAC address limit per port	10.1R1	9.0R1	11.3R1	10.2R1	12.2R1	11.3R2	10.3R1	12.3R2
MAC limiting per port and per VLAN (VLAN membership MAC limit)	12.3R1	12.3R1	12.3R1	12.3R1	12.3R1	12.3R1	12.3R1	12.3R2
MAC move limiting	10.1R1	9.0R2	11.3R1	Not supported	Not supported	11.3R2	Not supported	Not supported

Table 29: Port Security Features on Switches by Junos OS Release (*continued*)

Feature	EX2200	EX3200, EX4200	EX3300	EX4500	EX4550	EX6200	EX8200	EX9200
Persistent MAC learning (sticky MAC)	11.4R1	11.4R1	12.3R1	11.4R1	12.2R1	11.4R1	11.4R1	Not supported
Persistent storage for DHCP snooping	10.1R1	9.4R1	11.3R1	12.1R1	12.2R1	11.3R2	10.3R1	Not supported
Self-signed digital certificates for enabling SSL services	11.1R1	11.1R1	Not supported	11.1R1	12.2R1	12.1R1	11.1R1	Not supported
Static ARP support	10.1R1	9.0R2	11.3R1	10.2R1	12.2R1	11.3R2	9.4R1	Not supported

Table 30: Routing Policy and Packet Filtering Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4500	EX4550	EX6200	EX8200	EX9200
Dynamic allocation of TCAM memory to firewall filters	10.1R1	10.0R1	11.3R1	10.2R1	12.2R1	Not supported	10.3R1	See Table 31 on page 39 for a list of EX9200 routing policy and firewall filter features
Firewall filters and rate limiting	10.1R1	9.0R2	11.3R1	10.2R1	12.2R1	11.3R2	9.4R1	
For a list of supported firewall filter match conditions and actions, see Platform Support for Firewall Filter Match Conditions, Actions, and Action Modifiers on EX Series Switches.								
Firewall filters on LAGs	10.1R1	9.0R2	11.3R1	10.2R1	12.2R1	Not supported	10.0R1	
Firewall filters on the loopback interface	10.1R1	9.2R1	11.3R1	11.1R1	12.2R1	12.1R1	9.6R1	
For a list of supported firewall filter match conditions and actions on a loopback interface, see Support for Match Conditions and Actions for Loopback Firewall Filters on Switches.								
Firewall filters on the management interface	11.3R1	10.4R1	Not supported	10.4R1	12.2R1	12.1R1	10.4R1	

Table 30: Routing Policy and Packet Filtering Features on Switches by Junos OS Release (*continued*)

Feature	EX2200	EX3200, EX4200	EX3300	EX4500	EX4550	EX6200	EX8200	EX9200
Firewall filters on the virtual management interface	–	10.4R1 (EX4200 Virtual Chassis only)	Not supported	–	–	–	–	–
Firewall filters with IPv6	11.3R1	10.1R1	12.3R1	12.1R1	12.2R1	12.1R1	10.3R1	
Policing	10.1R1	9.0R2	11.3R1	10.2R1	12.2R1	11.3R2	9.4R1	

Table 31: Routing Policy and Firewall Filters on EX9200 Switches by Junos OS Release

Feature	Junos OS Release
Access and access-internal routes	12.3R2
Extension of numeric-range match conditions in firewall filters	12.3R2
Aggregate policer support for different family address types configured on the same interface	12.3R2
Authentication for BFD (MD5/SHA1)	12.3R2
BGP multipath link-bandwidth attribute	12.3R2
DHCP state persistence for DHCP relay agent	12.3R2
Dynamic configuration support for routing policies	12.3R2
Extended DHCP relay agent	12.3R2
Filter-based forwarding to a specific outgoing interface or destination IP address	12.3R2
Firewall filters within logical systems	12.3R2
IEEE 802.1p priority match conditions for Layer 2 VPN firewall filters	12.3R2
Filter-based forwarding to a specific outgoing interface or destination IP address	12.3R2
Layer 2 Gigabit Ethernet logical interface extended policing support	12.3R2
Layer 2 support for firewall filter match conditions	12.3R2
Load balancing of VPLS traffic	12.3R2
Option 60 support for extended DHCP relay agents	12.3R2

Table 31: Routing Policy and Firewall Filters on EX9200 Switches by Junos OS Release (*continued*)

Feature	Junos OS Release
Policers on physical interfaces	12.3R2
Firewall filters feature support	12.3R2
Support for policers that limit traffic on logical interfaces in ingress or egress directions	12.3R2
Support for policers that rate-limit based on a percentage of physical port speed on an interface	12.3R2
Support for the discard action for tricolor marking policers applied to firewall filters	12.3R2
Support for the prefix-list match condition for firewall filters for the VPLS protocol family	12.3R2
Support for enhanced policer statistics	12.3R2
Support for MAC address validation	12.3R2

Table 32: Spanning-Tree Protocols Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4500	EX4550	EX6200	EX8200	EX9200
BPDU protection for spanning-tree protocols	10.1R1	9.1R1	11.3R1	10.2R1	12.2R1	11.3R2	9.4R1	12.3R2
BPDU filter	12.2R1	12.2R1	12.2R1	12.2R1	12.2R1	12.2R1	12.2R1	12.3R2
Distributed periodic packet management (PPM) for Spanning Tree Protocols	Not supported	12.3R1	Not supported	Not supported	Not supported	Not supported	12.3R1	12.3R2
Loop protection for spanning-tree protocols	10.1R1	9.1R1	11.3R1	10.2R1	12.2R1	11.3R2	9.4R1	12.3R2
Root protection for spanning-tree protocols	10.1R1	9.1R1	11.3R1	10.2R1	12.2R1	11.3R2	9.4R1	12.3R2
Spanning tree: <ul style="list-style-type: none"> RSTP and VSTP concurrent configuration 	Not supported	10.2R1	12.3R1	10.2R1	12.2R1	11.3R2	10.2R1	

Table 32: Spanning-Tree Protocols Features on Switches by Junos OS Release (*continued*)

Feature	EX2200	EX3200, EX4200	EX3300	EX4500	EX4550	EX6200	EX8200	EX9200
Spanning tree:	10.1R1	9.0R2	11.3R1	10.2R1	12.2R1	11.3R2	9.4R1	12.3R2
<ul style="list-style-type: none"> Spanning Tree Protocol (STP) Rapid Spanning Tree Protocol (RSTP) Multiple Spanning Tree Protocol (MSTP) 								
Spanning tree:	10.1R1	9.4R1	11.3R1	10.2R1	12.2R1	11.3R2	9.6R1	12.3R2
<ul style="list-style-type: none"> VLAN Spanning Tree Protocol (VSTP) 								

Table 33: System Management Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4500	EX4550	EX6200	EX8200	EX9200
Autoinstallation of configuration files	10.1R1	9.4R1	11.3R1	10.2R1	12.2R1	11.3R2	Not supported	
Automatic software download	10.1R1	9.6R1	11.3R1	10.2R1	12.2R1	11.3R2	9.6R1	
Automatic repair of corrupted partition when booting from alternate partition	12.3R1	12.3R1	12.3R1	12.3R1	12.3R1	12.3R1	12.3R1	
Configuration rollback	10.1R1	9.0R2	11.3R1	10.2R1	12.2R1	11.3R2	9.4R1	
Zero Touch Provisioning (EZ Touchless Provisioning using DHCP)	12.2R1	12.2R1	12.2R1	12.2R1	12.2R1	Not supported	12.2R1	Not supported
J-Web interface, for switch configuration and management	10.1R1 (12.1R1 for EX2200-C switches)	9.0R2	12.1R1	10.2R1	12.2R1	12.1R1	9.4R1	Not supported
Junos Space Service Now support	12.3R1	12.3R1	12.3R1	12.3R1	12.3R1	12.3R1	12.3R1	Not supported

Table 33: System Management Features on Switches by Junos OS Release (*continued*)

Feature	EX2200	EX3200, EX4200	EX3300	EX4500	EX4550	EX6200	EX8200	EX9200
LCD panel management support	–	9.0R1	11.3R1	10.2R1	12.2R1	11.3R1	9.4R1	–
Online insertion and removal (OIR) of uplink modules	–	10.0R1	–	11.1R1	12.2R1	–	–	–

Table 34: User Interface and Configuration Features on EX9200 Switches by Junos OS Release

Feature	Junos OS Release
Device-initiated SSH connection (outbound SSH)	12.3R2
Dynamic IPv6 filters	12.3R2
Dynamic configuration of the switch advertisement protocol	12.3R2
Dynamic profiles support by extended DHCP local server and extended DHCP relay agent	12.3R2
Enhanced IPv6 statistics	12.3R2
Extended DHCP local server	12.3R2
IGMP dynamic profiles	12.3R2
Extended DHCP local server	12.3R2
Protection for device configuration	12.3R2
RADIUS MSCHAPv2 protocol support for administrator authentication, password aging, and update	12.3R2
Limit configuration command output	12.3R2
Remote tracing	12.3R2
Support for CLI edit mode wildcard range	12.3R2
Support for configuring ARP aging time for a logical interface	12.3R2
Support for configuring a proxy server for downloading licenses	12.3R2
Support for configuring time-based user access	12.3R2
Support for logical router system administrators	12.3R2

Table 35: VPN Features on EX9200 Switches by Junos OS Release

Feature	Junos OS Release
Aggregated Ethernet interfaces for VPLS routing instances	12.3R2
BGP autodiscovery for LDP VPLS (FEC 129)	12.3R2
Clearing MAC addresses for better convergence	12.3R2
Configurable label block sizes for VPLS	12.3R2
Disable TTL propagation behavior for the routes in a VRF routing instance	12.3R2
EXP-based traffic classification for VPLS	12.3R2
Enhanced show interface command for Layer 3 VPN functionality	12.3R2
Expanded interface support for the vrf-table-label statement	12.3R2
Extranet next-generation MVPN GRE tunnels for Layer 3 VPNs	12.3R2
GRE tunnels for Layer 3 VPNs Ignore MTU mismatch on Layer 2 circuits Integrated routing and bridging support for inter-AS VPLS between BGP-signaled VPLS and LDP-signaled VPLS LDP-based VPLS Label allocation and substitution policy	12.3R2
Ignore MTU mismatch on Layer 2 circuits	12.3R2
Integrated routing and bridging support for inter-AS VPLS between BGP-signaled VPLS and LDP-signaled VPLS	12.3R2
LDP-based VPLS	12.3R2
Label allocation and substitution policy	12.3R2
Layer 2 VPN multihoming	12.3R2
Layer 3 VPN BGP routes and labels	12.3R2
Layer 3 VPN localization	12.3R2
Load balancing and IP header filtering for Layer 3 VPNs	12.3R2
Local switching support for the ignore-encapsulation-mismatch statement	12.3R2
Multipath load balancing for EBGp and IBGP VPNs	12.3R2
Multiple logical trunk interfaces per physical interface	12.3R2
Multiprotocol BGP-based multicast VPN	12.3R2
NTP support for IPv4 VRF and IPv6 VRF	12.3R2

Table 35: VPN Features on EX9200 Switches by Junos OS Release (*continued*)

Feature	Junos OS Release
Nonstop active routing support for Layer 3 VPNs	12.3R2
PIM source-specific multicast (PIM-SSM) provider tunnel support added to Multiprotocol BGP-based multicast VPNs	12.3R2
Point-to-multipoint LSP support for VPLS	12.3R2
Point-to-multipoint LSP support for multicast VPNs	12.3R2
Proxy BGP route target filtering	12.3R2
Static VPLS	12.3R2
Static route target filtering	12.3R2
Support for autorp, BSR, PIM dense mode and mtrace for next-generation multicast VPNs	12.3R2
VLAN range for Layer 2 VPN	12.3R2
VPLS automatic site ID	12.3R2
VPLS automatic site ID for nonstop active routing	12.3R2
VPLS ping	12.3R2
VPLS trunk interfaces	12.3R2
eBGP and iBGP load-balancing support for MVPN and PIM	12.3R2

**Related
Documentation**

- EX Series Virtual Chassis Software Features Overview
- EX2200 Switches Hardware Overview
- EX3200 Switches Hardware Overview
- EX3300 Switches Hardware Overview
- EX4200 Switches Hardware Overview
- EX4500 Switches Hardware Overview
- EX4550 Switches Hardware Overview
- EX6210 Switch Hardware Overview
- EX8208 Switch Hardware Overview
- EX8216 Switch Hardware Overview
- EX9208 Switch Hardware Overview
- Line Card Model and Version Compatibility in an EX6200 Switch

- Line Card Model and Version Compatibility in an EX8200 Switch
- Line Card Model and Version Compatibility in an EX9200 Switch
- XRE200 External Routing Engine Hardware Overview
- Layer 3 Protocols Supported on EX Series Switches
- Layer 3 Protocols Not Supported on EX Series Switches

PART 2

Configuration

- [Configuration Statements on page 49](#)

CHAPTER 2

Configuration Statements

- [\[edit protocols bgp\] Hierarchy Level on page 49](#)

[\[edit protocols bgp\] Hierarchy Level](#)

Several statements in the **[edit protocols mpls]** hierarchy are valid at numerous locations within it. To make the complete hierarchy easier to read, the repeated statements are listed in “[Common BGP Family Options](#)” on page 49 and that section is referenced at the appropriate locations in “[Complete \[edit protocols bgp\] Hierarchy](#)” on page 50.

- [Common BGP Family Options on page 49](#)
- [Complete \[edit protocols bgp\] Hierarchy on page 50](#)

Common BGP Family Options

This section lists statements that are valid at the following hierarchy levels, and is referenced at those levels in “[Complete \[edit protocols bgp\] Hierarchy](#)” on page 50 instead of the statements being repeated.

- **[edit protocols bgp family inet (any | flow | labeled-unicast | multicast | unicast)]**
- **[edit protocols bgp family inet6 (any | labeled-unicast | multicast | unicast)]**
- **[edit protocols bgp family (inet-mdt | inet-mvpn | inet6-mvpn | l2vpn) signaling]**
- **[edit protocols bgp family inet-vpn (any | flow | multicast | unicast)]**
- **[edit protocols bgp family inet6-vpn (any | multicast | unicast)]**
- **[edit protocols bgp family iso-vpn unicast]**

The common BGP family options are as follows:

```
accepted-prefix-limit {  
    maximum number;  
    teardown <percentage> <idle-timeout (forever | minutes)>;  
}  
damping;  
loops number;  
prefix-limit {  
    maximum number;  
    teardown <percentage> <idle-timeout (forever | minutes)>;
```

```
}
rib-group group-name;
topology name {
  community {
    target identifier;
  }
}
```

Complete [edit protocols bgp] Hierarchy

The statement hierarchy listed in this section can also be included at the [edit logical-systems *logical-system-name*] hierarchy level.

```
protocols {
  bgp {
    disable;
    accept-remote-nexthop;
    advertise-external <conditional>;
    advertise-from-main-vpn-tables;
    advertise-inactive;
    (advertise-peer-as | no-advertise-peer-as);
    authentication-algorithm (aes-128-cmac-96 | hmac-sha-1-96 | md5);
    authentication-key key;
    authentication-key-chain key-chain;
    bfd-liveness-detection {
      authentication {
        algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
          meticulous-keyed-sha-1 | simple-password);
        key-chain key-chain-name;
        loose-check;
      }
      detection-time {
        threshold milliseconds;
      }
      holddown-interval milliseconds;
      minimum-interval milliseconds;
      minimum-receive-interval milliseconds;
      multiplier number;
      no-adaptation;
      session-mode (automatic | multihop | single-hop);
      transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
      }
      version (1 | automatic);
    }
    cluster cluster-identifier;
    damping;
    description text-description;
    export [ policy-names ];
    family family-name {
      ... the family subhierarchies appear after the main [edit protocols bgp] hierarchy ...
    }
    graceful-restart {
      disable;
      restart-time seconds;
    }
  }
}
```

```

    stale-routes-time seconds;
}
group group-name {
    ... the group subhierarchy appears after the main [edit protocols bgp] hierarchy ...
}
hold-time seconds;
idle-after-switch-over (seconds | forever);
import [ policy-names ];
include-mp-next-hop;
ipsec-sa ipsec-sa;
keep (all | none);
local-address address;
local-as autonomous-system <loops number> < alias> <private>;
local-interface interface-name;
local-preference local-preference;
log-updown;
metric-out (metric | igp (delay-med-update | offset) | minimum-igp offset);
mtu-discovery;
multihop {
    no-nexthop-change;
    ttl ttl-value;
}
no-aggregator-id;
no-client-reflect;
out-delay seconds;
outbound-route-filter {
    bgp-orf-cisco-mode;
    prefix-based {
        accept {
            inet;
            inet6;
        }
    }
}
passive;
path-selection {
    always-compare-med;
    as-path-ignore;
    cisco-non-deterministic;
    external-router-id;
    med-plus-igp {
        igp-multiplier number;
        med-multiplier number;
    }
}
peer-as autonomous-system;
preference preference;
remove-private;
tcp-mss segment-size;
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
vpn-apply-export;
}

```

```
bgp {
  family inet {
    (any | multicast) {
      ... statements in Common BGP Family Options on page 49 ...
    }
    flow {
      ... statements in Common BGP Family Options on page 49 PLUS ...
      no-validate [ validation-procedure-names ];
    }
    labeled-unicast {
      ... statements in Common BGP Family Options on page 49 PLUS ...
      add-path {
        receive;
        send {
          path-count number;
          prefix-policy [ policy-names ];
        }
      }
      aggregate-label {
        community community-name;
      }
      aigp [disable];
      explicit-null connected-only;
      per-group-label;
      per-prefix-label;
      resolve-vpn;
      rib (inet.3 | inet6.3);
      traffic-statistics {
        file filename <files number> <size maximum-file-size> <world-readable |
          no-world-readable>;
        interval seconds;
      }
    }
  }
  unicast {
    ... statements in Common BGP Family Options on page 49 PLUS ...
    add-path {
      receive;
      send {
        path-count number;
        prefix-policy [ policy-names ];
      }
    }
    topology name {
      community target identifier;
    }
  }
}

bgp {
  family inet6 {
    (any | multicast) {
      ... statements in Common BGP Family Options on page 49 ...
    }
    labeled-unicast {
```



```

... statements in Common BGP Family Options on page 49 PLUS ...
add-path {
    receive;
    send {
        path-count number;
        prefix-policy [ policy-names ];
    }
}
aggregate-label {
    community community-name;
}
aigp [disable];
explicit-null;
per-group-label;
traffic-statistics {
    file filename <files number> <size maximum-file-size> <world-readable |
    no-world-readable>;
    interval seconds;
}
}
unicast {
    ... statements in Common BGP Family Options on page 49 PLUS ...
    topology name {
        community target identifier;
    }
}
}
}

bgp {
    family (inet-mdt | inet-mvpn | inet6-mvpn | l2vpn) {
        signaling {
            ... statements in Common BGP Family Options on page 49 ...
        }
    }
}

bgp {
    family inet-vpn {
        (any | multicast | unicast) {
            ... statements in Common BGP Family Options on page 49 PLUS ...
            aggregate-label <community community-name>;
        }
        flow {
            ... statements in Common BGP Family Options on page 49 ...
        }
    }
}

bgp {
    family inet6-vpn {
        (any | multicast | unicast) {
            ... statements in Common BGP Family Options on page 49 PLUS ...
            aggregate-label <community community-name>;
        }
    }
}

```

```

}

bgp {
  family iso-vpn {
    unicast {
      ... statements in Common BGP Family Options on page 49 PLUS ...
      aggregate-label <community community-name>;
    }
  }
}

bgp {
  family route-target {
    accepted-prefix-limit {
      maximum number;
      teardown <percentage> <idle-timeout (forever | minutes)>;
    }
    advertise-default;
    external-paths number;
    prefix-limit {
      maximum number;
      teardown <percentage> <idle-timeout (forever | minutes)>;
    }
    proxy-generate <route-target-policy route-target-policy-name>;
  }
}

bgp {
  group group-name {
    ... same statements as at the [edit protocols bgp] hierarchy level PLUS ...
    allow [ all ip-prefix</prefix-length> ];
    as-override;
    multipath <multiple-as>;
    neighbor address {
      ... the neighbor subhierarchy appears after the main [edit protocols bgp group
        group-name] hierarchy ...
    }
    type (external | internal);
    ... BUT NOT ...
    disable; # NOT valid at this level
    group group-name { ... } # NOT valid at this level
    path-selection { ... } # NOT valid at this level
  }

  group group-name {
    neighbor address {
      ... same statements as at the [edit protocols bgp] hierarchy level PLUS ...
      as-override;
      multipath <multiple-as>;
      ... BUT NOT ...
      disable; # NOT valid at this level
      group group-name { ... } # NOT valid at this level
      neighbor address { ... } # NOT valid at this level
      path-selection { ... } # NOT valid at this level
    }
  }
}

```

```
}
}
```

- Related Documentation**
- Notational Conventions Used in Junos OS Configuration Hierarchies
 - [edit protocols] Hierarchy Level

accept-remote-nexthop

Syntax accept-remote-nexthop;

Hierarchy Level [edit logical-systems *logical-system-name* protocols bgp],
 [edit logical-systems *logical-system-name* protocols bgp group *group-name*],
 [edit logical-systems *logical-system-name* protocols bgp group *group-name* neighbor *address*],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name*],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* neighbor *address*],
 [edit protocols bgp],
 [edit protocols bgp group *group-name*],
 [edit protocols bgp group *group-name* neighbor *address*],
 [edit routing-instances *routing-instance-name* protocols bgp],
 [edit routing-instances *routing-instance-name* protocols bgp group *group-name*],
 [edit routing-instances *routing-instance-name* protocols bgp group *group-name* neighbor *address*]


Release Information Statement introduced in Junos OS Release 8.5.
 Statement introduced in Junos OS Release 9.0 for EX Series switches.
 Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Specify that a single-hop EBGP peer accepts a remote next hop with which it does not share a common subnet. Configure a separate import policy on the EBGP peer to specify the remote next hop. You cannot configure **multihop** and **accept-remote-nexthop** statements for the same EPBG peer.

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

- Related Documentation**
- Example: Configuring Single-Hop EBGP Peers to Accept Remote Next Hops
 - Understanding Route Advertisement
 - [multipath on page 113](#)

advertise-external

Syntax	<code>advertise-external {conditional};</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>neighbor-address</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.3.</p> <p>Statement introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Specify BGP to advertise the best external route into an IBGP mesh group, a route reflector cluster, or an AS confederation even if the best route is an internal route. In order to configure the advertise-external statement on a route reflector, you must disable intracluster reflection with the no-client-reflect statement. The advertise-external statement is supported at both the group and neighbor level. If you configure the statement at the neighbor level, you must configure it for all neighbors in a group. Otherwise, the group is automatically split into different groups.</p>
	<div>  <p>NOTE: When configuring the advertise-external statement for an AS confederation, it is recommended that EBGp peers belonging to different autonomous systems are configured in a separate EBGp peer group. This ensures consistency while BGP sends the best external route to peers in the configured peer group.</p> </div>
Options	<p>conditional—(Optional) Advertise the best external path only if the route selection process reaches the point at which the multiple exit discriminator (MED) metric is evaluated. The conditional option restricts advertisement to when the best external path and the active path are equal until the MED step of the route selection process. This implies that external routes with a longer AS path length than the active path, for instance, are not advertised. The criteria used for selecting the best external path is the same whether or not the conditional option is configured.</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"> • Example: Configuring BGP Route Advertisement • Understanding Route Advertisement

- [advertise-inactive on page 57](#)

advertise-inactive

Syntax	advertise-inactive;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Configure the routing table to export to BGP the best route learned by BGP even if Junos OS did not select this route to be an active route.</p> <p>One way to achieve multivendor compatibility is to include the advertise-inactive statement in the external BGP (EBGP) configuration. By default, BGP stores the route information it receives from update messages in the Junos OS routing table, and the routing table exports only active routes into BGP, which BGP then advertises to its peers. The advertise-inactive statement causes Junos OS to advertise the best BGP route that is inactive because of IGP preference. When you use the advertise-inactive statement, the Junos OS device uses, for example, the OSPF route for forwarding, and the other vendor's device uses the EBGP route for forwarding. However, from the perspective of an EBGP peer in a neighboring AS, both vendors' devices appear to behave the same way.</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"> • Example: Configuring the Preference Value for BGP Routes • Example: Configuring BGP Route Preference (Administrative Distance) • Understanding Route Advertisement • advertise-external on page 56



advertise-peer-as

Syntax	advertise-peer-as;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Disable the default behavior of suppressing AS routes.
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">• Example: Configuring BGP Route Advertisement• Understanding Route Advertisement• no-advertise-peer-as


aggregate-label

Syntax	aggregate-label { community <i>community-name</i> ; }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols bgp family inet labeled-unicast], [edit logical-systems <i>logical-system-name</i> protocols bgp family inet6 labeled-unicast], [edit logical-systems <i>logical-system-name</i> protocols bgp family inet-vpn unicast], [edit logical-systems <i>logical-system-name</i> protocols bgp family inet-vpn6 unicast], [edit protocols bgp family inet labeled-unicast], [edit protocols bgp family inet6 labeled-unicast], [edit protocols bgp family inet-vpn unicast], [edit protocols bgp family inet6-vpn unicast]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Specify matching criteria (in the form of a community) such that all routes which match are assigned the same VPN label, selected from one of the several routes in the set defined by this criteria. This reduces the number of VPN labels that the router must consider, and aggregates the received labels.
Options	community <i>community-name</i> —Specify the name of the community to which to apply the aggregate label.
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 Aggregate Labels for VPNs

allow

Syntax	<code>allow (all [<i>network/mask-length</i>]);</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>]
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 11.3 for the QFX Series.
Description	Implicitly configure BGP peers, allowing peer connections from any of the specified networks or hosts. To configure multiple BGP peers, configure one or more networks and hosts within a single allow statement or include multiple allow statements.
<div>  <p>NOTE: You cannot define a BGP group with dynamic peers with BGP authentication enabled.</p> </div>	
Options	<p>all—Allow all addresses, which is equivalent to 0.0.0.0/0 (or ::/0).</p> <p><i>network/mask-length</i>—IPv6 or IPv4 network number of a single address or a range of allowable addresses for BGP peers, followed by the number of significant bits in the subnet mask.</p> <div>  <p>NOTE: You cannot define a BGP group with dynamic peers with authentication enabled.</p> </div>
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"> neighbor on page 114

as-override

Syntax	as-override;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Compare the AS path of an incoming advertised route with the AS number of the BGP peer under the group and replace all occurrences of the peer AS number in the AS path with its own AS number before advertising the route to the peer.
<div>  <p>NOTE: The as-override statement is specific to a particular BGP group. This statement does not affect peers from the same remote AS configured in different groups.</p> </div>	
<p>Enabling the AS override feature allows routes originating from an AS to be accepted by a router residing in the same AS. Without AS override enabled, the routing device refuses the route advertisement once the AS path shows that the route originated from its own AS. This is done by default to prevent route loops. The as-override statement overrides this default behavior.</p> <p>Note that enabling the AS override feature may result in routing loops. Use this feature only for specific applications that require this type of behavior, and in situations with strict network control. One application is the IGP protocol between the provider edge routing device and the customer edge routing device in a virtual private network.</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"> • Example: Configuring a Layer 3 VPN with Route Reflection and AS Override • Understanding Route Advertisement • Junos OS VPNs Configuration Guide

authentication-algorithm

Syntax	<code>authentication-algorithm <i>algorithm</i>;</code>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> protocols ldp session <i>session-address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp session <i>session-address</i>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit protocols ldp session <i>session-address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols ldp session <i>session-address</i>]</pre>
Release Information	<p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced for BGP in Junos OS Release 8.0.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Configure an authentication algorithm type.
Options	<p><i>algorithm</i>—Specify one of the following types of authentication algorithms:</p> <ul style="list-style-type: none"> aes-128-cmac-96—Cipher-based message authentication code (AES128, 96 bits). hmac-sha-1-96—Hash-based message authentication code (SHA1, 96 bits). md5—Message digest 5.
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"> Understanding Route Authentication Example: Configuring Route Authentication for BGP

authentication-key (Protocols BGP)

Syntax	authentication-key <i>key</i> ;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name neighbor address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name neighbor address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp <i>group group-name</i>],</p> <p>[edit protocols bgp group <i>group-name neighbor address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name neighbor address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Configure an MD5 authentication key (password). Neighboring routing devices use the same password to verify the authenticity of BGP packets sent from this system.
Options	<i>key</i> —Authentication password. It can be up to 126 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").
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"> Example: Configuring Route Authentication for BGP

authentication-key-chain (Protocols BGP)

Syntax	<code>authentication-key-chain <i>key-chain</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.0.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Apply and enable an authentication keychain to the routing device. Note that the referenced key chain must be defined. When configuring the authentication key update mechanism for BGP, you cannot commit the 0.0.0.0/allow statement with authentication keys or key chains. The CLI issues a warning and fails to commit such configurations.
Options	key-chain —Authentication keychain name. It can be up to 126 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").
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"> • Example: Configuring Route Authentication for BGP • Example: Configuring BFD Authentication for Static Routes • Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols

autonomous-system

Syntax	<pre>autonomous-system <i>autonomous-system</i> <asdot-notation> <loops <i>number</i>> { independent-domain <no-attrset>; }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit routing-options]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>asdot-notation option introduced in Junos OS Release 9.3.</p> <p>asdot-notation option introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>no-attrset option introduced in Junos OS Release 10.4.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	<p>Specify the routing device's AS number.</p> <p>An autonomous system (AS) is a set of routing devices that are under a single technical administration and that generally use a single interior gateway protocol (IGP) and metrics to propagate routing information within the set of routing devices. An AS appears to other ASs to have a single, coherent interior routing plan and presents a consistent picture of what destinations are reachable through it. ASs are identified by a number that is assigned by the Network Information Center (NIC) in the United States (http://www.isi.edu).</p> <p>If you are using BGP on the routing device, you must configure an AS number.</p> <p>The AS path attribute is modified when a route is advertised to an EBGP peer. Each time a route is advertised to an EBGP peer, the local routing device prepends its AS number to the existing path attribute, and a value of 1 is added to the AS number.</p> <p>In Junos OS Release 9.1 and later, the numeric range is extended to provide BGP support for 4-byte AS numbers as defined in RFC 4893, <i>BGP Support for Four-octet AS Number Space</i>. RFC 4893 introduces two new optional transitive BGP attributes, AS4_PATH and AS4_AGGREGATOR. These new attributes are used to propagate 4-byte AS path information across BGP speakers that do not support 4-byte AS numbers. RFC 4893 also introduces a reserved, well-known, 2-byte AS number, AS 23456. This reserved AS number is called AS_TRANS in RFC 4893. All releases of Junos OS support 2-byte AS numbers.</p> <p>In Junos OS Release 9.3 and later, you can also configure a 4-byte AS number using the AS-dot notation format of two integer values joined by a period: <i><16-bit high-order value in decimal>.<16-bit low-order value in decimal></i>. For example, the 4-byte AS number of 65,546 in plain-number format is represented as 1.10 in the AS-dot notation format.</p>
Options	<p><i>autonomous-system</i>—AS number. Use a number assigned to you by the NIC.</p>

Range: 1 through 4,294,967,295 ($2^{32} - 1$) in plain-number format for 4-byte AS numbers

In this example, the 4-byte AS number 65,546 is represented in plain-number format:

```
[edit]
routing-options {
  autonomous-system 65546;
}
```

Range: 0.0 through 65535.65535 in AS-dot notation format for 4-byte numbers

In this example, 1.10 is the AS-dot notation format for 65,546:

```
[edit]
routing-options {
  autonomous-system 1.10;
}
```

Range: 1 through 65,535 in plain-number format for 2-byte AS numbers (this is a subset of the 4-byte range)

In this example, the 2-byte AS number 60,000 is represented in plain-number format:

```
[edit]
routing-options {
  autonomous-system 60000;
}
```

asdot-notation—(Optional) Display the configured 4-byte autonomous system number in the AS-dot notation format.

Default: Even if a 4-byte AS number is configured in the AS-dot notation format, the default is to display the AS number in the plain-number format.

loops number—(Optional) Specify the number of times detection of the AS number in the AS_PATH attribute causes the route to be discarded or hidden. For example, if you configure **loops 1**, the route is hidden if the AS number is detected in the path one or more times. This is the default behavior. If you configure **loops 2**, the route is hidden if the AS number is detected in the path two or more times.

Range: 1 through 10

Default: 1



.....

NOTE: When you specify the same AS number in more than one routing instance on the local routing device, you must configure the same number of loops for the AS number in each instance. For example, if you configure a value of 3 for the loops statement in a VRF routing instance that uses the same AS number as that of the master instance, you must also configure a value of 3 loops for the AS number in the master instance.

Use the independent-domain option if the loops statement must be enabled only on a subset of routing instances.

.....

The remaining statement is explained separately.

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">• Examples: Configuring External BGP Peering• Examples: Configuring Internal BGP Peering• 4-Byte Autonomous System Numbers Overview in the <i>Using 4-Byte Autonomous System Numbers in BGP Networks Technology Overview</i>• Juniper Networks Implementation of 4-Byte Autonomous System Numbers in the <i>Using 4-Byte Autonomous System Numbers in BGP Networks Technology Overview</i>• Configuring 4-Byte Autonomous System Numbers in the <i>Using 4-Byte Autonomous System Numbers in BGP Networks Technology Overview</i>

bfd-liveness-detection (Protocols BGP)

Syntax `bfd-liveness-detection {`
 `authentication {`
 `algorithm algorithm-name;`
 `key-chain key-chain-name;`
 `loose-check;`
 `}`
 `detection-time {`
 `threshold milliseconds;`
 `}`
 `hold-down-interval milliseconds;`
 `minimum-interval milliseconds;`
 `minimum-receive-interval milliseconds;`
 `multiplier number;`
 `no-adaptation;`
 `session-mode (automatic | multihop | single-hop);`
 `transmit-interval {`
 `minimum-interval milliseconds;`
 `threshold milliseconds;`
 `}`
 `version (1 | automatic);`
`}`

Hierarchy Level `[edit logical-systems logical-system-name protocols bgp],`
`[edit logical-systems logical-system-name protocols bgp group group-name],`
`[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],`
`[edit logical-systems logical-system-name routing-instances routing-instance-name protocols`
`bgp],`
`[edit logical-systems logical-system-name routing-instances routing-instance-name protocols`
`bgp group group-name],`
`[edit logical-systems logical-system-name routing-instances routing-instance-name protocols`
`bgp group group-name neighbor address],`
`[edit protocols bgp],`
`[edit protocols bgp group group-name],`
`[edit protocols bgp group group-name neighbor address],`
`[edit routing-instances routing-instance-name protocols bgp],`
`[edit routing-instances routing-instance-name protocols bgp group group-name],`
`[edit routing-instances routing-instance-name protocols bgp group group-name neighbor`
`address]`

Release Information Statement introduced in Junos OS Release 8.1.
 Statement introduced in Junos OS Release 9.0 for EX Series switches.
detection-time threshold and **transmit-interval threshold** options introduced in Junos OS Release 8.2
 Support for logical routers introduced in Junos OS Release 8.3.
 Support for IBGP and multihop EBGP sessions introduced in Junos OS Release 8.3.
holddown-interval statement introduced in Junos OS Release 8.5. You can configure this statement only for EBGP peers at the **[edit protocols bgp group *group-name* neighbor *address*]** hierarchy level.
no-adaptation statement introduced in Junos OS Release 9.0.
 Support for BFD authentication introduced in Junos OS Release 9.6.

Support for BFD on IPv6 interfaces with BGP introduced in Junos OS Release 11.2.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Configure bidirectional failure detection (BFD) timers and authentication for BGP.

For IBGP and multihop EBGP support, configure the **bfd-liveness-detection** statement at the global **[edit bgp protocols]** hierarchy level. You can also configure IBGP and multihop support for a routing instance or a logical system.

Options **authentication algorithm** *algorithm-name* (Optional)—Configure the algorithm used to authenticate the specified BFD session: **simple-password**, **keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, **meticulous-keyed-sha-1**.

authentication key-chain *key-chain-name* (Optional)—Associate a security key with the specified BFD session using the name of the security keychain. The keychain name must match one of the keychains configured in the **authentication-key-chains key-chain** statement at the **[edit security]** hierarchy level.

authentication loose-check—(Optional) Configure loose authentication checking on the BFD session. Use only for transitional periods when authentication may not be configured at both ends of the BFD session.

detection-time threshold *milliseconds* (Optional)—Configure a threshold. When the BFD session detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

holddown-interval *milliseconds* (Optional)—Configure an interval specifying how long a BFD session must remain up before a state change notification is sent. When you configure the hold-down interval for the BFD protocol for EBGp, the BFD session is unaware of the BGP session during this time. In this case, if the BGP session goes down during the configured hold-down interval, BFD already assumes it is down and does not send a state change notification. The **holddown-interval** statement is supported only for EBGp peers at the **[edit protocols bgp group group-name neighbor address]** hierarchy level. If the BFD session goes down and then comes back up during the configured hold-down interval, the timer is restarted. You must configure the hold-down interval on both EBGp peers. If you configure the hold-down interval for a multihop EBGp session, you must also configure a local IP address by including the **local-address** statement at the **[edit protocols bgp group group-name]** hierarchy level.

Range: 0 through 255,000

Default: 0

minimum-interval *milliseconds* (Required)—Configure the minimum intervals at which the local routing device transmits hello packets and then expects to receive a reply from a neighbor with which it has established a BFD session. This value represents the minimum interval at 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 in the range from 1 through 255,000 milliseconds. Optionally, instead of using this statement, you can specify the minimum transmit and receive intervals separately (using the **minimum-receive-interval** and **transmit-interval** statements).

Range: 1 through 255,000

minimum-receive-interval *milliseconds* (Optional)—Configure only 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.

Range: 1 through 255,000

multiplier *number* (Optional)—Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.

Range: 1 through 255

Default: 3

no-adaptation (Optional)—Configure BFD sessions not to adapt to changing network conditions. We recommend that you not disable BFD adaptation unless it is preferable to not to have BFD adaptation enabled in your network.

transmit-interval threshold *milliseconds* (Optional)—Configure a threshold. When the BFD session transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent. The interval threshold must be greater than the minimum transmit interval.

Range: 0 through 4,294,967,295 ($2^{32} - 1$)

transmit-interval minimum-interval *milliseconds* (Optional)—Configure only the minimum interval at which the local routing device transmits hello packets to a neighbor with which it has established a BFD session.

Range: 1 through 255,000

version (Optional)—Configure the BFD version to detect.

Range: 1 or **automatic** (autodetect the BFD version)

Default: **automatic**

The remaining statements are explained separately.


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

Related Documentation	<ul style="list-style-type: none">• Example: Configuring BFD for Static Routes• Example: Configuring BFD Authentication for Static Routes• Example: Configuring BFD on Internal BGP Peer Sessions• Example: Configuring BFD Authentication for BGP• Understanding BFD for BGP
------------------------------	---

bgp

Syntax	bgp { ... }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit protocols], [edit routing-instances <i>routing-instance-name</i> protocols]
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 11.3 for the QFX Series.
Description	Enable BGP on the routing device or for a routing instance.
Default	BGP 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">• BGP Configuration Guide

bgp-orf-cisco-mode

Syntax	<code>bgp-orf-cisco-mode;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp outbound-route-filter],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> outbound-route-filter],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> outbound-route-filter],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp outbound-route-filter],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> outbound-route-filter],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> outbound-route-filter],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options outbound-route-filter],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options outbound-route-filter],</p> <p>[edit protocols bgp outbound-route-filter],</p> <p>[edit protocols bgp group <i>group-name</i> outbound-route-filter],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> outbound-route-filter],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp outbound-route-filter],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> outbound-route-filter],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> outbound-route-filter],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options outbound-route-filter],</p> <p>[edit routing-options outbound-route-filter]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.2.</p> <p>Statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Support for the BGP group and neighbor hierarchy levels introduced in Junos OS Release 9.2.</p> <p>Support for the BGP group and neighbor hierarchy levels introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	Enable interoperability with routing devices that use the vendor-specific outbound route filter compatibility code of 130 and code type of 128.
	<div>  <p>NOTE: To enable interoperability for all BGP peers configured on the routing device, include the statement at the [edit routing-options outbound-route-filter] hierarchy level.</p> </div>
Default	Disabled
Required Privilege Level	routing—To view this statement in the configuration.
Level	routing-control—To add this statement to the configuration.

**Related
Documentation**

- [Example: Configuring BGP Prefix-Based Outbound Route Filtering](#)

cluster

Syntax	<code>cluster <i>cluster-identifier</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Specify the cluster identifier to be used by the route reflector cluster in an internal BGP group.



CAUTION:

If you configure both route reflection and VPNs on the same routing device, the following modifications to the route reflection configuration cause current BGP sessions to be reset:

- Adding a cluster ID—If a BGP session shares the same AS number with the group where you add the cluster ID, all BGP sessions are reset regardless of whether the BGP sessions are contained in the same group.
- Creating a new route reflector—If you have an IBGP group with an AS number and create a new route reflector group with the same AS number, all BGP sessions in the IBGP group and the new route reflector group are reset.



NOTE: If you change the address family specified in the [edit protocols bgp family] hierarchy level, all current BGP sessions on the routing device are dropped and then reestablished.

Options	<i>cluster-identifier</i> —IPv6 or IPv4 address to use as the cluster identifier.
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">• Example: Configuring BGP Route Reflectors• Understanding External BGP Peering Sessions• no-client-reflect on page 118

damping (Protocols BGP)

Syntax	damping;
Hierarchy Level	<pre> [edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp family <i>family</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp family <i>family</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp family <i>family</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp family <i>family</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family <i>family</i>], [edit protocols bgp], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> family <i>family</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i> family <i>family</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp family <i>family</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> family <i>family</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>] [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family <i>family</i>] </pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Support for flap damping at the address family level introduced in Junos OS Release 12.2.</p>
Description	<p>Enable route flap damping. BGP route flapping describes the situation in which BGP systems send an excessive number of update messages to advertise network reachability information. Flap damping reduces the number of update messages sent between BGP</p>

peers, thereby reducing the load on these peers, without adversely affecting the route convergence time for stable routes.

You typically apply flap damping to external BGP (EBGP) routes (that is, to routes in different ASs). You can also apply it within a confederation, between confederation member ASs. Because routing consistency within an AS is important, do not apply flap damping to internal BGP (IBGP) routes. (If you do, it is ignored.) The exception to this rule is when flap damping is applied at the address family level. When you apply flap damping at the address family level, it works for both IBGP and EBGP.

Default Flap damping is disabled on the routing device.

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

Related Documentation

- Examples: Configuring BGP Flap Damping
- Example: Configuring BGP Route Flap Damping Based on the MBGP MVPN Address Family

description (Protocols BGP)

Syntax	<code>description text-description;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Provide a description of the global, group, or neighbor configuration. If the text includes one or more spaces, enclose it in quotation marks (" "). The text is displayed in the output of the show command and has no effect on the configuration.
Options	text-description —Text description of the configuration. It is limited to 255 characters.
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"> BGP Configuration Guide

disable (Protocols BGP)

Syntax	disable;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp]
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 11.3 for the QFX Series.
Description	Disable BGP on the system.
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">• BGP Configuration Guide

explicit-null (Protocols BGP)

Syntax	explicit-null;
Hierarchy Level	<pre> [edit logical-systems <i>logical-system-name</i> protocols mpls], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>family</i> inet labeled-unicast], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>family</i> inet6 labeled-unicast], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> <i>family</i> inet labeled-unicast], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> <i>family</i> inet6 labeled-unicast], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> <i>family</i> inet labeled-unicast], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> <i>family</i> inet6 labeled-unicast], [edit logical-systems <i>logical-system-name</i> protocols ldap], [edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols bgp <i>family</i> inet labeled-unicast], [edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols bgp <i>family</i> inet6 labeled-unicast], [edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols bgp group <i>group-name</i> <i>family</i> inet labeled-unicast], [edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols bgp group <i>group-name</i> <i>family</i> inet6 labeled-unicast], [edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> <i>family</i> inet labeled-unicast], [edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> <i>family</i> inet6 labeled-unicast], [edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols ldap], [edit protocols mpls], [edit protocols bgp <i>family</i> inet labeled-unicast], [edit protocols bgp <i>family</i> inet6 labeled-unicast], [edit protocols bgp group <i>group-name</i> <i>family</i> inet labeled-unicast], [edit protocols bgp group <i>group-name</i> <i>family</i> inet6 labeled-unicast], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i> <i>family</i> inet labeled-unicast], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i> <i>family</i> inet6 labeled-unicast], [edit protocols ldap], [edit routing-instances <i>instance-name</i> protocols bgp <i>family</i> inet labeled-unicast], [edit routing-instances <i>instance-name</i> protocols bgp <i>family</i> inet6 labeled-unicast], [edit routing-instances <i>instance-name</i> protocols bgp group <i>group-name</i> <i>family</i> inet labeled-unicast], [edit routing-instances <i>instance-name</i> protocols bgp group <i>group-name</i> <i>family</i> inet6 labeled-unicast], [edit routing-instances <i>instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> <i>family</i> inet labeled-unicast], [edit routing-instances <i>instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> <i>family</i> inet6 labeled-unicast], [edit routing-instances <i>instance-name</i> protocols ldap] </pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
Description	Advertise label 0 to the egress routing device of an LSP.

Default	If you do not include the explicit-null statement in the configuration, label 3 (implicit null) is advertised.
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">Advertising Explicit Null Labels to BGP Peers

export (Protocols BGP)

Syntax	<code>export [<i>policy-names</i>];</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols bgp],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i></code> <code> <i>neighbor address</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> bgp],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> bgp <i>group group-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> bgp <i>group group-name</i> <i>neighbor address</i>],</code> <code>[edit protocols bgp],</code> <code>[edit protocols bgp <i>group group-name</i>],</code> <code>[edit protocols bgp <i>group group-name</i> <i>neighbor address</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i></code> <code> <i>neighbor address</i>]</code>
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 11.3 for the QFX Series.
Description	Apply one or more policies to routes being exported from the routing table into BGP.
Options	<i>policy-names</i> —Name of one or more policies.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">Understanding Route AdvertisementRouting Policy Configuration Guideimport on page 95

family (Protocols BGP)

```
Syntax  family {
    (inet | inet6 | inet-vpn | inet6-vpn | iso-vpn) {
        (any | flow | labeled-unicast | multicast | unicast) {
            accepted-prefix-limit {
                maximum number;
                teardown <percentage-threshold> idle-timeout (forever | minutes);
            }
            add-path {
                send {
                    path-count number;
                    prefix-policy [ policy-names ];
                }
                receive;
            }
            algp [disable];
            loops number;
            prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
            }
            protection;
            rib-group group-name;
            topology name {
                community {
                    target identifier;
                }
            }
        }
        flow {
            no-validate policy-name;
        }
        labeled-unicast {
            accepted-prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
            }
            aggregate-label {
                community community-name;
            }
            explicit-null {
                connected-only;
            }
            prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
            }
            resolve-vpn;
            rib (inet.3 | inet6.3);
            rib-group group-name;
            traffic-statistics {
                file filename <world-readable | no-world-readable>;
                interval seconds;
            }
        }
    }
}
```

```
    }
  }
  route-target {
    accepted-prefix-limit {
      maximum number;
      proxy-generate <route-target-policy route-target-policy-name>;
      teardown <percentage> <idle-timeout (forever | minutes)>;
    }
    advertise-default;
    external-paths number;
    prefix-limit {
      maximum number;
      teardown <percentage> <idle-timeout (forever | minutes)>;
    }
  }
  (inet-mdt | inet-mvpn | inet6-mvpn | l2vpn) {
    signaling {
      accepted-prefix-limit {
        maximum number;
        teardown <percentage-threshold> idle-timeout (forever | minutes);
      }
      add-path {
        send {
          path-count number;
          prefix-policy [ policy-names ];
        }
        receive;
      }
      aigp [disable];
      damping;
      loops number;
      prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
      }
      rib-group group-name;
    }
  }
}
```


Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>inet-mvpn and inet6-mvpn statements introduced in Junos OS Release 8.4.</p> <p>inet-mdt statement introduced in Junos OS Release 9.4.</p> <p>Support for the loops statement introduced in Junos OS Release 9.6.</p>
Description	<p>Enable multiprotocol BGP (MP-BGP) by configuring BGP to carry network layer reachability information (NLRI) for address families other than unicast IPv4, to specify MP-BGP to carry NLRI for the IPv6 address family, or to carry NLRI for VPNs.</p>

- Options**
- any**—Configure the family type to be both unicast and multicast.
 - inet**—Configure NLRI parameters for IPv4.
 - inet6**—Configure NLRI parameters for IPv6.
 - inet-mdt**—Configure NLRI parameters for the multicast distribution tree (MDT) subaddress family identifier (SAFI) for IPv4 traffic in Layer 3 VPNs.
 - inet-mvpn**—Configure NLRI parameters for IPv4 for multicast VPNs.
 - inet6-mvpn**—Configure NLRI parameters for IPv6 for multicast VPNs.
 - inet-vpn**—Configure NLRI parameters for IPv4 for Layer 3 VPNs.
 - inet6-vpn**—Configure NLRI parameters for IPv6 for Layer 3 VPNs.
 - iso-vpn**—Configure NLRI parameters for IS-IS for Layer 3 VPNs.
 - l2vpn**—Configure NLRI parameters for IPv4 for MPLS-based Layer 2 VPNs and VPLS.
 - labeled-unicast**—Configure the family type to be labeled-unicast. This means that the BGP peers are being used only to carry the unicast routes that are being used by labeled-unicast for resolving the labeled-unicast routes. This statement is supported only with **inet** and **inet6**.
 - multicast**—Configure the family type to be multicast. This means that the BGP peers are being used only to carry the unicast routes that are being used by multicast for resolving the multicast routes.
 - unicast**—Configure the family type to be unicast. This means that the BGP peers only carry the unicast routes that are being used for unicast forwarding purposes. The default family type is **unicast**.

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 | <ul style="list-style-type: none">• autonomous-system on page 65• local-as on page 102• Understanding Multiprotocol BGP |
|------------------------------|---|

graceful-restart (Protocols BGP)

Syntax	<pre>graceful-restart { disable; restart-time <i>seconds</i>; stale-routes-time <i>seconds</i>; }</pre>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i> <i>neighbor address</i>], [edit protocols bgp], [edit protocols bgp <i>group group-name</i>], [edit protocols bgp <i>group group-name neighbor address</i>]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Enable graceful restart for BGP. Graceful restart allows a routing device undergoing a restart to inform its adjacent neighbors and peers of its condition. Graceful restart is disabled by default.</p> <p>To configure the duration of the BGP graceful restart period, include the restart-time statement at the [edit protocols bgp graceful-restart] hierarchy level. To set the length of time the router waits to receive messages from restarting neighbors before declaring them down, include the stale-routes-time statement at the [edit protocols bgp graceful-restart] hierarchy level.</p>



NOTE: If you configure graceful restart after a BGP session has been established, the BGP session restarts and the peers negotiate graceful restart capabilities.

For graceful restart to function properly, graceful restart must be enabled at the **[edit routing-instance *instance-name* routing-options]** or **[edit routing-options]** hierarchy level as well as in the protocol level.

For example:

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

Graceful restart is enabled both at the **[edit routing-options]** hierarchy level, as well as at the routing protocol level. If graceful restart is not configured in both sections, the peer might have its route removed after a restart, which is not the intended behavior.

The remaining statements are explained separately.

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

Related Documentation	<ul style="list-style-type: none">• Configuring Graceful Restart Options for BGP• Configuring Graceful Restart for QFabric Systems• Junos OS High Availability Configuration Guide
------------------------------	--

group (Protocols BGP)

```

Syntax  group group-name {
        advertise-inactive;
        allow [ network/mask-length ];
        authentication-key key;
        cluster cluster-identifier;
        damping;
        description text-description;
        export [ policy-names ];
        family {
            (inet | inet6 | inet-vpn | inet6-vpn | l2-vpn) {
                (any | multicast | unicast | signaling) {
                    accepted-prefix-limit {
                        maximum number;
                        teardown <percentage> <idle-timeout (forever | minutes)>;
                    }
                    add-path {
                        send {
                            path-count number;
                            prefix-policy [ policy-names ];
                        }
                        receive;
                    }
                    aigp [disable];
                    damping;
                    prefix-limit {
                        maximum number;
                        teardown <percentage> <idle-timeout (forever | minutes)>;
                    }
                    rib-group group-name;
                    topology name {
                        community {
                            target identifier;
                        }
                    }
                }
            }
            flow {
                no-validate policy-name;
            }
            labeled-unicast {
                accepted-prefix-limit {
                    maximum number;
                    teardown <percentage> <idle-timeout (forever | minutes)>;
                }
                explicit-null {
                    connected-only;
                }
                prefix-limit {
                    maximum number;
                    teardown <percentage> <idle-timeout (forever | minutes)>;
                }
                resolve-vpn;
                rib inet.3;
            }
        }
    }

```

```

        rib-group group-name;
    }
}
route-target {
    accepted-prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
    }
    advertise-default;
    external-paths number;
    prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
    }
}
}
hold-time seconds;
import [ policy-names ];
ipsec-sa ipsec-sa;
keep (all | none);
local-address address;
local-as autonomous-system <private>;
local-preference local-preference;
log-updown;
metric-out metric;
multihop <ttl-value>;
multipath {
    multiple-as;
}
no-aggregator-id;
no-client-reflect;
out-delay seconds;
passive;
peer-as autonomous-system;
preference preference;
remove-private;
tcp-mss segment-size;
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
type type;
neighbor address {
    ... peer-specific-options ...
}
}

```

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

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 11.3 for the QFX Series.
Description	<p>Define a BGP peer group. BGP peer groups share a common type, peer autonomous system (AS) number, and cluster ID, if present. To configure multiple BGP groups, include multiple group statements.</p> <p>By default, the group's options are identical to the global BGP options. To override the global options, include group-specific options within the group statement.</p> <p>The group statement is one of the statements you must include in the configuration to run BGP on the routing device.</p> <p>Each group must contain at least one peer.</p>
Options	<p>group-name—Name of the BGP group.</p> <p>The remaining statements are explained separately.</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">• BGP Configuration Guide

hold-time (Protocols BGP)

Syntax	<code>hold-time seconds;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols bgp],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i></code> <code> <i>neighbor address</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> bgp],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> bgp <i>group group-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> bgp <i>group group-name neighbor address</i>],</code> <code>[edit protocols bgp],</code> <code>[edit protocols bgp <i>group group-name</i>],</code> <code>[edit protocols bgp <i>group group-name neighbor address</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i></code> <code> <i>neighbor address</i>]</code>
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 11.3 for the QFX Series.
Description	<p>Specify the hold-time value to use when negotiating a connection with the peer. The hold-time value is advertised in open packets and indicates to the peer the length of time that it should consider the sender valid. If the peer does not receive a keepalive, update, or notification message within the specified hold time, the BGP connection to the peer is closed and routing devices through that peer become unavailable.</p> <p>The hold time is three times the interval at which keepalive messages are sent.</p> <p>BGP on the local routing device uses the smaller of either the local hold-time value or the peer's hold-time value received in the open message as the hold time for the BGP connection between the two peers.</p> <p>Starting in Junos OS Release 12.3, the BGP hold-time value can be zero (0). This implies that the speaker does not expect keepalive messages from its peer to maintain the BGP session. When negotiating between two peers, if one side requests a nonzero hold time and the other requests a zero hold time, the negotiation settles on the nonzero value and keepalive intervals are determined accordingly. Both sides must be set to zero for keepalive messages to stop being sent.</p>
Options	seconds —Hold time. Range: 10 through 65,535 seconds (or 0 for infinite hold time) Default: 90 seconds



TIP: When you set a hold-time value of 1 through 19 seconds, we recommend that you also configure the BGP `precision-timers` statement. The `precision-timers` statement ensures that if scheduler slip messages occur, the routing device continues to send keepalive messages. When the `precision-timers` statement is included, keepalive message generation is performed in a dedicated kernel thread, which helps to prevent BGP session flaps.

Required Privilege Level	routing—To view this statement in the configuration.
	routing-control—To add this statement to the configuration.
Related Documentation	• BGP Messages Overview
	• <code>precision-timers</code>

idle-after-switch-over

Syntax	<code>idle-after-switch-over (forever seconds);</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>]
Release Information	Statement introduced in Junos OS Release 9.5. Statement introduced in Junos OS Release 9.5 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Configure the routing device so that it does not automatically reestablish BGP peer sessions after a nonstop active routing (NSR) switchover. This feature is particularly useful if you are using dynamic routing policies because the dynamic database is not synchronized with the backup Routing Engine when NSR is enabled.
Options	forever —Do not reestablish a BGP peer session after an non-stop routing switchover until the clear bgp neighbor command is issued. seconds —Do not reestablish a BGP peer session after an non-stop routing switchover until after the specified period. Range: 1 through 4,294,967,295 ($2^{32} - 1$)
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">Preventing Automatic Reestablishment of BGP Peer Sessions After NSR SwitchoversRouting Policy Configuration GuideJunos OS High Availability Configuration Guide

import (Protocols BGP)

Syntax	<code>import [<i>policy-names</i>];</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name neighbor address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name neighbor address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp <i>group group-name</i>],</p> <p>[edit protocols bgp <i>group group-name neighbor address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name neighbor address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Apply one or more routing policies to routes being imported into the Junos OS routing table from BGP.
Options	<i>policy-names</i> —Name of one or more policies.
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"> • Example: Configuring BGP Interactions with IGP • Understanding Route Advertisement • Importing and Exporting Routes • Routing Policy Configuration Guide • export on page 82

include-mp-next-hop

Syntax	include-mp-next-hop;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Enable multiprotocol updates to contain next-hop reachability information.
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">Examples: Configuring Multiprotocol BGP

keep

Syntax	keep (all none);
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Specify whether routes learned from a BGP peer are retained in the routing table even if they contain an AS number that was exported from the local AS.
Default	If you do not include this statement, most routes are retained in the routing table.
Options	<p>all—Retain all routes.</p> <p>none—Retain none of the routes. When keep none is configured for the BGP session and the inbound policy changes, Junos OS forces readvertisement of the full set of routes advertised by the peer.</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"> Understanding Route Advertisement out-delay on page 121

labeled-unicast (Protocols BGP)

Syntax	<pre> labeled-unicast { accepted-prefix-limit { maximum <i>number</i>; teardown <<i>percentage</i>> <idle-timeout (forever <i>minutes</i>)>; } aggregate-label { community <i>community-name</i>; } explicit-null { connected-only; } prefix-limit { maximum <i>number</i>; teardown <<i>percentage</i>> <idle-timeout (forever <i>minutes</i>)>; } resolve-vpn; rib (inet.3 inet6.3); rib-group <i>group-name</i>; } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp family (inet inet6)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> family (inet inet6)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family (inet inet6)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp family (inet inet6)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> family (inet inet6)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family (inet inet6)],</p> <p>[edit protocols bgp family (inet inet6)],</p> <p>[edit protocols bgp group <i>group-name</i> family (inet inet6)],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> family (inet inet6)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp family (inet inet6)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> family (inet inet6)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family (inet inet6)]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Configure the family type to be labeled-unicast.</p> <p>The remaining statements are explained separately.</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

- Examples: Configuring Multiprotocol BGP

local-address (Protocols BGP)

Syntax	<code>local-address address;</code>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i> <i>neighbor address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name neighbor address</i>], [edit protocols bgp], [edit protocols bgp <i>group group-name</i>], [edit protocols bgp <i>group group-name neighbor address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i> <i>neighbor address</i>]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Specify the address of the local end of a BGP session. This address is used to accept incoming connections to the peer and to establish connections to the remote peer. When none of the operational interfaces are configured with the specified local address, a session with a BGP peer is placed in the idle state.</p> <p>You generally configure a local address to explicitly configure the system's IP address from BGP's point of view. This IP address can be either an IPv6 or IPv4 address. Typically, an IP address is assigned to a loopback interface, and that IP address is configured here.</p> <p>For internal BGP (IBGP) peering sessions, generally the loopback interface (lo0) is used to establish connections between the IBGP peers. The loopback interface is always up as long as the device is operating. If there is a route to the loopback address, the IBGP peering session stays up. If a physical interface address is used instead and that interface goes up and down, the IBGP peering session also goes up and down. Thus, the loopback interface provides fault tolerance in case the physical interface or the link goes down, if the device has link redundancy.</p> <p>When a device peers with a remote device's loopback interface address, the local device expects BGP update messages to come from (be sourced by) the remote device's loopback interface address. The local-address statement enables you to specify the source information in BGP update messages. If you omit the local-address statement, the expected source of BGP update messages is based on the device's source address selection rules, which normally result in the egress interface address being the expected source of update messages. When this happens, the peering session is not established because a mismatch exists between the expected source address (the egress interface</p>

of the peer) and the actual source (the loopback interface of the peer). To ensure that the expected source address matches the actual source address, specify the loopback interface address in the **local-address** statement.



NOTE: Although a BGP session can be established when only one of the paired routing devices has **local-address** configured, we strongly recommend that you configure **local-address** on both paired routing devices for IBGP and multihop EBGP sessions. The **local-address** statement ensures that deterministic fixed addresses are used for the BGP session end-points.

If you include the **default-address-selection** statement in the configuration, the software chooses the system default address as the source for most locally generated IP packets. For protocols in which the local address is unconstrained by the protocol specification, for example IBGP and multihop EBGP, if you do not configure a specific local address when configuring the protocol, the local address is chosen using the same methods as other locally generated IP packets.

Default If you do not configure a local address, BGP uses the routing device's source address selection rules to set the local address.

Options **address**—IPv6 or IPv4 address of the local end of the connection.

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

Related Documentation

- Example: Configuring Internal BGP Peering Sessions on Logical Systems
- Example: Configuring Internal BGP Peer Sessions
- Understanding Internal BGP Peering Sessions
- router-id

local-as

Syntax	<code>local-as <i>autonomous-system</i> <loops <i>number</i>> <private alias> <no-prepend-global-as>;</code>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group</i> <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp], [edit protocols bgp <i>group</i> <i>group-name</i>], [edit protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>alias option introduced in Junos OS Release 9.5.</p> <p>no-prepend-global-as option introduced in Junos OS Release 9.6.</p>
Description	<p>Specify the local autonomous system (AS) number. An AS is a set of routing devices that are under a single technical administration and generally use a single interior gateway protocol (IGP) and metrics to propagate routing information within the set of routing devices.</p> <p>Internet service providers (ISPs) sometimes acquire networks that belong to a different AS. When this occurs, there is no seamless method for moving the BGP peers of the acquired network to the AS of the acquiring ISP. The process of configuring the BGP peers with the new AS number can be time-consuming and cumbersome. In this case, it might not be desirable to modify peer arrangements or configuration. During this kind of transition period, it can be useful to configure BGP-enabled devices in the new AS to use the former AS number in BGP updates. This former AS number is called a <i>local</i> AS.</p>



NOTE: If you are using BGP on the routing device, you must configure an AS number before you specify the local as number.

In Junos OS Release 9.1 and later, the AS numeric range in plain-number format is extended to provide BGP support for 4-byte AS numbers, as defined in RFC 4893, *BGP Support for Four-octet AS Number Space*.

In Junos OS Release 9.3 and later, you can also configure a 4-byte AS number using the AS-dot notation format of two integer values joined by a period: *<16-bit high-order value in decimal>.<16-bit low-order value in decimal>*. For

example, the 4-byte AS number of 65546 in plain-number format is represented as 1.10 in the AS-dot notation format.

Options **alias**—(Optional) Configure the local AS as an alias of the global AS number configured for the router at the **[edit routing-options]** hierarchy level. As a result, a BGP peer considers any local AS to which it is assigned as equivalent to the primary AS number configured for the routing device. When you use the **alias** option, only the AS (global or local) used to establish the BGP session is prepended in the AS path sent to the BGP neighbor.

autonomous-system—AS number.

Range: 1 through 4,294,967,295 ($2^{32} - 1$) in plain-number format

Range: 0.0 through 65535.65535 in AS-dot notation format

loops number—(Optional) Specify the number of times detection of the AS number in the AS_PATH attribute causes the route to be discarded or hidden. For example, if you configure **loops 1**, the route is hidden if the AS number is detected in the path one or more times. This is the default behavior. If you configure **loops 2**, the route is hidden if the AS number is detected in the path two or more times.



NOTE: If you configure the local AS values for any BGP group, the detection of routing loops is performed using both the AS and the local AS values for all BGP groups.

If the local AS for the EBGP or IBGP peer is the same as the current AS, do not use the **local-as** statement to specify the local AS number.

When you configure the local AS within a VRF, this impacts the AS path loop-detection mechanism. All of the **local-as** statements configured on the device are part of a single AS domain. The AS path loop-detection mechanism is based on looking for a matching AS present in the domain.

Range: 1 through 10

Default: 1

no-prepend-global-as—(Optional) Specify to strip the global AS and to prepend only the local AS in AS paths sent to external peers.

private—(Optional) Configure to use the local AS only during the establishment of the BGP session with a BGP neighbor but to hide it in the AS path sent to external BGP peers. Only the global AS is included in the AS path sent to external peers.



NOTE: The **private** and **alias** options are mutually exclusive. You cannot configure both options with the same **local-as** statement.

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

Related Documentation

- Examples: Configuring BGP Local AS
- Example: Configuring a Local AS for EBGp Sessions
- [autonomous-system on page 65](#)
- [family on page 83](#)

local-interface (IPv6)

Syntax `local-interface interface-name;`

Hierarchy Level [edit logical-systems *logical-system-name* protocols bgp [group](#) *group-name* [neighbor](#) *ipv6-link-local-address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* [neighbor](#) *ipv6-link-local-address*],
[edit protocols bgp group *group-name* [neighbor](#) *ipv6-link-local-address*],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* [neighbor](#) *ipv6-link-local-address*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Specify the interface name of the EBGp peer that uses IPv6 link-local addresses. This peer is link-local in scope.

Options *interface-name*—Interface name of the EBGp IPv6 peer.

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

Related Documentation

- Example: Configuring Internal BGP Peering Sessions on Logical Systems
- Example: Configuring Internal BGP Peer Sessions
- Example: Configuring External BGP on Logical Systems with IPv6 Interfaces
- Understanding Internal BGP Peering Sessions

local-preference

Syntax	<code>local-preference local-preference;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Modify the value of the LOCAL_PREF path attribute, which is a metric used by IBGP sessions to indicate the degree of preference for an external route. The route with the highest local preference value is preferred.</p> <p>The LOCAL_PREF path attribute always is advertised to internal BGP peers and to neighboring confederations. It is never advertised to external BGP peers.</p>
Default	If you omit this statement, the LOCAL_PREF path attribute, if present, is not modified.
Options	<p>local-preference—Preference to assign to routes learned from BGP or from the group or peer.</p> <p>Range: 0 through 4,294,967,295 ($2^{32} - 1$)</p> <p>Default: If the LOCAL_PREF path attribute is present, do not modify its value. If a BGP route is received without a LOCAL_PREF attribute, the route is handled locally (it is stored in the routing table and advertised by BGP) as if it were received with a LOCAL_PREF value of 100. By default, non-BGP routes that are advertised by BGP are advertised with a LOCAL_PREF value of 100.</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"> Example: Configuring the Local Preference Value for BGP Routes Understanding Internal BGP Peering Sessions

- [preference on page 126](#)

log-updown (Protocols BGP)

Syntax	log-updown;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>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 11.3 for the QFX Series.</p>
Description	<p>Specify to generate a log a message whenever a BGP peer makes a state transition. Messages are logged using the system logging mechanism located at the [edit system syslog] hierarchy level.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Example: Preventing BGP Session Resets• Junos OS System Basics Configuration Guide• traceoptions on page 133

metric-out (Protocols BGP)

Syntax	<code>metric-out (<i>metric</i> minimum-igp <i>offset</i> igp (delay-med-update <i>offset</i>);</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name neighbor address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name neighbor address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp <i>group group-name</i>],</p> <p>[edit protocols bgp <i>group group-name neighbor address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name neighbor address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Option delay-med-update introduced in Junos OS Release 9.0.</p>
Description	<p>Specify the metric for all routes sent using the multiple exit discriminator (MED, or MULTI_EXIT_DISC) path attribute in update messages. This path attribute is used to discriminate among multiple exit points to a neighboring AS. If all other factors are equal, the exit point with the lowest metric is preferred.</p> <p>You can specify a constant metric value by including the metric option. For configurations in which a BGP peer sends third-party next hops that require the local system to perform next-hop resolution—IBGP configurations, configurations within confederation peers, or EBGP configurations that include the multihop command—you can specify a variable metric by including the minimum-igp or igp option.</p> <p>You can increase or decrease the variable metric calculated from the IGP metric (either from the igp or minimum-igp statement) by specifying a value for offset. The metric is increased by specifying a positive value for offset, and decreased by specifying a negative value for offset.</p> <p>In Junos OS Release 9.0 and later, you can specify that a BGP group or peer not advertise updates for the MED path attributes used to calculate IGP costs for BGP next hops unless the MED is lower. You can also configure an interval to delay when MED updates are sent by including the med-igp-update-interval minutes statement at the [edit routing-options] hierarchy level.</p>
Options	<p>delay-med-update—Specify that a BGP group or peer configured with the metric-out igp statement not advertise MED updates unless the current MED value is lower than</p>

the previously advertised MED value, or another attribute associated with the route has changed, or the BGP peer is responding to a refresh route request.



NOTE: You cannot configure the `delay-med-update` statement at the global BGP level.

igp—Set the metric to the most recent metric value calculated in the IGP to get to the BGP next hop. Routes learned from an EBGP peer usually have a next hop on a directly connected interface and thus the IGP value is equal to zero. This is the value advertised.

metric—Primary metric on all routes sent to peers.

Range: 0 through 4,294,967,295 ($2^{32} - 1$)

Default: No metric is sent.

minimum-igp—Set the metric to the minimum metric value calculated in the IGP to get to the BGP next hop. If a newly calculated metric is greater than the minimum metric value, the metric value remains unchanged. If a newly calculated metric is lower, the metric value is lowered to that value. When you change a neighbor's export policy from any configuration to a configuration that sets the minimum IGP offset on an exported route, the advertised MED is not updated if the value would increase as a result, even if the previous configuration does not use a minimum IGP-based MED value. This behavior helps to prevent unnecessary route flapping when an IGP cost changes, by not forcing a route update if the metric value increases past the previous lowest known value.

offset—Increases or decreases the metric by this value.

Range: -2^{31} through $2^{31} - 1$

Default: None

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">• Example: Associating the MED Path Attribute with the IGP Metric and Delaying MED Updates• Examples: Configuring BGP MED• Example: Configuring the MED Attribute Directly• Understanding the MED Attribute• med-igp-update-interval
------------------------------	--


mtu-discovery

Syntax	mtu-discovery;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Configure TCP path maximum transmission unit (MTU) discovery.</p> <p>TCP path MTU discovery enables BGP to automatically discover the best TCP path MTU for each BGP session. In Junos OS, TCP path MTU discovery is disabled by default for all BGP neighbor sessions.</p> <p>When MTU discovery is disabled, TCP sessions that are not directly connected transmit packets of 512-byte maximum segment size (MSS). These small packets minimize the chances of packet fragmentation at a device along the path to the destination. However, because most links use an MTU of at least 1500 bytes, 512-byte packets do not result in the most efficient use of link bandwidth. For directly connected EBGP sessions, MTU mismatches prevent the BGP session from being established. As a workaround, enable path MTU discovery within the EBGP group.</p> <p>Path MTU discovery dynamically determines the MTU size on the network path between the source and the destination, with the goal of avoiding IP fragmentation. Path MTU discovery works by setting the Don't Fragment (DF) bit in the IP headers of outgoing packets. When a device along the path has an MTU that is smaller than the packet, the device drops the packet. The device also sends back an ICMP Fragmentation Needed (Type 3, Code 4) message that contains the device's MTU, thus allowing the source to reduce its path MTU appropriately. The process repeats until the MTU is small enough to traverse the entire path without fragmentation.</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**

- [Example: Limiting TCP Segment Size for BGP](#)
- [Configuring the Junos OS for IPv6 Path MTU Discovery](#)
- [Configuring the Junos OS for Path MTU Discovery on Outgoing GRE Tunnel Connections](#)

multihop

Syntax	<pre>multihop { no-nexthop-change; ttl <i>ttl-value</i>; }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>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 11.3 for the QFX Series.</p>
Description	<p>Configure an EBGp multihop session.</p> <p>For Layer 3 VPNs, you configure the EBGp multihop session between the PE and CE routing devices. This allows you to configure one or more routing devices between the PE and CE routing devices.</p> <p>An external confederation peer is a special case that allows unconnected third-party next hops. You do not need to configure multihop sessions explicitly in this particular case because multihop behavior is implied.</p> <p>If you have external BGP confederation peer-to-loopback addresses, you still need the multihop configuration.</p>
	<div>  <p>NOTE: You cannot configure the <code>accept-remote-nexthop</code> statement at the same time.</p> </div>
Default	<p>If you omit this statement, all EBGp peers are assumed to be directly connected (that is, you are establishing a nonmultihop, or “regular,” BGP session), and the default time-to-live (TTL) value is 1.</p>

The remaining statements are explained separately.

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

Related Documentation	<ul style="list-style-type: none">• Example: Configuring EBGp Multihop Sessions• Configuring EBGp Multihop Sessions Between PE and CE Routers in Layer 3 VPNs• accept-remote-nextthop on page 55• no-nextthop-change• ttl
------------------------------	---

multipath (Protocols BGP)

Syntax	<pre> multipath { multiple-as; vpn-unequal-cost equal-external-internal; } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Allow load sharing among multiple EBGP paths and multiple IBGP paths. A path is considered a BGP equal-cost path (and will be used for forwarding) if a tie-break is performed. The tie-break is performed after the BGP route path selection step that chooses the next-hop path that is resolved through the IGP route with the lowest metric. All paths with the same neighboring AS, learned by a multipath-enabled BGP neighbor, are considered.</p>
Options	<p>multiple-as—Disable the default check requiring that paths accepted by BGP multipath must have the same neighboring AS.</p> <p>vpn-unequal-cost equal-external-internal—Enable load-balancing in a Layer 3 VPN with unequal cost paths.</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"> Understanding BGP Path Selection Example: Load Balancing BGP Traffic

neighbor (Protocols BGP)

```
Syntax  neighbor address {
    accept-remote-nexthop;
    advertise-external <conditional>;
    advertise-inactive;
    (advertise-peer-as | no-advertise-peer-as);
    as-override;
    authentication-algorithm algorithm;
    authentication-key key;
    authentication-key-chain key-chain;
    cluster cluster-identifier;
    damping;
    description text-description;
    export [ policy-names ];
    family {
        (inet | inet6 | inet-mvpn | inet6-mpvn | inet-vpn | inet6-vpn | iso-vpn | l2-vpn) {
            (any | flow | multicast | unicast | signaling) {
                accepted-prefix-limit {
                    maximum number;
                    teardown <percentage> <idle-timeout (forever | minutes)>;
                }
                damping;
                prefix-limit {
                    maximum number;
                    teardown <percentage> <idle-timeout (forever | minutes)>;
                }
                rib-group group-name;
                topology name {
                    community {
                        target identifier;
                    }
                }
            }
        }
        flow {
            no-validate policy-name;
        }
        labeled-unicast {
            accepted-prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
            }
            aggregate-label {
                community community-name;
            }
            explicit-null {
                connected-only;
            }
            prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
            }
            resolve-vpn;
            rib inet.3;
        }
    }
}
```

```

    rib-group group-name;
    topology name {
        community {
            target identifier;
        }
    }
}
route-target {
    advertise-default;
    external-paths number;
    accepted-prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
    }
    prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
    }
}
signaling {
    prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
    }
}
}
graceful-restart {
    disable;
    restart-time seconds;
    stale-routes-time seconds;
}
hold-time seconds;
import [ policy-names ];
ipsec-sa ipsec-sa;
keep (all | none);
local-address address;
local-as autonomous-system <private>;
local-interface interface-name;
local-preference preference;
log-updown;
metric-out (metric | minimum-igp <offset> | igp <offset>);
mtu-discovery;
multihop <ttl-value>;
multipath {
    multiple-as;
}
no-aggregator-id;
no-client-reflect;
out-delay seconds;
passive;
peer-as autonomous-system;
preference preference;
tcp-mss segment-size;
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
}

```

```
    flag flag <flag-modifier> <disable>;  
  }  
  vpn-apply-export;  
}
```

Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>]
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 11.3 for the QFX Series.
Description	<p>Explicitly configure a neighbor (peer). To configure multiple BGP peers, include multiple neighbor statements.</p> <p>By default, the peer's options are identical to those of the group. You can override these options by including peer-specific option statements within the neighbor statement.</p> <p>The neighbor statement is one of the statements you can include in the configuration to define a minimal BGP configuration on the routing device. (You can include an allow all statement in place of a neighbor statement.)</p>
Options	<p>address—IPv6 or IPv4 address of a single peer.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• BGP Configuration Guide

no-aggregator-id

Syntax	no-aggregator-id;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Prevent different routing devices within an AS from creating aggregate routes that contain different AS paths.</p> <p>Junos OS performs route aggregation, which is the process of combining the characteristics of different routes so that only a single route is advertised. Aggregation reduces the amount of information that BGP must store and exchange with other BGP systems. When aggregation occurs, the local routing device adds the local AS number and the router ID to the aggregator path attribute. The no-aggregator-id statement causes Junos OS to place a 0 in the router ID field and thus eliminate the possibility of having multiple aggregate advertisements in the network, each with different path information.</p>
Default	If you omit this statement, the router ID is included in the BGP aggregator path attribute.
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"> BGP Messages Overview

no-client-reflect

Syntax	no-client-reflect;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Disable intracluster route redistribution by the system acting as the route reflector. Include this statement when the client cluster is fully meshed to prevent the sending of redundant route advertisements. Route reflection provides a way to decrease BGP control traffic and minimizing the number of update messages sent within the AS.
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">• Example: Configuring BGP Route Reflectors• cluster on page 75

no-validate

Syntax	<code>no-validate <i>policy-name</i>;</code>
Hierarchy Level	<code>[edit protocols bgp group <i>group-name</i> <i>family</i> (inet inet flow)],</code> <code>[edit protocols bgp group <i>group-name</i> neighbor address <i>family</i> (inet inet flow)],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <i>family</i> (inet inet flow)],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor address <i>family</i> (inet inet flow)]</code>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>When BGP is carrying flow-specification network layer reachability information (NLRI) messages, the no-validate statement omits the flow route validation procedure after packets are accepted by a policy.</p> <p>The receiving BGP-enabled device accepts a flow route if it passes the following criteria:</p> <ul style="list-style-type: none"> • The originator of a flow route matches the originator of the best match unicast route for the destination address that is embedded in the route. • There are no more specific unicast routes, when compared to the destination address of the flow route, for which the active route has been received from a different next-hop autonomous system. <p>The first criterion ensures that the filter is being advertised by the next-hop used by unicast forwarding for the destination address embedded in the flow route. For example, if a flow route is given as 10.1.1.1, proto=6, port=80, the receiving BGP-enabled device selects the more specific unicast route in the unicast routing table that matches the destination prefix 10.1.1.1/32. On a unicast routing table containing 10.1/16 and 10.1.1/24, the latter is chosen as the unicast route to compare against. Only the active unicast route entry is considered. This follows the concept that a flow route is valid if advertised by the originator of the best unicast route.</p> <p>The second criterion addresses situations in which a given address block is allocated to different entities. Flows that resolve to a best-match unicast route that is an aggregate route are only accepted if they do not cover more specific routes that are being routed to different next-hop autonomous systems.</p> <p>You can bypass the validation process and use your own specific import policy. To disable the validation procedure and use an import policy instead, include the no-validate statement in the configuration.</p> <p>Flow routes configured for VPNs with family inet-vpn are not automatically validated, so the no-validate statement is not supported at the <code>[edit protocols bgp group <i>group-name</i> family <i>inet-vpn</i>]</code> hierarchy level. No validation is needed if the flow routes are configured locally between devices in a single AS.</p>

Options *policy-name*—Import policy to match NLRI messages.


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

Related • Example: Configuring Flow Routes
Documentation

out-delay

Syntax	<code>out-delay seconds;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Specify how long a route must be present in the Junos OS routing table before it is exported to BGP. Use this time delay to help bundle routing updates.</p> <p>When configured, the out-delay value displays as Outbound Timer when using show bgp group or show bgp group neighbor commands.</p>
Default	If you omit this statement, routes are exported to BGP immediately after they have been added to the routing table.
Options	<p>seconds—Output delay time.</p> <p>Range: 0 through 65,535 seconds</p> <p>Default: 0 seconds</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"> Understanding Route Advertisement

outbound-route-filter

Syntax	<pre> outbound-route-filter { bgp-orf-cisco-mode; prefix-based { accept { (inet inet6); } } } </pre>
Hierarchy Level	<pre> [edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>] </pre>
Release Information	<p>Statement introduced in Junos OS Release 9.2.</p> <p>Statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Configure a BGP peer to accept outbound route filters from a remote peer.
Options	<p>accept—Specify that outbound route filters from a BGP peer be accepted.</p> <p>inet—Specify that IPv4 prefix-based outbound route filters be accepted.</p> <p>inet6—Specify that IPv6 prefix-based outbound route filters be accepted.</p>
	<div>  <p>NOTE: You can specify that both IPv4 and IPv6 outbound route filters be accepted.</p> </div>
	<p>prefix-based—Specify that prefix-based filters be accepted.</p> <p>The bgp-orf-cisco-mode statement is explained separately.</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**
- Example: Configuring BGP Prefix-Based Outbound Route Filtering

passive (Protocols BGP)

Syntax	passive;
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group</i> <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group</i> <i>group-name</i> <i>neighbor</i> <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i> <i>neighbor</i> <i>address</i>], [edit protocols bgp], [edit protocols bgp <i>group</i> <i>group-name</i>], [edit protocols bgp <i>group</i> <i>group-name</i> <i>neighbor</i> <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i> <i>neighbor</i> <i>address</i>]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Configure the routing device so that active open messages are not sent to the peer. Once you configure the routing device to be passive, the routing device will wait for the peer to issue an open request before a message is sent.
Default	If you omit this statement, all explicitly configured peers are active, and each peer periodically sends open requests until its peer responds.
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"> • Example: Preventing BGP Session Flaps When VPN Families Are Configured

peer-as (Protocols BGP)

Syntax	<code>peer-as <i>autonomous-system</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Specify the neighbor (peer) autonomous system (AS) number.</p> <p>For EBGP, the peer is in another AS, so the AS number you specify in the peer-as statement must be different from the local router's AS number, which you specify in the autonomous-system statement. For IBGP, the peer is in the same AS, so the two AS numbers that you specify in the autonomous-system and peer-as statements must be the same.</p> <p>The AS numeric range in plain-number format has been extended in Junos OS Release 9.1 to provide BGP support for 4-byte AS numbers, as defined in RFC 4893, <i>BGP Support for Four-octet AS Number Space</i>. RFC 4893 introduces two new optional transitive BGP attributes, AS4_PATH and AS4_AGGREGATOR. These new attributes are used to propagate 4-byte AS path information across BGP speakers that do not support 4-byte AS numbers. RFC 4893 also introduces a reserved, well-known, 2-byte AS number, AS 23456. This reserved AS number is called AS_TRANS in RFC 4893. All releases of the Junos OS support 2-byte AS numbers.</p> <p>In Junos OS Release 9.2 and later, you can also configure a 4-byte AS number using the AS-dot notation format of two integer values joined by a period: <i><16-bit high-order value in decimal>.<16-bit low-order value in decimal></i>. For example, the 4-byte AS number of 65,546 in plain-number format is represented as 1.10 in the AS-dot notation format.</p> <p>With the introduction of 4-byte AS numbers, you might have a combination of routers that support 4-byte AS numbers and 2-byte AS numbers. For more information about what happens when establishing BGP peer relationships between 4-byte and 2-byte capable routers, see the following topics:</p>

- Establishing a Peer Relationship Between a 4-Byte Capable Router and a 2-Byte Capable Router Using a 2-Byte AS Number in the *Using 4-Byte Autonomous System Numbers in BGP Networks Technology Overview*.
- Establishing a Peer Relationship Between a 4-Byte Capable Router and a 2-Byte Capable Router Using a 4-Byte AS Number in the *Using 4-Byte Autonomous System Numbers in BGP Networks Technology Overview*.

Options *autonomous-system*—AS number.

Range: 1 through 4,294,967,295 ($2^{32} - 1$) in plain-number format for 4-byte AS numbers

Range: 1 through 65,535 in plain-number format for 2-byte AS numbers (this is a subset of the 4-byte range)

Range: 0.0 through 65535.65535 in AS-dot notation format for 4-byte AS numbers

Required Privilege routing—To view this statement in the configuration.

Level routing-control—To add this statement to the configuration.

Related Documentation

- 4-Byte Autonomous System Numbers Overview in the [Using 4-Byte Autonomous System Numbers in BGP Networks Technology Overview](#)

- Juniper Networks Implementation of 4-Byte Autonomous System Numbers in the [Using 4-Byte Autonomous System Numbers in BGP Networks Technology Overview](#)

preference (Protocols BGP)

Syntax	<code>preference <i>preference</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Specify the preference for routes learned from BGP.</p> <p>At the BGP global level, the preference statement sets the preference for routes learned from BGP. You can override this preference in a BGP group or peer preference statement.</p> <p>At the group or peer level, the preference statement sets the preference for routes learned from the group or peer. Use this statement to override the preference set in the BGP global preference statement when you want to favor routes from one group or peer over those of another.</p>
Options	<p>preference—Preference to assign to routes learned from BGP or from the group or peer.</p> <p>Range: 0 through 4,294,967,295 ($2^{32} - 1$)</p> <p>Default: 170 for the primary preference</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"> • local-preference on page 105 • Example: Configuring the Preference Value for BGP Routes

prefix-limit

Syntax	<pre>prefix-limit { maximum <i>number</i>; teardown <percentage> <idle-timeout (forever minutes)>; }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp family (inet inet6) (any flow labeled-unicast multicast unicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> family (inet inet6) (any flow labeled-unicast multicast unicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family (inet inet6) (any flow labeled-unicast multicast unicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp family (inet inet6) (any flow labeled-unicast multicast unicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> family (inet inet6) (any flow labeled-unicast multicast unicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family (inet inet6) (any flow labeled-unicast multicast unicast)],</p> <p>[edit protocols bgp family (inet inet6) (any flow labeled-unicast multicast unicast)],</p> <p>[edit protocols bgp group <i>group-name</i> family (inet inet6) (any labeled-unicast multicast unicast)],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> family (inet inet6) (any flow labeled-unicast multicast unicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp family (inet inet6) (any flow labeled-unicast multicast unicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> family (inet inet6) (any flow labeled-unicast multicast unicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family (inet inet6) (any flow labeled-unicast multicast unicast)]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Limit the number of prefixes received on a BGP peer session and a rate-limit logging when injected prefixes exceed a set limit.
Options	<p>maximum <i>number</i>—When you set the maximum number of prefixes, a message is logged when that number is exceeded.</p> <p>Range: 1 through 4,294,967,295 ($2^{32} - 1$)</p> <p>teardown <percentage>—If you include the teardown statement, the session is torn down when the maximum number of prefixes is reached. If you specify a percentage, messages are logged when the number of prefixes exceeds that percentage. After the session is torn down, it is reestablished in a short time unless you include the idle-timeout statement. Then the session can be kept down for a specified amount of time, or forever. If you specify forever, the session is reestablished only after you issue a clear bgp neighbor command.</p> <p>Range: 1 through 100</p>

idle-timeout (**forever** | *timeout-in-minutes*)—(Optional) If you include the **idle-timeout** statement, the session is torn down for a specified amount of time, or forever. If you specify a period of time, the session is allowed to reestablish after this timeout period. If you specify **forever**, the session is reestablished only after you intervene with a **clear bgp neighbor** command.

Range: 1 through 2400

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

Related Documentation	<ul style="list-style-type: none">• accepted-prefix-limit• Understanding Multiprotocol BGP
------------------------------	---

remove-private

Syntax `remove-private all replace nearest;`

Hierarchy Level [edit logical-systems *logical-system-name* protocols bgp],
 [edit logical-systems *logical-system-name* protocols bgp **group** *group-name*],
 [edit logical-systems *logical-system-name* protocols bgp **group** *group-name* neighbor *address*],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp *group* *group-name* **neighbor** *address*],
 [edit protocols bgp],
 [edit protocols bgp **group** *group-name*],
 [edit protocols bgp *group* *group-name* **neighbor** *address*],
 [edit routing-instances *routing-instance-name* protocols bgp],
 [edit routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
 [edit routing-instances *routing-instance-name* protocols bgp *group* *group-name* **neighbor** *address*]

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 11.3 for the QFX Series.

Description When advertising AS paths to remote systems, have the local system strip private AS numbers from the AS path. The numbers are stripped from the AS path starting at the left end of the AS path (the end where AS paths have been most recently added). The routing device stops searching for private ASs when it finds the first nonprivate AS or a peer's private AS. If the AS path contains the AS number of the external BGP (EBGP) neighbor, BGP does not remove the private AS number.



NOTE: As of Junos OS 10.0R2 and higher, if there is a need to send prefixes to an EBGP peer that has an AS number that matches an AS number in the AS path, consider using the `as-override` statement instead of the `remove-private` statement.

The operation takes place after any confederation member ASs have already been removed from the AS path, if applicable.

The Junos OS recognizes the set of AS numbers that is considered private, a range that is defined in the Internet Assigned Numbers Authority (IANA) assigned numbers document.

The set of reserved AS numbers is in the range from 64,512 through 65,535.

Options **all**—Remove all private AS numbers from the original path. Do not stop the process of removing private AS numbers, even if a public AS number is encountered.

nearest—When you use the **all** and **replace** options, choose the last (right-most) public AS number encountered in the original AS path for the replacement value, as the AS path is processed from left to right. If no public AS number is encountered, the default replacement value is used. (See the **replace** option for information about the default replacement value.)

replace—When you use the **all** option, instead of removing private AS numbers, perform a replace operation. The default replacement value for the private AS number is the local AS number at the BGP group level for the BGP peer. If you are unsure about the replacement value, check the local AS value displayed in the output of the **show bgp group group-name** command.

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

Related Documentation	<ul style="list-style-type: none">• Example: Removing Private AS Numbers from AS Paths
------------------------------	--


rib-group (Protocols BGP)

Syntax	<code>rib-group <i>group-name</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp family inet (labeled-unicast unicast multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> family inet (labeled-unicast unicast multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family inet (labeled-unicast unicast multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp family inet (labeled-unicast unicast multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> family inet (labeled-unicast unicast multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family inet (labeled-unicast unicast multicast)],</p> <p>[edit protocols bgp family inet (labeled-unicast unicast multicast)],</p> <p>[edit protocols bgp group <i>group-name</i> family inet (labeled-unicast unicast multicast)],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> family inet (labeled-unicast unicast multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp family inet (labeled-unicast unicast multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> family inet (labeled-unicast unicast multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family inet (labeled-unicast unicast multicast)]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Add unicast prefixes to unicast and multicast tables.
Options	<i>group-name</i> —Name of the routing table group. The name must start with a letter and can include letters, numbers, and hyphens. You generally specify only one routing table group.
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"> • Example: Exporting Specific Routes from One Routing Table Into Another Routing Table • Example: Importing Direct and Static Routes Into a Routing Instance • Understanding Multiprotocol BGP

tcp-mss (Protocols BGP)

Syntax	<code>tcp-mss <i>segment-size</i>;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols bgp],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor</code> <code> <i>neighbor-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> bgp],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> bgp group <i>group-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> bgp group <i>group-name</i> neighbor <i>neighbor-name</i>],</code> <code>[edit protocols bgp],</code> <code>[edit protocol bgp group <i>group-name</i>],</code> <code>[edit protocols bgp group <i>group-name</i> neighbor <i>neighbor-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor</code> <code> <i>neighbor-name</i>]</code>
Release Information	Statement introduced in Junos OS Release 8.1. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Configure the maximum segment size (MSS) for the TCP connection for BGP neighbors. The MSS is only valid in increments of 2 KB. The value used is based on the value set, but is rounded down to the nearest multiple of 2048.
Options	<i>segment-size</i> —MSS for the TCP connection. Range: 1 through 4096
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">Example: Limiting TCP Segment Size for BGP

traceoptions (Protocols BGP)

Syntax	<pre> traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <flag-modifier> <disable>; } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>4byte-as statement introduced in Junos OS Release 9.2.</p> <p>4byte-as statement introduced in Junos OS Release 9.2 for EX Series switches.</p>
Description	Configure BGP protocol-level tracing options. To specify more than one tracing operation, include multiple flag statements.
<div>  <p>NOTE: The traceoptions statement is not supported on QFabric systems.</p> </div>	
Default	<p>The default BGP protocol-level tracing options are inherited from the routing protocols traceoptions statement included at the [edit routing-options] hierarchy level. The default group-level trace options are inherited from the BGP protocol-level traceoptions statement. The default peer-level trace options are inherited from the group-level traceoptions statement.</p>
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 BGP tracing output in the file bgp-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. If you specify a maximum number of files, you must also specify a maximum file size with the **size** option.

Range: 2 through 1000 files

Default: 10 files

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

BGP Tracing Flags

- **4byte-as**—4-byte AS events.
- **bfd**—BFD protocol events.
- **damping**—Damping operations.
- **graceful-restart**—Graceful restart events.
- **keepalive**—BGP keepalive messages. If you enable the the BGP **update** flag only, received keepalive messages do not generate a trace message.
- **nsr-synchronization**—Nonstop routing synchronization events.
- **open**—Open packets. These packets are sent between peers when they are establishing a connection.
- **packets**—All BGP protocol packets.
- **refresh**—BGP refresh packets.
- **update**—Update packets. These packets provide routing updates to BGP systems. If you enable only this flag, received keepalive messages do not generate a trace message. Use the **keepalive** flag to generate a trace message for keepalive messages.

Global Tracing Flags

- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations

Default: If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Routing protocol task processing
- **timer**—Routing protocol timer processing

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Provide detailed trace information.
- **filter**—Provide filter trace information. Applies only to **route** and **damping** tracing flags.
- **receive**—Trace the packets being received.
- **send**—Trace the packets being transmitted.

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. 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	<ul style="list-style-type: none"> • log-updown on page 106 statement • Understanding Trace Operations for BGP Protocol Traffic • Configuring OSPF Refresh and Flooding Reduction in Stable Topologies

type (Protocols BGP)

Syntax	<code>type type;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>]
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 11.3 for the QFX Series.
Description	<p>Specify the type of BGP peer group.</p> <p>When configuring a BGP group, you can indicate whether the group is an IBGP group or an EBGP group. All peers in an IBGP group are in the same AS, while peers in an EBGP group are in different ASs and normally share a subnet.</p>
Options	<p>type—Type of group:</p> <ul style="list-style-type: none">• external—External group, which allows inter-AS BGP routing• internal—Internal group, which allows intra-AS BGP routing
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">• BGP Configuration Guide

PART 3

Administration

- [Operational Commands on page 139](#)

CHAPTER 3

Operational Commands

clear bgp damping

Syntax	clear bgp damping <logical-system (all <i>logical-system-name</i>)> < <i>prefix</i> >
Syntax (EX Series Switch and QFX Series)	clear bgp damping < <i>prefix</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Clear BGP route flap damping information.
Options	none —Clear all BGP route flap damping information. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. <i>prefix</i> —(Optional) Clear route flap damping information for only the specified destination prefix.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show policy damping on page 172• show route damping
List of Sample Output	clear bgp damping on page 140
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear bgp damping user@host> clear bgp damping

clear bgp neighbor

Syntax	<pre>clear bgp neighbor <as <i>as-number</i>> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)> <neighbor> <soft soft-inbound> <soft-minimum-igp></pre>
Syntax (EX Series Switch and QFX Series)	<pre>clear bgp neighbor <as <i>as-number</i>> <instance <i>instance-name</i>> <neighbor> <soft soft-inbound> <soft-minimum-igp></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Perform one of the following tasks:</p> <ul style="list-style-type: none"> • Change the state of one or more BGP neighbors to IDLE. For neighbors in the ESTABLISHED state, this command drops the TCP connection to the neighbors and then reestablishes the connection. • (soft or soft-inbound keyword only) Reapply export policies or import policies, respectively, and send refresh updates to one or more BGP neighbors without changing their state.
Options	<p>none—Change the state of all BGP neighbors to IDLE.</p> <p>as <i>as-number</i>—(Optional) Apply this command only to neighbors in the specified autonomous system (AS).</p> <p>instance <i>instance-name</i>—(Optional) Apply this command only to neighbors for the specified routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>neighbor—(Optional) IP address of a BGP peer. Apply this command only to the specified neighbor.</p> <p>soft—(Optional) Reapply any export policies and send refresh updates to neighbors without clearing the state.</p> <p>soft-inbound—(Optional) Reapply any import policies and send refresh updates to neighbors without clearing the state.</p>

soft-minimum-igp—(Optional) Provides soft refresh of the outbound state when the interior gateway protocol (IGP) metric is reset.

Required Privilege Level clear

Related Documentation • [show bgp neighbor on page 153](#)

List of Sample Output [clear bgp neighbor on page 142](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

`clear bgp neighbor` user@host> clear bgp neighbor

clear bgp table

Syntax	<code>clear bgp table <i>table-name</i></code> <code><logical-system (all <i>logical-system-name</i>)></code>
Syntax (EX Series Switch and QFX Series)	<code>clear bgp table <i>table-name</i></code>
Release Information	Command introduced in Junos OS Release 9.0. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Request that BGP refresh routes in a specified routing table.
Options	<code>logical-system (all <i>logical-system-name</i>)</code> —(Optional) Perform this operation on all logical systems or on a particular logical system. <code>table-name</code> —Request that BGP refresh routes in the specified table.
Additional Information	In some cases, a prefix limit is associated with a routing table for a VPN instance. When this limit is exceeded (for example, because of a network misconfiguration), some routes might not be inserted in the table. Such routes need to be added to the table after the network issue is resolved. Use the clear bgp table command to request that BGP refresh routes in a VPN instance table.
Required Privilege Level	clear
List of Sample Output	clear bgp table private.inet.0 on page 144 clear bgp table inet.6 logical-system all on page 144 clear bgp table private.inet.6 logical-system ls1 on page 144 clear bgp table logical-system all inet.0 on page 144 clear bgp table logical-system ls2 private.inet.0 on page 144
Output Fields	This command produces no output.

Sample Output

clear bgp table
private.inet.0

user@host> clear bgp table private.inet.0

clear bgp table inet.6
logical-system all

user@host> clear bgp table inet.6 logical-system all

clear bgp table
private.inet.6
logical-system ls1

user@host> clear bgp table private.inet.6 logical-system ls1

clear bgp table
logical-system all
inet.0

user@host> clear bgp table logical-system all inet.0

clear bgp table
logical-system ls2
private.inet.0

user@host> clear bgp table logical-system ls2 private.inet.0

show bgp bmp

Syntax	show bgp bmp
Release Information	Command introduced in Junos OS Release 9.5. Command introduced in Junos OS Release 9.5 for EX Series switches.
Description	Display information about the BGP Monitoring Protocol (BMP).
Options	This command has no options.
Required Privilege Level	view
List of Sample Output	show bgp bmp on page 145
Output Fields	Table 36 on page 145 lists the output fields for the show bgp bmp command. Output fields are listed in the approximate order in which they appear.

Table 36: show bgp bmp Output Fields

Field Name	Field Description
BMP station address/port	IP address and port number of the monitoring station to which BGP Monitoring Protocol (BMP) statistics are sent.
BMP session state	Status of the BMP session: UP or DOWN .
Memory consumed by BMP	Memory used by the active BMP session.
Statistics timeout	Amount of time, in seconds, between transmissions of BMP data to the monitoring station.
Memory limit	Threshold, in bytes, at which the routing device stops collecting BMP data.
Memory-connect retry timeout	Amount of time, in seconds, after which the routing device attempts to resume a BMP session that was ended after the configured memory threshold was exceeded.

Sample Output

```

show bgp bmp

user@host> show bgp bmp
  BMP station address/port: 172.24.24.157+5454
  BMP session state: DOWN
  Memory consumed by BMP: 0
  Statistics timeout: 15
  Memory limit: 10485760
  Memory connect retry timeout: 600

```

show bgp group

Syntax	<code>show bgp group</code> <code><brief detail summary></code> <code><group-name></code> <code><exact-instance instance-name></code> <code><instance instance-name></code> <code><logical-system (all logical-system-name)></code> <code><rtf></code>
Syntax (EX Series Switch and QFX Series)	<code>show bgp group</code> <code><brief detail summary></code> <code><group-name></code> <code><exact-instance instance-name></code> <code><instance instance-name></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series. exact-instance option introduced in Junos OS Release 11.4.
Description	Display information about the configured BGP groups.
Options	none —Display group information about all BGP groups. brief detail summary —(Optional) Display the specified level of output. group-name —(Optional) Display group information for the specified group. exact-instance instance-name —(Optional) Display information for the specified instance only. instance instance-name —(Optional) Display information about BGP groups for all routing instances whose name begins with this string (for example, cust1 , cust11 , and cust111 are all displayed when you run the show bgp group instance cust1 command). The instance name can be master for the main instance, or any valid configured instance name or its prefix. logical-system (all logical-system-name) —(Optional) Perform this operation on all logical systems or on a particular logical system. rtf —(Optional) Display BGP group route targeting information.
Required Privilege Level	view
List of Sample Output	show bgp group on page 151 show bgp group brief on page 151 show bgp group detail on page 151 show bgp group rtf detail on page 152 show bgp group summary on page 152

Output Fields Table 37 on page 147 describes the output fields for the **show bgp group** command. Output fields are listed in the approximate order in which they appear.

Table 37: show bgp group Output Fields

Field Name	Field Description	Level of Output
Group Type or Group	Type of BGP group: Internal or External .	All levels
group-index	Index number for the BGP peer group. The index number differentiates between groups when a single BGP group is split because of different configuration options at the group and peer levels.	rtf detail
AS	AS number of the peer. For internal BGP (IBGP), this number is the same as Local AS .	brief detail none
Local AS	AS number of the local routing device.	brief detail none
Name	Name of a specific BGP group.	brief detail none
Index	Unique index number of a BGP group.	brief detail none
Flags	Flags associated with the BGP group. This field is used by Juniper Networks customer support.	brief detail none
Remove-private options	Options associated with the remove-private statement.	brief detail none
Holdtime	Maximum number of seconds allowed to elapse between successive keepalive or update messages that BGP receives from a peer in the BGP group, after which the connection to the peer is closed and routing devices through that peer become unavailable.	brief detail none
Export	Export policies configured for the BGP group with the export statement.	brief detail none
MED tracks IGP metric update delay	Time, in seconds, that updates to multiple exit discriminator (MED) are delayed. Also displays the time remaining before the interval is set to expire	All levels
Traffic Statistics Interval	Time between sample periods for labeled-unicast traffic statistics, in seconds.	brief detail none
Total peers	Total number of peers in the group.	brief detail none
Established	Number of peers in the group that are in the established state.	All levels

Table 37: show bgp group Output Fields (*continued*)

Field Name	Field Description	Level of Output
Active/Received/Accepted/Damped	<p>Multipurpose field that displays information about BGP peer sessions. The field's contents depend upon whether a session is established and whether it was established in the main routing device or in a routing instance.</p> <ul style="list-style-type: none"> If a peer is not established, the field shows the state of the peer session: Active, Connect, or Idle. If a BGP session is established in the main routing device, the field shows the number of active, received, accepted, and damped routes that are received from a neighbor and appear in the inet.0 (main) and inet.2 (multicast) routing tables. For example, 8/10/10/2 and 2/4/4/0 indicate the following: <ul style="list-style-type: none"> 8 active routes, 10 received routes, 10 accepted routes, and 2 damped routes from a BGP peer appear in the inet.0 routing table. 2 active routes, 4 received routes, 4 accepted routes, and no damped routes from a BGP peer appear in the inet.2 routing table. 	summary
ip-addresses	List of peers who are members of the group. The address is followed by the peer's port number.	All levels
Route Queue Timer	Number of seconds until queued routes are sent. If this time has already elapsed, this field displays the number of seconds by which the updates are delayed.	detail
Route Queue	Number of prefixes that are queued up for sending to the peers in the group.	detail
inet.number	<p>Number of active, received, accepted, and damped routes in the routing table. For example, inet.0: 7/10/9/0 indicates the following:</p> <ul style="list-style-type: none"> 7 active routes, 10 received routes, 9 accepted routes, and no damped routes from a BGP peer appear in the inet.0 routing table. 	none

Table 37: show bgp group Output Fields (*continued*)

Field Name	Field Description	Level of Output
Table inet.number	Information about the routing table. <ul style="list-style-type: none"> • Received prefixes—Total number of prefixes from the peer, both active and inactive, that are in the routing table. • Active prefixes—Number of prefixes received from the peer that are active in the routing table. • Suppressed due to damping—Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols. • Advertised prefixes—Number of prefixes advertised to a peer. • Received external prefixes—Total number of prefixes from the external BGP (EBGP) peers, both active and inactive, that are in the routing table. • Active external prefixes—Number of prefixes received from the EBGP peers that are active in the routing table. • Externals suppressed—Number of routes received from EBGP peers currently inactive because of damping or other reasons. • Received internal prefixes—Total number of prefixes from the IBGP peers, both active and inactive, that are in the routing table. • Active internal prefixes—Number of prefixes received from the IBGP peers that are active in the routing table. • Internals suppressed—Number of routes received from IBGP peers currently inactive because of damping or other reasons. • RIB State—Status of the graceful restart process for this routing table: BGP restart is complete, BGP restart in progress, VPN restart in progress, or VPN restart is complete. 	detail
Groups	Total number of groups.	All levels
Peers	Total number of peers.	All levels
External	Total number of external peers.	All levels
Internal	Total number of internal peers.	All levels
Down peers	Total number of unavailable peers.	All levels
Flaps	Total number of flaps that occurred.	All levels
Table	Name of a routing table.	brief , none
Tot Paths	Total number of routes.	brief , none
Act Paths	Number of active routes.	brief , none
Suppressed	Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols.	brief , none

Table 37: show bgp group Output Fields (*continued*)

Field Name	Field Description	Level of Output
History	Number of withdrawn routes stored locally to keep track of damping history.	brief, none
Damp State	Number of active routes with a figure of merit greater than zero, but lower than the threshold at which suppression occurs.	brief, none
Pending	Routes being processed by the BGP import policy.	brief, none
Group	Group the peer belongs to in the BGP configuration.	detail
Receive mask	Mask of the received target included in the advertised route.	detail
Entries	Number of route entries received.	detail
Target	Route target that is to be passed by route-target filtering. If a route advertised from the provider edge (PE) routing device matches an entry in the route-target filter, the route is passed to the peer.	detail
Mask	Mask which specifies that the peer receive routes with the given route target.	detail

Sample Output

show bgp group

```
user@host> show bgp group
Groups: 2 Peers: 2 External: 0 Internal: 2 Down peers: 1 Flaps: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending

inet.0
0 0 0 0 0 0

bgp.l3vpn.0
0 0 0 0 0 0

bgp.rtarget.0
2 0 0 0 0 0
```

show bgp group brief

```
user@host> show bgp group brief
Groups: 2 Peers: 2 External: 0 Internal: 2 Down peers: 1 Flaps: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending

inet.0
0 0 0 0 0 0

bgp.l3vpn.0
0 0 0 0 0 0

bgp.rtarget.0
2 0 0 0 0 0
```

show bgp group detail

```
user@host> show bgp group detail
Group Type: Internal AS: 1 Local AS: 1
Name: ibgp Index: 0 Flags: <Export Eval>
Holdtime: 0
Total peers: 3 Established: 0
22.0.0.2
22.0.0.8
22.0.0.5

Groups: 1 Peers: 3 External: 0 Internal: 3 Down peers: 3 Flaps: 3
Table bgp.l3vpn.0
Received prefixes: 0
Accepted prefixes: 0
Active prefixes: 0
Suppressed due to damping: 0
Received external prefixes: 0
Active external prefixes: 0
Externals suppressed: 0
Received internal prefixes: 0
Active internal prefixes: 0
Internals suppressed: 0
RIB State: BGP restart is complete
RIB State: VPN restart is complete
Table bgp.mdt.0
Received prefixes: 0
Accepted prefixes: 0
Active prefixes: 0
Suppressed due to damping: 0
Received external prefixes: 0
Active external prefixes: 0
```

```

Externals suppressed:      0
Received internal prefixes: 0
Active internal prefixes:  0
Internals suppressed:      0
RIB State: BGP restart is complete
RIB State: VPN restart is complete
Table VPN-A.inet.0
Received prefixes:         0
Accepted prefixes:         0
Active prefixes:           0
Suppressed due to damping: 0
Received external prefixes: 0
Active external prefixes:  0
Externals suppressed:      0
Received internal prefixes: 0
Active internal prefixes:  0
Internals suppressed:      0
RIB State: BGP restart is complete
RIB State: VPN restart is complete
Table VPN-A.mdt.0
Received prefixes:         0
Accepted prefixes:         0
Active prefixes:           0
Suppressed due to damping: 0
Received external prefixes: 0
Active external prefixes:  0
Externals suppressed:      0
Received internal prefixes: 0
Active internal prefixes:  0
Internals suppressed:      0
RIB State: BGP restart is complete
RIB State: VPN restart is complete

```

show bgp group rtf detail

```

user@host> show bgp group rtf detail
Group: internal (group-index: 0)
  Receive mask: 00000002
  Table: bgp.rtarget.0
    Target      Mask
    100:100/64  00000002
    200:201/64  (Group)
Group: internal (group-index: 1)
  Table: bgp.rtarget.0
    Target      Mask
    200:201/64  (Group)

```

show bgp group summary

```

user@host> show bgp group summary
Group      Type      Peers  Established  Active/Received/Accepted/Damped
ibgp       Internal  3      0
Groups: 1  Peers: 3   External: 0   Internal: 3   Down peers: 3   Flaps: 3
bgp.l3vpn.0 : 0/0/0/0 External: 0/0/0/0 Internal: 0/0/0/0
bgp.mdt.0   : 0/0/0/0 External: 0/0/0/0 Internal: 0/0/0/0
VPN-A.inet.0 : 0/0/0/0 External: 0/0/0/0 Internal: 0/0/0/0
VPN-A.mdt.0 : 0/0/0/0 External: 0/0/0/0 Internal: 0/0/0/0

```

show bgp neighbor

Syntax	<pre>show bgp neighbor <exact-instance <i>instance-name</i>> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)> <neighbor-address> <orf (detail <i>neighbor-address</i>)</pre>
Syntax (EX Series Switch and QFX Series)	<pre>show bgp neighbor <instance <i>instance-name</i>> <exact-instance <i>instance-name</i>> <neighbor-address> <orf (<i>neighbor-address</i> detail)</pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>orf option introduced in Junos OS Release 9.2.</p> <p>exact-instance option introduced in Junos OS Release 11.4.</p>
Description	Display information about BGP peers.
Options	<p>none—Display information about all BGP peers.</p> <p>exact-instance <i>instance-name</i>—(Optional) Display information for the specified instance only.</p> <p>instance <i>instance-name</i>—(Optional) Display information about BGP peers for all routing instances whose name begins with this string (for example, cust1, cust11, and cust111 are all displayed when you run the show bgp neighbor instance cust1 command).</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>neighbor-address—(Optional) Display information for only the BGP peer at the specified IP address.</p> <p>orf (detail <i>neighbor-address</i>)—(Optional) Display outbound route-filtering information for all BGP peers or only for the BGP peer at the specified IP address. The default is to display brief output. Use the detail option to display detailed output.</p>
Additional Information	For information about the local-address , nlri , hold-time , and preference statements, see the Junos OS Routing Protocols Configuration Guide.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear bgp neighbor on page 141

List of Sample Output

- [show bgp neighbor on page 161](#)
- [show bgp neighbor \(CLNS\) on page 161](#)
- [show bgp neighbor \(Layer 2 VPN\) on page 162](#)
- [show bgp neighbor \(Layer 3 VPN\) on page 164](#)
- [show bgp neighbor neighbor-address on page 165](#)
- [show bgp neighbor neighbor-address on page 165](#)
- [show bgp neighbor orf neighbor-address detail on page 166](#)

Output Fields Table 38 on page 154 describes the output fields for the **show bgp neighbor** command. Output fields are listed in the approximate order in which they appear.

Table 38: show bgp neighbor Output Fields

Field Name	Field Description
Peer	Address of the BGP neighbor. The address is followed by the neighbor port number.
AS	AS number of the peer.
Local	Address of the local routing device. The address is followed by the peer port number.
Type	Type of peer: Internal or External .
State	<p>Current state of the BGP session:</p> <ul style="list-style-type: none"> • Active—BGP is initiating a transport protocol connection in an attempt to connect to a peer. If the connection is successful, BGP sends an Open message. • Connect—BGP is waiting for the transport protocol connection to be completed. • Established—The BGP session has been established, and the peers are exchanging update messages. • Idle—This is the first stage of a connection. BGP is waiting for a Start event. • OpenConfirm—BGP has acknowledged receipt of an open message from the peer and is waiting to receive a keepalive or notification message. • OpenSent—BGP has sent an open message and is waiting to receive an open message from the peer.
Flags	<p>Internal BGP flags:</p> <ul style="list-style-type: none"> • Aggregate Label—BGP has aggregated a set of incoming labels (labels received from the peer) into a single forwarding label. • CleanUp—The peer session is being shut down. • Delete—This peer has been deleted. • Idled—This peer has been permanently idled. • ImportEval—At the last commit operation, this peer was identified as needing to reevaluate all received routes. • Initializing—The peer session is initializing. • SendRtn—Messages are being sent to the peer. • Sync—This peer is synchronized with the rest of the peer group. • TryConnect—Another attempt is being made to connect to the peer. • Unconfigured—This peer is not configured. • WriteFailed—An attempt to write to this peer failed.

Table 38: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
Last state	<p>Previous state of the BGP session:</p> <ul style="list-style-type: none"> • Active—BGP is initiating a transport protocol connection in an attempt to connect to a peer. If the connection is successful, BGP sends an Open message. • Connect—BGP is waiting for the transport protocol connection to be completed. • Established—The BGP session has been established, and the peers are exchanging update messages. • Idle—This is the first stage of a connection. BGP is waiting for a Start event. • OpenConfirm—BGP has acknowledged receipt of an open message from the peer and is waiting to receive a keepalive or notification message. • OpenSent—BGP has sent an open message and is waiting to receive an open message from the peer.
Last event	<p>Last activity that occurred in the BGP session:</p> <ul style="list-style-type: none"> • Closed—The BGP session closed. • ConnectRetry—The transport protocol connection failed, and BGP is trying again to connect. • HoldTime—The session ended because the hold timer expired. • KeepAlive—The local routing device sent a BGP keepalive message to the peer. • Open—The local routing device sent a BGP open message to the peer. • OpenFail—The local routing device did not receive an acknowledgment of a BGP open message from the peer. • RecvKeepAlive—The local routing device received a BGP keepalive message from the peer. • RecvNotify—The local routing device received a BGP notification message from the peer. • RecvOpen—The local routing device received a BGP open message from the peer. • RecvUpdate—The local routing device received a BGP update message from the peer. • Start—The peering session started. • Stop—The peering session stopped. • TransportError—A TCP error occurred.
Last error	<p>Last error that occurred in the BGP session:</p> <ul style="list-style-type: none"> • Cease—An error occurred, such as a version mismatch, that caused the session to close. • Finite State Machine Error—In setting up the session, BGP received a message that it did not understand. • Hold Time Expired—The session's hold time expired. • Message Header Error—The header of a BGP message was malformed. • Open Message Error—A BGP open message contained an error. • None—No errors occurred in the BGP session. • Update Message Error—A BGP update message contained an error.
Export	Name of the export policy that is configured on the peer.
Import	Name of the import policy that is configured on the peer.

Table 38: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
Options	Configured BGP options: <ul style="list-style-type: none"> • AddressFamily—Configured address family: inet or inet-vpn. • AuthKeyChain—Authentication key change is enabled. • DropPathAttributes—Certain path attributes are configured to be dropped from neighbor updates during inbound processing. • GracefulRestart—Graceful restart is configured. • HoldTime—Hold time configured with the hold-time statement. The hold time is three times the interval at which keepalive messages are sent. • IgnorePathAttributes—Certain path attributes are configured to be ignored in neighbor updates during inbound processing. • Local Address—Address configured with the local-address statement. • Multihop—Allow BGP connections to external peers that are not on a directly connected network. • NLRI—Configured MBGP state for the BGP group: multicast, unicast, or both if you have configured nlri any. • Peer AS—Configured peer autonomous system (AS). • Preference—Preference value configured with the preference statement. • Refresh—Configured to refresh automatically when the policy changes. • Rib-group—Configured routing table group.
Path-attributes dropped	Path attribute codes that are dropped from neighbor updates.
Path-attributes ignored	Path attribute codes that are ignored during neighbor updates.
Authentication key change	(appears only if the authentication-keychain statement has been configured) Name of the authentication keychain enabled.
Authentication algorithm	(appears only if the authentication-algorithm statement has been configured) Type of authentication algorithm enabled: hmac or md5 .
Address families configured	Names of configured address families for the VPN.
Local Address	Address of the local routing device.
Remove-private options	Options associated with the remove-private statement.
Holdtime	Hold time configured with the hold-time statement. The hold time is three times the interval at which keepalive messages are sent.
Flags for NLRI inet-label-unicast	Flags related to labeled-unicast: <ul style="list-style-type: none"> • TrafficStatistics—Collection of statistics for labeled-unicast traffic is enabled.

Table 38: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
Traffic statistics	Information about labeled-unicast traffic statistics: <ul style="list-style-type: none"> • Options—Options configured for collecting statistics about labeled-unicast traffic. • File—Name and location of statistics log files. • size—Size of all the log files, in bytes. • files—Number of log files.
Traffic Statistics Interval	Time between sample periods for labeled-unicast traffic statistics, in seconds.
Preference	Preference value configured with the preference statement.
Outbound Timer	Time for which the route is available in Junos OS routing table before it is exported to BGP. This field is displayed in the output only if the out-delay parameter is configured to a non-zero value.
Number of flaps	Number of times the BGP session has gone down and then come back up.
Peer ID	Router identifier of the peer.
Group index	Index number for the BGP peer group. The index number differentiates between groups when a single BGP group is split because of different configuration options at the group and peer levels.
Peer index	Index that is unique within the BGP group to which the peer belongs.
Local ID	Router identifier of the local routing device.
Local Interface	Name of the interface on the local routing device.
Active holdtime	Hold time that the local routing device negotiated with the peer.
Keepalive Interval	Keepalive interval, in seconds.
BFD	Status of BFD failure detection.
Local Address	Name of directly connected interface over which direct EBGP peering is established.
NLRI for restart configured on peer	Names of address families configured for restart.
NLRI advertised by peer	Address families supported by the peer: unicast or multicast .
NLRI for this session	Address families being used for this session.
Peer supports Refresh capability	Remote peer's ability to send and request full route table readvertisement (route refresh capability). For more information, see RFC 2918, <i>Route Refresh Capability for BGP-4</i> .
Restart time configured on peer	Configured time allowed for restart on the neighbor.

Table 38: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
Stale routes from peer are kept for	When graceful restart is negotiated, the maximum time allowed to hold routes from neighbors after the BGP session has gone down.
Peer does not support Restarter functionality	Graceful restart restarter-mode is disabled on the peer.
Peer does not support Receiver functionality	Graceful restart helper-mode is disabled on the peer.
Restart time requested by this peer	Restart time requested by this neighbor during capability negotiation.
Restart flag received from the peer	When this field appears, the BGP speaker has restarted (Restarting), and this peer should not wait for the end-of-rib marker from the speaker before advertising routing information to the speaker.
NLRI that peer supports restart for	Neighbor supports graceful restart for this address family.
NLRI peer can save forwarding state	Neighbor supporting this address family saves all forwarding states.
NLRI that peer saved forwarding for	Neighbor saves all forwarding states for this address family.
NLRI that restart is negotiated for	Router supports graceful restart for this address family.
NLRI of received end-of-rib markers	Address families for which end-of-routing-table markers are received from the neighbor.
NLRI of all end-of-rib markers sent	Address families for which end-of-routing-table markers are sent to the neighbor.
Peer supports 4 byte AS extension (peer-as 1)	Peer understands 4-byte AS numbers in BGP messages. The peer is running Junos OS Release 9.1 or later.
NLRIs for which peer can receive multiple paths	Appears in the command output of the local router if the downstream peer is configured to receive multiple BGP routes to a single destination, instead of only receiving the active route. Possible value is inet-unicast .
NLRIs for which peer can send multiple paths: inet-unicast	Appears in the command output of the local router if the upstream peer is configured to send multiple BGP routes to a single destination, instead of only sending the active route. Possible value is inet-unicast .

Table 38: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
Table inet.number	<p>Information about the routing table:</p> <ul style="list-style-type: none"> • RIB State—BGP is in the graceful restart process for this routing table: restart is complete or restart in progress. • Bit—Number that represents the entry in the routing table for this peer. • Send state—State of the BGP group: in sync, not in sync, or not advertising. • Active prefixes—Number of prefixes received from the peer that are active in the routing table. • Received prefixes—Total number of prefixes from the peer, both active and inactive, that are in the routing table. • Accepted prefixes—Total number of prefixes from the peer that have been accepted by a routing policy. • Suppressed due to damping—Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols.
Last traffic (seconds)	Last time any traffic was received from the peer or sent to the peer, and the last time the local routing device checked.
Input messages	Messages that BGP has received from the receive socket buffer, showing the total number of messages, number of update messages, number of times a policy is changed and refreshed, and the buffer size in octets. The buffer size is 16 KB.
Output messages	Messages that BGP has written to the transmit socket buffer, showing the total number of messages, number of update messages, number of times a policy is changed and refreshed, and the buffer size in octets. The buffer size is 16 KB.
Input dropped path attributes	<p>Information about dropped path attributes:</p> <ul style="list-style-type: none"> • Code—Path attribute code. • Count—Path attribute count.
Input ignored path attributes	<p>Information about ignored path attributes:</p> <ul style="list-style-type: none"> • Code—Path attribute code. • Count—Path attribute count.
Output queue	Number of BGP packets that are queued to be transmitted to a particular neighbor for a particular routing table. Output queue 0 is for unicast NLRIs, and queue 1 is for multicast NLRIs.
Trace options	Configured tracing of BGP protocol packets and operations.
Trace file	Name of the file to receive the output of the tracing operation.
Filter Updates rcv	<p>(orf option only) Number of outbound-route filters received for each configured address family.</p> <p>NOTE: The counter is cumulative. For example, the counter is increased after the remote peer either resends or clears the outbound route filtering prefix list.</p>

Table 38: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
Immediate	<p>(orf option only) Number of route updates received with the immediate flag set. The immediate flag indicates that the BGP peer should readvertise the updated routes.</p> <p>NOTE: The counter is cumulative. For example, the counter is increased after the remote peer either resends or clears the outbound route filtering prefix list.</p>
Filter	(orf option only) Type of prefix filter received: prefix-based or extended-community .
Received filter entries	(orf option only) List of received filters displayed.
seq	(orf option only) Numerical order assigned to this prefix entry among all the received outbound route filter prefix entries.
prefix	(orf option only) Address for the prefix entry that matches the filter.
minlength	(orf option only) Minimum prefix length, in bits, required to match this prefix.
maxlength	(orf option only) Maximum prefix length, in bits, required to match this prefix.
match	(orf option only) For this prefix match, whether to permit or deny route updates.

Sample Output

show bgp neighbor

```

user@host > show bgp neighbor
Peer: 10.255.7.250+179 AS 10 Local: 10.255.7.248+63740 AS 10
  Type: Internal State: Established Flags: <Sync>
  Last State: OpenConfirm Last Event: RecvKeepAlive
  Last Error: None
  Export: [ redist_static ]
  Options: <Preference LocalAddress PeerAS Refresh>
  Local Address: 10.255.7.248 Holdtime: 90 Preference: 170 Outbound Timer: 50
  Number of flaps: 0
  Peer ID: 10.255.7.250 Local ID: 10.255.7.248 Active Holdtime: 90
  Keepalive Interval: 30 Group index: 0 Peer index: 0
  BFD: disabled, down
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast
  NLRI for this session: inet-unicast
  Peer supports Refresh capability (2)
  Stale routes from peer are kept for: 300
  Peer does not support Restarter functionality
  NLRI that restart is negotiated for: inet-unicast
  NLRI of received end-of-rib markers: inet-unicast
  NLRI of all end-of-rib markers sent: inet-unicast
  Peer supports 4 byte AS extension (peer-as 10)
  Peer does not support Addpath
  Table inet.0 Bit: 10000
    RIB State: BGP restart is complete
    Send state: in sync
    Active prefixes: 1
    Received prefixes: 1
    Accepted prefixes: 1
    Suppressed due to damping: 0
    Advertised prefixes: 1
  Last traffic (seconds): Received 9 Sent 5 Checked 5
  Input messages: Total 36 Updates 2 Refreshes 0 Octets 718
  Output messages: Total 37 Updates 1 Refreshes 0 Octets 796
  Output Queue[0]: 0

Peer: 10.255.162.214+52193 AS 100 Local: 10.255.167.205+179 AS 100
  Type: Internal State: Established (route reflector client)Flags: <Sync>
  Last State: OpenConfirm Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress Cluster AddressFamily Rib-group Refresh>
  Address families configured: inet-unicast inet-vpn-unicast route-target
  Local Address: 10.255.167.205 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 10.255.162.214 Local ID: 10.255.167.205 Active Holdtime: 90
  Keepalive Interval: 30 Group index: 0 Peer index: 1

```

show bgp neighbor (CLNS)

```

user@host> show bgp neighbor
Peer: 10.245.245.1+179 AS 200 Local: 10.245.245.3+3770 AS 100
  Type: External State: Established Flags: <ImportEval Sync>
  Last State: OpenConfirm Last Event: RecvKeepAlive
  Last Error: None
  Options: <Multihop Preference LocalAddress HoldTime AddressFamily PeerAS
  Rib-group Refresh>
  Address families configured: iso-vpn-unicast
  Local Address: 10.245.245.3 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 10.245.245.1 Local ID: 10.245.245.3 Active Holdtime: 90

```

```

Keepalive Interval: 30          Peer index: 0
NLRI advertised by peer: iso-vpn-unicast
NLRI for this session: iso-vpn-unicast
Peer supports Refresh capability (2)
Table bgp.isovpn.0 Bit: 10000
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: in sync
  Active prefixes:              3
  Received prefixes:            3
  Suppressed due to damping:    0
  Advertised prefixes:          3
Table aaaa.iso.0
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: not advertising
  Active prefixes:              3
  Received prefixes:            3
  Suppressed due to damping:    0
Last traffic (seconds): Received 6    Sent 5    Checked 5
Input messages: Total 1736    Updates 4    Refreshes 0    Octets 33385
Output messages: Total 1738    Updates 3    Refreshes 0    Octets 33305
Output Queue[0]: 0
Output Queue[1]: 0

```

show bgp neighbor (Layer 2 VPN)

```

user@host> show bgp neighbor
Peer: 10.69.103.2      AS 65100 Local: 10.69.103.1      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: 10.69.103.1 Holdtime: 90 Preference: 170
Number of flaps: 0
Peer: 10.69.104.2      AS 65100 Local: 10.69.104.1      AS 65104
Type: External        State: Active          Flags: <ImportEval>
Last State: Idle      Last Event: Start
Last Error: None
Export: [ BGP-L-import ]
Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily PeerAS
Refresh>
Address families configured: inet-labeled-unicast
Local Address: 10.69.104.1 Holdtime: 90 Preference: 170
Number of flaps: 0
Peer: 10.255.14.182+179 AS 69    Local: 10.255.14.176+2131 AS 69
Type: Internal        State: Established    Flags: <ImportEval>
Last State: OpenConfirm Last Event: RecvKeepAlive
Last Error: None
Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily
Rib-group Refresh>
Address families configured: inet-vpn-unicast l2vpn
Local Address: 10.255.14.176 Holdtime: 90 Preference: 170
Number of flaps: 0
Peer ID: 10.255.14.182    Local ID: 10.255.14.176    Active Holdtime: 90
Keepalive Interval: 30
NLRI for restart configured on peer: inet-vpn-unicast l2vpn
NLRI advertised by peer: inet-vpn-unicast l2vpn
NLRI for this session: inet-vpn-unicast l2vpn
Peer supports Refresh capability (2)

```

```

Restart time configured on the peer: 120
Stale routes from peer are kept for: 300
Restart time requested by this peer: 120
NLRI that peer supports restart for: inet-vpn-unicast 12vpn
NLRI peer can save forwarding state: inet-vpn-unicast 12vpn
NLRI that peer saved forwarding for: inet-vpn-unicast 12vpn
NLRI that restart is negotiated for: inet-vpn-unicast 12vpn
NLRI of received end-of-rib markers: inet-vpn-unicast 12vpn
Table bgp.13vpn.0 Bit: 10000
  RIB State: BGP restart in progress
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          10
  Received prefixes:        10
  Suppressed due to damping: 0
Table bgp.12vpn.0 Bit: 20000
  RIB State: BGP restart in progress
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          1
  Received prefixes:        1
  Suppressed due to damping: 0
Table BGP-INET.inet.0 Bit: 30000
  RIB State: BGP restart in progress
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          2
  Received prefixes:        2
  Suppressed due to damping: 0
Table BGP-L.inet.0 Bit: 40000
  RIB State: BGP restart in progress
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          2
  Received prefixes:        2
  Suppressed due to damping: 0
Table LDP.inet.0 Bit: 50000
  RIB State: BGP restart is complete
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          1
  Received prefixes:        1
  Suppressed due to damping: 0
Table OSPF.inet.0 Bit: 60000
  RIB State: BGP restart is complete
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          2
  Received prefixes:        2
  Suppressed due to damping: 0
Table RIP.inet.0 Bit: 70000
  RIB State: BGP restart is complete
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          2
  Received prefixes:        2
  Suppressed due to damping: 0
Table STATIC.inet.0 Bit: 80000
  RIB State: BGP restart is complete
  RIB State: VPN restart in progress
  Send state: in sync

```

```

Active prefixes:          1
Received prefixes:       1
Suppressed due to damping: 0
Table L2VPN.l2vpn.0 Bit: 90000
RIB State: BGP restart is complete
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:          1
Received prefixes:       1
Suppressed due to damping: 0
Last traffic (seconds): Received 0    Sent 0    Checked 0
Input messages: Total 14    Updates 13    Refreshes 0    Octets 1053
Output messages: Total 3    Updates 0    Refreshes 0    Octets 105
Output Queue[0]: 0
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

```

show bgp neighbor (Layer 3 VPN)

```

user@host> show bgp neighbor
Peer: 4.4.4.4+179    AS 10045 Local: 5.5.5.5+1214    AS 10045
Type: Internal    State: Established    Flags: <ImportEval>
Last State: OpenConfirm    Last Event: RecvKeepAlive
Last Error: None
Export: [ match-all ] Import: [ match-all ]
Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily
Rib-group Refresh>
Address families configured: inet-vpn-unicast
Local Address: 5.5.5.5 Holdtime: 90 Preference: 170
Flags for NLRI inet-labeled-unicast: TrafficStatistics
Traffic Statistics: Options: all File: /var/log/bstat.log
                        size 131072 files 10

Traffic Statistics Interval: 60
Number of flaps: 0
Peer ID: 192.168.1.110    Local ID: 192.168.1.111    Active Holdtime: 90
Keepalive Interval: 30
NLRI for restart configured on peer: inet-vpn-unicast
NLRI advertised by peer: inet-vpn-unicast
NLRI for this session: inet-vpn-unicast
Peer supports Refresh capability (2)
Restart time configured on the peer: 120
Stale routes from peer are kept for: 300
Restart time requested by this peer: 120
NLRI that peer supports restart for: inet-vpn-unicast
NLRI peer can save forwarding state: inet-vpn-unicast
NLRI that peer saved forwarding for: inet-vpn-unicast
NLRI that restart is negotiated for: inet-vpn-unicast
NLRI of received end-of-rib markers: inet-vpn-unicast
NLRI of all end-of-rib markers sent: inet-vpn-unicast
Table bgp.l3vpn.0 Bit: 10000
RIB State: BGP restart is complete
RIB State: VPN restart is complete
Send state: in sync
Active prefixes:          2
Received prefixes:       2
Suppressed due to damping: 0
Table vpn-green.inet.0 Bit: 20001

```



```

RIB State: BGP restart is complete
RIB State: VPN restart is complete
Send state: in sync
Active prefixes:          2
Received prefixes:        2
Suppressed due to damping: 0
Last traffic (seconds): Received 15   Sent 20   Checked 20
Input messages:  Total 40   Updates 2   Refreshes 0   Octets 856
Output messages: Total 44   Updates 2   Refreshes 0   Octets 1066
Output Queue[0]: 0
Output Queue[1]: 0
Trace options: detail packets
Trace file: /var/log/bgpr.log size 131072 files 10

```

show bgp neighbor neighbor-address

```

user@host> show bgp neighbor 192.168.1.111
Peer: 10.255.245.12+179 AS 35 Local: 10.255.245.13+2884 AS 35
Type: Internal State: Established (route reflector client)Flags: <Sync>
Last State: OpenConfirm Last Event: RecvKeepAlive
Last Error: None
Options: <Preference LocalAddress HoldTime Cluster AddressFamily Rib-group
Refresh>
Address families configured: inet-vpn-unicast inet-labeled-unicast
Local Address: 10.255.245.13 Holdtime: 90 Preference: 170
Flags for NLRI inet-vpn-unicast: AggregateLabel
Flags for NLRI inet-labeled-unicast: AggregateLabel
Number of flaps: 0
Peer ID: 10.255.245.12 Local ID: 10.255.245.13 Active Holdtime: 90
Keepalive Interval: 30
BFD: disabled
NLRI advertised by peer: inet-vpn-unicast inet-labeled-unicast
NLRI for this session: inet-vpn-unicast inet-labeled-unicast
Peer supports Refresh capability (2)
Restart time configured on the peer: 300
Stale routes from peer are kept for: 60
Restart time requested by this peer: 300
NLRI that peer supports restart for: inet-unicast inet6-unicast
NLRI that restart is negotiated for: inet-unicast inet6-unicast
NLRI of received end-of-rib markers: inet-unicast inet6-unicast
NLRI of all end-of-rib markers sent: inet-unicast inet6-unicast
Table inet.0 Bit: 10000
RIB State: restart is complete
Send state: in sync
Active prefixes: 4
Received prefixes: 6
Suppressed due to damping: 0
Table inet6.0 Bit: 20000
RIB State: restart is complete
Send state: in sync
Active prefixes: 0
Received prefixes: 2
Suppressed due to damping: 0
Last traffic (seconds): Received 3 Sent 3 Checked 3
Input messages:  Total 9 Updates 6 Refreshes 0 Octets 403
Output messages: Total 7 Updates 3 Refreshes 0 Octets 365
Output Queue[0]: 0
Output Queue[1]: 0
Trace options: detail packets
Trace file: /var/log/bgpr size 131072 files 10

```

```

user@host> show bgp neighbor 192.168.4.222

```

show bgp neighbor neighbor-address

```

Peer: 192.168.4.222+4902 AS 65501 Local: 192.168.4.221+179 AS 65500
Type: External      State: Established      Flags: <Sync>
Last State: OpenConfirm  Last Event: RecvKeepAlive
Last Error: Cease
Export: [ export-policy ] Import: [ import-policy ]
Options: <Preference HoldTime AddressFamily PeerAS PrefixLimit Refresh>
Address families configured: inet-unicast inet-multicast
Holdtime: 60000 Preference: 170
Number of flaps: 4
Last flap event: RecvUpdate
Error: 'Cease' Sent: 5 Recv: 0
Peer ID: 10.255.245.6      Local ID: 10.255.245.5      Active Holdtime: 60000
Keepalive Interval: 20000      Peer index: 0
BFD: disabled, down
Local Interface: fxp0.0
NLRI advertised by peer: inet-unicast inet-multicast
NLRI for this session: inet-unicast inet-multicast
Peer supports Refresh capability (2)
Table inet.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          8
  Received prefixes:        10
  Accepted prefixes:        10
  Suppressed due to damping: 0
  Advertised prefixes:      3
Table inet.2 Bit: 20000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          0
  Received prefixes:        0
  Accepted prefixes:        0
  Suppressed due to damping: 0
  Advertised prefixes:      0
Last traffic (seconds): Received 357 Sent 357 Checked 357
Input messages: Total 4 Updates 2 Refreshes 0 Octets 211
Output messages: Total 4 Updates 1 Refreshes 0 Octets 147
Output Queue[0]: 0
Output Queue[1]: 0
Trace options: all
Trace file: /var/log/bgp size 10485760 files 10

```

show bgp neighbor orf neighbor-address detail

```

user@host > show bgp neighbor orf 192.168.165.56 detail
Peer: 192.168.165.56+179 Type: External
Group: ext1

inet-unicast
  Filter updates rcv:          1 Immediate:          1
  Filter: prefix-based receive
  Received filter entries:
    seq 1: prefix 2.2.2.2/32: minlen 32: maxlen 32: match deny:

inet6-unicast
  Filter updates rcv:          0 Immediate:          1
  Filter: prefix-based receive
  Received filter entries:
    *.*

```

show bgp summary

Syntax	<pre>show bgp summary <exact-instance <i>instance-name</i>> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)></pre>
Syntax (EX Series Switch and QFX Series)	<pre>show bgp summary <exact-instance <i>instance-name</i>> <instance <i>instance-name</i>></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>exact-instance option introduced in Junos OS Release 11.4.</p>
Description	Display BGP summary information.
Options	<p>none—Display BGP summary information for all routing instances.</p> <p>exact-instance <i>instance-name</i>—(Optional) Display information for the specified instance only.</p> <p>instance <i>instance-name</i>—(Optional) Display information for all routing instances whose name begins with this string (for example, cust1, cust11, and cust111 are all displayed when you run the show bgp summary instance cust1 command). The instance name can be master for the main instance, or any valid configured instance name or its prefix.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	<p>show bgp summary (When a Peer Is Not Established) on page 170</p> <p>show bgp summary (When a Peer Is Established) on page 170</p> <p>show bgp summary (CLNS) on page 170</p> <p>show bgp summary (Layer 2 VPN) on page 170</p> <p>show bgp summary (Layer 3 VPN) on page 171</p>
Output Fields	<p>Table 39 on page 167 describes the output fields for the show bgp summary command. Output fields are listed in the approximate order in which they appear.</p>

Table 39: show bgp summary Output Fields

Field Name	Field Description
Groups	Number of BGP groups.
Peers	Number of BGP peers.

Table 39: show bgp summary Output Fields (*continued*)

Field Name	Field Description
Down peers	Number of down BGP peers.
Table	Name of routing table.
Tot Paths	Total number of paths.
Act Paths	Number of active routes.
Suppressed	Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols.
History	Number of withdrawn routes stored locally to keep track of damping history.
Damp State	Number of routes with a figure of merit greater than zero, but still active because the value has not reached the threshold at which suppression occurs.
Pending	Routes in process by BGP import policy.
Peer	Address of each BGP peer. Each peer has one line of output.
AS	Peer's AS number.
InPkt	Number of packets received from the peer.
OutPkt	Number of packets sent to the peer.
OutQ	Number of BGP packets that are queued to be transmitted to a particular neighbor. It normally is 0 because the queue usually is emptied quickly.
Flaps	Number of times the BGP session has gone down and then come back up.
Last Up/Down	Last time since the neighbor transitioned to or from the established state.

Table 39: show bgp summary Output Fields (*continued*)

Field Name	Field Description
State #Active /Received/Accepted /Damped	<p>Multipurpose field that displays information about BGP peer sessions. The field's contents depend upon whether a session is established and whether it was established on the main routing device or in a routing instance.</p> <ul style="list-style-type: none"> If a peer is not established, the field shows the state of the peer session: Active, Connect, or Idle. In general, the Idle state is the first stage of a connection. BGP is waiting for a Start event. A session can be idle for other reasons as well. The reason that a session is idle is sometimes displayed. For example: Idle (Removal in progress) or Idle (LicenseFailure). If a BGP session is established on the main routing device, the field shows the number of active, received, accepted, and damped routes that are received from a neighbor and appear in the inet.0 (main) and inet.2 (multicast) routing tables. For example, 8/10/10/2 and 2/4/4/0 indicate the following: <ul style="list-style-type: none"> 8 active routes, 10 received routes, 10 accepted routes, and 2 damped routes from a BGP peer appear in the inet.0 routing table. 2 active routes, 4 received routes, 4 accepted routes, and no damped routes from a BGP peer appear in the inet.2 routing table. If a BGP session is established in a routing instance, the field indicates the established (Establ) state, identifies the specific routing table that receives BGP updates, and shows the number of active, received, and damped routes that are received from a neighbor. For example, Establ VPN-AB.inet.0: 2/4/0 indicates the following: <ul style="list-style-type: none"> The BGP session is established. Routes are received in the VPN-AB.inet.0 routing table. The local routing device has two active routes, four received routes, and no damped routes from a BGP peer. <p>When a BGP session is established, the peers are exchanging update messages.</p>

Sample Output

show bgp summary (When a Peer Is Not Established)

```
user@host> show bgp summary
Groups: 2 Peers: 4 Down peers: 1
Table          Tot Paths  Act Paths Suppressed  History  Damp State   Pending
inet.0          6          4          0          0          0          0
Peer           AS      InPkt    OutPkt    OutQ    Flaps Last Up/Dwn
State|#Active/Received/Damped...
10.0.0.3        65002      86       90       0        2      42:54 0/0/0

0/0/0
10.0.0.4        65002      90       91       0        1      42:54 0/2/0

0/0/0
10.0.0.6        65002      87       90       0        3          3 Active
10.1.12.1       65001      89       89       0        1      42:54 4/4/0

0/0/0
```

show bgp summary (When a Peer Is Established)

```
user@host> show bgp summary
Groups: 1 Peers: 3 Down peers: 0
Table          Tot Paths  Act Paths Suppressed  History  Damp State   Pending
inet.0          6          4          0          0          0          0
Peer           AS      InPkt    OutPkt    OutQ    Flaps Last Up/Dwn
State|#Active/Received/Damped...
10.0.0.2        65002    88675    88652     0        2      42:38 2/4/0

0/0/0
10.0.0.3        65002    54528    54532     0        1     2w4d22h 0/0/0

0/0/0
10.0.0.4        65002    51597    51584     0        0     2w3d22h 2/2/0

0/0/0
```

show bgp summary (CLNS)

```
user@host> show bgp summary
Groups: 1 Peers: 1 Down peers: 0
Peer           AS      InPkt    OutPkt    OutQ    Flaps Last Up/Dwn
State|#Active/Received/Damped...
10.245.245.1    200     1735     1737     0        0     14:26:12 Establ
  bgp.isovpn.0: 3/3/0
  aaaa.iso.0: 3/3/0
```

show bgp summary (Layer 2 VPN)

```
user@host> show bgp summary
Groups: 1 Peers: 5 Down peers: 0
Table          Tot Paths  Act Paths Suppressed  History  Damp State   Pending
bgp.l2vpn.0      1          1          0          0          0          0
inet.0           0          0          0          0          0          0
Peer           AS      InPkt    OutPkt    OutQ    Flaps Last Up/Dwn
State|#Active/Received/Damped...
10.255.245.35   65299      72       74       0        1     19:00 Establ
  bgp.l2vpn.0: 1/1/0
  frame-vpn.l2vpn.0: 1/1/0
10.255.245.36   65299    2164    2423     0        4     19:50 Establ
  bgp.l2vpn.0: 0/0/0
  frame-vpn.l2vpn.0: 0/0/0
10.255.245.37   65299      36       37       0        4     17:07 Establ
```

```

inet.0: 0/0/0
10.255.245.39 65299 138 168 0 6 53:48 Estab1
bgp.12vpn.0: 0/0/0
frame-vpn.12vpn.0: 0/0/0
10.255.245.69 65299 134 140 0 6 53:42 Estab1
inet.0: 0/0/0

```

show bgp summary (Layer 3 VPN)

```

user@host> show bgp summary
Groups: 2 Peers: 2 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
bgp.13vpn.0 2 2 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Damped...
10.39.1.5 2 21 22 0 0 6:26 Estab1
VPN-AB.inet.0: 1/1/0
10.255.71.15 1 19 21 0 0 6:17 Estab1
bgp.13vpn.0: 2/2/0
VPN-A.inet.0: 1/1/0
VPN-AB.inet.0: 2/2/0
VPN-B.inet.0: 1/1/0

```

show policy damping

Syntax	show policy damping <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switch and QFX Series)	show policy damping
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Display information about BGP route flap damping parameters.
Options	<p>none—Display information about BGP route flap damping parameters.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Additional Information	In the output from this command, figure-of-merit values correlate with the probability of future instability of a routing device. Routes with higher figure-of-merit values are suppressed for longer periods of time. The figure-of-merit value decays exponentially over time. A figure-of-merit value of zero is assigned to each new route. The value is increased each time the route is withdrawn or readvertised, or when one of its path attributes changes.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • “Configuring BGP Flap Damping Parameters” in the Routing Policy Configuration Guide • clear bgp damping on page 140 • show route damping
List of Sample Output	show policy damping on page 173
Output Fields	Table 40 on page 172 describes the output fields for the show policy damping command. Output fields are listed in the approximate order in which they appear.

Table 40: show policy damping Output Fields

Field Name	Field Description
Halflife	Decay half-life, in minutes. The value represents the period during which the accumulated figure-of-merit value is reduced by half if the route remains stable. If a route has flapped, but then becomes stable, the figure-of-merit value for the route decays exponentially. For example, for a route with a figure-of-merit value of 1500, if no incidents occur, its figure-of-merit value is reduced to 750 after 15 minutes and to 375 after another 15 minutes.

Table 40: show policy damping Output Fields (*continued*)

Field Name	Field Description
Reuse merit	Figure-of-merit value below which a suppressed route can be used again. A suppressed route becomes reusable when its figure-of-merit value decays to a value below a reuse threshold, and the route once again is considered usable and can be installed in the forwarding table and exported from the routing table.
Suppress/cutoff merit	Figure-of-merit value above which a route is suppressed for use or inclusion in advertisements. When a route's figure-of-merit value reaches a particular level, called the cutoff or suppression threshold, the route is suppressed. When a route is suppressed, the routing table no longer installs the route into the forwarding table and no longer exports this route to any of the routing protocols.
Maximum suppress time	Maximum hold-down time, in minutes. The value represents the maximum time that a route can be suppressed no matter how unstable it has been before this period of stability.
Computed values	<ul style="list-style-type: none"> • Merit ceiling—Maximum merit that a flapping route can collect. • Maximum decay—Maximum decay half-life, in minutes.

Sample Output

```

show policy damping      user@host> show policy damping
                          Default damping information:
                          Halflife: 15 minutes
                          Reuse merit: 750 Suppress/cutoff merit: 3000
                          Maximum suppress time: 60 minutes
                          Computed values:
                            Merit ceiling: 12110
                            Maximum decay: 6193
                          Damping information for "standard-damping":
                          Halflife: 10 minutes
                          Reuse merit: 4000 Suppress/cutoff merit: 8000
                          Maximum suppress time: 30 minutes
                          Computed values:
                            Merit ceiling: 32120
                            Maximum decay: 12453

```

